

This is a “core dump” of potential questions for Midterm 2. This should give you a good idea of the *types* of questions that we will ask on the exam, but the actual exam questions may or may not appear in this handout. This list intentionally includes a few questions that are too long or difficult for exam conditions. In particular, there are several multipart questions where we would ask only one or two parts on an exam.

Questions from Jeff’s past exams are labeled with the semester then they were used. Questions from this semester’s labs are labeled **⟨⟨Lab⟩⟩**.

Solving every problem in this handout is *not* the best way to study for the exam. Memorizing the solutions to every problem in this handout is the *absolute worst* way to study for the exam.

What we recommend instead is to work on a *sample* of the problems. Choose one or two problems at random from each section and try to solve them from scratch under exam conditions—by yourself, in a quiet room, with a 30-minute timer, *without* your notes, *without* the internet, and if possible, even without your cheat sheet. If you’re comfortable solving a few problems in a particular section, you’re probably ready for that type of problem on the exam. Move on to the next section.

If you find yourself getting stuck on a problem, try to figure out *why* you’re stuck. Do you understand the problem statement? Are you stuck on choosing the right high-level approach, or are you stuck on the details? For recursion/dynamic programming: Are you solving the right recursive generalization of the stated problem? For greedy algorithms: Are you sure a greedy algorithm is a good idea? [*Hint: Probably not.*] For graph algorithms: Are you aiming for the right problem? Are you using the right graph?

Discussing problems with other people (in your study groups, in the review sessions, in office hours, or on Piazza) and/or looking up old solutions can be *extremely* helpful, but **only after** you have (1) made a good-faith effort to solve the problem on your own, and (2) you have either a candidate solution or some idea about where you’re getting stuck.

When you do discuss problems with other people, remember that your goal is *not* merely to “understand” the solution to any particular problem, but to become more comfortable with solving a certain *type* of problem on your own. If you can identify specific steps that you find problematic, read more *about those steps*, focus your practice *on those steps*, and try to find helpful information *about those steps* to write on your cheat sheet. Then work on the next problem!

Recursion and Dynamic Programming

Elementary Recursion/Divide and Conquer

1. **⟨⟨Lab⟩⟩**

- (a) Suppose $A[1..n]$ is an array of n distinct integers, sorted so that $A[1] < A[2] < \dots < A[n]$. Each integer $A[i]$ could be positive, negative, or zero. Describe a fast algorithm that either computes an index i such that $A[i] = i$ or correctly reports that no such index exists..
- (b) Now suppose $A[1..n]$ is a sorted array of n distinct **positive** integers. Describe an even faster algorithm that either computes an index i such that $A[i] = i$ or correctly reports that no such index exists. [*Hint: This is really easy.*]

2. **⟨⟨Lab⟩⟩** Suppose we are given an array $A[1..n]$ such that $A[1] \geq A[2]$ and $A[n-1] \leq A[n]$. We say that an element $A[x]$ is a **local minimum** if both $A[x-1] \geq A[x]$ and $A[x] \leq A[x+1]$. For example, there are exactly six local minima in the following array:

9	7	7	2	1	3	7	5	4	7	3	3	4	8	6	9
	▲			▲				▲		▲	▲			▲	

Describe and analyze a fast algorithm that returns the index of one local minimum. For example, given the array above, your algorithm could return the integer 5, because $A[5]$ is a local minimum. [*Hint: With the given boundary conditions, any array must contain at least one local minimum. Why?*]

3. **⟨⟨Lab⟩⟩**

- (a) Suppose you are given two sorted arrays $A[1..n]$ and $B[1..n]$ containing distinct integers. Describe a fast algorithm to find the median (meaning the n th smallest element) of the union $A \cup B$. For example, given the input

$$A[1..8] = [0, 1, 6, 9, 12, 13, 18, 20] \quad B[1..8] = [2, 4, 5, 8, 17, 19, 21, 23]$$

your algorithm should return the integer 9. [*Hint: What can you learn by comparing one element of A with one element of B ?*]

- (b) Now suppose you are given two sorted arrays $A[1..m]$ and $B[1..n]$ and an integer k . Describe a fast algorithm to find the k th smallest element in the union $A \cup B$. For example, given the input

$$A[1..8] = [0, 1, 6, 9, 12, 13, 18, 20] \quad B[1..5] = [2, 5, 7, 17, 19] \quad k = 6$$

your algorithm should return the integer 7.

4. **⟨⟨HW⟩⟩** You are a visitor at a political convention (or perhaps a faculty meeting) with n delegates; each delegate is a member of exactly one political party. It is impossible to tell which political party any delegate belongs to; in particular, you will be summarily ejected from the convention if you ask. However, you can determine whether any pair of delegates

belong to the *same* party or not simply by introducing them to each other—members of the same party always greet each other with smiles and friendly handshakes; members of different parties always greet each other with angry stares and insults.

- (a) Suppose more than half of the delegates belong to the same political party. Describe an efficient algorithm that identifies all members of this majority party.
- (b) Now suppose exactly k political parties are represented at the convention and one party has a *plurality*: more delegates belong to that party than to any other. Present a practical procedure to pick out the people from the plurality political party as parsimoniously as possible. (Please.)
5. **⟨⟨F14, S14⟩⟩** An array $A[0..n-1]$ of n distinct numbers is **bitonic** if there are unique indices i and j such that $A[(i-1) \bmod n] < A[i] > A[(i+1) \bmod n]$ and $A[(j-1) \bmod n] > A[j] < A[(j+1) \bmod n]$. In other words, a bitonic sequence either consists of an increasing sequence followed by a decreasing sequence, or can be circularly shifted to become so. For example,

4	6	9	8	7	5	1	2	3
---	---	---	---	---	---	---	---	---

 is bitonic, but

3	6	9	8	7	5	1	2	4
---	---	---	---	---	---	---	---	---

 is *not* bitonic.

Describe and analyze an algorithm to find the index of the *smallest* element in a given bitonic array $A[0..n-1]$ in $O(\log n)$ time. You may assume that the numbers in the input array are distinct. For example, given the first array above, your algorithm should return 6, because $A[6] = 1$ is the smallest element in that array.

Dynamic Programming

1. **⟨⟨Lab⟩⟩** Describe and analyze efficient for the following problems.
 - (a) Given an array $A[1..n]$ of integers, compute the length of a longest **increasing** subsequence of A . A sequence $B[1..l]$ is *increasing* if $B[i] > B[i - 1]$ for every index $i \geq 2$.
 - (b) Given an array $A[1..n]$ of integers, compute the length of a longest **decreasing** subsequence of A . A sequence $B[1..l]$ is *decreasing* if $B[i] < B[i - 1]$ for every index $i \geq 2$.
 - (c) Given an array $A[1..n]$ of integers, compute the length of a longest **alternating** subsequence of A . A sequence $B[1..l]$ is *alternating* if $B[i] < B[i - 1]$ for every even index $i \geq 2$, and $B[i] > B[i - 1]$ for every odd index $i \geq 3$.
 - (d) Given an array $A[1..n]$ of integers, compute the length of a longest **convex** subsequence of A . A sequence $B[1..l]$ is *convex* if $B[i] - B[i - 1] > B[i - 1] - B[i - 2]$ for every index $i \geq 3$.
 - (e) Given an array $A[1..n]$, compute the length of a longest **palindrome** subsequence of A . Recall that a sequence $B[1..l]$ is a *palindrome* if $B[i] = B[l - i + 1]$ for every index i .

2. **⟨⟨S16⟩⟩** After the Revolutionary War, Alexander Hamilton's biggest rival as a lawyer was Aaron Burr. (Sir!) In fact, the two worked next door to each other. Unlike Hamilton, Burr cannot work non-stop; every case he tries exhausts him. The bigger the case, the longer he must rest before he is well enough to take the next case. (Of course, he is willing to wait for it.) If a case arrives while Burr is resting, Hamilton snatches it up instead.

Burr has been asked to consider a sequence of n upcoming cases. He quickly computes two arrays $profit[1..n]$ and $skip[1..n]$, where for each index i ,

- $profit[i]$ is the amount of money Burr would make by taking the i th case, and
- $skip[i]$ is the number of consecutive cases Burr must skip if he accepts the i th case. That is, if Burr accepts the i th case, he cannot accept cases $i + 1$ through $i + skip[i]$.

Design and analyze an algorithm that determines the maximum total profit Burr can secure from these n cases, using his two arrays as input.

3. **⟨⟨S14⟩⟩** Recall that a *palindrome* is any string that is the same as its reversal. For example, **I**, **DAD**, **HANNAH**, **AIBOHPHOBIA** (fear of palindromes), and the empty string are all palindromes.
 - (a) Describe and analyze an algorithm to find the length of the longest substring (not *subsequence*!) of a given input string that is a palindrome. For example, **BASEESAB** is the longest palindrome substring of **BUBBASEESABANANA** ("Bubba sees a banana."). Thus, given the input string **BUBBASEESABANANA**, your algorithm should return the integer 8.

- (b) Any string can be decomposed into a sequence of palindrome substrings. For example, the string **BUBBASEESABANANA** can be broken into palindromes in the following ways (and many others):

$$\begin{aligned} & \text{BUB} + \text{BASEESAB} + \text{ANANA} \\ & \text{B} + \text{U} + \text{BB} + \text{A} + \text{SEES} + \text{ABA} + \text{NAN} + \text{A} \\ & \text{B} + \text{U} + \text{BB} + \text{A} + \text{SEES} + \text{A} + \text{B} + \text{ANANA} \\ & \text{B} + \text{U} + \text{B} + \text{B} + \text{A} + \text{S} + \text{E} + \text{E} + \text{S} + \text{A} + \text{B} + \text{A} + \text{N} + \text{A} + \text{N} + \text{A} \end{aligned}$$

Describe and analyze an algorithm to find the smallest number of palindromes that make up a given input string. For example, if your input is the string **BUBBASEESA-BANANA**, your algorithm should return the integer 3.

4. A *shuffle* of two strings X and Y is formed by interspersing the characters into a new string, keeping the characters of X and Y in the same order. For example, the string **BANANAANANAS** is a shuffle of the strings **BANANA** and **ANANAS** in several different ways.

$$\text{BANANA}_{\text{ANANAS}} \quad \text{BAN}_{\text{ANA}}\text{ANA}_{\text{NAS}} \quad \text{B}_{\text{AN}}\text{AN}_{\text{A}}\text{A}_{\text{NA}}\text{NA}_{\text{S}}$$

Similarly, the strings **PRODGYRNAMAMMIINCG** and **DYPRONGARMAMMICING** are both shuffles of **DYNAMIC** and **PROGRAMMING**:

$$\text{PRO}_{\text{D}}\text{G}_{\text{Y}}\text{R}_{\text{NAM}}\text{AMMI}_{\text{I}}\text{N}_{\text{C}}\text{G} \quad \text{DY}_{\text{PRO}}\text{N}_{\text{G}}\text{A}_{\text{R}}\text{M}_{\text{AMMI}}\text{I}_{\text{C}}\text{ING}$$

Describe and analyze an efficient algorithm to determine, given three strings $A[1..m]$, $B[1..n]$, and $C[1..m+n]$, whether C is a shuffle of A and B .

5. Suppose you are given a sequence of non-negative integers separated by $+$ and \times signs; for example:

$$2 \times 3 + 0 \times 6 \times 1 + 4 \times 2$$

You can change the value of this expression by adding parentheses in different places. For example:

$$\begin{aligned} 2 \times (3 + (0 \times (6 \times (1 + (4 \times 2)))))) &= 6 \\ (((((2 \times 3) + 0) \times 6) \times 1) + 4) \times 2 &= 80 \\ ((2 \times 3) + (0 \times 6)) \times (1 + (4 \times 2)) &= 108 \\ (((2 \times 3) + 0) \times 6) \times ((1 + 4) \times 2) &= 360 \end{aligned}$$


Describe and analyze an algorithm to compute, given a list of integers separated by $+$ and \times signs, the smallest possible value we can obtain by inserting parentheses.

Your input is an array $A[0..2n]$ where each $A[i]$ is an integer if i is even and $+$ or \times if i is odd. Assume any arithmetic operation in your algorithm takes $O(1)$ time.

6. Suppose you are given an array $A[1..n]$ of numbers, which may be positive, negative, or zero, and which are *not* necessarily integers.

- (a) Describe and analyze an algorithm that finds the largest sum of elements in a contiguous subarray $A[i..j]$.
- (b) Describe and analyze an algorithm that finds the largest *product* of elements in a contiguous subarray $A[i..j]$.

For example, given the array $[-6, 12, -7, 0, 14, -7, 5]$ as input, your first algorithm should return the integer 19, and your second algorithm should return the integer 504.



For the sake of analysis, assume that comparing, adding, or multiplying any pair of numbers takes $O(1)$ time.

[Hint: Problem (a) has been a standard computer science interview question since at least the mid-1980s. You can find many correct solutions on the web; the problem even has its own [Wikipedia page](#)! But at least in 2016, a significant fraction of the solutions I found on the web for problem (b) were either significantly slower than necessary or actually incorrect. Remember that the product of two negative numbers is positive.]

7. Suppose you are given three strings $A[1..n]$, $B[1..n]$, and $C[1..n]$.
- (a) Describe and analyze an algorithm to find the length of the longest common subsequence of all three strings. For example, given the input strings

$$A = \text{AxxBxxCDxEF}, \quad B = \text{yyABCDyEyFy}, \quad C = \text{zAzzBCDzEFz},$$

your algorithm should output the number 6, which is the length of the longest common subsequence **ABCDEF**.

- (b) Describe and analyze an algorithm to find the length of the shortest common supersequence of all three strings. For example, given the input strings

$$A = \text{AxxBxxCDxEF}, \quad B = \text{yyABCDyEyFy}, \quad C = \text{zAzzBCDzEFz},$$

your algorithm should output the number 21, which is the length of the shortest common supersequence **zyxAzzxBxxCDxyzEyFzy**.

8. (a) Suppose we are given a set L of n line segments in the plane, where each segment has one endpoint on the line $y = 0$ and one endpoint on the line $y = 1$, and all $2n$ endpoints are distinct. Describe and analyze an algorithm to compute the largest subset of L in which no pair of segments intersects.
- (b) Suppose we are given a set L of n line segments in the plane, where each segment has one endpoint on the line $y = 0$ and one endpoint on the line $y = 1$, and all $2n$ endpoints are distinct. Describe and analyze an algorithm to compute the largest subset of L in which *every* pair of segments intersects.

9. Suppose you are given an $m \times n$ bitmap, represented by an array $M[1..n, 1..n]$ of 0s and 1s. A *solid block* in M is a subarray of the form $M[i..i', j..j']$ containing only 1-bits. A solid block is square if it has the same number of rows and columns.
- (a) Describe an algorithm to find the maximum area of a solid *square* block in M in $O(n^2)$ time.
- (b) Describe an algorithm to find the maximum area of a solid block in M in $O(n^3)$ time. [Hint: Try for $O(n^4)$ first.]

10. **⟨⟨F14⟩⟩** The new swap-puzzle game *Candy Swap Saga XIII* involves n cute animals numbered 1 through n . Each animal holds one of three types of candy: circus peanuts, Heath bars, and Cioccolateria Gardini chocolate truffles. You also have a candy in your hand; at the start of the game, you have a circus peanut.

To earn points, you visit each of the animals in order from 1 to n . For each animal, you can either keep the candy in your hand or exchange it with the candy the animal is holding.

- If you swap your candy for another candy of the *same* type, you earn one point.
- If you swap your candy for a candy of a *different* type, you lose one point. (Yes, your score can be negative.)
- If you visit an animal and decide not to swap candy, your score does not change.

You *must* visit the animals in order, and once you visit an animal, you can never visit it again.

Describe and analyze an efficient algorithm to compute your maximum possible score. Your input is an array $C[1..n]$, where $C[i]$ is the type of candy that the i th animal is holding.

11. **⟨⟨F14⟩⟩** Farmers Boggis, Bunce, and Bean have set up an obstacle course for Mr. Fox. The course consists of a row of n booths, each with an integer painted on the front with bright red paint, which could be positive, negative, or zero. Let $A[i]$ denote the number painted on the front of the i th booth. Everyone has agreed to the following rules:
- At each booth, Mr. Fox *must* say either “Ring!” or “Ding!”.
 - If Mr. Fox says “Ring!” at the i th booth, he earns a reward of $A[i]$ chickens. (If $A[i] < 0$, Mr. Fox pays a penalty of $-A[i]$ chickens.)
 - If Mr. Fox says “Ding!” at the i th booth, he pays a penalty of $A[i]$ chickens. (If $A[i] < 0$, Mr. Fox earns a reward of $-A[i]$ chickens.)
 - Mr. Fox is forbidden to say the same word more than three times in a row. For example, if he says “Ring!” at booths 6, 7, and 8, then he *must* say “Ding!” at booth 9.
 - All accounts will be settled at the end; Mr. Fox does not actually have to carry chickens through the obstacle course.
 - If Mr. Fox violates any of the rules, or if he ends the obstacle course owing the farmers chickens, the farmers will shoot him.

Describe and analyze an algorithm to compute the largest number of chickens that Mr. Fox can earn by running the obstacle course, given the array $A[1..n]$ of booth numbers as input.

Greedy Algorithms

Remember that you will receive **zero** points for a greedy algorithm, even if it is perfectly correct, unless you also give a formal proof of correctness.

1. **⟨⟨Lab⟩⟩** Recall the class scheduling problem described in lecture on Tuesday. We are given two arrays $S[1..n]$ and $F[1..n]$, where $S[i] < F[i]$ for each i , representing the start and finish times of n classes. Your goal is to find the largest number of classes you can take without ever taking two classes simultaneously. We showed in class that the following greedy algorithm constructs an optimal schedule:

Choose the course that *ends first*, discard all conflicting classes, and recurse.

But this is not the only greedy strategy we could have tried. For each of the following alternative greedy algorithms, either prove that the algorithm always constructs an optimal schedule, or describe a small input example for which the algorithm does not produce an optimal schedule. Assume that all algorithms break ties arbitrarily (that is, in a manner that is completely out of your control).

[Hint: Exactly three of these greedy strategies actually work.]

- (a) Choose the course x that *ends last*, discard classes that conflict with x , and recurse.
 - (b) Choose the course x that *starts first*, discard all classes that conflict with x , and recurse.
 - (c) Choose the course x that *starts last*, discard all classes that conflict with x , and recurse.
 - (d) Choose the course x with *shortest duration*, discard all classes that conflict with x , and recurse.
 - (e) Choose a course x that *conflicts with the fewest other courses*, discard all classes that conflict with x , and recurse.
 - (f) If no classes conflict, choose them all. Otherwise, discard the course with *longest duration* and recurse.
 - (g) If no classes conflict, choose them all. Otherwise, discard a course that *conflicts with the most other courses* and recurse.
 - (h) Let x be the class with the *earliest start time*, and let y be the class with the *second earliest start time*.
 - If x and y are disjoint, choose x and recurse on everything but x .
 - If x completely contains y , discard x and recurse.
 - Otherwise, discard y and recurse.
 - (i) If any course x completely contains another course, discard x and recurse. Otherwise, choose the course y that *ends last*, discard all classes that conflict with y , and recurse.
2. **⟨⟨S14⟩⟩** Binaria uses coins whose values are $1, 2, 4, \dots, 2^k$, the first k powers of two, for some integer k . As in most countries, Binarian shopkeepers always make change using the following greedy algorithm:

```
MAKECHANGE(N):
  if  $N = 0$ 
    say "Thank you, come again!"
  else
     $c \leftarrow$  largest coin value such that  $c \leq N$ 
    give the customer one  $c$  cent coin
    MAKECHANGE( $N - c$ )
```

For example, to make 37 cents in change, the shopkeeper would give the customer one 32 cent coin, one 4 cent coin, and one 1 cent coin, and then say "Thank you, come again!" (For purposes of this problem, assume that every shopkeeper has an unlimited supply of each type of coin.)



Prove that this greedy algorithm always uses the smallest possible number of coins. [Hint: Prove that the greedy algorithm uses at most one coin of each denomination.]

- Let X be a set of n intervals on the real line. We say that a set P of points *stabs* X if every interval in X contains at least one point in P . Describe and analyze an efficient algorithm to compute the smallest set of points that stabs X . Assume that your input consists of two arrays $L[1..n]$ and $R[1..n]$, representing the left and right endpoints of the intervals in X . If you use a greedy algorithm, don't forget to **prove** that it is correct.

Graph Algorithms

Sanity Check

1. **⟨⟨S14, F14⟩⟩** Indicate the following structures in the example graphs below. If the requested structure does not exist, just write NONE. To indicate a subgraph, draw over the entire edge with a heavy black line; your answer should be visible from across the room. *[On an exam, we would only ask about one graph, we would not ask for all these structures, and we would give you several copies of the graph on which to mark your answers.]*




- (a) A depth-first spanning tree rooted at node s .
- (b) A breadth-first spanning tree rooted at node s .
- (c) A shortest-path tree rooted at node s .
- (d) The set of all vertices reachable from node c . (Circle each vertex.)
- (e) The set of all vertices from which node c is reachable. (Circle each vertex.)
- (f) The strongly connected components. (Circle each strongly connected component.)
- (g) The shortest directed cycle.
- (h) A topological order. (List the vertices in order.)
- (i) A walk from s to d with the maximum number of edges.
- (j) A walk from s to d with the largest total weight.

Reachability/Connectivity/Traversal

1. **⟨⟨F14⟩⟩** Suppose you are given a directed graph $G = (V, E)$ and two vertices s and t . Describe and analyze an algorithm to determine if there is a walk in G from s to t (possibly repeating vertices and/or edges) whose length is divisible by 3.

For example, given the graph below, with the indicated vertices s and t , your algorithm should return TRUE, because the walk $s \rightarrow w \rightarrow y \rightarrow x \rightarrow s \rightarrow w \rightarrow t$ has length 6.



[Hint: Build a (different) graph.]

2. **⟨⟨Lab⟩⟩** *Snakes and Ladders* is a classic board game, originating in India no later than the 16th century. The board consists of an $n \times n$ grid of squares, numbered consecutively from 1 to n^2 , starting in the bottom left corner and proceeding row by row from bottom to top, with rows alternating to the left and right. Certain pairs of squares, always in different rows, are connected by either “snakes” (leading down) or “ladders” (leading up). Each square can be an endpoint of at most one snake or ladder.

100	99	98	97	96	95	94	93	92	91
81	82	83	84	85	86	87	88	89	90
80	79	78	77	76	75	74	73	72	71
61	62	63	64	65	66	67	68	69	70
60	59	58	57	56	55	54	53	52	51
41	42	43	44	45	46	47	48	49	50
40	39	38	37	36	35	34	33	32	31
21	22	23	24	25	26	27	28	29	30
20	19	18	17	16	15	14	13	12	11
1	2	3	4	5	6	7	8	9	10

A typical Snakes and Ladders board.

Upward straight arrows are ladders; downward wavy arrows are snakes.


You start with a token in cell 1, in the bottom left corner. In each move, you advance your token up to k positions, for some fixed constant k (typically 6). If the token ends the move at the *top* end of a snake, you **must** slide the token down to the bottom of that snake. If the token ends the move at the *bottom* end of a ladder, you **may** move the token up to the top of that ladder.

Describe and analyze an algorithm to compute the smallest number of moves required for the token to reach the last square of the grid.

3. Let G be a connected undirected graph. Suppose we start with two coins on two arbitrarily chosen vertices of G , and we want to move the coins so that they lie on the same vertex using as few moves as possible. At every step, each coin *must* move to an adjacent vertex.

- (a) **⟨⟨Lab⟩⟩** Describe and analyze an algorithm to compute the minimum number of steps to reach a configuration where both coins are on the same vertex, or to report correctly that no such configuration is reachable. The input to your algorithm consists of a graph $G = (V, E)$ and two vertices $u, v \in V$ (which may or may not be distinct).
 - (b) Now suppose there are three coins. Describe and analyze an algorithm to compute the minimum number of steps to reach a configuration where both coins are on the same vertex, or to report correctly that no such configuration is reachable. [Hint: Some people already considered this problem in lab.]
 - (c) Finally, suppose there are *forty-two* coins. Describe and analyze an algorithm to determine whether it is possible to move all 42 coins to the same vertex. Again, *every* coin must move at *every* step. For full credit, your algorithm should run in $O(V + E)$ time.
4. A graph (V, E) is bipartite if the vertices V can be partitioned into two subsets L and R , such that every edge has one vertex in L and the other in R .
- (a) Prove that every tree is a bipartite graph.
 - (b) Describe and analyze an efficient algorithm that determines whether a given undirected graph is bipartite.
5. **⟨⟨F14⟩⟩** A **number maze** is an $n \times n$ grid of positive integers. A token starts in the upper left corner; your goal is to move the token to the lower-right corner. On each turn, you are allowed to move the token up, down, left, or right; the distance you may move the token is determined by the number on its current square. For example, if the token is on a square labeled 3, then you may move the token three steps up, three steps down, three steps left, or three steps right. However, you are never allowed to move the token off the edge of the board.


Describe and analyze an efficient algorithm that either returns the minimum number of moves required to solve a given number maze, or correctly reports that the maze has no solution. For example, given the maze shown below, your algorithm would return the number 8.



A 5×5 number maze that can be solved in eight moves.

6. **Kaniel Dane** is a solitaire puzzle played with two tokens on an $n \times n$ square grid. Some squares of the grid are marked as *obstacles*, and one grid square is marked as the *target*. In each turn, the player can move either of the tokens from its current position as far as possible upward, downward, right, or left, stopping just before the token hits (1) the edge

of the board, (2) an obstacle square, or (3) the other token. The goal is to move either of the tokens onto the target square.



An instance of the Kaniel Dane puzzle that can be solved in six moves.
Circles indicate the initial token positions; black squares are obstacles; the center square is the target.

Describe and analyze an algorithm to determine whether an instance of this puzzle is solvable. Your input consist of the integer n , a list of obstacle locations, the target location, and the initial locations of the tokens. The output of your algorithm is a single boolean: `TRUE` if the given puzzle is solvable and `FALSE` otherwise. The running time of your algorithm should be a small polynomial in n .

Depth-First Search, Dags, Strong Connectivity

1. **⟨⟨Lab⟩⟩** Inspired by an earlier question, you decided to organize a Snakes and Ladders competition with n participants. In this competition, each game of Snakes and Ladders involves three players. After the game is finished, they are ranked first, second and third. Each player may be involved in any (non-negative) number of games, and the number needs not be equal among players.

At the end of the competition, m games have been played. You realized that you had forgotten to implement a proper rating system, and therefore decided to produce the overall ranking of all n players as you see fit. However, to avoid being too suspicious, if player A ranked better than player B in any game, then A must rank better than B in the overall ranking.

You are given the list of players involved and the ranking in each of the m games. Describe and analyze an algorithm to produce an overall ranking of the n players that satisfies the condition, or correctly reports that it is impossible.

2. Let G be a directed acyclic graph with a unique source s and a unique sink t .
 - (a) A *Hamiltonian path* in G is a directed path in G that contains every vertex in G . Describe an algorithm to determine whether G has a Hamiltonian path.
 - (b) Suppose the vertices of G have weights. Describe an efficient algorithm to find the path from s to t with maximum total weight.
 - (c) Suppose we are also given an integer ℓ . Describe an efficient algorithm to find the maximum-weight path from s to t , such that the path contains at most ℓ edges. (Assume there is at least one such path.)
 - (d) Suppose several vertices in G are marked *essential*, and we are given an integer k . Design an efficient algorithm to determine whether there is a path from s to t that passes through at least k essential vertices.
 - (e) Suppose the vertices of G have integer labels, where $label(s) = -\infty$ and $label(t) = \infty$. Describe an algorithm to find the path from s to t with the maximum number of edges, such that the vertex labels define an increasing sequence.
 - (f) **⟨⟨Lab⟩⟩** Describe an algorithm to compute the number of distinct paths from s to t in G . (Assume that you can add arbitrarily large integers in $O(1)$ time.)
3. Suppose you are given a directed graph G in which **every edge has negative weight**, and a source vertex s . Describe and analyze an efficient algorithm that computes the shortest path distances from s to every other vertex in G . Specifically, for every vertex t :
 - If t is not reachable from s , your algorithm should report $dist(t) = \infty$.
 - If the shortest-path distance from s to t is not well-defined because of negative cycles, your algorithm should report $dist(t) = -\infty$.
 - If neither of the two previous conditions applies, your algorithm should report the correct shortest-path distance from s to t .

[Hint: First think about graphs where the first two conditions never happen.]

Shortest Paths

1. **⟨⟨F14⟩⟩** Let G be a directed graph with weighted edges, and let s be a vertex of G . Suppose every vertex $v \neq s$ stores a pointer $pred(v)$ to another vertex in G . Describe and analyze an algorithm to determine whether these predecessor pointers correctly define a single-source shortest path tree rooted at s . Do **not** assume that G has no negative cycles.
2. **⟨⟨F14⟩⟩** Suppose we are given an undirected graph G in which every *vertex* has a positive weight.
 - (a) Describe and analyze an algorithm to find a *spanning tree* of G with minimum total weight. (The total weight of a spanning tree is the sum of the weights of its vertices.)
 - (b) Describe and analyze an algorithm to find a *path* in G from one given vertex s to another given vertex t with minimum total weight. (The total weight of a path is the sum of the weights of its vertices.)

3. **⟨⟨S14, Lab⟩⟩** You just discovered your best friend from elementary school on Twitbook. You both want to meet as soon as possible, but you live in two different cities that are far apart. To minimize travel time, you agree to meet at an intermediate city, and then you simultaneously hop in your cars and start driving toward each other. But where *exactly* should you meet?

You are given a weighted graph $G = (V, E)$, where the vertices V represent cities and the edges E represent roads that directly connect cities. Each edge e has a weight $w(e)$ equal to the time required to travel between the two cities. You are also given a vertex p , representing your starting location, and a vertex q , representing your friend's starting location.

Describe and analyze an algorithm to find the target vertex t that allows you and your friend to meet as quickly as possible.

4. There are n galaxies connected by m intergalactic teleport-ways. Each teleport-way joins two galaxies and can be traversed in both directions. Also, each teleport-way e has an associated cost of $c(e)$ dollars, where $c(e)$ is a positive integer. A teleport-way can be used multiple times, but the toll must be paid every time it is used.

Judy wants to travel from galaxy s to galaxy t , but teleportation is not very pleasant and she would like to minimize the number of times she needs to teleport. However, she wants the total cost to be a multiple of five dollars, because carrying small change is not pleasant either.

- (a) Describe and analyze an algorithm to compute the smallest number of times Judy needs to teleport to travel from galaxy s to galaxy t while the total cost is a multiple of five dollars.
- (b) Solve part (a), but now assume that Judy has a coupon that allows her to use one teleport-way for free.

[Hint: No, this is **not** the same Intergalactic Judy problem that you saw in lab.]

5. **⟨⟨Lab⟩⟩** A *looped tree* is a weighted, directed graph built from a binary tree by adding an edge from every leaf back to the root. Every edge has non-negative weight.




Figure 1. A looped tree.

- (a) How much time would Dijkstra's algorithm require to compute the shortest path between two vertices u and v in a looped tree with n nodes?
- (b) Describe and analyze a faster algorithm.
6. After graduating you accept a job with Aerophobes-Я-Us, the leading traveling agency for people who hate to fly. Your job is to build a system to help customers plan airplane trips from one city to another. All of your customers are afraid of flying (and by extension, airports), so any trip you plan needs to be as short as possible. You know all the departure and arrival times of all the flights on the planet.
- Suppose one of your customers wants to fly from city X to city Y . Describe an algorithm to find a sequence of flights that minimizes the *total time in transit*—the length of time from the initial departure to the final arrival, including time at intermediate airports waiting for connecting flights.
7. When there is more than one shortest path from one node s to another node t , it is often convenient to choose a shortest path with the fewest edges; call this the *best* path from s to t . Suppose we are given a directed graph G with positive edge weights and a source vertex s in G . Describe and analyze an algorithm to compute best paths in G from s to every other vertex.