

PROOF BY CONTRADICTION: \leftrightarrow **PROOF BY SWAPPING:**

- ASSUME G.C. not optimal
- let opt. soln differ from G.C. & 1st choice
- show we can construct new opt. soln. matching G.C.
- contradiction!

GREEDY ALG: best option doesn't block optimality

- make 1st decision locally optimally, & at least 1 global optimal soln will be consistent
- order input in some way then apply G.C. repeat

STEPS:

greedy choice	greedy remaining	choice problem	repeated up
---------------	------------------	----------------	-------------

ind. step

1) **GREEDY CHOICE:** Schedule earliest finishing time 1st

2) **GCP:** at least 1 optimal soln containing greedy choice

3) **PROOF OF GCP:** compare some optimal soln that doesn't make G.C. prove you can make it at least as optimal by swapping to make G.C.

methods of prove

- proof by contradiction \rightarrow preserves feasibility (satisfies constraints)
- proof by swapping \rightarrow since... it covers all...

4) **PROOF OF ALG BY IND:** argue that remaining problem is of some form (self-reduction) and that solving smaller problems optimally & adding 1st choice back in gives optimal soln. \rightarrow recursively repeating G.C. self reduction

5) **GIVE ALG & RUNTIME (IMPLEMENT G.C.)**

+3

i:3 o:2 u:1
i: 1
o: 00
u: 01

HUFFMANS ALG: encode letters as binary strings

GREEDY CHOICE: pick smallest frequencies as siblings

- prefix-free (no code word is prefix of others)
- minimize total weighted path length

COMMON GREEDY PATTERNS:

- interval scheduling max # non-overlapping \rightarrow G.C: earliest finish time
- Dijkstras pick closest unvisited node
- coin change use largest coin \leq remaining amt.
- scheduling w/ deadlines highest profit & place job in latest possible slot

COMPLEXITY CLASS: collection of decision problems w/ some property

① **P:** problems solvable in $O(nk)$ for constant k where n in input space

② **NP:** problems verifiable in $O(n^k)$ for some constant k .

- VERIFIER $V(x, c)$: poly-time alg that takes instance x & certificate c then returns Y/N

NP-HARD: B is NP-hard if for all problems A \in NP, $A \leq B$

NP-COMPLETE: NP-hard & in NP

- If $A \leq B$ & B is in P/NP/EXPTIME, then so is A
- If $A \leq B$ & A is in NP-hard, then so is B.

REDUCTIONS:

- convert one problem into another
- A reduce to B \rightarrow B at least as hard as A
- $A \leq B$ if can solve B in poly, can solve A
- used to compare difficulty of problems
- are transitive!

Ex: Partition \leq Subset Sum

POLYNOMIAL TIME

- running time = # of bits of input size
- Ex: array of n #, each up to 1 mil. poly depends on n

PSEUDOPOLY TIME

- polynomial in numeric value of inputs
- time $\leq O(n^k)$ deg. poly in input size \leq int inputs
- Ex: array of n #, each up to 1 mil. pseudo poly depends on time prop. to 1 mil. which can be exponential.

HALTING PROBLEM:

in: (P, x) alg, input
out: yes if P halts on x
no otherwise

\rightarrow NP-hard & undecidable (not in NP)

TOTALITY:

in: P
out: yes if P halts on all inputs no o.w.

thm: totality is undecidable

PF: Halt \leq Totality
assume alg T decides totality
construct H, a decider for Halt

$H(P, x):$
1) define $A(y) :=$ ignore input. run P on x . output yes (as long as no ∞ loop)
2) return $T(Q)$ (whether P halts)

(a) Prove that ACCEPT is recognizable. That is, prove that there exists an algorithm B such that $B(\mathcal{A}, x) = \text{YES}$ iff $\text{ACCEPT}(\mathcal{A}, x) = \text{YES}$.

RECOGNIZABLE
Solution: $B(\mathcal{A}, x)$ runs $\mathcal{A}(x)$ and outputs the result (if any). Then $B(\mathcal{A}, x) = \text{YES}$ iff $\mathcal{A}(x) = \text{YES}$ iff $\text{ACCEPT}(\mathcal{A}, x) = \text{YES}$.

(b) Prove that ACCEPT is undecidable by reducing from HALT.

Solution:

UNDECIDABLE

Input is (\mathcal{A}, x)

Create a new algorithm \mathcal{A}' , which on input y , runs $\mathcal{A}(y)$ and then outputs YES.

Output is (\mathcal{A}', x)

Now $\mathcal{A}'(x) = \text{YES}$ iff $\mathcal{A}(x)$ halts, i.e. $\text{ACCEPT}(\mathcal{A}', x) = \text{HALT}(\mathcal{A}, x)$. HALT is undecidable, and $\text{HALT} \leq \text{ACCEPT}$, so ACCEPT is undecidable also.

CORRECTNESS:

- by GCP, we assume optimal soln contains the G.C.
- after taking G.C., remaining problem is smaller but of the same type.
- by strong induction, let's prove this greedy alg is optimal.

define what you're inducting on

base case

1H (assume greedy alg produces opt. soln for ... fewer than n)
Inductive step

V1: Recursion

$\max\{\text{coins}[\dots]\}$
 $\{A[\dots] + \text{coins}(S[\dots])\}$
many overlapping subp.

V3: Pure DP

- start from smaller problems
- build in reverse topological order

V2: Memoization (top down)

• store answers in DAA

• recursive calls don't repeat work

DYNAMIC PROGRAMMING: overlapping subproblems

SUBPROBLEM: memo definition (domain, dimensions) $M[i] = \dots$
 $M[i] = \dots$ # dims depends on states (weight, len, ind)

RELATIONSHIP: DP transition (recursion equation)
 $M[i] = \max\{\dots\}$

TOPOLOGICAL SORT: argue graph formed in P is acyclic, & table can be filled iteratively in reverse topo. order

BASE CASES: initialize

OUTPUT: which entry of M to output
 $M[\dots]$

TIME: runtime. often product of (# entries in M) \times (work per entry)
 $O(\dots)$

DP SP:

DAG SSPP:
 $S: x(v) = \text{dist. to } v$ (edges)
 $R: x(v) = \min\{x(u) + w(u, v)\}$
 $T: \text{rev. topo. order of } G$
 $B: x(S, 0) = 0$
 $O: \text{whole table}$
 $T: O(|V| + |E|)$

COMMON STATES:
• post resource
• L+R endpt
• vertex + # edges

COMMON DP PATTERNS:

- knapsack $d[i][j] = \text{best among 1st } i \text{ items} \rightarrow \text{NP-complete (open q)}$
- weighted intervals DP on jobs sorted by finish times
- edit distance $i[i][j] = \text{cost to convert 1st } i \text{ chars to 1st } j \text{ chars}$
- LCS grid $i[i][j] = \text{LCS of prefixes } i \leq i \text{ & } j$

• LIS $i[i] = \text{longest incr. ending } @ i$

• grid path from top-left or bottom right. DP cell = best from neighbors

• DAG SP/LP-topo/rev. topo-order, recurrence on neighbors

• trees combine child subtrees

• coin change min #coins/ count ways to make amount

• sandwich cutting $T(l) = \text{max val cutting sandwich of len } l$

• subset sum decision problem. $T(l, i) = T$ if $\text{Subset } A[i...n] \text{ sum to } l$. not poly-time - input size isn't l , it's 2^n (bits).

if l is huge (2^n), $O(nl)$ is actually exponential (NP complete), pseudopoly

• S-t reachability in NP & not known to be NP-complete

NP PROOF:

- certificate
- verifier: returns Y/N based on input & certificate
- argue poly RT.

NP-hard PROOF:

- we reduce from \dots , which is NP-complete
- common choices: 3-SAT, 3-COLOR, CLIQUE
- construct reduction
- correctness argument \rightarrow show YES maps to YES, poly time RT argument NO maps to NO (both directions)

4. Consider two decision problems A and B. The problem $A \cup B$ asks whether its input is a YES instance of A or a YES instance of B . Similarly, the problem $A \cap B$ asks whether its input is a YES instance of A and a YES instance of B .

Circle all necessarily true statements: **NP & P**

- (a) If $A, B \in P$, then $A \cup B \in P$. YES if poly time solvable, can just solve both
- (b) If $A, B \in P$, then $A \cap B \in P$. YES solve both
- (c) If $A, B \in NP$, then $A \cup B \in NP$. YES certificate tells us which problem to use & has certificate for that problem.
- (d) If $A, B \in NP$, then $A \cap B \in NP$. YES

Recall that in PARTITION, we are given a list of numbers A and asked whether it can be partitioned into two lists with the same sum. This problem is NP-complete.

We can define a decision problem 0-1 KNAPSACK as follows: we are given a capacity S , a target value V , and a list of items, which each have a size s_i and a value v_i . We are asked whether there's a subset of items with total size at most S and total value at least V . (The "0-1" in the name comes from the fact that each item can be taken zero or one times, but not multiple times or a fractional value.)

Prove that 0-1 KNAPSACK is NP-hard by describing a reduction from PARTITION to 0-1 KNAPSACK. **NP-HARD: KNAPSACK** partition \leq knapsack

Solution: Given an instance A of PARTITION, we construct an instance of 0-1 KNAPSACK. For each element a_i of A , we create an item with $s_i = v_i = a_i$. Let $T = \sum a_i$ be the sum of all of the numbers. Our capacity and target value are $S = V = T/2$.

If we started with a YES instance of PARTITION, then there is a solution to the instance of 0-1 KNAPSACK: take the items corresponding to either set in a valid partition, which have total value and total weight $T/2$.

Conversely, if this is a YES instance of 0-1 KNAPSACK, then we can find a partition: one side is the elements of A corresponding to items taken in the solution to 0-1 KNAPSACK, and the other side is everything else. Both of these sets have sum $T/2$.

Incidentally, 0-1 KNAPSACK is in NP, so it is also NP-Complete.

NP-COMPLETE EXS:

- clique
- partition

• 0-1 knapsack
• 3-colorability
• 3SAT
• subset sum

(k) [4 points] If $P = NP$, then which of the following must be true?

Select all correct answers.

- There is an NP-complete problem that can be solved in polynomial time.
- Every NP-complete problem can be solved in polynomial time.
- There is an NP-hard problem that can be solved in polynomial time.
- Every NP-hard problem can be solved in polynomial time.

EQUIV:

in: P, Q

out: yes if P, Q output yes on same inputs

no O.W.

thm: Equiv is undecidable

PF: Halt \leq Equiv

assume alg. E decides Equiv

H(P, x):

- 1) define $Q(y)$: ignore input, run $P(\approx)$, output yes
- 2) define $R(y)$: ignore input, output yes" always
- 3) return $E(Q, R)$

Partition

• Input: Sequence of n positive integers $A = \{a_1, \dots, a_n\}$.

• Output: Is there a bipartition of A into two subsets with equal sum?

• To solve Partition, we can reduce it to Subset Sum. We efficiently convert instances of Partition into instances of Subset Sum, and then we can use any algorithm for Subset Sum to solve them.

Subset Sum

NP Complete

• Input: Sequence of n positive integers $A = \{a_1, \dots, a_n\}$, another integer L .

• Output: Is there a subset of A that sums exactly to L (i.e., $\exists A' \subseteq A$ s.t. $\sum_{a \in A'} a = L$)

• This is a decision problem. Answer is YES or NO, TRUE or FALSE

• Example: $A = \{2, 5, 7, 8, 9\}$, $L = 21$ is YES ($L = 5 + 7 + 9$ is a valid solution, or witness)

LONGEST INCREASING SUBSEQUENCE
LIS DP
 $[8, 2, 10, 6, 7, 9] \rightarrow [5, 6, 7, 9]$
 $S: T(i, j) = \text{len of LIS of } A[i:j] \text{ w/ } (0 \leq i \leq j) \text{ numbers larger than } A[i:j]$
 $R: T(i, j) = \max \{ T(i+1, j), T(i, j-1) \text{ or } A[i:j] > A[i:j-1] \}$
 $O: T(0, -1)$
Alternate:
 $S: T(i) = \text{len of LIS of } A[i:n] \text{ using } A[i:j]$
 $R: T(i) = \max \{ T(j), A[i] < A[j] \}$
SRBTBOT for the coin row problem
COIN ROW
Input: an array A denoting the values of the coins.

 S Let $M(i)$ be the maximum value of coins we can pick from $A[i:n]$, for $i \in \{0, \dots, n-1\}$.

 R $M(i) = \max\{M(i+1), A[i] + M(i+2)\}$, for $i = \{0, \dots, n-1\}$.

 T Each $M(i)$ depends on $M(i')$ with larger index $i' > i$. Therefore, we can compute $M(i)$ in decreasing order of i .

 B $M(i) = 0$ if $i \geq n$ (i.e. if $A[i:n]$ is empty).

 O Output is $M(0)$. Note, however, that we may want the actual list of coins, which we will discuss next.

 T There are n values $M(i)$. Each takes $O(1)$ to compute for a total of $O(n)$.

 S $T(k, u, v)$ = minimum weight of a path from u to v through vertices in $[k+1] \cup \{u, v\}$.

The notation $[k+1]$ means $0, 1, \dots, k$.

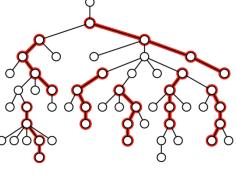
APSP: FLOYD WARSHALL
 R We have the choice of using node k or not. A shortest path that uses k as an intermediate node is formed of a shortest path from u to k using only $[k]$, i.e., $0, 1, \dots, k-1$, as possible intermediary nodes plus a shortest path from k to v also restricted to $[k]$ as possible intermediary nodes. Therefore,

$$T(k, u, v) = \min \begin{cases} T(k-1, u, k) + T(k-1, k, v) \\ T(k-1, u, v) \end{cases}$$

 T Decreasing k .

 B $T(-1) = w$ (i.e. $T(-1, u, v) = w(u, v)$ for all u and v). There are no intermediate vertices between u and v .

 O $T(|V|, \cdot, \cdot)$ contains all pairs of shortest paths without any restrictions.

 T The table T has size $(|V|+1) \times |V| \times |V| \in O(|V|^3)$. Each entry takes $O(1)$ work for a total of $O(|V|^3)$.

GREEDY
 (a) Describe, analyze, and prove correctness of an $O(n)$ -time greedy algorithm to compute the maximum number of disjoint paths that can fit in the tree. Your algorithm is given a rooted tree T and an integer k as input, and it should output the largest possible number of disjoint k -edge paths directed towards the root in T . Do not assume that T is a binary tree. For example, given the tree above as input, your algorithm should return 8. Note that you only need to return the number of paths and not the location of the paths.

 S Prove the greedy choice property you rely on, and argue the correctness and runtime of your algorithm.

 $a)$ $GC: \text{consider node } x \text{ w/ height } k$
 $\text{use a path starting from leaf } x \text{ and ending at } x$
 $GC_P: \text{there is at least 1 optimal soln containing } GC$
 $PF \text{ of } GC_P: \text{consider optimal soln } S \text{ that doesn't make } GC$
 $\text{if } S \text{ doesn't use } x \text{ at all, we can add new path}$
 $\text{if } S \text{ does use } x, \text{ can shift path down}$
 $Alg:$
 $\text{recursively compute height of each vertex}$
 $\text{when node } x \text{ has height } k$
 $1. \text{ add } 1 \text{ to counter}$
 $2. \text{ remove } x \text{ \& its subtree from tree}$
 $O(n) \text{ post order traversal}$
 $PF \text{ by induction}$
 $BC: \text{height} \leq k$
 $Ind \text{ Step: consider opt. soln}$
 $\text{by } GC, \text{ can assume } S \text{ makes same 1st step as } G$
 $\text{say we removed node } x \text{ to get } T$
 $\text{we can consider our alg. to have run on } T \text{ up to this point}$
 $\text{by IH: } 16|12|15|1 \rightarrow 16|15|1$
 $\text{greedy does as good as optimal!}$
 (b) [15 points] Describe modifications to your algorithm in Part (a) so that it runs in $O(n \log k)$ time. Analyze its runtime, but you need not prove correctness. You need not re-state your SRBTBOT other than steps that you modify (if applicable). Correct $O(nk)$ algorithm will be eligible for partial credit.

 $Solution:$ Store $X(j)$ in a max-heap cross-linked with a DAA. At each iteration of the R step, we access the max m of the heap and compute $X(i) = \lfloor m/2 \rfloor + p_i$. We then insert $X(i)$ into index i in the DAA, insert it in the heap, and remove $X(i-k)$ from the heap (if it exists).

OR **MAX HEAP/AVL**
 $Store X(j)$ in an AVL sequence augmented with the max in each subtree. At each iteration of the R step, we access the root augmentation m and compute $X(i) = \lfloor m/2 \rfloor + p_i$. We then append the new $X(i)$ and, if there are more than k elements, delete the first.

 $\text{For } Y, \text{ we instead use a min-heap or min augmentation, compute } Y(i) = 2m, \text{ and store } Y(i) - p_i.$
 $\text{Let } L \text{ be a decision language. Prove that } L \text{ is decidable iff both } L \text{ and its complement } \bar{L} \text{ are recognizable.}$
DECIDABLE
 $Solution:$ The forward direction is easy; a decider for \bar{L} is also a recognizer for L , and flipping its output produces a decider (and hence recognizer) for L .

 $\text{Conversely, suppose } L \text{ and its complement are both recognizable, say by algorithms } A \text{ and } B$
 $\text{respectively. We decide } L \text{ as follows:}$
 $\text{Input is } x$
 $\text{Run } A(x) \text{ and } B(x) \text{ in parallel}$
 $\text{If } A(x) \text{ says YES, output YES}$
 $\text{If } B(x) \text{ says YES, output NO}$

**can't loop forever b/c together they cover all inputs
will eventually say YES from one of them.**

 $L \text{ and } \bar{L} \text{ partition all inputs, so one of } A(x) \text{ or } B(x) \text{ outputs YES. The above algorithm is therefore a decider. It outputs YES iff } A \text{ does, so it decides } L.$
 \downarrow
 $\text{Each city has a } \text{disapproval score}, \text{ which is an integer. The disapproval scores may be zero, negative, or positive. The disapproval score for the } j^{\text{th}} \text{ city on row } i \text{ is } A[i][j]. \text{ For example, if } n = 4, \text{ the disapproval scores may be as shown above.}$
 $\text{Over all possible journeys, Sriini wants to know the } \text{maximum product} \text{ of disapproval scores of the cities he visits. In our example above for the particular values of } A, \text{ the maximum disapproval score corresponds to the path marked and has value } (-1)(-2)(-2)(-1) = 4.$
 $\text{Design an } O(n^2) \text{-time Dynamic Programming algorithm that returns the cities on the path with maximum product of disapproval scores. You may assume that all arithmetic operations take constant time.}$
DP: T_{\max} & T_{\min}
 $\text{As with the arithmetic parenthesization problem in R19, we don't know when we want to maximize the product of a subpath and when we want to minimize it. So we'll do both.}$
 $\text{Subproblems: For } 0 \leq j \leq i \leq n, \text{ let } T_{\max}(i, j) \text{ and } T_{\min}(i, j) \text{ be the maximum and minimum (respectively) possible products of disapproval ratings of southward journeys starting at } C[i][j].$
 $\text{Note that we allow } i \text{ to be } n, \text{ which is a fake city south of all the real cities---this makes our base case simpler, since we can consider the 'journey' that starts at such a fake city and does nothing to have product 1.}$
 $\text{Relate: There are two places we could go from } C[i][j]. \text{ Whether we need to maximize or minimize the product after leaving } C[i][j] \text{ depends on the sign of } A[i][j]. \text{ We get the following recurrences.}$

$$T_{\max}(i, j) = \begin{cases} A[i][j] \cdot \max(T_{\max}(i+1, j), T_{\max}(i+1, j+1)) & \text{if } A[i][j] > 0 \\ 0 & \text{if } A[i][j] = 0 \\ A[i][j] \cdot \min(T_{\min}(i+1, j), T_{\min}(i+1, j+1)) & \text{if } A[i][j] < 0 \end{cases}$$

$$T_{\min}(i, j) = \begin{cases} A[i][j] \cdot \min(T_{\min}(i+1, j), T_{\min}(i+1, j+1)) & \text{if } A[i][j] > 0 \\ 0 & \text{if } A[i][j] = 0 \\ A[i][j] \cdot \max(T_{\max}(i+1, j), T_{\max}(i+1, j+1)) & \text{if } A[i][j] < 0 \end{cases}$$

 $\text{It would also work to take the max or min of all four possible products of } A[i][j] \text{ and a relevant subproblem.}$
 $\text{Topological order: } T_{\max}(i, j) \text{ and } T_{\min}(i, j) \text{ only depend on subproblems } T_{\max}(i', j') \text{ and } T_{\min}(i', j') \text{ with } i' > i.$
 $\text{Base case: Our recurrence breaks down in the last row (of fake cities) when } i = n. \text{ There we have } T_{\max}(n, j) = T_{\min}(n, j) = 1 \text{ for all } j.$
 $\text{Original problem: The maximum possible product is } T_{\max}(0, 0). \text{ We can follow use parent pointers to construct the path, by recording which option was the best for each subproblem.}$
 $\text{Time: It takes } O(1) \text{ time to solve each of } O(n^2) \text{ subproblems, so the runtime is } O(n^2).$
Problem 1. [20 points] DP (part)
 $\text{iven a string } S \text{ of length } n \text{ where each character is one of the 26 English characters, design an } O(n) \text{ namic programming algorithm to compute the length of the longest strictly-increasing sub-sequence in } S \text{ in alphabetical order.}$
 $\text{ake sure to use the SRBTBOT framework.}$
DP: state=constant
uition: Treat the characters as integers (1-26).

 S Let $T(i, k)$ be the length of a LIS in $S[i:n]$ with start value $\geq k$. ($0 \leq i \leq n, 1 \leq k \leq 27$)

$$R \quad T(i, k) = \max \begin{cases} T(i+1, k) \\ 1 + T(i+1, S[i] + 1) \quad \text{if } S[i] \geq k \end{cases}$$

T Decreasing k .

 $B \quad T(n, k) = T(i, 27) = 0 \text{ for all } i, k.$
 $O \quad T(0, 1).$
 T The table T has size $27(n+1) = O(n)$. Each entry takes $O(1)$ to compute for a total of $O(n)$.

 (d) [10 points] Give a full greedy algorithm to solve the problem. Prove correctness and analyze runtime.

Solution: Algorithm A :

GREEDY: big proof w/ ind/contra.
 $1. \text{ Order the intervals by } b_i$
 $O(n \log n)$
 $2. \text{ Set } S \leftarrow \emptyset \text{ and } b \leftarrow \infty$
 $O(1)$
 $3. \text{ Scan the intervals in order:}$
 $\text{When processing } [a_i, b_i], \text{ add } b_i \text{ to } S \text{ and set } b \leftarrow b_i \text{ if } a_i > b.$
 $O(n)$
 $4. \text{ Output } S$
 $O(1)$
 $\text{Total runtime: } O(n \log n). \text{ To prove correctness, we use strong induction on } n.$
 $1. \text{ Let } GREEDY \text{ be the output by } A, \text{ and } S \text{ be an optimal set of points that contains the greedy choice } b_k \text{ (which exists by (b)).}$
 $2. \text{ For contradiction, assume that } |GREEDY| > |S|.$
 $3. \text{ Let } T' \text{ be the set of intervals not covered by } b_k, \text{ i.e. intervals that start after } b_k.$
 $4. \text{ Note that after inserting } b_k \text{ into } S, \text{ ignores all intervals covered by } b_k, \text{ and the repeats the process on } T', \text{ solely for which } A \text{ would output } GREEDY \setminus \{b_k\}.$
 $5. \text{ By the induction hypothesis, } A \text{ is optimal for } T' \text{ as } |T'| < n \text{ and hence}$

$$|GREEDY \setminus \{b_k\}| \leq |S \setminus \{b_k\}|.$$

 $6. \text{ Therefore, } |GREEDY| \leq |S|, \text{ a contradiction.}$
Problem 5. [45 points] Reduction (3 parts)
 $Erik, Mohsen, and Brynnor want to celebrate the end of the semester by throwing a party. However, among their mutual friends, there are certain pairs of friends that cause trouble together.$
 $The Professors still want to celebrate, so they come up with a plan: host 3 parties (i.e. Erik, Mohsen, and Brynnor each throw their own party, so the troublemakers can go to separate parties and everyone can still celebrate). **REDUCTION**$
 $We define INVITE-FRIENDS as the decision problem: given a list of n friends and a list of m pairs of friends that cause trouble together, is it possible to invite everyone to one of the three parties.$
 $\text{For example, if the Professors are friends with [Maggie, Selina, Daph, Ragulan] and the trouble making pairs are [Maggie, Selina], [Selina, Daph], [Daph, Ragulan], [Selina, Ragulan], one solution to the problem is: Maggie and Ragulan could go to Erik's party, Daph could go to Mohsen's party, and Selina could go to Brynnor's party. Unfortunately, INVITE-FRIENDS is NP-complete. Poor Professors! To console them, you will write a great proof arguing that the INVITE-FRIENDS is NP-complete.$
 $\text{Unfortunately, INVITE-FRIENDS is NP-complete. Poor Professors! To console them, you will write a great proof arguing that the INVITE-FRIENDS is NP-complete.}$
 (a) [20 points] Show that INVITE-FRIENDS \in NP.

Solutions: To show that INVITE-FRIENDS \in NP, we show that we can verify a certificate in polynomial time.

 \bullet Let the certificate be three lists of friends: one list of attendees for each party.

 \bullet Verifier: For each party and all pairs of friends invited to that party, check if any pair will cause trouble. If any pair of friends at the same party cause trouble, return NO. Otherwise, return YES.

 \bullet If the professors have n friends, the certificate will be of size $O(n)$.

 \bullet Runtime for the verifier algorithm: For each of the $O(n^2)$ pairs of friends at the same party, scan the $O(m)$ list of friends that cause trouble. Thus, the verifier takes $O(n^2m)$, which is polynomial in the size of the instance of the problem (which is $\Theta(n+m)$).

 (b) [20 points] Show that INVITE-FRIENDS is NP-hard by reducing from 3-COLOR, an NP-complete problem.

 $3-COLOR$ is the following decision problem: Given a graph G , is it possible to color the vertices of G using 3 colors, such that neighboring vertices are not colored the same color? **3-COLOR INVITE FRIENDS**
 $Solution:$ In order to show INVITE-FRIENDS is NP hard, we must transform an arbitrary instance of 3-COLOR into an instance of INVITE-FRIENDS.

 $\text{Given graph } G = (V, E), \text{ we can let each vertex } u, v \text{ be a friend in INVITE-FRIENDS. Each edge } (u, v) \text{ can represent a pair of friends that cause trouble in INVITE-FRIENDS. So a valid coloring in 3-COLOR becomes a valid party assignment in INVITE-FRIENDS. Thus, we have transformed an instance of 3-COLOR into an instance of INVITE-FRIENDS.}$
 $\text{It takes } O(|V| + |E|) \text{ to iterate through the vertices and edges, so this reduction takes polynomial time.}$
 $\text{Thus, because 3-COLOR is NP-complete, and we can reduce 3-COLOR to INVITE-FRIENDS in poly-$
 $\text{nomial time, INVITE-FRIENDS is NP-hard.}$
 $\text{Each city has a } \text{disapproval score}, \text{ which is an integer. The disapproval scores may be zero, negative, or positive. The disapproval score for the } j^{\text{th}} \text{ city on row } i \text{ is } A[i][j]. \text{ For example, if } n = 4, \text{ the disapproval scores may be as shown above.}$
 $\text{Over all possible journeys, Sriini wants to know the } \text{maximum product} \text{ of disapproval scores of the cities he visits. In our example above for the particular values of } A, \text{ the maximum disapproval score corresponds to the path marked and has value } (-1)(-2)(-2)(-1) = 4.$
 $\text{Design an } O(n^2) \text{-time Dynamic Programming algorithm that returns the cities on the path with maximum product of disapproval scores. You may assume that all arithmetic operations take constant time.}$
(a) Given strings A , B and C , design a DP algorithm that checks in $O(|A||B||C|)$ time whether there is a way to merge A and B to produce C . You may assume that $|C| = |A| + |B|$. **PSETS**
 $\text{For example, if } A = abab, B = aabb \text{ and } C = abaabb + b \text{ where the first and last substrings come from } A \text{ and the middle one comes from } B.$
 $\text{Make sure to use the SRBTBOT framework.}$
Solution:
 $\text{We assume that } |A| = n, |B| = m \text{ and } |C| = n + m.$
 $\text{1. Subproblems Let } T(i, j) = \text{TRUE iff. one can merge } A[i:j] \text{ and } B[j:n] \text{ into } C[i:j+n]. \text{ Index } i \text{ runs from } 0, \dots, n \text{ and index } j \text{ runs from } 0, \dots, m.$
 $\text{2. Relate We consider the case where the next letter will come from } A \text{ and } B \text{ by considering incrementing either the pointer at } i \text{ or the pointer at } j, \text{ respectively.}$
 $T(i, j) = \text{V} \left\{ \begin{array}{l} T(i+1, j) \text{ if } C[i:j] == A[i:j] \\ T(i, j+1) \text{ if } C[i:j] == B[j:n] \end{array} \right.$
 $\text{We also handle edge cases—if one of our conditional statements requires checking an index out of range, we assume it to be FALSE.}$
 $\text{3. Topo. Order by } i + j$
 $\text{4. Base } T(0, 0) = \text{TRUE.}$
 $\text{5. Original } T(0, 0).$
 $\text{6. Time the table } T \text{ has size } (n+1)(m+1) \in O(nm). \text{ each entry takes constant time to compute, for a total of } O(nm).$
 (a) Frieren has a row of treasure chests numbered 1 through n . There is an unknown s such that chests 1 through s are safe, and the rest are mimics. Frieren can open chest i and learn whether or not it is safe by spending a positive amount m_i of mana.

 $\text{Frieren must compute } s \text{ without opening more than } m \text{ mimics, given } k \text{ and all of the } m_i \text{ inputs. Frieren will use the most efficient algorithm, in terms of the worst-case total mana cost. For example, if } k = 1, \text{ then the only correct algorithm is to open all chests in increasing order of } i, \text{ which costs } \sum m_i.$
 $\text{Design a DP that computes the worst-case mana cost of Frieren's algorithm. Your DP should run in time } O(n^3k).$
 $\text{Use the SRBTBOT framework, but you need not prove correctness beyond that.}$
 $\text{Hint: Discuss with friends or ask course staff to make sure you have the right subproblem definition!}$
Solution: Give as hint: $M(a, b, c)$ is the worst-case cost of Frieren's algorithm on input $(c, m[a : b])$
 $S \quad M(a, b, c) \text{ is the worst-case cost of Frieren's algorithm on input } (c, m[a : b]), \text{ where } 0 \leq c \leq k \text{ and } 0 \leq a \leq b \leq n$
 $R \quad M(a, b, c) = \min_{a \leq c \leq b} (m_q + \max\{M(a, q, c-1), M(q+1, b, c)\})$
 $T \quad \text{Decreasing } b - a$
 $B \quad M(a, b, a) = \infty \text{ if } b > a; M(a, a, c) = 0$

CATTU TEST 2

SIMPLE GRAPH:

- no duplicate edges and no self loops

SIMPLE PATH:

- no edge or node duplicates (no loops)

ADJACENCY LIST: $\Theta(|V| + |E|)$ - DAA, hash, LL

ADJACENCY MATRIX: $\Theta(|V|^2)$

REDUCTION: alg. for transferring one problem to another.

TURING A STB iff A is solved using B as atomic subroutine:

- poly time upper bound: cook red.

MANY-ONE: A \leq_{MB} iff 3 some f converts input of A to input of B

- if f poly time: KARP so $A \leq_{P} B$ means that f [inputs to A] \leq_{P} [inputs to B]
- if A reduces to B, CANT assume the converse
- A \leq_{B} in time f(n) & B solved in g(n), then A can be solved in f(n) · g(nf(n))

WEIGHTED GRAPHS:

GRAPH DUPLICATION: EX: only want to return even # paths

can use BFS to track even/odd





DUMMY NODES: insert nodes to match weight

must be positive weight & finite!

$0 \xrightarrow{3} 0 \xrightarrow{1} 0 \rightarrow 0 \rightarrow 0 \rightarrow 0$ becomes $0 \rightarrow 0 \rightarrow 0 \rightarrow 0 \rightarrow 0$

DIJKSTRA: $O(|V| \log |V| + |E|)$ SSSP on + weight graph

- distance estimates $d[v] = \infty$ except source $d[s] = 0$
- consider nodes in increasing distance
- for all outgoing edges, relax them.
- RELAX: if $d[v] > d[u] + w(u, v)$, set $d[v] = d[u] + w(u, v)$
- # want to relax each edge exactly once
- # doesn't work w/ negative edge weights $d = \text{pred.}$
- at end, $d(v) = d^*(v)$ \leftarrow path relaxation lemma

DIJKSTRA DIFFERENT METRICS:

- min, +: $[\infty, \infty]$ \rightarrow minimizes sum of edge weights along path. SSSP
- min, x: $[\infty, \infty]$ \rightarrow minimizes product of edge weights along path
- min, max: $(-\infty, \infty)$ \rightarrow minimizes largest edge along path
- max, min: $(-\infty, \infty)$ \rightarrow maximizes smallest edge along path
- max, x: $[0, \infty]$ \rightarrow maximizes product of edge weights (all < 1)

for paths:

TRIANGLE INEQ: $d(a, c) \leq d(a, b) + d(b, c)$

FOR EDGE: $w(a, c) > w(a, b) + w(b, c)$

DAG SP: $O(|V| + |E|)$ acyclic + g - edge weights

- start distances from ∞
- get topological order w/ full DFS
- relax outgoing edges of each node in this order.
- if $d[v] > d[u] + w(u, v)$, set $d[v] = d[u] + w(u, v)$ (triangle inequality)

BF & DAG SP - DIFFERENT METRICS:

- min, +: $(-\infty, \infty)$ \rightarrow minimizes sum of edge weights
- min, x: $[0, \infty)$

JOHNSON'S:

ALL PAIRS SHORTEST PATH: 'distance b/t any 2"

GOAL: output a $|V| \times |V|$ table $\rightarrow |E| \leq 2(|V|^2)$

- if DAG: run DAG SP from each node $O(|V||E| + |V|^2)$
- if non-neg. weights: Dijkstras from each node $O(|V||E| + |V|^2 \log |V|)$

→ make weights non-negative while preserving SP?

① $\phi(v)$ is some function on v for each v, subtract $\phi(v)$ from all incoming edges add $\phi(v)$ to all outgoing edges

all paths from s to t change only by $\phi(s) - \phi(t)$

② set $\phi(v) = \text{SSSP } d^*(v)$ from some start node new weight for (u, v) : $w'(u, v) = w(u, v) + d^*(u) - d^*(v)$ non-neg.

SSSP DISTANCE from ANY...: SUPERNODE

- add supernode w/ edges to all nodes. run BF from it. $O(|V||E|)$
- once reweighted, Dijkstras from every node. $O(|V||E| + |V|^2 \log |V|)$

allows for disconnected nodes to be calc.

OUTPUT: min distance from every $u \rightarrow v$ in graph (often represented as matrix)

FLOYD-WARSHALL: $O(|V|^3)$ - APSP

- DP alg. finding SP b/t all pairs of vertices in weighted graph
- undirected & directed, handles neg. edge weights
- doesn't handle neg. cycles

GRAPH: $G = (V, E)$ is set of pairs of vertices V & edges (pairs of vertices) $E \subseteq V \times V \rightarrow |E| = O(|V|^2)$

adj⁺(u) = set of outgoing edges from u

adj⁻(u) = set of incoming edges

$\deg^+(u) = |\text{adj}^+(u)|$ $\deg^-(u) = |\text{adj}^-(u)|$

$\sum_{v \in V} \deg^*(u) = |E|$

ADJACENCY LIST: $\Theta(|V| + |E|)$ - DAA, hash, LL

ADJACENCY MATRIX: $\Theta(|V|^2)$

PATH: sequence of vertices connected by edges

$d^*(u, v)$: shortest path from $u \rightarrow v$ (∞ if no path)

Single-pair-reachability (G, s, t): True if path $s \rightarrow t$

Single-pair-shortest-path (G, s, t): SP & distance $d^*(s, t)$

Single-source-shortest-path (G, s): SP tree & dist. for all V starting @ s .

SHORTEST PATH TREE:

- contains one path from s to every vertex v reachable from s .
- parent array (pred.)
- can convert bit SPT & $d^*(s, -)$ in linear time.

UNDIRECTED GRAPHS:

CONNECTED COMPONENTS: $u \rightarrow v$ path b/t vertices in same CC

a-b 

BFS for each node & find CCs $\rightarrow O(|V| + |E|)$ (touch each node once)

EQUAL WEIGHT GRAPHS:

- # for directed & undirected
- BFS: $O(|V| + |E|)$ DFS: $O(|E|)$

FULL DFS: run DFS from every unexplored vertex in G until all are explored. $O(|V| + |E|)$

- go in layers
- not SP
- solves SPR
- return shortest reachability distances from start node to each reachable nodes from node (shortest, start node in layers)
- retrieves set of reachable nodes from start node

FINISHING ORDER: mark in order DFS run has explored V & all its neighbors. parent FO is always $>$ than children (child finishes first)

CYCLE DETECTION (full dfs):

- run full DFS, check active vs. finished for nodes
- if an edge goes back to a visited but unfinished node, that node is ancestor in DFS & forms cycle.

DIRECTED GRAPHS:

STRONGLY CONNECTED COMPS (SCCS): 3 u-v & v-u path for vertices in directed graph

a-b 

CONDENSATION GRAPH: src sink src
 $a \rightarrow b \leftarrow c \rightarrow d \rightarrow e$

V = SCCs of graph

E = 3 edge from vertex in C_i to vertex in G .

CG must be DAGS! use Kosaraju-Shavir to make SCCs

KOSARJU-SHAVIR: finds SCCs $O(|V| + |E|)$

- run full DFS on G & record finishing times $f[v]$ for each v in reverse $f[V]$ order.
- for each v in rev. finishing order:
 - set current = v
 - if v = unvisited, run DFS starting from v . for any vertex u w/o leader, set $\text{leader}[u] = v$ (grouping SCCs)

intuition: hierarchy of SCCs

← starts from 3rd, then 2, 1.

TOPOLOGICAL SORT: order of nodes s.t. $(u, v) \in E$, then u comes before v in order.

- full DFS, topo order = reverse finishing time.
- # must be acyclic, directed graphs

CHANGEABLE PQS: $O(|V| \log |V| + |E|)$

- build \rightarrow generating RT for Dijkstra
- insert
- delete min $O(\log n)$
- decrease key(i, new.key) $O(1) \leftarrow$ Fib. heap
 - change val $i \rightarrow$ & heapify-up
 - store every $v \in V$ keyed by distance estimate
 - perform one delete min for every node and one decrease key for every edge

interface:

PRIORITY QUEUE:

- build(A)
- insert(x)
- delete-min() / delete-min() $O(\log n)$

implementation:

HEAP: AVL:

- $O(n)$ $O(n \log n)$
- $O(\log n)$ $O(\log n)$

COMPACT BINARY TREE: all nodes as far left as possible

- has 1:1 correspondence w/ array implementation
- n nodes
- parent x 's children:
 - LEFT(x) = $2x$
 - RIGHT(x) = $2x + 1$
 - PARENT(x) = $\lfloor \frac{x}{2} \rfloor$

HEAPS: compact binary tree implementation of PQ I.P. ✓

MAX HEAP PROPERTY: for $x \in \text{Tree}$, $x \geq \text{left}(x), x \geq \text{right}(x)$

→ MIN HEAP: opposite (parent = min elt.)

OPERATIONS:

- find: $O(n) \rightarrow$ not sorted
- heapify-up(x): $O(\log n) \rightarrow$ swap x up tree while bigger than parent
- heapify-down(x): $O(\log n) \rightarrow$ swap x down while smaller than child

PQ OPERATIONS (w/ HEAP):

- insert(x): $O(\log n)$ put x @ end, heapify-up(x)
- delete-min(): $O(\log n)$ swap root & last elt in array, delete last, heapify-down from root
- build: $O(n)$ heapify every item down starting from leaves S: x

SORTING {

- HEAP SORT:** $O(n \log n)$ I.P. ✓
- build a heap (n = size of heap)
- repeatedly remove max elt.

BELLMON-FORD: $O(|V| \cdot |E|)$ + g - weight

- any SP takes at most $|V| - 1$ edges
- make $|V|$ layers & run DAG SP to look at last layer distance
- "how far we can get after # edges"

NEGATIVE CYCLE DETECTION:

- BF & add another layer
- anything that has decr. at added layer is in negative cycle. can ID by tracing parent pointers
- DFS all points reachable to cycles & remove them.

≤ 0 edges

≤ 1 edge

≤ 2 edges

≤ 3 edges

HANDLING NEGATIVE CYCLES:

- compute SCCs & condensation graph $O(|V| + |E|)$
- within each SCC in original graph, run BF to learn which SCCs have negative cycles. $O(|V| \cdot |E|)$
- weigh condensation graph edges s.t. outgoing edges from a negative cycle SCC have weight -1, o.w. edges have weight 0.
- get DAG APSP in condensation graph $d^*(A, B)$
- set $d^*(A, A) = -1$ if A contains neg. cycle
- create G' by removing all negative SCCs
- run Johnson's on G' to get $d^*(u, v)$
- go through every u, v let A, B be their SCCs if $d^*(A, B) < 0$, $d^*(u, v) = -\infty$ o.w. $d^*(u, v) < d^*(u, v)$

REC + PTS:

5. Given an unweighted graph $G = (V, E)$ in which some edges are red and some are blue, find a path from s to t with the minimal number of red edges.

Solution: We combine BFS and DFS. The main idea is that using a DEQUE, we can BFS w.r.t. red edges while simultaneously condensing blue edges using DFS.

6FS + DFS

- Initialize a DEQUE $Q = [s]$, $d(s) = 0$, and $P(s) = \perp$
- While Q is non-empty:
 - Remove the first element u from Q
 - For each unvisited out-neighbor v of u :
 - set $P(v) = u$
 - if (u, v) is red, set $d(v) = d(u) + 1$ and append v to Q
 - if (u, v) is blue, instead set $d(v) = d(u)$ and prepend v to Q
 - Follow parent pointer from t and output the resulting path

Correctness and runtime proofs are as DFS and BFS.

3. There are n lock boxes and m keys. Each box has a distinct lock, so each key can open exactly one box. There's at least one copy of each box's key, but for some boxes there may be multiple copies of the key. Some people put all keys in the boxes and locked them up, but luckily they made a note of which keys are stored in each box. Keys and boxes are numbered so that we know which box is opened by each key. Some boxes contain no keys while others contain multiple keys. Boxes can also be forced open with a rusty crowbar. Design an algorithm to find the smallest set of boxes that you need to force open in order to open all the boxes. **K5, CG + INDEGREES** **O(n+m)**

Solution: Define a graph $G = (V, E)$ such that there is a node in V for each lock box and there is a directed edge $(u \rightarrow v) \in E$ between two lock boxes u and v if lock box u contains the key to lock box v .

The key idea is that if there is a path $a \rightarrow b \rightarrow c$ in G , then opening box a gives us the key to b , and opening box b gives us the key to c . So if we open a box, we get access to every box reachable from it in G . As a consequence, opening a box allows us to open all other boxes in its strongly connected component.

This last fact suggests we consider the condensation graph G_C . If a strongly connected component x of G is a *source* in G_C , meaning there are no edges into x , then we must force open some box in x . This gives as access to every box in x , as well as every box in every SCC reachable in (G_C) from x .

After forcing open one box in every source SCC, we're done: if a vertex in G_C isn't a source, we can follow edges backwards until reaching a source, so every vertex is 'downstream' of a source.

So this question is equivalent to finding the number of source SCCs in G_C . To do this, we construct G_C , and then for each edge $x \rightarrow y$ in G_C we mark y as not a source. The vertices that do not get marked are the sources. This takes $O(n + m)$ time.

1. Totodile has k PP to spend and wants to go to Tangela Island. Design an $O(k|E|)$ time algorithm to decide whether Totodile can reach Tangela Island. **DUNNY NOOBS**, **EPS**

Solution: This reduces to SPSP by using a variant of graph duplication that subdivides edges instead of duplicating the original vertices. We assume WLOG that G is connected; or we can prune unreachable vertices in $O(|E|)$ time using e.g. DFS.

• Construct a graph G_1 by subdividing every edge (u, v) into a $c_{u,v}$ -edge chain of edges if $c_{u,v} \leq k$ (instead remove the edge if $c_{u,v} > k$).

• BFS in G_1 from Shamouti Island to compute the distance d to Tangela Island.

• Output $d \leq k$.

2. Ducklett has k PP to spend and wants to go to Tarocco Island. Ducklett can either swim across a route (u, v) by spending $c_{u,v}$ PP, or fly over a route by spending 1 PP, regardless of $c_{u,v}$. However, after flying, Ducklett must rest and cannot fly again until after swimming across another route. Design an $O(k|E|)$ time algorithm to decide whether Ducklett can reach Tarocco Island. **GRAPH DUPLICATION** **O(k|E|)**

Solution: To solve this problem, we combine the ideas of the previous two parts with some graph duplication ideas. We construct a G_2 with two vertices per island:

- (a) vertex i_F corresponding to being at island i with Ducklett being rested
- (b) vertex i_R corresponding to being at island i after flying

Then for each route $[u, v]$, add the following edges:



and if the cost of the route is at most k PP, we also add following additional edges:

- $c_{u,v}$ -edge directed chain of edges from u_F to v_F
- $c_{u,v}$ -edge directed chain of edges from u_F to v_R
- $c_{u,v}$ -edge directed chain of edges from u_R to v_F
- $c_{u,v}$ -edge directed chain of edges from u_R to v_R

The number of edges and vertices in this graph is still $\Theta(k|E|)$, so constructing this graph takes $\Theta(k|E|)$ time. Then any sequence of routes from Shamouti Island to Tarocco Island in this graph corresponds to visiting a sequence of islands without flying twice in a row. As before, we can run BFS from Shamouti Island (starting from the vertex that is allowed to fly to other islands) to find the distance d to Tarocco Island (either vertex in $\Theta(k|E|)$ time). Ducklett can reach Tarocco Island iff $d \leq k$.

3. CIA officer Mary Cathouse needs to drive to meet with an informant across an unwelcome city. Some roads in the city are equipped with government surveillance cameras, and Mary will be detained if cameras from more than one road observe her car on the way to her informant. Mary has a map describing the length of each road and knows which roads have surveillance cameras. Help Mary find the shortest drive to reach her informant, being seen by at most one surveillance camera along the way. **DIESTRA + DUPLICATION**

Solution: Construct a graph having two vertices $(v, 0)$ and $(v, 1)$ for every road intersection v within the city. Vertex $(v, 0)$ represents arriving at intersection v having already been spotted by exactly 1 camera(s). For each road from intersection v to u add two directed edges from $(v, 0)$ to $(u, 0)$ and from $(v, 1)$ to $(u, 1)$ if traveling on the road will not be visible by a camera, and add one directed edge from $(v, 0)$ to $(u, 1)$ traveling on the road will be visible by a camera. If a road $v \rightarrow u$ is visible by k cameras, then $c_{v,u} = k$. If $v \rightarrow u$ is not visible by a camera, then $c_{v,u} = 0$. In the constructed graph, will be a path visible by at most one camera. Let n be the number of road intersections and m be the number of roads in the network. Assuming lengths of roads are positive, use Djikstra's algorithm to find the shortest path $\Theta(m + n \log m)$ time using a Fibonacci Heap for Djikstra's priority queue.

4. Ash is trying to cycle from Pallet Town to Viridian City without destroying Misty's bike. For every trail e in the area, he knows the probability $p(e)$ of destroying the bike if he cycles along e . Help Ash find the safest path to Viridian City (the path that minimizes his probability of destroying the bike). Assume that all probabilities are independent and that arithmetic operations take constant time. **DIESTRA, MAX X, O(1)**

Solution: Construct a graph having a vertex v for every trail intersection, and weight each edge (v, e) with $p(e) = 1 - p(e)$. Run Djikstra using $(0, 1)$, max, \perp , 0 instead of $[0, \infty]$, min, \perp , ∞ , respectively. If N is the number of intersections and M the number of trails in the graph, then Djikstra will take $O(M + N \log M)$ time using a Fibonacci Heap.

5. Problem 2, [15 points] **Tricky Tolls** (1 part)

1. You are traveling from Syroco to Throntos, and you wish to pay off your fellow classmates. Your map shows that the students are traveling through n cities, connected by m two-way roads, and each road charges a toll s_i except for the one road you own between Hamilton and Burlington. You are allowed to set your toll price to be any positive integer you'd like.

2. Design an algorithm that returns the highest toll price you could set so that the MIT students will still choose to take your road, or state that no toll exists. The students always pick the cheapest route. In case of a tie, the students will pick the route that does not use your road. Analyze the runtime of your algorithm. You don't need to prove correctness.

REWEIGHT, BF

Solution: Build a graph having a vertex w for every trail intersection, and weight each edge (w, e) with $p(e) = 1 - p(e)$. Run Djikstra using $(0, 1)$, max, \perp , 0 instead of $[0, \infty]$, min, \perp , ∞ , respectively. If N is the number of intersections and M the number of trails in the graph, then Djikstra will take $O(M + N \log M)$ time using a Fibonacci Heap.

6. Problem 2, [15 points] **Tricky Tolls** (1 part)

1. You are traveling from Syroco to Throntos, and you wish to pay off your fellow classmates. Your map shows that the students are traveling through n cities, connected by m two-way roads, and each road charges a toll s_i except for the one road you own between Hamilton and Burlington. You are allowed to set your toll price to be any positive integer you'd like.

2. Design an algorithm that returns the highest toll price you could set so that the MIT students will still choose to take your road, or state that no toll exists. The students always pick the cheapest route. In case of a tie, the students will pick the route that does not use your road. Analyze the runtime of your algorithm. You don't need to prove correctness.

ONE UNIQUE EDGE

2 BFS

Solution: Here, s, t, \perp can be computed using a BFS on G_{Toll} .

BFS of G w/o road \rightarrow 0, (s, t)

BFS of G w/ road \rightarrow 0, (s, t)

Output \perp , (now much we can change)

Solution: Let $G = (V, E)$ be a directed graph. Say that vertices $u, v \in V$ are semi-connected if G contains either a u -path or a v -path to both u and v .

Problem 3. Hide and Seek **SKITTY**

Solution: Hoothoot wants to evade Skitty as long as possible. A naive attempt is to compute the vertex x that is reachable from y and far from v_S as possible, go to x , and wait. Describe justifying your answer.

Problem 4. [16 points] Prismatic Paths (1 part)

Solution: Let $G = (V, E)$ be a positive-weighted, undirected graph, and let $s, t \in V$. Suppose each edge of G has one of seven colors, and every vertex has exactly one incident edge of each color. Design an $O(|V| \log |V|)$ algorithm to compute the shortest $s-t$ path that contains edges of at least three different colors.

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state \perp that collapses all levels of all these colors. (Note this avoidable.)

Solution: We duplicate G 30 times to keep track of the 29 sets of at most two colors, plus a success state $\perp</$

CATTU
6.121 1

KARATSUBA:
A = multiply(x_{10}, y_{10})
B = multiply(x_{10}, y_{10})
C = multiply($x_{10} + x_{11}, y_{10} + y_{11}$)

$T(n) = 3T\left(\frac{n}{2}\right) + C_2 n$
 $\Theta(n \log_2^2 n)$

$\log x^n < n^x < n!$
poly exp fact

INTERFACES:

STACK: LIFO	QUEUE: FIFO	DEQUEUE: double-end queue	SET: key-value pairs
push(x)	enqueue(x)	push-front	build
pop()	dequeue(x)	push-back	find(k)
peek()	peek()	pop-front	unique keys
isempty()	len()	pop-back	multiset: can be non-unique (chaining)

SEQUENCE: maintain order; only have indices

IMPLEMENTATIONS:

DOUBLY LINKED:	LINKED LIST:
• pointer-based	• x_1, x_2, \dots, x_n
• x_1, x_2, \dots, x_n	• x_1, x_2, \dots, x_n
• can track tail	• x_1, x_2, \dots, x_n

SORTING:

INSERTION $\Theta(n^2)$	SELECTION $\Theta(n^2)$	MERGE SORT $\Theta(n \log n)$	DIVIDE & CONQUER $\Theta(n \log n)$
• comparing $S: \checkmark$	• handpicking $S: \times$	• stable: prefer left arr. 8 6 7 5 3 0 9 div. 8 cong. 0 7 8 0 3 5 9 S: \checkmark 0 3 5 6 7 8 9 I.P: \times $T(n) = \Theta(n) + 2T(n/2)$	• base case: if $n = 1$: • if $A[m] \dots$, recurse on $A[m:]$ • else, recurse on $A[:m]$ • prove correctness on B.C., both sides, no overlap
$S: \checkmark$	$S: \times$		
$I.P: \checkmark$	$I.P: \checkmark$		

DIRECT-ACCESS ARRAY: $\Theta(u)$

- all keys distinct $S: \checkmark$
- u = possible range $I.P: \times$

$\begin{bmatrix} 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \\ 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \end{bmatrix}$

AVL PROPERTY: V_{NET} , $l_{SKEW}(n) \leq 1$ $h = \Theta(\log n)$ \times I.P.

- at most 2 rot. to fix violations $find(x) = \Theta(n)$
- AVL sequence: maintains order of Elts. nodes contain height(size of subtree)
- AVL set: BST property (comparison model)

AUGMENTATIONS: dynamic order stats, computed for nodes from itself & left/right

ORDER-STATISTIC TREE: os-select(T, i), os-rank(T, x)

- T.size = T.left.size + T.right.size + 1
- i'th order: return element w/ i'th smallest key
- rank: give element + return position i in ascending ord

INTERVAL TREES: interval-search($T[a, b]$)

- find node who's interval intersects i if exists
- T.low & T.high
- keyed by max T.high go left if T.left max $\geq a$
- while node T doesn't overlap, traverse down tree.
- if left.max $\geq a$ go left if left.max $< T.a$, skip.

bst:

- distinct keys
- all keys on left subtree $<$ node
- all keys on right subtree $>$ node
- uses comparison model
- ins/find/del $\in \Theta(n)$, height $\in \Theta(n)$

MASTER'S THM: let $T(n) = aT\left(\frac{n}{b}\right) + f(n)$, $b > 1, a \geq 1$

$O(n)$ work done \leq **CASE 1:** if $f(n) \in O(n^{\log_b a - \epsilon})$, for some $\epsilon > 0$, $T(n) \in \Theta(n^{\log_b a})$

$O(n \log n)$ evenly dist. = **CASE 2:** if $f(n) \in \Theta(n^{\log_b a})$, $T(n) \in \Theta(n^{\log_b a} \cdot \log n)$

$O(n^2)$ work done @ leaves \geq **CASE 3:** if $f(n) \in \Omega(n^{\log_b a + \epsilon})$ & $aT\left(\frac{n}{b}\right) \leq c f(n)$, $T(n) \in \Theta(f(n))$ for some $\epsilon > 0$

ASYMPTOTICS:

$f \in \Theta(g)$	$f \in g$	$f \in O(g) \leq$	$f \in o(g) <$	$f \in \omega(g) >$	$f \in \Omega(g) \geq$
$f \in O(g) \&$	$\frac{f(n)}{n \rightarrow \infty} = 1$	$3c > 0 \exists n_0 \geq 0$	$\frac{f(n)}{n \rightarrow \infty} = 0$	$\frac{f(n)}{n \rightarrow \infty} = \infty$	$g \in O(f)$
$f \in \Omega(g)$	$\frac{f(n)}{n \rightarrow \infty} = \infty$	$\forall n_0 \exists n \geq n_0$ $f(n) \leq c g(n)$	$\frac{f(n)}{n \rightarrow \infty} = \infty$	$\frac{f(n)}{n \rightarrow \infty} = \infty$	

AMORTIZED Cost (T): if any sequence of m ops. (starting from empty d.s.) takes $\leq M$ # decr. greatly when potential func $\Phi(A) \geq 0$ for all ops. from $A_i \rightarrow A_{i+1}$ that take time t_i . A.C. = $t_i + \Phi(A_{i+1}) - \Phi(A_i)$

Ex: $\Phi(A) = c \lceil n - \frac{m}{2} \rceil$ $c = \text{const.}$, $n = \# \text{ Elts.}$, $m = \text{capacity}$

WORLDRAM: math model, steps ops. WORD: W bits, fit in 2 reg./mem. if $* < 2^W$ LOAD: $R[i] \leftarrow M[R[i]]$, STORE: $M[R[i]] \leftarrow R[i]$

$O \rightarrow O$
 $W \rightarrow S2$
 $O \otimes \Omega \leftrightarrow \Theta$

RECURRENCES: inductive sequence

$T(n) = aT\left(\frac{n}{b}\right) + f(n)$

PLUG & CHUG:

$T(n) = aT\left(\frac{n}{b}\right) + f(n)$

$T\left(\frac{n}{b}\right) = a\left(aT\left(\frac{n}{b^2}\right) + f\left(\frac{n}{b}\right)\right) + f\left(\frac{n}{b}\right)$

\vdots

$= a^k T\left(\frac{n}{b^k}\right) + -n^? \sum_{i=1}^k \frac{1}{b^i}$ get closed form

$= \dots \text{ Simplify}$

let $k = \log_b n$ and we know $b^k = n$

$T(n) = a^{\log_b n} T(1) + \dots$

(if prove): let $P(n) := T(n) = \dots$ (induction)

LOG PROPERTIES:

$\log_a b = \frac{\log_b b}{\log_b a}$

$\log(x^z) = z \log x$

$\log(ab) = \log a + \log b$

$\log_b a = \frac{1}{\log_a b}$

DERIVATIVES:

$\frac{d}{dx}(\log_a x) = \frac{1}{x \ln a}$

$\frac{d}{dx}(\ln x) = \frac{1}{x}$

$\frac{d}{dx}(a^x) = a^x \ln a$

IMPLEMENTATIONS:

LINKED LIST:

- pointer-based
- x_1, x_2, \dots, x_n
- x_1, x_2, \dots, x_n
- ins/del in $\Theta(1)$
- isempty in $\Theta(1)$

STABLE: preserve order: Elts. $w/ =$ keys

IN PLACE: no added memory

DYNAMIC ARRAY: expands/shrinks

- expands when doubling $2x$
- when $cn/4$ Elts, halve & move Elts.
- avg. 2 series of expensive ops. over series of cheap ones
- good for get-at(i) $\rightarrow \Theta(1)$

HASHING: SUHA

- lookup by key
- pairwise indep.
- chain len $\Theta(n \alpha + 1)$, $\alpha = \frac{n}{m}$
- $n = \text{total items inserted}$
- $m = \text{len of array}$
- $L = \text{bucket size}$
- collisions \rightarrow chaining (points to DAA)

$P[h(k_j) = h(k_{\text{new}})] = \frac{1}{m}$ (suha)

$E[L] = 1 + \Pr[L = 1] + \frac{n}{m} = 1 + \alpha E[1]$

LOAD FACT. α : avg. # items/slot (n/m)

RADIX SORT: $\Theta(cn)$, $\Theta(n \log n)$ or $\Theta(n \log \log n)$

- must be bounded
- $b = \# \text{ buckets}$
- tuple sort w/ aux. counting sort
- if every key $c \in \mathbb{N}$ \rightarrow linear $\Theta(cn)$
- whole #s $\in \mathbb{N}^d$ $\# \text{ digits} = c$ buckets
- if n words of len d in base b , each round $\Theta(n+b)$, total time $\Theta(n+b)d$
- Sort LSB \rightarrow MSB

STABILITY: $\Theta(n^2)$

COUNTING SORT: $\Theta(n+k)$

- keys don't have to be distinct - have a sequence within list
- use queue, etc. for stable
- $S: \checkmark$
- $I.P: \times$
- $K = \text{bucket size}$
- $n = \text{len arr.}$, $[a_1, a_2] \text{ Pseudo poly!!}$

TUPLE SORT USING _____:

- need auxiliary algo.
- $\Theta(n + n \log n)$ \leftarrow counting S.
- linear if $\log n \leq \Theta(1)$
- $[1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106, 107, 108, 109, 110, 111, 112, 113, 114, 115, 116, 117, 118, 119, 120, 121, 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132, 133, 134, 135, 136, 137, 138, 139, 140, 141, 142, 143, 144, 145, 146, 147, 148, 149, 150, 151, 152, 153, 154, 155, 156, 157, 158, 159, 160, 161, 162, 163, 164, 165, 166, 167, 168, 169, 170, 171, 172, 173, 174, 175, 176, 177, 178, 179, 180, 181, 182, 183, 184, 185, 186, 187, 188, 189, 190, 191, 192, 193, 194, 195, 196, 197, 198, 199, 200, 201, 202, 203, 204, 205, 206, 207, 208, 209, 210, 211, 212, 213, 214, 215, 216, 217, 218, 219, 220, 221, 222, 223, 224, 225, 226, 227, 228, 229, 230, 231, 232, 233, 234, 235, 236, 237, 238, 239, 240, 241, 242, 243, 244, 245, 246, 247, 248, 249, 250, 251, 252, 253, 254, 255, 256, 257, 258, 259, 260, 261, 262, 263, 264, 265, 266, 267, 268, 269, 270, 271, 272, 273, 274, 275, 276, 277, 278, 279, 280, 281, 282, 283, 284, 285, 286, 287, 288, 289, 290, 291, 292, 293, 294, 295, 296, 297, 298, 299, 300, 301, 302, 303, 304, 305, 306, 307, 308, 309, 310, 311, 312, 313, 314, 315, 316, 317, 318, 319, 320, 321, 322, 323, 324, 325, 326, 327, 328, 329, 330, 331, 332, 333, 334, 335, 336, 337, 338, 339, 340, 341, 342, 343, 344, 345, 346, 347, 348, 349, 350, 351, 352, 353, 354, 355, 356, 357, 358, 359, 360, 361, 362, 363, 364, 365, 366, 367, 368, 369, 370, 371, 372, 373, 374, 375, 376, 377, 378, 379, 380, 381, 382, 383, 384, 385, 386, 387, 388, 389, 390, 391, 392, 393, 394, 395, 396, 397, 398, 399, 400, 401, 402, 403, 404, 405, 406, 407, 408, 409, 410, 411, 412, 413, 414, 415, 416, 417, 418, 419, 420, 421, 422, 423, 424, 425, 426, 427, 428, 429, 430, 431, 432, 433, 434, 435, 436, 437, 438, 439, 440, 441, 442, 443, 444, 445, 446, 447, 448, 449, 450, 451, 452, 453, 454, 455, 456, 457, 458, 459, 460, 461, 462, 463, 464, 465, 466, 467, 468, 469, 470, 471, 472, 473, 474, 475, 476, 477, 478, 479, 480, 481, 482, 483, 484, 485, 486, 487, 488, 489, 490, 491, 492, 493, 494, 495, 496, 497, 498, 499, 500, 501, 502, 503, 504, 505, 506, 507, 508, 509, 510, 511, 512, 513, 514, 515, 516, 517, 518, 519, 520, 521, 522, 523, 524, 525, 526, 527, 528, 529, 530, 531, 532, 533, 534, 535, 536, 537, 538, 539, 540, 541, 542, 543, 544, 545, 546, 547, 548, 549, 550, 551, 552, 553, 554, 555, 556, 557, 558, 559, 560, 561, 562, 563, 564, 565, 566, 567, 568, 569, 570, 571, 572, 573, 574, 575, 576, 577, 578, 579, 580, 581, 582, 583, 584, 585, 586, 587, 588, 589, 589, 590, 591, 592, 593, 594, 595, 596, 597, 598, 599, 600, 601, 602, 603, 604, 605, 606, 607, 608, 609, 610, 611, 612, 613, 614, 615, 616, 617, 618, 619, 620, 621, 622, 623, 624, 625, 626, 627, 628, 629, 630, 631, 632, 633, 634, 635, 636, 637, 638, 639, 640, 641, 642, 643, 644, 645, 646, 647, 648, 649, 650, 651, 652, 653, 654, 655, 656, 657, 658, 659, 660, 661, 662, 663, 664, 665, 666, 667, 668, 669, 669, 670, 671, 672, 673, 674, 675, 676, 677, 678, 679, 680, 681, 682, 683, 684, 685, 686, 687, 688, 689, 689, 690, 691, 692, 693, 694, 695, 696, 697, 698, 699, 699, 700, 701, 702, 703, 704, 705, 706, 707, 708, 709, 709, 710, 711, 712, 713, 714, 715, 716, 717, 718, 719, 719, 720, 721, 722, 723, 724, 725, 726, 727, 728, 729, 729, 730, 731, 732, 733, 734, 735, 736, 737, 738, 739, 739, 740, 741, 742, 743, 744, 745, 746, 747, 748, 749, 749, 750, 751, 752, 753, 754, 755, 756, 757, 758, 759, 759, 760, 761, 762, 763, 764, 765, 766, 767, 768, 769, 769, 770, 771, 772, 773, 774, 775, 776, 777, 778, 779, 779, 780, 781, 782, 783, 784, 785, 786, 787, 788, 789, 789, 790, 791, 792, 793, 794, 795, 796, 797, 798, 799, 799, 800, 801, 802, 803, 804, 805, 806, 807, 808, 809, 809, 810, 811, 812, 813, 814, 815, 816, 817, 817, 818, 819, 820, 821, 822, 823, 824, 825, 826, 827, 828, 829, 829, 830, 831, 832, 833, 834, 835, 836, 837, 838, 839, 839, 840, 841, 842, 843, 844, 845, 846, 847, 848, 849, 849, 850, 851, 852, 853, 854, 855, 856, 857, 858, 859, 859, 860, 861, 862, 863, 864, 865, 866, 867, 868, 869, 869, 870, 871, 872, 873, 874, 875, 876, 877, 878, 879, 879, 880, 881, 882, 883, 884, 885, 886, 887, 888, 889, 889, 890, 891, 892, 893, 894, 895, 896, 897, 898, 899, 899, 900, 901, 902, 903, 904, 905, 906, 907, 908, 909, 909, 910, 911, 912, 913, 914, 915, 916, 917, 918, 919, 919, 920, 921, 922, 923, 924, 925, 926, 927, 928, 929, 929, 930, 931, 932, 933, 934, 935, 936, 937, 938, 939, 939, 940, 941, 942, 943, 944, 945, 946, 947, 948, 949, 949, 950, 951, 952, 953, 954, 955, 956, 957, 958, 959, 959, 960, 961, 962, 963, 964, 965, 966, 967, 968, 969, 969, 970, 971, 972, 973, 974, 975, 976, 977, 978, 979, 979, 980, 981, 982, 983, 984, 985, 986, 987, 988, 989, 989, 990, 991, 992, 993, 994, 995, 996, 997, 998, 999, 999, 1000, 1001, 1002, 1003, 1004, 1005, 1006, 1007, 1008, 1009, 1009, 1010, 1011, 1012, 1013, 1014, 1015, 1016, 1017, 1018, 1019, 1019, 1020, 1021, 1022, 1023, 1024, 1025, 1026, 1027, 1028, 1029, 1029, 1030, 1031, 1032, 1033, 1034, 1035, 1036, 1037, 1038, 1039, 1039, 1040, 1041, 1042, 1043, 1044, 1045, 1046, 1047, 1048, 1049, 1049, 1050, 1051, 1052, 1053, 1054, 1055, 1056, 1057, 1058, 1059, 1059, 1060, 1061, 1062, 1063, 1064, 1065, 1066, 1067, 1068, 1069, 1069, 1070, 1071, 1072, 1073, 1074, 1075, 1076, 1077, 1078, 1079, 1079, 1080, 1081, 1082, 1083, 1084, 1085, 1086, 1087, 1087, 1088, 1089, 1089, 1090, 1091, 1092, 1093, 1094, 1095, 1095, 1096, 1097, 1098, 1099, 1099, 1100, 1101, 1102, 1103, 1104, 1105, 1106, 1107, 1108, 1109, 1109, 1110, 1111, 1112, 1113, 1114, 1115, 1116, 1117, 1118, 1119, 1119, 1120, 1121, 1122, 1123, 1124, 1125, 1126, 1127, 1128, 1129, 1129, 1130, 1131, 1132, 1133, 1134, 1135, 1136, 1137, 1138, 1139, 1139, 1140, 1141, 1142, 1143, 1144, 1145, 1146, 1147, 1148, 1149, 1149, 1150, 1151, 1152, 1153, 1154, 1155, 1156, 1156, 1157, 1158, 1159, 1159, 1160, 1161, 1162, 1163, 1164, 1165, 1166, 1167, 1167, 1168, 1169, 1169, 1170, 1171, 1172, 1173, 1174, 1175, 1176, 1177, 1178, 1179, 1179, 1180, 1181, 1182, 1183, 1184, 1185, 1186, 1187, 1187, 1188, 1189, 1189, 1190, 1191, 1192, 1193, 1194, 1195, 1195, 1196, 1197, 1197, 1198, 1199, 1199, 1200, 1201, 1202, 1203, 1204, 1205, 1206, 1207, 1207, 1208, 1209, 1209, 1210, 1211, 1212, 1213, 1214, 1215, 1216, 1217, 1217, 1218, 1219, 1219, 1220, 1221, 1222, 1223, 1224, 1225, 1226, 1227, 1227, 1228, 1229, 1229, 1230, 1231, 1232, 1233, 1234, 1235, 1236, 1237, 1238, 1239, 1239, 1240, 1241, 1242, 1243, 1244, 1245, 1246, 1247, 1247, 1248, 1249, 1249, 1250, 1251, 1252, 1253, 1254, 1255, 1256, 1257, 1257, 1258, 1259, 1259, 1260, 1261, 1262, 1263, 1264, 1265, 1266, 1267, 1267, 1268, 1269, 1269, 1270, 1271, 1272, 1273, 1274, 1275, 1276, 1277, 1277, 1278, 1279, 1279, 1280, 1281, 1282, 1283, 1284, 1285, 1286, 1287, 1287, 1288, 1289, 1289, 1290, 1291, 1292, 1293, 1294, 1295, 1296, 1296, 1297, 1298, 1298, 1299, 1300, 1301, 1302, 1303, 1304, 1305, 1306, 1307, 130$

(b) Describe (with proof) an augmentation from which CLOSEST() can be computed in constant time.

Hint: Use an ordered triple that captures the difference between the examples you described above. $\text{CLOSEST} = \min + \text{diff. bit} + 2 \text{ bits}$.

Solution: In the given solution, the *minimum* value in the right subtree changed without changing any augmentation values. Generally, the closest pair can be in the left subtree, be in the right subtree, or contain the root value. In order to identify the last case, we want to know the minimum and maximum values of each subtree. The triple $(\min, \text{CLOSEST}, \max)$ is a valid augmentation that can be maintained via:

$$\begin{aligned} \min(T) &= \min(\min(T.\text{LEFT}), T.\text{ITEM}) \\ \max(T) &= \max(\max(T.\text{RIGHT}), T.\text{ITEM}) \\ \text{CLOSEST}(T) &= \min \begin{cases} \text{CLOSEST}(T.\text{LEFT}) \\ \text{CLOSEST}(T.\text{RIGHT}) \\ |T.\text{ITEM} - \max(T.\text{LEFT})| \\ |T.\text{ITEM} - \min(T.\text{RIGHT})| \end{cases} \end{aligned}$$

Problem 2. [10 points] Pidgey (1 part)

Help Pidgey design an algorithm that takes as input an array A of n integers, and outputs a pair of indices $i \neq j$ (if they exist) such that $A[i] - A[j]$ is a multiple of $\lceil \log n \rceil$. If there are many such pairs, you may output any one of them. Briefly justify correctness. Analyze the runtime, including proving matching upper and lower asymptotic bounds. Classify the runtime as worst-case, expected, and/or amortized.

Choose your own adventure:

■ For full credit, your algorithm must run in $O(\log n)$ time.

□ For up to 5 points, your algorithm must run in $O(n)$ time.

PIGEONHOLE + DAA

Solution: Create a DAA that stores i at location $\text{rem}(A[i], \lceil \log n \rceil)$. Output the colliding indices from the first collision. $A[i] \rightarrow i$ using the hash function $h(k) = \text{rem}(k, \lceil \log n \rceil)$. Output the colliding indices from the first collision.

Correctness follows from the definition of modular equivalence and the fact that $\text{rem}(a, c) = \text{rem}(b, c) \iff a \equiv b \pmod{c}$.

Every insertion takes worst-case $O(1)$ time because there is at most one collision by construction. By the pigeonhole principle, the first collision is found within the first $\lceil \log n \rceil + 1$ elements, so runtime is worst case $O(\log n)$. This bound is tight, witnessed by the input array $[1, 2, \dots, n]$, for which the collision between 1 and $\lceil \log n \rceil + 1$ is found after $\log n$ insertions.

(b) [16 points] Suppose instead that Riolu knows that there exists some $\sigma < n^{1210}$ such that her friend can be partitioned into n σ -pairs with one left over. Design a worst-case $O(n)$ algorithm to find the unpaired Pokémon's aura, given as input an array A of the $2n + 1$ Pokémon's auras. Note that σ is unknown. Prove your algorithm correct, and analyze its runtime. You may assume a correct solution to part (a), even if you did not solve it yourself.

Choose your own adventure:

RAIDX SORT

■ For full credit, your algorithm must output the unpaired aura.

□ For up to 8 points, your algorithm need only decide whether or not the largest aura is unpaired.

Solution: We reduce to part (a).

1. If any aura is larger than n^{1210} , return it.
2. Radix sort A .
3. Check if for every index $i \in [1, n]$, $A[i] + A[-i]$ is the same (and $< n^{1210}$); if so, return $A[0]$.
4. Repeat the previous step on the reverse of A .
5. Return the result of part (a), with $\sigma = A[0] + A[-1]$.

Proof of correctness: Auras are positive, so every paired aura must be less than $\sigma < n^{1210}$. Hence step 1 is correct or a no-op. We now condition on whether the unpaired aura is the smallest, largest, or neither.

- The smallest aura is unpaired. In this case, because A is sorted, $A[i]$ and $A[-i]$ must form a σ -pair for every i . This is detected by step 3.
- The largest aura is unpaired. This case is symmetric to the previous and is detected by step 4.
- The smallest and largest auras are both paired. In this case, they must be paired with each other, so we know σ and can reduce to part (a).

Runtime analysis: Steps 1, 3, 4 are linear scans. If step 1 does not return, then auras are polynomial, and step 2 takes linear time. Step 5 takes linear time by assumption. Total runtime is linear.

Problem 4. [20 points] Happy Go Lucky (1 part)

Blissey's Pokémon friends are at integer locations along a line. Blissey is keeping them happy by placing Lucky Eggs. Each Lucky Egg e is represented as a pair of integers (a, b) with $a \leq b$, and it makes all Pokémon at locations x with $a \leq x \leq b$ happy. Blissey's friends may accidentally break the Lucky Eggs, which then lose their effect.

Blissey would like to know who is happy. Design a data structure that supports the following operations:

- PLACE(e): Add a Lucky Egg e to the data structure
- BREAK(e): Delete a Lucky Egg e from the data structure
- HAPPY(x): Output TRUE iff the Pokémon at location x is currently happy

AVL BST, AUGM.

Your data structure must use $O(n)$ space, and all operations should have worst-case $O(\log n)$ runtime, where n is the number of Lucky Eggs currently in existence. You need not prove runtime or correctness.

Choose your own adventure:

■ For full credit, implement the operations as given above.

□ For up to 10 points, assume that there exists d such that all Lucky Eggs have the form $(a, a + d)$. Your runtime and space requirements must not depend on d .

Solution: Keep all Lucky Eggs in an AVL BST keyed lexicographically by either (a, b) or $(a, -b)$; that is, they are primarily keyed by a , and if multiple eggs have the same a , ties are broken by some ordering of b . Augment each node with the maximum b in each subtree, computed by $\text{node}.maxb} = \max(\text{node}.item.b, \text{node}.left.maxb, \text{node}.right.maxb}$. PLACE and BREAK are standard add and delete operations.

HAPPY(x): Recursively search from the root. At each node, if $node$ is null, output FALSE. Otherwise:

1. If $node.item.a > x$, recurse on $node.left$ and output its value. (The only eggs that can possibly have $a > x$ are all in the left subtree.)
2. Else, $node.item.a \leq x$:
 - (a) If $node.item.b \geq x$, output TRUE.
 - (b) If $node.item.b \geq x$, output TRUE. (All eggs in the left subtree have $a \leq x$, and this condition means at least one of them have $b \geq x$.)
 - (c) Otherwise, recurse on $node.right$ and output its value. (The only eggs that can possibly have $b \geq x$ are all in the right subtree.)

- (b) Suppose instead that every Pokémon will train for more than one hour, i.e. $b_i - a_i > 1$. (You may no longer assume that a_i is an integer.) Give an $O(n)$ time algorithm to solve the problem. You may assume a correct solution to part (a), even if you didn't manage to solve it yourself. $I_1 = [1, 2, 4, 5] \rightarrow c_1 = 2 \rightarrow [1, 2, 2, 5, 4, 5]$

Solution: Reduction to (a) use counting sort on right & left respectively, then merge unions

1. For every i , there is an integer c_i such that $a_i < c_i < b_i$.
2. Compute $\bigcup\{[n - c_i, n - a_i]\}$ using part (a).
3. Replace each interval $[x, y]$ in the union resulting from step 2 with $[n - y, n - x]$.
4. Compute $\bigcup\{[c_i, b_i]\}$ using part (a).
5. Merge the two lists into a single sorted list S .
6. Iterate over S , merging each interval with the previous interval if they overlap.
7. Return S .

Proof of Correctness: We are computing the function $f(\mathcal{I}) = \mathcal{J}$, where \mathcal{I} and \mathcal{J} are finite sets of intervals with the same union, and \mathcal{J} is minimal. Because c_i is defined as an integer, we know that $n - c_i$ is also an integer, and step 2 correctly finds the union of $[n - c_i, n - a_i]$ from the correctness proof of part (a).

To show that steps 2 and 3 correctly output the union of $[a_i, c_i]$, we observe that both steps apply the transformation $g(\mathcal{I}) = \{[n - x : x \in I] : I \in \mathcal{I}\}$, which is self-inverting and commutes with f . (1)

Step 4 is correct because c_i is an integer and thus correctness from part (a) applies. The remaining steps are correct per the proof from (a).

Aside: This level of detail is not needed, but for a full proof of (1), it suffices to show that $g(f(g(\{[p, q], [r, s]\}))) = f(\{[p, q], [r, s]\})$ for any p, q, r, s , as this can then be applied inductively.

WLOG, assume that $p < r$. Then, we have two cases:

- If $q < r$ then the two intervals are disjoint. Similarly, $n - r < n - q$ and $n - r < n - p$, so $[n - q, n - p], [n - s, n - r]$ are also disjoint. On disjoint

• $\text{find_best}(m)$: Find the highest rated player whose cost is between $m/2$ and m (inclusive), or report that no such player exists

All of insert , delete , and find_best must still have $O(\log n)$ runtime.

Hint: There are two very similar solutions to part (a), but one of them can be adapted much more naturally to solve part (b) as well. ***KEY BY COST**

Solution: The first solution above can be adapted; the second should not. We keep three augmentations: the **best player** in each subtree, the **min key** in each subtree, and the **max key** in each subtree. Both of insert and delete are unchanged. We adapt find_best : **AUGMENTATION/RANGE QUERY**

1. If the tree is empty, the min key is greater than m , or the max key is less than $m/2$, then output ⊥
2. If the min and max keys are both in the interval $[m/2, m]$, then output the best augmentation
3. If the root cost is greater than m , then recurse on the left
4. If the root cost is less than $m/2$, then recurse on the right
5. Otherwise, recurse on both children, and output whichever of these two calls or the root player has the best rating.

Correctness: Step 1 outputs ⊥ if there are no costs in the range $[m/2, m]$. Step 2 executes iff all costs are in the range $[m/2, m]$, in which case they are all eligible. In this case, the best augmentation is the correct output by definition. Steps 3 and 4 execute iff the root is not eligible, in which case one subtree is also ineligible, and recursion on the other subtree gives the correct output. Step 5 is a catch-all and is correct by casework on where the correct output is.

Runtime: The steps are mutually exclusive. Steps 1-2 return immediately, and Steps 3-4 make a single recursive call on a subtree. By the guards on the previous steps, Step 5 executes iff the root cost and at most one extremum are in the range $[m/2, m]$. This means at least one of $m/2$ or m is strictly between the min and max keys. Suppose Step 5 executes on two subtrees T_1 and T_2 , neither of which is an ancestor of the other. By the BST property, their least common ancestor has a key between them, so WLOG we have $\min T_1 < m/2 \leq T_1.\text{root} \leq \max T_1 \leq \min T_2 \leq T_2.\text{root} \leq m < \max T_2$. Therefore the right child of T_1 and the left child of T_2 both return immediately. This means there are at most four recursive calls at every level of the recursion tree, giving $O(\log n)$ runtime.

Problem 3. Long Jump **2 AVL TREES**

Henry is unsatisfied with the current long jump rules and wants to give athletes prizes according to his own set of rules. He arrives at the competition with a bag full of Snickers bars to give the athletes. Whenever he's feeling generous, he selects a minimum distance threshold t . An athlete is eligible to receive this prize if their most recent jump was at least distance t . He gives the prize to the eligible athlete who jumped most recently, or to Srinivasa if no athlete is eligible. Athletes have no limits to their number of attempts.

Henry needs a data structure that can support the following operations:

- JUMP(a, d): Record that an athlete with the name a just achieved a long jump distance of d
- PRIZE(t): Output the name of the prize winner, given distance threshold t

Describe a data structure that uses $O(n)$ space and implements both of the above operations in $O(\log n)$ time, where n is the total number of athletes who have participated so far. Briefly justify correctness and analyze runtime. You do not need to prove space complexity or runtime lower bounds. You need not analyze data structures or algorithms presented in class, but you must describe and analyze any modifications that you make. Assume that the JUMP operation is executed at the time of the recorded jump, but Henry cannot tell what this time is.

Solution: We store RECORDS as tuples (a, d) , representing that the athlete with name a achieved a jump of distance d .

We keep two AVL trees:

- In DISTANCES, we store a SEQUENCE of $n+1$ RECORDS, in order of time. Initially, the only RECORD is (Srinivasa, 0). We add the augmentation MAX(T), which is the maximum RECORD in T 's subtree, comparing by distance. We compute this by comparing the root RECORD's size with the augmentations of both children.
- In ATHLETES, we store a set mapping each athlete's name a to the node in DISTANCES containing a record (a, d) .

Operations are implemented as follows:

```

if MAX(A) >= t:
  go right
  prize(c) = (b, t)
  M = 7
  M = 7
  (a, 7) ←
  (b, 7) ←
  Remove a from ATHLETES, and remove the removed node from DISTANCES
  Add (a, d) to the end of DISTANCES
  Add (a, d) to the end of DISTANCES and in ATHLETES map a to the new node
  PRIZE(t): (defined recursively and wrapped, starting with DISTANCES)
    - If T's right child augmentation is at least t, recurse on right
    - Else if T's root RECORD has distance at least t, output it
    - Else recurse on left
  
```

Runtime analysis: We store two AVL trees of size $n+1$, for a total of $O(n)$ space. Augmentation runtime is immediate from definition. JUMPing performs at most four AVL insertions / deletions, each of which takes $O(\log n)$ time. PRIZE follows at most a path from the root of DISTANCES to a leaf, which takes $O(\log n)$ time.

Proof of correctness: DISTANCES stores every Athlete's most recently jump attempt, plus a dummy for Srinivasa. We observe that WLOG these are the *only* jumps that have been attempted, as no others affect prize distribution.

Correctness of MAX augmentation is immediate from definition.

When JUMP is called, the new RECORD is necessarily the most recent. It should therefore be added at the end of DISTANCES and replace any RECORD from the same Athlete.

Correctness of PRIZE follows from induction and the traversal order of DISTANCES, noting that c is always "eligible". Assume that T contains an eligible Athlete and that PRIZE executes correctly on smaller trees. We condition on where the prize winner is.

- If the athlete is in the right subtree, then they are also the last eligible Athlete in the right subtree, so are found by the first recursive call.
- If they are at the root, then all Athletes later in traversal order, i.e. in the right subtree, are ineligible. The maximum size appearing in the right subtree is less than t , so the first recursive call isn't executed, and the root is returned.
- If they are in the left subtree, then they are the last eligible record in the left subtree. All Athletes later in traversal order, i.e. the root and everyone in the right subtree, are ineligible. The else clause is executed, and recursion on the left finds the correct prize winner.

Problem 1. Counting Sheep

Mareep is implementing a counter that stores a natural number k (initially 0) in base- n using an array A of length n . The counter has two operations:

- increment(): increments k by one
- get(): outputs and resets k

Describe how to implement both operations, and prove that they take amortized constant time.

Solution: For increment, we start by incrementing $A[0]$. Each time $A[0]$ overflows, we reset it and increment $A[1] + 1$. If $A[0] - 1$ overflows, double n , reallocate A , and set $A[0] = 1$.

For get, output A , reset $n = 2$, and reallocate A .

To analyze runtime, we want to figure out a good potential function. The first observation is that increment becomes expensive when $n - 1$ appears many times in A , but it also gets rid of those appearances. Therefore these values should contribute to the potential. The second observation is that get must always use linear work, but this work cannot exceed the number of earlier increments.

We define $\phi(A) = (m + k)$, where k is a constant and m is the number of indices i such that $A[i] = n - 1$. (Note that $\phi(A) = c(2m + n)$ will also work.)

The work done for increment is $O(m' + 1)$, where m' is the length of the longest prefix of A whose values are all $n - 1$. The potential increases by $O(1)$ and decreases by $O(m')$. By choosing c to be sufficiently large, the potential decrease can be made to dominate the work done, so the amortized cost is $O(1)$.

The work done for get is $\Theta(n)$, and the potential decrease is $2^{O(n)}$, so there is no amortized cost.

(b) Assuming unique node values, give a worst-case $\Theta(n^2)$ algorithm to determine Bulbasaur's PreOrder traversal given Ivysaur's InOrder and Venusaur's PostOrder traversals of a tree T . Prove your algorithm correct and observe that it proves that the PreOrder traversal is uniquely determined by the InOrder and PostOrder traversals.

PreOrder ← Post, In

Solution: We know that the root is always the last item in the PostOrder traversal. Next, we know that the InOrder traversal has the structure **Left_SubTree, Root, Right_SubTree**, so we can search the InOrder for the root. We know that the left subtree contains all nodes to the left of the root in the InOrder traversal and the right subtree contains all nodes to its right. Similarly, all left nodes appear before any of the right nodes in the PostOrder. So, once we find the index of the root in the PostOrder to have that split into left and right subtrees.

Once we have the traversals split into the two subtrees, we recurse and run the same procedure again, first with the left subtree, and then with the right subtree.

We then form the PreOrder traversal by placing first the root, then the PreOrder of the left subtree and finally the PreOrder of the right subtree.

Correctness: We are correctly obtaining the root node in each recursive call as the root node will always be the last item in the PostOrder traversal. We are also splitting at the correct node to determine left and right subtrees since each node value is unique.

Runtime: The worst case runtime is $O(n^2)$ since each node will be the root exactly once, and each time a node is the root we have to search for it in the InOrder, which takes at most $O(n)$ time (the length of the list of nodes we have to search through).

We are able to show this bound is tight by considering the case where the InOrder and PostOrder are the same, making a tree that is a chain of left children. This exhibits worst case behavior since each time we have to search for the root we must iterate through all of the remaining nodes in the traversal, giving a runtime of $\sum_{i=0}^{n-1} (n - i) = \Theta(n^2)$.

Problem 2. Rocket Recruits **BS, SORTING, PAIRS**

Giovanni has recruited n new Grunts for Team Rocket. For each $i \in \{0, \dots, n-1\}$, Grunt i has k_i Pokémons. All of the k_i are positive and distinct.

Giovanni wants to pair them into teams with exactly p Pokémons between them. A pair of Grunts (x, y) , with $x < y$, such that $k_x + k_y = p$ is called a *battle pair*.

Giovanni wants to find two quantities: the number B of battle pairs he can make, and the battle pair (x^*, y^*) for which $x^* + y^*$ is minimal.¹

(a) Describe an algorithm to find B and (x^*, y^*) with a worst-case runtime of $O(n \log n)$. Prove your algorithm correct, and analyze its runtime.

Solution: We present two solutions. The first is simpler, and uses binary search to find a valid y that pairs with each x after sorting the list of Pokémons quantities. The second solution improves on the binary search step, replacing it with a two-finger algorithm to find all battle pairs in linear time. While the second algorithm is more efficient in practice, both have the same asymptotic runtime $\Theta(n \log n)$ due to the sorting step being the bottleneck.

For all parts of the question, we denote "record" (x^*, y) as incrementing the running count B , and then updating the running optimum $(x^*, y^*) \leftarrow (x, y)$ if $x < x^*$.

Algorithm 1 (Sort + Binary Search): **Sort, ←, k ← k + 1**

1. Initialize the running count $B = 0$ and running optimum $(x^*, y^*) = (n, n)$.
2. Keep track of the original indices of each element in $\{k_i\}$ before sorting. This can be done by creating a new sequence $\{c_i\}$, and setting $c_i = (k_i, i)$ for each $0 \leq i \leq n - 1$.
3. Sort $\{c_i\}$ by their Pokémon quantities (the first element k_i of each tuple), using a $\Theta(n \log n)$ sorting algorithm such as merge sort.
4. For each $x \in [0, n - 1]$, binary search on $\{c_i\}$ to find a tuple (k_y, y) with $k_y = p - k_x$. If there exists such a y , and $x < y$, record (x, y) as a battle pair.
5. Return S and (x^*, y^*) .

Proof of correctness. It is easy to see that if we record exactly the set of all battle pairs once each, the correctness of S and (x^*, y^*) follows naturally. We show that this is indeed what we record in Step 4.

Any battle pair (x, y) must satisfy $k_y = p - k_x$, so when iterating through x , the binary search will correctly find the unique y with Pokémon quantity $p - k_x$, and then record (x, y) . Likewise, we can show that any pair that's not a battle pair will not be found and recorded. Thus, we record all and only each battle pair exactly once.

Problem 4. Find the Missing No.

Given an array A of strictly increasing integers of length n , and a number s , your task is to design an algorithm to find the smallest integer larger than or equal to s not in A .

(a) Design an algorithm to find the missing integer in time $O(\log n)$. You must describe your algorithm in English.

Solution: 1 (Modification of binary search): We modify the comparison step. When A is not empty, we find the index of the middle element i and check if $A[i] < s$. If so, we want to recurse on the right side of the list. If not, we check if there are consecutive numbers from s to $A[i]$ by finding the difference between $A[i]$ and s . Let's call this k . Next, we check if $A[i] - k = s$. If so, we know that there are consecutive numbers from s to $A[i]$, so we recurse on the right side of the list, setting s to $A[i] + 1$. If not, we recurse on the left side of A .

(b) Prove your algorithm correct.

Solution: 1: We proceed by strong induction on n . Our hypothesis is that our solution finds the smallest integer larger than or equal to s not in A correctly. The base case is when $n = 0$. In this case, we have an empty array and therefore, we return s , which is the smallest integer larger than or equal to s not in A . For the inductive step, let's assume that our hypothesis holds for all arrays of size $\text{len}(A) < n$ and let's show that for arrays of size n the algorithm is still correct. There are three cases.

1. The middle element is less than s : Since we have a strictly increasing array, any missing element in the list on the left side is smaller than s , which means we can disregard the left side. According to our algorithm, we'll recurse on the right side of the array, and as the right side of the array will have size less than n , our inductive hypothesis ensures that the algorithm will run correctly on this side.

2. The middle element is larger than s and there is a consecutive sequence of integers from s to the middle element: This means that there is a missing number on the left side of the list that is larger than s . Thus, we should search in the left side, which is what our algorithm does. Since the left side of the array has size less than n , our algorithm will reach the correct output based on our inductive hypothesis.

3. The middle element is larger than s but there is not a consecutive sequence from s to the middle element: This means that there is a missing number on the left side of the list that is larger than s . Thus, we should search in the left side, which is what our algorithm does. Since the left side of the array has size less than n , our algorithm will reach the correct output based on our inductive hypothesis.

5. We saw how to create an iterator for an AVL tree in which the NEXT operation takes amortized constant time but worst-case logarith