

Write your answers in the separate answer booklet.

You have 180 minutes (after you get the answer booklet) to answer seven questions.

1. For each statement below, there are two boxes in the answer booklet labeled “Yes” and “No”. Check “Yes” if the statement is *always* true and “No” otherwise, and give a *brief* (at most one short sentence) explanation of your answer. **Assume $P \neq NP$.** If there is any other ambiguity or uncertainty about an answer, check “No”. For example:

- $x + y = 5$

 Yes

 No

Suppose $x = 3$ and $y = 4$.

- 3SAT can be solved in polynomial time.

 Yes

 No

3SAT is NP-hard.

- If $P = NP$, then Emily can beat Super Mario Brothers in four minutes and 53 seconds.

 Yes

 No

The hypothesis is false, so the implication is true.

Read each statement *very* carefully; some of these are deliberately subtle!

- (a) Which of the following statements are true for *every* language $L \subseteq \{0, 1\}^*$?

- $(L^*)^*$ is infinite.
- If L is decidable then its complement \bar{L} is undecidable.
- $\{\langle M \rangle \mid M \text{ accepts } L\}$ is undecidable.
- Either L is finite or L is NP-hard. (The *or* is inclusive.)
- Either L has an infinite fooling set or $L \in P$. (The *or* is inclusive.)

- (b) Consider the following pair of languages:

- TREE = $\{G \mid G \text{ is a connected undirected graph with no cycles}\}$
- HAMPATH = $\{G \mid G \text{ is an undirected graph that contains a Hamiltonian path}\}$

(For concreteness, assume that in both of these languages, graphs are represented by their adjacency matrices.) Which of the following statements are true, assuming $P \neq NP$?

- TREE is NP-hard.
- $\text{TREE} \cap \text{HAMPATH}$ is NP-hard.
- $\text{TREE} \cup \text{HAMPATH}$ is NP-hard.
- HAMPATH is undecidable.
- A reduction from TREE to HAMPATH would imply $P = NP$.

Problems 2–7 appear on the next three pages.

2. Suppose we want to prove that the following language is undecidable.

$$\text{CHALMERS} := \{ \langle M \rangle \mid M \text{ accepts both STEAMED and HAMS} \}$$

- (a) Professor Skinner suggests a reduction from the standard halting language

$$\text{HALT} := \{ \langle \langle M \rangle, w \rangle \mid M \text{ halts on inputs } w \}.$$

Specifically, Professor Skinner claims that if there is a Turing machine SUPERNINTENDO that decides the language CHALMERS, then the following algorithm decides HALT.

```

DECIDEHALT( $\langle \langle M \rangle, w \rangle$ ):
  Encode the following Turing machine:
  AURORABOREALIS( $x$ ):
    if  $x = \text{STEAMED}$  or  $x = \text{HAMS}$  or  $x = \text{UTICA}$ 
      run  $M$  on input  $w$ 
      return FALSE  $\langle\langle \text{reject } x \rangle\rangle$ 
    else
      return TRUE  $\langle\langle \text{accept } x \rangle\rangle$ 
  return SUPERNINTENDO( $\langle \text{AURORABOREALIS} \rangle$ )

```


Show that Professor Skinner's reduction is incorrect by describing a Turing machine (or an algorithm) M and string w such that DECIDEHALT($\langle M \rangle, w$) either never returns or returns the wrong answer.

- (b) Correctly **prove** that CHALMERS is undecidable.
3. Let T be a *full* binary tree, meaning that every node has either two children or no children.
- Recall that the *height* of a vertex v in T is the length of the longest path in T from v down to a leaf. In particular, every leaf of T has height zero.
 - A vertex v is *AVL-balanced* if v is a leaf, or if the heights of v 's children differ by at most 1. (You might recall from CS 225 that an *AVL-tree* is a binary search tree in which every vertex is AVL-balanced.)

Describe and analyze an algorithm to compute the number of AVL-balanced vertices in T .

Problems 4–7 appear on the next two pages.

4. Submit a solution to *exactly one* of the following problems.
- (a) A *Hamiltonian bicycle* in a graph G is a pair of simple cycles in G , with identical lengths, such that every vertex of G lies on exactly one of the two cycles.



A Hamiltonian bicycle in the 6×6 grid graph.

Prove that determining whether a given undirected graph G has a Hamiltonian bicycle is NP-complete.

- (b) A *clique-partition* of a graph $G = (V, E)$ is a partition of the vertices V into disjoint subsets $V_1 \cup V_2 \cup \dots \cup V_k$, such that for each index i , every pair of vertices in subset V_i is connected by an edge in G . The *size* of a clique partition is the number of subsets V_i .

Prove that determining whether a given undirected graph G has a clique-partition of size 3 is NP-complete.

In fact, both of these problems are NP-complete, but we only want a proof for one of them. Don't forget to tell us which problem you've chosen!

5. Suppose we are given a directed graph $G = (V, E)$, where every edge $e \in E$ has a *positive* weight $w(e)$, along with two vertices s and t .
- (a) Suppose each *vertex* of G is colored either orange, green, or purple. Describe and analyze an algorithm to find the shortest walk from s to t in G that never visits two consecutive *vertices* with the same color.
- (b) Now suppose each *edge* of G is colored either orange, green, or purple. Describe and analyze an algorithm to find the shortest walk from s to t in G that never traverses two consecutive *edges* with the same color.

Problems 5 and 6 appear on the next page.

6. *Vankin's Mile* is a solitaire game played on an $n \times n$ square grid. The player starts by placing a token on any square of the grid. Then on each turn, the player moves the token either one square to the right or one square down. The game ends when player moves the token off the edge of the board. Each square of the grid has a numerical value, which could be positive, negative, or zero. The player starts with a score of zero; whenever the token lands on a square, the player adds its value to his score. The object of the game is to score as many points as possible.

For example, given the grid below, we can score $7 - 2 + 3 + 5 + 6 - 4 + 8 + 0 = 23$ points by following the path on the left, or we can score $8 - 4 + 1 + 5 + 1 - 4 + 8 = 15$ points by following the path on the right.

-1	7	-2	10	-5
8	-4	3	-6	0
5	1	5	6	-5
-7	-4	1	-4	8
7	1	-9	4	0

-1	7	-2	10	-5
8	-4	3	-6	0
5	1	5	6	-5
-7	-4	1	-4	8
7	1	-9	4	0

- (a) Describe and analyze an efficient algorithm to compute the maximum possible score for a game of *Vankin's Mile*, given the $n \times n$ array of values as input.
- (b) A variant called *Vankin's Niknav* adds an additional constraint to *Vankin's Mile*: *The sequence of values that the token touches must be a **palindrome***. Thus, the example path on the right is valid, but the example path on the left is not. Describe and analyze an efficient algorithm to compute the maximum possible score for an instance of *Vankin's Niknav*, given the $n \times n$ array of values as input.
7. (a) Let L_a denote the set of all strings $w \in \{0, 1, 2\}^*$ such that $\#(1, w) + 2 \cdot \#(2, w)$ is divisible by 3. For example, L_a contains the strings **0012** and **20210202** and the empty string ϵ , but L_a does not include the strings **121** or **0122210**.
Describe a DFA or NFA that accepts L_a . (You do not need to prove that your answer is correct.)
- (b) Let L_b denote the set of all strings $w \in \{0, 1, 2\}^*$ such that no two symbols appear the same number of times, or in other words, the integers $\#(0, w)$ and $\#(1, w)$ and $\#(2, w)$ are all different. For example, L_b contains the strings **110212** and **20220**, but L_b does not include the strings **01212** or **2120210** or the empty string ϵ .
Prove that L_b is not a regular language.