1. *1d-Clickomania.* Consider the following game on a sequence of black-white tiles: In each step, you are allowed to choose any maximal run of *at least two* single-colored tiles, and remove them from the sequence. The rest of the tiles on the two sides will rejoin into a single sequence. If there is a way to remove all the tiles from the sequence, then we say that the sequence is *winnable*. For example,

 $\blacksquare \blacksquare \Box \Box \Box \blacksquare \blacksquare \Box \to \blacksquare \blacksquare \Box \Box \Box \Box \to \blacksquare \blacksquare \to \varepsilon$

is winnable, while $\blacksquare \square \square \square \square \square \blacksquare$ is not.

Prove that the following language is not regular:

$$L = \left\{ w \in \{\blacksquare, \Box\}^* : w \text{ is winnable} \right\}.$$

Solution: We prove that *L* is not regular by describing a fooling set for *L* of infinite size. Let

$$F := \left\{ (\blacksquare \Box)^i : \text{for all integer } i \right\}.$$

Take two arbitrary elements in *F*, say without loss of generality $x := (\blacksquare \square)^i$ and $y := (\blacksquare \square)^j$ for i < j. Let $z := (\square \blacksquare)^i$.

- On one hand, we have $xz = (\square \square)^i (\square \square)^i = \varepsilon$ based on the rule. This shows that xz is winnable.
- On the other hand, we have yz = (■□)^j(□■)ⁱ. As we attempt to simplify the tile sequence, at any moment there is only one monochromatic run with exactly two tiles. This implies the simplification process is unique and yz = (■□)^{j−i}, which is not winnable because there are no further tiles can be removed.

This implies that *F* is a fooling set for *L* of infinite size.

Rubric: Standard 5-point grading scaled to 10 points. 6 points if the fooling set is correct.

2. *Fault-tolerant computing.* Imagine you are given a good-old Turing machine with a single-tape and a single-head, but there is one catch: exactly one of the cells on the tape is *faulty*. Any symbol written on that cell might turn into other symbols unexpectedly, while the reading head is away. Worst of all, there are no ways to tell if the cell is faulty by examination, even after the cell has turned a symbol into another. (We assume that the faulty cell is not within the input.)

Prove that the faulty Turing machine can still simulate a standard Turing machine.

Solution: We describe how to emulate a standard Turing machine using a faulty Turing machine. The idea is to duplicate each symbol three times and read three consecutive cells at a time; now a majority vote always returns the correct symbol.

More precisely, we encode the symbol of the standard TM M by a duplicating the symbol three times in the faulty TM M'. ($\langle 0 \rangle = 000$, $\langle 1 \rangle = 111$, and $\langle \Box \rangle = \Box \Box \Box$). The input of M' is also encoded in the same fashion from the input of M. For each transition of M (that is, read a symbol, write a symbol, and move the head to an adjacent cell), the new machine M' read the corresponding three symbols, and treat the majority of the symbols as the symbol read. Then M' writes down three identical copies of the same symbol as required, and move the head three steps to the left or right.

To see the correctness of the emulation, because there is at most one faulty cell across the tape, the majority vote within the three cells will always return the correct symbol and thus M' correctly simulates the computation of M.

Rubric: Standard 5-point grading scaled to 10 points. Maximum 2 points if the majority vote idea is missing.

3. *Regular language equivalence.* Let L_0 be an arbitrary regular language over alphabet Σ . Let D_0 be a DFA deciding L_0 , given by some encoding $\langle D_0 \rangle$. Prove that the following language can be decidable by a Turing machine in polynomial time:

$$L_1 := \left\{ \langle D \rangle : L(D) = L(D_0) \right\},\$$

where L(D) denote the language recognized by DFA D.

Solution: To show that L_1 can be solved in polynomial time, we prove that one can test whether the exclusive-or between L(D) and $L(D_0)$ is empty on a given encoding of some DFA $\langle D \rangle$.

Testing if $\langle D \rangle$ encodes a valid DFA can be done in linear time. We construct DFA D' recognizing the language $L(D) \oplus L(D_0)$ — the exclusive-or between L(D) and $L(D_0)$ — using product construction.

- For each state q in D and q' in D_0 , construct a state (q,q') in D'.
- The alphabet of D' is the union of the alphabet sets of D and D_0 .
- For each transition $q \xrightarrow{a} r$ in D and $q' \xrightarrow{a} r'$ in D_0 for some identical symbol a, add a transition $(q, q') \xrightarrow{a} (r, r')$ in D'.
- The starting state in D' is (s, s') where s/s' is the starting state of D/D_0 , respectively.
- The accepting states of D' are the *exclusive-or* of those in D and D_0 ; in notation,

 $\left\{ (q,q'): \begin{array}{c} q \text{ is accepting in } D \text{ but } q' \text{ is not accepting in } D_0 \\ \text{or } q \text{ is not accepting in } D \text{ but } q' \text{ is accepting in } D_0 \end{array} \right\}.$

To decide if L(D') is empty, we can perform a graph-traversal from the starting state of D' and see if any accepting state can be reached. We accept $\langle D \rangle$ if the language L(D') is empty. Overall the running time is proportional to the size of D' which is polynomial in $|\langle D \rangle|$ because for a fixed machine D_0 , DFA D' has size linear in the size of D.

Rubric: Standard 5-point grading scaled to 10 points. Maximum 6 points if the construction of exclusive-or DFA is missing. Maximum 4 points if the algorithm is testing whether the *two DFAs* are equivalent. 4. *Inception.* Let M_1 be a TM deciding the language L_1 from the previous problem (given as some encoding $\langle M_1 \rangle$). Prove that the following language is undecidable:

$$L_2 := \left\{ \langle M \rangle : L(M) = L(M_1) \right\},\$$

where L(M) denote the language accepted by TM M.

Solution: We prove that *L*₂ is undecidable by reduction from the HALTING problem:

HALTING

Input: An encoding of a Turing machine (M), and a string w. *Output:* Does M run on input w halt?

Assume for contradiction that L_2 can be decided by some TM M_2 . We now construction another Turing machine *H* that decides HALTING, thus deriving a contradiction.



 $\frac{M_1'(x):}{\text{run } M \text{ on } w}$ return $M_1(x)$

To prove that *H* decides the HALTING problem,

- If $(\langle M \rangle, w)$ is a yes-instance of HALTING, the constructed machine M'_1 always terminates when running M(w) and then mimics the behavior of M_1 on the same input x, which implies that $L(M'_1) = L(M_1)$. Therefore M_2 on input $\langle M'_1 \rangle$ will answer *yes*, implies that H accepts $(\langle M \rangle, w)$.
- If $(\langle M \rangle, w)$ is a no-instance of HALTING, inside machine M'_1 the simultation M(w) runs forever, which implies that $L(M'_1) = \emptyset \neq L(M_1)$ because $L(M_1)$ is not empty $(\langle D_0 \rangle$ must be in $L(M_1)$ from Q3). Therefore M_2 on input $\langle M'_1 \rangle$ will answer *no*, implies that H rejects $(\langle M \rangle, w)$.

Notice that we never really execute M'_1 as a program but only as input to M_2 , so H always terminates.

Rubric: Standard 5-point grading scaled to 10 points. Maximum 4 points if the reduction is incorrect or works for the wrong problem.

5. Consider the following problem:

ZeroSumSet
• Input: A set X of n integers in $[-N N] := \{-N,, 0, 1,, N\}$.
• Output: Is there a nonempty subset S of X such that numbers in S sum
up to 0?

Prove that ZEROSUMSET is NP-complete.

Solution: We prove that ZEROSUMSET is NP-hard by reduction from the NP-hard problem SUBSETSUM. Given a set of positive integers Y in [1..N] and a target t, the SUBSETSUM problem asks if there is a subset R of Y such that the numbers in R sum up to t.

We first construct an instance *X* of ZEROSUMSET as follows. We set $X := Y \cup \{-t\}$, that is, adding -t as another integer into *Y* to form *X*. The construction can clearly be done in polynomial time.

To prove that the reduction is correct:

- If (Y, t) is a *yes*-instance for SUBSETSUM, we know there is a subset R of Y such that the numbers in R sum up to t. Now the subset $S := R \cup \{-t\}$ of X is nonempty and must sum up to zero, and thus X is a yes-instance for ZEROSUMSET.
- If (*Y*, *t*) is a *no*-instance for SUBSETSUM, assume for contradiction that there is a subset *S* of *X* such that the numbers in *S* sum up to zero. Because every integer other than −*t* in *X* came from *Y* which is positive, set *S* must contain −*t* to sum up to zero. Now removing −*t* from *S* to form subset *R* of *Y*; the numbers in *R* sum up to *t*. Thus (*Y*, *t*) is a yes-instance for ZEROSUMSET, a contradiction.

Rubric: Standard 5-point grading scaled to 10 points. Maximum 4 points if instance construction does not work.

Watch out if you are reducing from a version of subset sum where negative integers are allowed; in such case simply setting $S := R \cup \{-t\}$ is insufficient (the no-instance fails). One has to offset every integer in Y by a large number K, and pair it with a copy of K; then change the new integer added to $-(t + |Y| \cdot K)$.

Solution: (*Alternative solution.*) We prove that ZEROSUMSET is NP-hard by reduction from the standard NP-hard problem 3SAT.

Given a 3CNF-formula ϕ for 3SAT, we construct an instance *X* of ZEROSUMSET as follows. Let x_1, \ldots, x_n be variables of ϕ and C_1, \ldots, C_m be clauses of ϕ . We add integers of the form

$$\sum_{1 \le i \le n} \alpha_i \cdot 10^{i+m} + \sum_{1 \le j \le m} \beta_j \cdot 10^j$$

to *X*, represented as a vector $(\alpha_1, \ldots, \alpha_n, \beta_1, \ldots, \beta_m)$. Based on the following construction no sum of digits will ever overflow to other digits, and thus we can safely treat each base-10 integer as a vector.

- For each variable x_i , add integers a_i and $\overline{a_i}$ to X, corresponding to literals x_i and $\overline{x_i}$:
 - $\alpha_i = 1$ and $\alpha_{i'} = 0$ for every other i'.
 - $\beta_i = 1$ if clause C_i contains literal $x_i/\overline{x_i}$ and 0 otherwise, respectively.
- For each clause C_j , add two identical integers c_j whose only nonzero digit is at $\beta_j = -1$.
- Add the integer -T = (-1, ..., -1, -3, ..., -3) to *X*.

This finishes the description of the (multi-)set *X*. To turn *X* into an actual set (without duplicate elements), it is sufficient to add different infinitesimals to each element and also include the negative of all the infinitesimals. The construction can be carried out in time polynomial to the size of the formula ϕ .

To prove that the reduction is correct:

- If ϕ is a *yes*-instance, we pick the numbers in *X* that correspond to the literals set to *true* in a fixed satisfying assignment of ϕ , together with -T and enough c_j s equal to one less than the number of true literals in each clause C_j , so that the total sum is zero. Exactly one literal per variable is picked, so each α_i digit sums to zero. Because the assignment is satisfying, at least one literal in each clause is true, and thus we have enough c_j s so that the β_j digit sums up to 3, which is than cancelled by the -3 at the same digit in -T; thus each β_j also sums to zero.
- If ϕ is a *no*-instance, any assignment of ϕ has at least one clause with all literals being false. Any zero-sumset *S* of *X* we choose must contain -T as it is the only number with negative values at digit α_i s. Because each α_i sums to zero, exactly one a_i or $\overline{a_i}$ is chosen. Such choice give rise to an assignment of ϕ ; let C_j denote the unsatisfied clause. Because -T is chosen, the sum of digit β_j of the rest of $S \setminus \{-T\}$ has to be 3, which is impossible as no chosen numbers a_i have positive β_j and there are only two c_j s.

This shows that an instance ϕ is satisfiable if and only if there is a non-empty subset of *X* summing up to zero.