

Time limit: 50 minutes.

Maximum score: 200 points.

Instructions: For this test, you work in teams of eight to solve a multi-part, proof-oriented question.

Problems that use the words “compute” or “list” only call for an answer; no explanation or proof is needed. Unless otherwise stated, all other questions require explanation or proof. Answers for problems should be written on sheets of scratch paper, clearly labeled, with every problem *on its own sheet*. If you have multiple pages for a problem, number them and write the total number of pages for the problem (e.g. 1/2, 2/2).

Write the name of your team on every submitted page. Only submit one set of solutions for the team. Do not turn in any scratch work. After the test, put the sheets you want graded into your packet. If you do not have your packet, ensure your sheets are labeled *extremely clearly* and stack the loose sheets neatly.

In your solution for a given problem, you may cite the statements of earlier problems (but not later ones) without additional justification, even if you haven’t solved them.

The problems are ordered by content, NOT DIFFICULTY. It is to your advantage to attempt problems from throughout the test.

No calculators.

Introduction

This Power Round develops the many and varied properties of the Thue-Morse sequence, an infinite sequence of 0s and 1s which starts $0, 1, 1, 0, 1, 0, 0, 1, \dots$ and appears in a remarkable number of different contexts in recreational and research mathematics. We will see applications to geometry, probability, game theory, combinatorics, algebra, and fractals. Nevertheless, we won't even come close to exhausting the amusing and useful properties of this sequence, some of which require mathematics beyond our scope to discuss.

Remark: Regardless of which problem you decide to work on, it is recommended that you read Problem 1 first to become familiar with the definitions.

Remark 2: The following problems rely heavily on the technique of proof by induction. If you are not yet comfortable with induction, we have copies of an introduction available for you to consult—ask your proctor.

Defining the Thue-Morse sequence

The first sign that there's something special about the Thue-Morse sequence is that it's hard to make up your mind about how to define it, because there are numerous very different-looking definitions which all turn out to be equivalent. In this problem, we work through a few of these definitions and determine that each of them gives the same result. We refer to the n th term of the Thue-Morse sequence by t_n , starting with t_0, t_1, t_2, \dots

1. (a) [3] Our first definition is a simple recursive one. The zeroth term of the Thue-Morse sequence is $t_0 = 0$. For n a nonnegative integer, after the first 2^n terms of the Thue-Morse sequence (including the zeroth term) have been specified, construct the next 2^n terms by taking the first 2^n terms, replacing each 0 by a 1, and replacing each 1 by a 0 (simultaneously). (This is called “bitwise negation”.) Therefore, we have $t_1 = 1$, and the next two terms are $t_2 = 1, t_3 = 0$. The zeroth through fifteenth terms (leaving out the commas, as we will often do for convenience) are 0110100110010110.

Write down (no justification required) the 16th through 31st terms.

- (b) [6] Our second definition is direct. The Thue-Morse sequence is the sequence $\{t_n\}$ ($n = 0, 1, \dots$) where t_n is 1 if the number of ones in the binary (base-2) expansion of n is odd and 0 if the number of ones in the binary expansion of n is even. For example, 5 is 101_2 in base 2, which has two ones, so $t_5 = 0$.

Prove that this definition gives the same sequence as the one from part (a).

- (c) [6] Our third definition is recursive again, but uses a different recursion. The Thue-Morse sequence is the sequence $\{t_n\}$ satisfying $t_0 = 0$, $t_{2n} = t_n$, and $t_{2n+1} = 1 - t_n$.

Prove that this definition is equivalent to either of the first two definitions.

- (d) [6] Our fourth definition is by a certain algorithm (known as a Lindenmeyer system). We start with the single digit 0 (call this stage zero). At each stage, we take the digits we already have, replace each 0 by a 01, and replace each 1 by a 10 (simultaneously). So stage one is 01, stage two is 0110, and so on. The Thue-Morse sequence is the sequence $\{t_n\}$ whose first 2^n terms are the digits from stage n .

Prove that this definition is equivalent to any of the first three definitions. (Note that as stated, it is not clear that this definition is even coherent, since it redefines each term over and over again. Your job is to show that it nevertheless uniquely defines each term as the corresponding term of the Thue-Morse sequence as given by parts (a)-(c).)

Now we derive a few simple properties of the Thue-Morse sequence, just to play with it some more.

2. (a) [5] Prove that the string $t_0 t_1 \cdots t_{2^{2n}-1}$ is a palindrome for all $n \geq 0$. (Recall that a palindrome is a string of digits which reads the same forward and backward.)
- (b) Let A be the set of all nonnegative integers n such that $t_n = 0$. Let $n \oplus m$ denote the *binary xor* of n and m . (To compute the binary xor of n and m , we write both n and m in binary, then add them without carrying. For example, if $n = 5$ and $m = 13$, then $n = 101_2$ and $m = 1101_2$, so $n \oplus m = 1000_2 = 8$.)
 - (i.) [1] Compute $14 \oplus 23$.
 - (ii.) [5] Prove that if n and m are both in A , then $n \oplus m$ is also in A .
- (c) [6] Prove that given any finite string $X = t_a t_{a+1} \cdots t_b$ of consecutive terms from the Thue-Morse sequence, there exists a number n_X such that *every* string of n_X consecutive terms $t_{k+1} t_{k+2} \cdots t_{k+n_X}$ from the sequence must contain X .
- (d) [6] Given a finite or infinite string T of 0s and 1s, let $f(T)$ be the string created by simultaneously replacing each 0 by a 01 and each 1 by a 10. For example, if $T = 001$, then $f(T) = 010110$. Note that we previously saw this procedure in problem 1, part d. A *fixed point* of f is an infinite string T such that $f(T) = T$. Prove that f has exactly two fixed points: the Thue-Morse sequence $\{t_n\}$, and its bitwise negation (meaning the sequence constructed from $\{t_n\}$ by replacing each 0 with a 1 and each 1 with a 0).

Greedy Galois Games

Time for some probability and game theory. Alice and Bob are in a duel where in each round (beginning with round 0), one duelist fires a shot at the other, hitting them with a success probability of p . The first person to fire a successful shot wins. They want to choose the shooter each round in a way that's fair—just switching back and forth after every shot wouldn't be fair, since we can see intuitively that whoever goes first is more likely to win. Also, they're both terrible at aiming, so p is very low, though positive. What do they do?

They come up with the following idea: Alice shoots first. Then, Bob shoots as many times as is necessary for his win probability to meet or exceed that of Alice's win probability so far. Then, Alice starts shooting again, again taking as many turns as is necessary for her win probability to meet or exceed that of Bob's win probability. And so on (if at any point, they have the same probability of winning, we let the person who was not shooting in the previous round shoot in the next round).

For example, suppose $p = 1/3$. Alice shoots during round 0, after which her win probability is $1/3$ and Bob's win probability is 0. Bob shoots during round 1. For Bob to win during round 1, Alice has to miss in round 0, which happens with probability $2/3$, and Bob has to hit in round 1, which happens with probability $1/3$. So after round 1, Bob's win probability is $(2/3)(1/3) = 2/9$, which is still less than Alice's win probability of $1/3$. Therefore, Bob shoots again in round 2. By the same logic, his overall win probability after round 2 is $(2/3)(1/3) + (2/3)(2/3)(1/3) = 10/27$, which is now higher than $1/3$. So Alice gets to shoot in round 3. And so on.

Let $P(A)$ be Alice's overall win probability after a given round, and $P(B)$ be Bob's win probability. We summarize the above information in the following table:

Round #	Shooter	$P(A)$	$P(B)$
0	Alice	1/3	0
1	Bob	1/3	2/9
2	Bob	1/3	10/27
3	Alice	?	10/27
4	?	?	?

3. (a) (i.) [2] Fill in the question marks in the above table (no justification required).
(ii.) [3] Fill in the same table for $p = 1/4$ instead of $1/3$ (no justification required).
(b) [6] Let $q = 1 - p$. Let $\{a_n\}$ be the sequence such that $a_n = -1$ if Alice shoots in round n and $a_n = 1$ if Bob shoots in round n . Let $P(A_n)$ be Alice's overall win probability after round n , and $P(B_n)$ Bob's overall win probability after round n . Finally, let

$$f_n(x) = a_n \left(\sum_{j=0}^n a_j x^j \right).$$

Prove that

$$a_{n+1} = \begin{cases} -a_n & \text{if } f_n(q) \geq 0, \\ a_n & \text{otherwise.} \end{cases}$$

- (c) [3] Prove that regardless of the value of p , we always have $a_0 = -1, a_1 = 1, a_2 = 1$.
(d) [3] Determine, with proof, all values of p such that $a_3 = -1$.

Our goal is now to prove that as p gets close to 0, or equivalently as q gets close to 1, the pattern of who shoots who becomes more and more like the Thue-Morse sequence, in the following sense. Recall that we define a_n to be -1 if Alice shoots in round n and 1 if Bob shoots in round n , and that $\{t_n\}$ is the Thue-Morse sequence. Let $\{t'_n\}$ be the sequence such that $t'_n = -1$ if $t_n = 0$ and $t'_n = 1$ if $t_n = 1$. That is, $\{t'_n\}$ is basically also the Thue-Morse sequence, just using -1 and 1 instead of 0 and 1 , since that's more convenient for our current application. We're going to show that as p gets close to 0, more and more of the first few terms of $\{a_n\}$ equal the first few terms of $\{t'_n\}$.

4. (a) [8] Prove that for each $n \in \mathbb{N}$, there is an $\epsilon > 0$ such that the sequence a_0, a_1, \dots, a_n is the same for all $q \in (1 - \epsilon, 1)$. Intuitively, this shows that as the success probability p nears zero, more and more of the first few terms of a_n stabilize and become fixed. (Hint: start with your solution to Problem 3).
(b) (i) [3] Prove that for any m , we have $\sum_{i=0}^{2m+1} t'_i = 0$.
(ii) [7] Suppose that there exists $\epsilon > 0$ such that for all $q \in (1 - \epsilon, 1)$, $a_i = t'_i$ for $0 \leq i \leq 2m$. Prove that then there is an $\epsilon' > 0$ such that $a_{2m+1} = -a_{2m}$ for all $q \in (1 - \epsilon', 1)$.
(c) [6] Suppose that there exists $\epsilon > 0$ such that for all $q \in (1 - \epsilon, 1)$, $a_i = t'_i$ for $0 \leq i \leq 2m + 1$. Prove that when $q \in (1 - \epsilon, 1)$, $f_{2m+1}(q) = (q - 1)f_m(q^2)$.
(d) [6] Prove that for each $n \in \mathbb{N}$, there is an $\epsilon > 0$ such that the sequence a_0, a_1, \dots, a_n is the same as the sequence t'_0, t'_1, \dots, t'_n for all $q \in (1 - \epsilon, 1)$. (This demonstrates the claim we made in the paragraph before this problem.)

Pattern avoidance

Now we develop and prove some more complicated but really cool properties of the Thue-Morse sequence. The goal of the next problem is to prove that no string of consecutive terms in the Thue-Morse sequence repeats itself three times consecutively. That is, the Thue-Morse sequence contains no *cubes*, where a cube is a nonempty string of consecutive terms which looks like www , where w is any string of 0s and 1s (for example, 001001001 is a cube with $w = 001$). As in some previous problems, we will leave out the commas between terms for convenience.


5. (a) [3] Of course, the simplest cubes are 000 and 111. Prove directly that in the Thue-Morse sequence, there are never three consecutive 0s or three consecutive 1s. (You may leave this part blank and receive full credit for it, but *only* if you receive full credit on the entire rest of this problem.)
- (b) We define an *overlapping factor* to be a nonempty string x of consecutive terms which begins with a string w of length shorter than x , and ends with the same string w , such that the two occurrences of w overlap in at least one term. For example, $x = 11011011$ is an overlapping factor because it both begins and ends with $w = 11011$, and the two instances of 11011 overlap by two terms (the middle two 1s).
 - (i.) [3] Prove that if a sequence contains a cube, then it also contains an overlapping factor.
 - (ii.) [8] Prove that if a sequence contains an overlapping factor, then it also contains an overlapping factor of the form $avava$, where a is a single term and v is a (possibly empty) string of terms.
- (c) [5] Suppose that $x = a_0a_1 \cdots a_{2n-1}$ where each a_i is either 0 or 1 and each string $a_{2i}a_{2i+1}$ is either 01 or 10. Prove that it is not possible to write $0x0$ or $1x1$ in the form $b_0b_1 \cdots b_{2n+1}$ where each b_j is either 0 or 1 and each string $b_{2i}b_{2i+1}$ is either 01 or 10.
- (d) Given a string T of 0s and 1s, let $f(T)$ be the function from problem 2, part d—that is, the string created by simultaneously replacing each 0 by a 01 and each 1 by a 10.
 - (i.) [6] Suppose $f(T) = xavavay$ where a is a single term (0 or 1) and x, v, y are strings of 0s and 1s. Prove that v consists of an odd number of terms.
 - (ii.) [7] Prove that if $f(T)$ contains an overlapping factor, then T also contains an overlapping factor.
 - (iii.) [3] Prove that the Thue-Morse sequence contains no overlapping factors, and therefore no cubes.

Miscellaneous


Just for fun, here are a few more cute and unexpected things you can do with the Thue-Morse sequence.

6. (a) The Koch snowflake is a well-known fractal that is constructed over iterations as follows. Our initial “snowflake”, the zeroth iteration, is just a straight line segment.

In the first iteration, we take the middle third of the line segment, draw an equilateral triangle using that middle third as a base, and then erase the middle third, resulting in the following figure.



In the second iteration, we take every line segment in the above figure and repeat the same procedure: replacing the middle third of the line segment with the other two sides of the outward-facing equilateral triangle that has that middle third as a base.



In general, we create the n th iteration of the Koch snowflake by taking each line segment in the $(n - 1)$ th iteration and replacing the middle third by a “corner” in the shape of an equilateral triangle, in the same way as before.

- (i) [2] Draw (no justification required) the third iteration of the Koch snowflake.
- (ii) [7] A turtle reads the Thue-Morse sequence t_0, t_1, \dots and decides to crawl according to the sequence, as follows. At the n th step, if $t_n = 0$, it will crawl forward one unit and then turn 60 degrees to the left. If instead $t_n = 1$, it will turn 180 degrees (without moving). Prove that after 2^{2n+1} steps (that is, after following the sequence from t_0, t_1, \dots up to $t_{2^{2n+1}-1}$), the turtle will have traced out the n th iteration of the Koch snowflake. (Of course, we are ignoring the scale of the resulting snowflake here; we are only interested in its shape.)
- (b) [9] Let $N = 2^{n+1}$. Let A_N be the set of integers i in $\{0, 1, \dots, N - 1\}$ such that $t_i = 0$, and let B_N be the set of integers j in $\{0, 1, \dots, N - 1\}$ such that $t_j = 1$. Prove that

$$\sum_{i \in A_N} i^k = \sum_{j \in B_N} j^k$$

for all integers k from 1 to n . (This is a special case of the *Prouhet-Tarry-Escott problem*.)

- (c) [11] As in the discussion after Problem 4, let $\{t'_n\}$ be the Thue-Morse sequence using $-1, 1$ instead of $0, 1$. Prove that

$$\left(\frac{1}{2}\right)^{t'_0} \left(\frac{3}{4}\right)^{t'_1} \left(\frac{5}{6}\right)^{t'_2} \cdots = \prod_{n=0}^{\infty} \left(\frac{2n+1}{2n+2}\right)^{t'_n} = \sqrt{2}.$$