The Space Complexity of Mirror Games

Sumegha Garg
Princeton University, Princeton, USA
sumeghag@cs.princeton.edu

Jon Schneider
Google Research, New York, USA
jschnei@google.com

Abstract

We consider the following game between two players Alice and Bob, which we call the mirror game. Alice and Bob take turns saying numbers belonging to the set $\{1, 2, \ldots, N\}$. A player loses if they repeat a number that has already been said. Otherwise, after N turns, when all the numbers have been spoken, both players win. When N is even, Bob, who goes second, has a very simple (and memoryless) strategy to avoid losing: whenever Alice says x, respond with $N+1-x$. The question is: does Alice have a similarly simple strategy to win that avoids remembering all the numbers said by Bob?

The answer is no. We prove a linear lower bound on the space complexity of any deterministic winning strategy of Alice. Interestingly, this follows as a consequence of the Eventown-Oddtown theorem from extremal combinatorics. We additionally demonstrate a randomized strategy for Alice that wins with high probability that requires only $\tilde{O}(\sqrt{N})$ space (provided that Alice has access to a random matching on K_N).

We also investigate lower bounds for a generalized mirror game where Alice and Bob alternate saying 1 number and b numbers each turn (respectively). When $1+b$ is a prime, our linear lower bounds continue to hold, but when $1+b$ is composite, we show that the existence of a $o(N)$ space strategy for Bob (when $N \not\equiv 0 \mod (1+b)$) implies the existence of exponential-sized matching vector families over $\mathbb{Z}_{1+b}^N$.

2012 ACM Subject Classification Theory of computation → Interactive computation

Keywords and phrases Mirror Games, Space Complexity, Eventown-Oddtown

Digital Object Identifier 10.4230/LIPIcs.ITCS.2019.36

Related Version `https://arxiv.org/abs/1710.02898`

Acknowledgements We would like to thank Mark Braverman, Sivakanth Gopi, and Jieming Mao for helpful advice and discussions.

1 Introduction

1.1 The Mirror Game and Mirror Strategies

Consider the following simple game. Alice and Bob take turns saying numbers belonging to the set $\{1, 2, \ldots, N\}$. If either player says a number that has previously been said, they lose. Otherwise, after N turns, all the numbers in the set have been spoken aloud, and both players win. Alice says the first number.

If N is even, there is a very simple and computationally efficient strategy that allows Bob to win this game, regardless of Alice's strategy: whenever Alice says x, Bob replies with $N+1-x$. This is an example of a *mirror strategy* (and for this reason, we refer to the game above as the *mirror game*). Mirror strategies are an effective tool for figuring out who wins

10th Innovations in Theoretical Computer Science (ITCS 2019).
Editor: Avrim Blum; Article No. 36; pp. 36:1–36:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

in a variety of combinatorial games (for example, two-pile Nim [2]). More practically, mirror strategies can be applied when playing more complex games, such as chess or go (to varying degrees of success). From a computational perspective, mirror strategies are interesting as they require very limited computational resources - most mirror strategies can be described via a simple transformation of the preceding action.

Returning to the game above, this leads to the following natural question: does Alice have a simple strategy to avoid losing when N is even? Since both players have access to the same set of actions, one may be tempted to believe that the answer is yes - in fact, if N is odd, then Alice can start by saying the number N and then adopt the mirror strategy for Bob described above for a set of $N-1$ elements. However, when N is even, the mirror strategy as stated does not work.

To answer the above question, we need to formalize what we mean by simple; after all, Alice has plenty of strategies such as "say the smallest number which has not yet been said". One useful metric of simplicity is the metric of *space complexity*, the amount of memory a player requires to implement their strategy (we formalize this in Section 2). Note that Bob needs only $O(\log N)$ bits of memory to implement his mirror strategy, whereas the naive strategy for Alice (remembering everything) requires $O(N)$ bits.

In this paper, we show that this gap is necessary; any successful, deterministic strategy for Alice requires at least $\Omega(N)$ bits of memory:

▶ **Theorem 1** (Restatement of Theorem 11). *If N is even, then any winning strategy for Alice in the mirror game requires at least* $(\log_2 5 - 2)N - o(N)$ *bits of space.*

1.2 Eventown and Oddtown

While many tools exist in the computer science literature for showing space lower bounds (e.g. communication complexity, information theory, etc.), one interesting feature of this problem absent from many others is that any proof of Theorem 1 *must depend crucially on the parity* of N.

In the study of set families in extremal combinatorics, an "Oddtown" is a collection of subsets of $\{1, 2, \dots N\}$ where every subset has even cardinality but each pair of distinct subsets has an intersection of odd cardinality. Likewise, an "Eventown" is a collection of subsets of $\{1, 2, \dots, N\}$ where every subset has even cardinality but each pair of distinct subsets has an intersection of even cardinality. In 1969, Berlekamp [3], answering a question of Erdős, showed that while there exist Eventowns containing up to $2^{N/2}$ subsets, any Oddtown contains at most N subsets.

It turns out that this exponential gap between the size of Oddtowns and the size of Eventowns is directly responsible for the exponential gap between the space complexity of Alice's strategy and the space complexity of Bob's strategy! (One way to see this connection is that, just as Bob's $O(\log N)$-space strategy involves pairing up the numbers of $\{1, 2, \dots, N\}$, one way to construct an Eventown of size $2^{N/2}$ is to perform a similar pairing, and then consider all subsets formed by unions of pairs).

The Eventown-Oddtown theorem figures into our proof of Theorem 1 in the following way. At a given turn, for each possible state of memory of Alice, label the state with the possible subsets of numbers that could have possibly been said before this turn. We show (Lemma 13) that these subsets must form an Oddtown, or else Bob has some strategy that can force Alice to lose. Since there are a large (exponential) number of total possible subsets (Corollary 10), and since each Oddtown contains at most N subsets, this implies that Alice's memory must be large.

1.3 Randomized Strategies for Alice

A natural followup to Theorem 1 is whether these lower bounds continue to hold if Alice instead uses a randomized strategy, which only needs to succeed with high probability.

We provide some evidence to show that this might not be the case. We demonstrate an $O(\sqrt{N}\log^2 N)$-space algorithm for Alice that succeeds with probability $1 - O(1/N)$, as long as Alice is provided access to a uniformly chosen perfect matching on the complete graph on N vertices (K_N) (Theorem 18). In addition, even with $O(\log N)$-space (and access to a uniformly chosen perfect matching on K_N), Alice can guarantee success with probability $\Omega(1/N)$.

In both of these algorithms, Alice attempts the mirror strategy of Bob, hoping that Bob does not choose the number with no match. In the $O(\sqrt{N}\log^2 N)$ algorithm, Alice decreases her probability of failure by maintaining a set of $\tilde{O}(\sqrt{N})$ possible "backup" points she can switch to if Bob identifies the unmatched point. This lets Alice survive until turn $N-\tilde{O}(\sqrt{N})$, whereupon Alice can reconstruct the remaining $O(\sqrt{N})$ elements by maintaining $O(\sqrt{N})$ power sums during the computation.

Since Alice cannot store a perfect matching on K_N in $o(N)$ space, this unfortunately does not give any $o(N)$-space strategies for Alice. In addition, our lower bound techniques from Theorem 1 fail to give non-trivial guarantees in the randomized model. Demonstrating non-trivial algorithms or non-trivial lower bounds for the general case is an interesting open problem.

1.4 General Mirror Games

It is possible to generalize the mirror game presented earlier to the case where in each turn Alice says a numbers and Bob says b numbers. We refer to this game as the (a, b)-*mirror game*.

By a generalization of the Oddtown theorem that works modulo prime p, our proof of Theorem 1 immediately carries over to show that if N is not divisible by p, then any winning strategy for Bob in the $(1, p-1)$-mirror game requires $\Omega(N)$ space (Theorem 20).

The natural generalization of the Oddtown theorem is known not to hold modulo composite m. Grolmusz showed [9] that if m is composite, there exists a set family of quasipolynomial size such that every set has cardinality divisible by m, but the intersection of any two distinct sets has cardinality not divisible by m. The best known upper bounds for the size of such a set family are of the form $2^{O(N)}$, and improving these to any bound of the form $2^{o(N)}$ is an important open problem, with applications in coding theory to the construction of matching vector families [5, 4].

As long as the size of a modulo m Oddtown is bounded above by $2^{o(N)}$, the proof of Theorem 20 still achieves an $\Omega(N)$ lower bound on the space complexity of Bob's strategy. It follows that finding any $o(N)$-space winning strategy for Bob in the $(1, m-1)$ mirror game implies the existence of a $2^{\Omega(N)}$-size modulo m Oddtown, and hence similarly sized matching vector families (Theorem 22). Admittedly, Bob may not have an $o(N)$-space winning strategy in this game for unrelated reasons; it is interesting whether it is possible to show a converse result, constructing a low-space strategy for Bob from any $2^{\Omega(N)}$-size modulo m Oddtown.

For other (a, b)-mirror games, we understand much less about the complexity of winning strategies (if $a+b$ is prime, then depending on the residue of N modulo $a+b$ it is occasionally possible to show an $\Omega(N)$ lower bound on the lower bound of a winning strategy).

2 Definitions and Preliminaries

2.1 The Mirror Game

The *mirror game on N elements* as discussed before is a game between two players, Alice and Bob. Alice and Bob take turns (Alice going first) saying a number in $[N]$ ($[N]$ indicates the set $\{1, 2, \ldots, N\}$). If a player says a number that has previously been said (either by the other player or by themselves) they lose and the other player wins. In addition, after N successful turns (so no numbers in $[N]$ remain unsaid), both players are declared winners.

With unbounded memory, it is clear that both players can easily win this game. We therefore restrict our attention to strategies with bounded memory. A *strategy computable in memory m* for Alice in the mirror game is defined by an initial memory state $s_0 \in [2^m]$ and a pair of transition functions $f : [N] \times [2^m] \times [N] \to [2^m]$ and $g : [2^m] \to [N]$ computable in $SPACE(m)$ (see [1] for an introduction to bounded space complexity). The function f takes in the previous reply $b \in [N]$ of Bob, Alice's current memory state $s \in [2^m]$, and the current turn $t \in [N]$, and returns Alice's new memory state $s' \in [2^m]$; the function g takes in Alice's current memory state $s \in [2^m]$ (after updating it based on Bob's move), and outputs her next move a. We say a strategy for Alice is a *winning strategy* if Alice is guaranteed to win regardless of Bob's choice of actions.

By removing the constraint that the transition function f is computable in $SPACE(m)$, we obtain a larger class of strategies, which we refer to as strategies *weakly computable in memory m*. Our lower bounds in Theorems 1 and 20 continue to hold for this larger class of strategies.

2.2 Eventown and Oddtown

In this section we review the known literature on the Eventown-Oddtown problem. Note that while in the introduction we only defined the terms "Eventown" and "Oddtown", there are actually four different classes of set system depending on the parity of cardinalities of the subsets and the parity of the cardinality of the intersection.

▶ **Definition 2.** A collection of subsets $\mathcal{F} \subseteq [N]$ forms an *(Odd, Even)-town* of sets if:

1. For every $F \in \mathcal{F}$, $|F| \equiv 1 \mod 2$.
2. For every $F_1 \neq F_2 \in \mathcal{F}$, $|F_1 \cap F_2| = 0 \mod 2$.

We define *(Odd, Odd)-towns*, *(Even, Odd)-towns*, and *(Even, Even)-towns* similarly.

Note that there exist (Even, Even)-towns and (Odd, Odd)-towns containing exponentially (in N) many sets; one simple construction of an (Even, Even)-town is to partition the ground set into pairs (possibly with a leftover element), and consider all sets formed by taking unions of these pairs. In contrast, (Even, Odd)-towns and (Odd, Even)-towns each contain at most N sets.

▶ **Lemma 3.** *Any (Odd, Even)-town $\mathcal{F}$ has size at most N i.e. $|\mathcal{F}| \leq N$.* [3, 12]

For completeness, we give the proof below.

Proof. Let $\mathcal{F}$ be $\{F_1, F_2, ..., F_k\}$. Let $v_i \in \{0,1\}^N$ be the characteristic function of F_i i.e. $v_i[l] = 1 \iff l \in F_i$. We know that $v_i^T v_j \mod 2 = 1$ for $i = j$ and 0 otherwise. We claim that $v_1, v_2, ..., v_k$ are linearly independent over $\mathbb{F}_2$ (where $\mathbb{F}_2$ is the field with 2 elements) and hence, $k \leq N$. We prove this claim by contradiction. If they are not linearly independent, then $\exists \lambda_1, \lambda_2, ..., \lambda_k \in \{0,1\}$ not all zero such that $\sum_{i=1}^k \lambda_i v_i = 0 \mod 2$. Without loss of generality, assume $\lambda_1 \neq 0$. As $v_1^T(\sum_{i=1}^k \lambda_i v_i) = \lambda_1$ over $\mathbb{F}_2$, this implies $\lambda_1 = 0$; a contradiction. ◀

▶ **Corollary 4.** *Any (Even, Odd)-town contains at most N sets.*

Proof. We prove an upper bound of $N+1$ sets by reduction to Lemma 3. Given an (Even, Odd)-town $\mathcal{B} = \{B_1, ..., B_m\}$, construct an (Odd, Even)-town in space $[N+1]$ from sets $B_1 \cup \{N+1\}, B_2 \cup \{N+1\}, ..., B_m \cup \{N+1\}$. Using Lemma 3, we get $m \leq N+1$. For a stronger upper bound of N, refer to [3, 12]. ◀

By adapting the above linear algebraic arguments to the field $\mathbb{F}_p$ (p is prime and $\mathbb{F}_p$ refers to field over p elements), it is possible to show similar upper bounds on the size of set families with cardinality constraints modulo p. We will use the following lemma, due to Frankl and Wilson.

▶ **Lemma 5** ([8])**.** *Let p be a prime and L be a set of s integers. Let $B_1, B_2, ..., B_m \subseteq [N]$ be a family of subsets such that:*

1. $|B_i| \mod p \notin L$, $\forall i \in [m]$.
2. $|B_i \cap B_j| \mod p \in L$, $\forall i \neq j$.

Then

$$m \leq \sum_{i=0}^{s} \binom{N}{s}$$

We call such family of subsets a (p, L)-Modtown.

Interestingly, by a result of Grolmusz, there is no straightforward generalization of these results modulo a composite number k.

▶ **Theorem 6.** *Let k be a positive integer with $r > 1$ different prime divisors. Then there exists a $c > 0$, such that for every N, there exists a family of subsets $B_1, B_2, \ldots, B_m$ such that:*

1. $m \geq \exp(c(\log N)^r/(\log\log N)^{r-1})$,
2. $|B_i| \equiv 0 \bmod k$, $\forall i$,
3. $|B_i \cap B_j| \not\equiv 0 \bmod k$, $\forall i \neq j$.

Proof. See [9]. ◀

This type of set family is captured in coding theory by the definition of a matching vector family.

▶ **Definition 7.** A *matching vector family* [4] over $\mathbb{Z}_m^n$ of size t is a pair of ordered lists $U = \{u_1, u_2, ..., u_t\}$ and $V = \{v_1, v_2, ..., v_t\}$ where $u_i, v_j \in \mathbb{Z}_m^n$ such that for all i, $u_i^T v_i = 0$, and for all $i \neq j$, $u_i^T v_j \neq 0$.

By taking $u_i = v_i$ to be the characteristic vector of B_i above, it is clear that a family of subsets of size m in Theorem 6 gives rise to a Matching Vector family of size m. Matching vector families have deep applications to many problems in coding theory, such as private information retrieval and locally decodable codes [6, 7]. Understanding the maximum possible size of a matching vector family is an important open problem in coding theory.

3 Alice Requires Linear Space

In this section we prove that Alice requires linear space to win the mirror game when N is even. To do this, we will show that if Alice is at a specific memory state, then the possible sets of numbers that have been said till that point in time must form an (Even, Odd)-town (and hence there are at most N such sets for any memory state).

Before we get into the main proof, we will define and lower-bound the size of what we call a "covering collection" of subsets (this will later allow us to lower bound the total number of possible sets of numbers said by round t ($2t$ turns)).

▶ **Definition 8.** A collection of subsets $\mathcal{C}$ of $[N]$ is (p, r)*-covering* if:

1. each $S \in \mathcal{C}$ has $|S| = pr$.
2. for every $T \subset [N]$ with $|T| = r$, there exists a $S \in \mathcal{C}$ with $T \subset S$.

▶ **Lemma 9.** *Every (p, r)-covering collection $\mathcal{C}$ has size at least $\frac{\binom{N}{r}}{\binom{pr}{r}}$.*

Proof. Every set in this collection contains at most $\binom{pr}{r}$ sets T with cardinality r. There are $\binom{N}{r}$ possible sets T. ◀

When $p = 2$, it turns out that the lower bound in Lemma 9 is maximized when $r = N/5$.

▶ **Corollary 10.** *When $r = N/5$, every $(2, r)$-covering collection has size at least $2^{(\log_2 5-2)N-o(N)}$.*

We now prove our main theorem.

▶ **Theorem 11.** *If N is even, then any winning strategy in the mirror game for Alice requires at least $(\log_2 5 - 2)N - o(N)$ bits of space.*

Proof. Fix a winning strategy for Alice. Assume this strategy uses m bits of memory, and thus has $M = 2^m$ distinct states of memory.

Call a subset S of $[N]$ r*-occurring* if it is possible that immediately after turn $2r$ (r rounds), the set of numbers that have been said is equal to S. Let $\mathcal{S}_r$ be the collection of all r-occurring sets. Before diving into the main proof, for any fixed deterministic strategy of Alice, we prove a lower bound on $\mathcal{S}_r$ i.e. the number of different subsets of numbers that could have been said in the first $2r$ turns over various strategies of Bob.

▶ **Lemma 12.** *$\mathcal{S}_r$ is $(2, r)$-covering.*

Proof. Since $2r$ numbers have been said immediately after turn $2r$, every set in $\mathcal{S}_r$ has cardinality $2r$. We must show that for any $T \subset [N]$ with $|T| = r$, that there exists an S in $\mathcal{S}_r$ with $T \subset S$.

Consider the following strategy for Bob: "say the smallest number in T which has not yet been said". Note that if Bob follows this strategy, then the set of numbers said by turn $2r$ must contain the entire set T. This set belongs to $\mathcal{S}_r$, and it follows that $\mathcal{S}_r$ is $(2, r)$-covering. ◀

Write $N = 2n$, and fix a value $r \in [n]$. For a memory state x out of the M possible memory states and an r-occurring set S, label x with S if it is possible that Alice is at memory state x when the set of numbers that have been said is equal to S. Each state of memory may be labeled with several or none r-occurring sets, but each r-occurring set must exist as a label to some state of memory (by definition). Let $\mathcal{U}_x$ be the collection of r-occurring labels for memory state x. We want to upper bound the size of $\mathcal{U}_x$. Following lemma along with (Even, Odd)-town Lemma (Corollary 4) helps us in doing exactly that.

▶ **Lemma 13.** *If S_1 and S_2 belong to $\mathcal{U}_x$, then $|S_1 \setminus S_2|$ is odd.*

Proof. Let $D = S_1 \Delta S_2$, let $D_1 = S_1 \setminus S_2$, and let $D_2 = S_2 \setminus S_1$. Assume to the contrary that $|D_1|$ is even. Note then that $|D_2|$ is also even (since $|S_1| = |S_2| = 2r$) and that $|D| = 2|D_1|$. We'll consider two possible cases for the state of the game after turn $2r$: 1. Alice is at state x, and the set of numbers that have been said is S_1, and 2. Alice is at state x, and the set of numbers that have been said is S_2.

Consider the following strategy that Bob can play in either of these cases: "say the smallest number that has not been said and is not in D". We claim that if Bob uses this strategy, Alice will be the first person (after turn $2r$) to say an element of D. Note that Bob will not say an element of D until after turn $2n - |D|/2$; this is since:

1. If we are in case 1, then all of the numbers in D_1 have been said but none of the numbers in D_2 have been said. There are therefore $|D_2| = |D|/2$ numbers in D which have not been said, so Bob can avoid saying an element of D till turn $2n - |D|/2$.
2. Likewise, if we are in case 2, the argument proceeds symmetrically.

On the other hand, if no element of D is spoken by either player between turn $2r$ and turn $2n - |D|/2$, then at turn $2n - |D|/2 + 1$, the only remaining elements belong to D (the set of remaining elements is either D_1 or D_2, depending on which case we are in). If $|D_1| = |D|/2$ is even, then it is Alice's turn to speak at turn $2n - |D|/2 + 1$. It follows that Alice will be the first person after turn r to say an element of D.

Let y_1 be the memory state of Alice when she first speaks an element of D in case 1, and define y_2 similarly. We claim that $y_1 = y_2$. Indeed, since Alice's strategy is deterministic and starts from x in both cases 1 and 2, Bob's strategy plays identically in both case 1 and case 2 until an element of D has been spoken. It follows that Alice must speak the same element of D in both cases. But if this element is in D_1, and they are in case 1, then this element has already been said before; similarly, if this element is in D_2, and they are in case 2, then this element has also been said before. Regardless of which element in D Alice speaks at this state, there is some case where she loses, which contradicts the fact that Alice's strategy is successful. It follows that $|D_1|$ must be odd, as desired. ◀

▶ **Claim 14.** $|\mathcal{U}_x| \leq N$

Proof. We claim the sets in $\mathcal{U}_x$ form an (Even, Odd)-town, from which this conclusion follows (Corollary 4). Each set in $\mathcal{U}_x$ has cardinality $2r$, so all sets have even cardinality. By Lemma 13, any pair of distinct sets $S_1, S_2 \in \mathcal{U}_x$ has odd $|S_1 \setminus S_2|$. Note that since $|S_1 \cap S_2| = |S_1| - |S_1 \setminus S_2|$, and since $|S_1|$ is even, it follows that $|S_1 \cap S_2|$ is odd, so any pair of sets have an odd cardinality intersection. ◀

Choose $r = N/5$. By Corollary 10, $\mathcal{S}_r$ must have cardinality at least $2^{(\log_2 5 - 2)N - o(N)}$. Since each element in $\mathcal{S}_r$ belongs to at least one collection $\mathcal{U}_x$, and since each collection $\mathcal{U}_x$ has cardinality at most N (Claim 14), the number M of memory states x is at least $2^{(\log_2 5-2)N-o(N)}/(N+1) = 2^{(\log_2 5-2)N-o(N)}$. It follows that $m = \log M \geq (\log_2 5 - 2)N - o(N)$, as desired. This proves Theorem 11. ◀

4 Randomized Strategies for Alice

In this section, we consider randomized strategies for Alice. A *randomized strategy computable in memory m* is defined similarly as in Section 2, with the exception that the transition function no longer needs to be deterministic and need only be computable in $RSPACE(m)$;

i.e. it must be computable by a space m Turing machine with access to a $poly(N, m)$-sized read-once random tape. We say a randomized strategy wins with probability p if Alice wins with probability at least p against every possible strategy for Bob.

Unfortunately, we do not know any randomized strategies with sublinear memory which win with high probability. We therefore relax the definition of randomized strategy above and consider randomized strategies where Alice has oracle access to a perfect matching M chosen uniformly at random from all perfect matchings on K_N where K_N is the complete graph on N vertices (we assume here that N is even). Calling this oracle with $x \in [N]$ returns x's match $M(x)$ in this matching, and M may be called arbitrarily many times during the computation of f.

One way to interpret the following upper bounds is as a source of difficulty for proving strong lower bounds for randomized strategies (of the form "you need linear space to succeed with high probability"), since a randomized strategy with memory m with access to a random matching can be viewed as a convex combination of deterministic strategies weakly computable in memory m. This means that any attempt to extend aforementioned strong lower bounds against randomized strategies must fail against this stronger class of randomized strategies.

Another way to interpret this model of computation is as a specific case of the setting where Alice has arbitrary read access to her random tape (as opposed to read-once access). It is known [11] that having arbitrary read access to randomness is more powerful than read-once access as long as certain probabilistic space classes do not collapse. Proving a strong lower bound for randomized strategies would provide more evidence for this separation (this time in the setting of streaming games).

We begin by showing that with only logarithmic space (and access to a random perfect matching), Alice can already win with probability $\Omega(1/N)$. In contrast, without access to a random matching, we know no logarithmic space strategy that succeeds with better than an exponentially small probability.

▶ **Theorem 15.** *When N is even, there exists a randomized strategy (with access to a uniformly random perfect matching M on K_N) for Alice with space complexity $O(\log N)$ which succeeds with probability $\Omega(1/N)$.*

Proof. Consider the following strategy for Alice. She begins by sampling a uniform element x from $[N]$. On her first turn, Alice says x. On subsequent turns, if Bob has just previously said y, Alice replies with $M(y)$.

Note that with this strategy, Alice wins playing against Bob if the last number said by Bob on turn N is $M(x)$. This follows since M is a matching, so there is no other number $y \neq M(x)$ that Bob can say where $M(y)$ has also already been said. It follows that if Bob says $M(x)$ at turn N, then Alice is guaranteed to win.

What is the probability Bob says $M(x)$ before turn N? We will show it is less than $1 - 1/N$. To do this, we first claim that any winning strategy for Bob (i.e. any strategy that doesn't repeat previously said elements) has the same probability of saying $M(x)$ before turn N. This follows from symmetry: since a uniformly random perfect matching conditioned on containing a submatching is still a uniformly random perfect matching on the remaining vertices, at any point in the protocol, if $M(x)$ has not been said yet, it has an equally likely chance of being any of the unsaid elements. Therefore, without loss of generality, assume Bob is playing according to the strategy where each turn he says the smallest number that has not been said so far.

Now, note that if $M(x) = N$, Bob will not say $M(x)$ before turn N (there will always be a smaller unsaid element). But this happens with probability $1/N$, and therefore Alice succeeds with probability at least $1/N$. ◀

We will next show how to extend this idea to construct an $\tilde{O}(\sqrt{N})$ strategy for Alice which succeeds with high probability. To do this, we will need the following folklore result on determining missing elements from a set in a streaming setting.

▶ **Definition 16.** The "missing k numbers problem" is a streaming problem where a subset S of cardinality $N - k$ is chosen from $[N]$, and Alice is shown the elements of S one at a time, in some order. Alice's goal is to output the set $[N] \setminus S$.

▶ **Lemma 17.** *There exists an $O(k \log N)$-space deterministic algorithm for the missing k numbers problem.*

Proof. See [10]. For completeness, we include the proof here.

Choose a prime $q \in (N, 2N]$. We will perform all subsequent computations over the field $\mathbb{F}_q$. For a subset T of $[N]$, define $p_i(T) = \sum_{x \in T} x^i$. We first claim that if $|T| \leq k$, then Alice can recover the elements of T given the values of $p_i(T)$ for $1 \leq i \leq k$.

To show this, define $e_i(T)$ to be the value of ith elementary symmetric polynomial over the elements of T (that is, $e_i(T) = \sum_{T' \subseteq T, |T'| = i} \prod_{x \in T'} x$). Newton's identities allow us to compute (in space $O(\log N)$) the values of $e_i(T)$ for $1 \leq i \leq k$ from the values of $p_i(T)$ for $1 \leq i \leq k$ (since $k < N < q$, all of these operations are valid over $\mathbb{F}_q$). Now, each element x of T is a root of the polynomial

$$x^k - e_1(T)x^{k-1} + e_2(T)x^{k-2} - \cdots + (-1)^k e_k(T) = 0.$$

We can factor this in space $O(\log N)$ by trying all possible $x \in [N]$, and therefore recover the original set T.

To solve the missing k numbers problem, Alice first computes $p_i(S)$ for each $1 \leq i \leq k$ using $O(k \log N)$ space, updating each count every time she receives a new element from S. This allows her to compute $p_i([N] \setminus S)$ via $p_i([N] \setminus S) = p_i([N]) - p_i(S)$. Finally, since $|[N] \setminus S| = k$, she can use the previous algorithm to recover the set $[N] \setminus S$. ◀

▶ **Theorem 18.** *When N is even, there exists a randomized strategy (with access to a uniformly random perfect matching M on K_N) for Alice with space complexity $O(\sqrt{N} \log^2 N)$ which succeeds with probability $1 - O(1/N)$.*

Proof. Consider the following strategy for Alice. Alice begins by sampling $r = \sqrt{N}$ distinct elements $X = \{x_1, x_2, \ldots, x_r\}$ uniformly at random from all $\binom{N}{r}$ subsets of r elements of N (and stores them in memory).

At every point in the game, Alice uses her $O(\sqrt{N} \log^2 N) = O(r \log^2 N)$ bits of memory to keep track of: i) which elements of X have already been said, and ii) the first $k = r \log N$ power sums of the elements said so far, modulo some prime $q \in (N, 2N]$ (as per the proof of Lemma 17).

Alice begins by choosing a random element of X and saying it. Her algorithm in later rounds is defined as follows:

1. If less than or equal to k elements remain unchosen (i.e., this is turn $N - k$ or later), Alice uses the $O(k \log n)$-space algorithm from Lemma 17 to compute the remaining unsaid elements. She then says one of these elements, chosen arbitrarily.
2. Otherwise, if Bob has just said y, Alice checks if $M(y)$ belongs to X.
   - If it does not belong to X, Alice says $M(y)$.

- If $M(y)$ belongs to X but has not been said yet, Alice says $M(y)$ and marks down $M(y)$ as having been said.
- If $M(y)$ belongs to X and has already been said, Alice chooses an element of X uniformly at random from the elements of X which have not yet been said (if there are no such elements, Alice gives up and says a random element). She then marks this element as having been said.

Let T be the random variable denoting the first turn in which all of the elements of X have been said. We first claim that if $T \geq N - k$, then Alice succeeds. Indeed, since the algorithm of Lemma 17 always succeeds, Alice is guaranteed to succeed if she makes it to turn $N - k$. On the other hand, the only way for Alice to fail before turn $N - k$ is if all of the elements of X have already been said, and Alice is forced to say a random element.

We now claim that with high probability, $T \geq N - k$. To see this, note that, similarly as in the proof of Theorem 15, any non-trivial strategy of Bob will give rise to the same distribution over T. This follows from the fact that at any point in the protocol, each unsaid element has the same probability of belonging to the set X.

We can therefore assume that Bob is playing under the strategy where on his turn, he says the smallest number that has not been said so far. Note that under this strategy, Bob will only say numbers less than or equal to i on turn i; it follows that if there exists any x_i such that $x_i \geq N - k$, then also $T \geq N - k$.

We can thus compute

$$\begin{aligned}
\Pr[T < N - k] &\leq \Pr[\forall i,\ x_i \leq N - k] \\
&= \frac{\binom{N-k}{r}}{\binom{N}{r}} \\
&\leq \left(\frac{N-k}{N}\right)^{r} \\
&= \left(1 - \frac{\log N}{\sqrt{N}}\right)^{\sqrt{N}} \\
&= O(1/N).
\end{aligned}$$

It follows that Alice wins with probability at least $1 - O(\frac{1}{N})$, as desired. ◀

Note that the same upper bound holds when Alice does not have an oracle access to a random matching but Bob's strategy is to choose a random number in $[N]$ in every turn.

5 The (a, b)-Mirror Game

Consider the following generalization of the mirror game, where Alice says a numbers each turn, and Bob says b numbers each turn. We call this new game the (a, b)-mirror game.

As with the regular $(1, 1)$-mirror game, mirror strategies exist for the class of $(1, b)$-mirror games (and similarly, the $(a, 1)$-mirror games).

▶ **Theorem 19.** *If N is divisible by $b + 1$, then Bob has a winning strategy computable in $O(\log N)$ memory for the $(1, b)$-mirror game.*

Proof. Divide the N elements of $[N]$ into $N/(b+1)$ consecutive $(b+1)$-tuples. Any time Alice says an element in a $(b+1)$-tuple, Bob says all the remaining elements in that $(b+1)$-tuple. ◀

Unlike in the $(1, 1)$-mirror game, we cannot rule out the existence of low space winning strategies for other games and other choices of N. In this section, we summarize what we know about low-space winning strategies for (a, b)-mirror games.

5.1 $a + b$ is Prime

We begin by considering the set of (a, b)-mirror games where $a + b$ is a prime, p. The Frankl-Wilson Lemma (Lemma 5) allows us to extend some of our proof techniques from Theorem 11 to this case.

In the case where $a = 1$, we have the following analogue to Theorem 11, fully characterizing the N where Bob has a low space winning strategy (Theorem 19) and where Bob requires linear space.

▶ **Theorem 20.** *Let p be a prime. If N is not divisible by p, then any winning strategy for Bob in the $(1, p - 1)$-mirror game requires $\Omega(N)$ space.*

Proof. See Appendix A. ◀

When $a > 1$, our proof techniques no longer allow us completely characterize the set of N where Bob requires $\Omega(N)$ space to win. Instead, we only have the following partial characterization.

▶ **Theorem 21.** *Let p be a prime. If $a + b = p$, and $N \bmod p \in \{a, a + 1, \ldots, p - 1\}$, then any winning strategy for Bob in the (a, b)-mirror game requires $\Omega(N)$ space.*

Proof. See Appendix A. ◀

5.2 $a + b$ is Composite

The failure of the Frankl-Wilson lemma (Lemma 5) to hold modulo composite numbers prevents us from directly applying our proof techniques in this case. Gromulz's construction (Theorem 6) shows that there exist (m, L)-Modtown families containing a superpolynomial in N number of sets when m is composite.

However, a sufficiently small superpolynomial upper bound on the size of such a family would still allow us to prove the analogue of Theorem 20. In fact, any upper bound of F on the size of such a family leads to a lower bound of $\log(2^{\Omega(n)}/F) = \Omega(n) - \log F$ on Bob's memory. This implies the following "contrapositive" to Theorem 20:

▶ **Theorem 22.** *If Bob has a $o(N)$-space winning strategy in a $(1, m - 1)$ game when $N \equiv k \not\equiv 0 \mod m$, then there exists $2^{\Omega(N)}$ sized matching vector families over $\mathbb{Z}_m^N$ (for sufficiently large N).*

Proof. See Appendix A. ◀

6 Open Problems

Our understanding of the space complexity aspects of mirror games is still very rudimentary. We conclude by listing some open problems we find interesting.

1. **When do low-space winning strategies exist?** Does either Alice or Bob ever have a low-space winning strategy for any (a, b) game when a and b are both larger than 1 (e.g., the $(2, 2)$ game)? Are there any cases where the best deterministic winning strategy

has space complexity strictly between $O(\log N)$ and $O(N)$ (e.g. $O(\sqrt{N})$), or does the best winning strategy always fit into one of these two extremes? Is it ever the case that both Alice and Bob simultaneously have $O(\log N)$ winning strategies?

2. **Beating low-space with low-space.** In order to show that Alice needs $\Omega(N)$ space, our adversarial strategy for Bob also requires $\Omega(N)$ space. Given a deterministic low-memory m strategy for Bob, is there a low-memory strategy for Alice which wins against it?
3. **Lower bounds for randomized strategies.** Can we prove any sort of lower bound against randomized strategies that win with high probability? What about against randomized strategies with additional power, such as access to a uniformly chosen matching or with multiple read access to the random tape? Is our upper bound of $\tilde{O}(\sqrt{N})$ tight in the latter contexts?
4. **Composite $a + b$.** Is it possible to show some sort of converse to Theorem 22, that any large enough matching vector family gives rise to a low space strategy for Bob? Or is it possible to show linear space lower bounds for composite $a + b$ via some other approach?
5. **More general mirror games.** One of the great successes of the theory of combinatorial games [2] is that its techniques apply to essentially all sequential two-player games with perfect information. Is there a more general class of games beyond the family of (a, b)-mirror games with similar space complexity phenomena? One possible other family of such games with mirror strategies comes from a generalization of the combinatorial game "Cram", where Alice and Bob take turn placing dominoes (more generally, an element of some set of "symmetric" polyominos) onto a w by h grid (more generally, some high dimensional grid, or any "symmetric" subset of a high-dimensional grid). For example, the original mirror game can be thought of as Cram on a 1 by N board, where players take turn placing unit squares. Can we say anything about the space complexity of playing Cram, or its generalizations?

References

1 Sanjeev Arora and Boaz Barak. *Computational complexity: a modern approach.* Cambridge University Press, 2009.

2 Elwyn R Berlekamp, John Horton Conway, and Richard K Guy. *Winning ways for your mathematical plays*, volume 3. AK Peters Natick, 2003.

3 ER Berlekamp. On subsets with intersections of even cardinality. *Canad. Math. Bull*, 12(4):471–477, 1969.

4 Abhishek Bhowmick, Zeev Dvir, and Shachar Lovett. New bounds for matching vector families. *SIAM Journal on Computing*, 43(5):1654–1683, 2014.

5 Zeev Dvir, Parikshit Gopalan, and Sergey Yekhanin. Matching vector codes. *SIAM Journal on Computing*, 40(4):1154–1178, 2011.

6 Zeev Dvir and Sivakanth Gopi. 2-Server PIR with Subpolynomial Communication. *Journal of the ACM (JACM)*, 63(4):39, 2016.

7 Klim Efremenko. 3-query locally decodable codes of subexponential length. *SIAM Journal on Computing*, 41(6):1694–1703, 2012.

8 Peter Frankl and Richard M. Wilson. Intersection theorems with geometric consequences. *Combinatorica*, 1(4):357–368, 1981.

9 Vince Grolmusz. Superpolynomial size set-systems with restricted intersections mod 6 and explicit Ramsey graphs. *Combinatorica*, 20(1):71–86, 2000.

10 S Muthukrishnan. Data Streams: Algorithms and Applications (Foundations and Trends in Theoretical Computer Science). *Foundations and Trends in Theoretical Computer Science*, 2005.

11 Noam Nisan. On read-once vs. multiple access to randomness in logspace. In *Structure in Complexity Theory Conference, 1990, Proceedings., Fifth Annual*, pages 179–184. IEEE, 1990.

12 Tibor Szabó. http://discretemath.imp.fu-berlin.de/DMII-2011-12/linalgmethod.pdf, 2011.

A Omitted Proofs

Proof of Theorem 20

Theorem 20 is a special case of Theorem 21 setting $a = 1$. Hence, we go on to prove that.

Proof of Theorem 21

Proof. We proceed similarly to the proof of Theorem 11. As before, call a subset S of $[N]$ r-occurring if it is possible that immediately after turn $2r$, the set of numbers that have been said is equal to S. Let $\mathcal{S}_r$ be the collection of all r-occurring sets. As before, $\mathcal{S}_r$ is (p, r)-covering and we can choose r such that $|\mathcal{S}_r| \geq 2^{\Omega(N)}$.

Write $N = pn + k$, $k \in \{a, a + 1, ..., p - 1\}$, and fix a value $r \in [n]$. For a memory state x out of the M possible memory states of Bob and an r-occurring set S, label x with S if it is possible that Bob is at memory state x when the set of numbers that have been said is equal to S. Each state of memory may be labelled with several r-occurring sets or none, but each r-occurring set must exist as a label to some state of memory (by definition). Let $\mathcal{U}_x$ be the collection of r-occurring labels for memory state x.

▶ **Lemma 23.** *If S_1 and S_2 belong to $\mathcal{U}_x$, then* $|S_1 \setminus S_2| \bmod p \in \{k-a+1, ..., k-1, k, .., p-1\}$.

Proof. Let $D = S_1 \Delta S_2$, let $D_1 = S_1 \setminus S_2$, and let $D_2 = S_2 \setminus S_1$. Assume to the contrary that $|D_1| \equiv k_1 \mod p$ where $k_1 \in \{0, 1, .., k - a\}$. Note then that $|D_2| = |D_1|$ (since $|S_1| = |S_2|$) and that $|D| = 2|D_1|$. We'll consider two possible cases for the state of the game after turn $2r$ (each turn, either Alice says a numbers or Bob says b numbers): 1. Bob is at state x, and the set of numbers that have been said is S_1, and 2. Bob is at state x, and the set of numbers that have been said is S_2.

Consider the following strategy that Alice can play in either of these cases: "say the smallest number that has not been said that is not in D". We claim that if Alice uses this strategy, Bob will be the first person (after turn $2r$) to say an element of D. Note that Alice will not say an element of D until turn $2\lfloor \frac{N-|D_1|}{p} \rfloor + 1$; this is since:

1. If we are in case 1, then all of the numbers in D_1 have been said but none of the numbers in D_2 have been said. There are therefore $|D_2| = |D_1|$ numbers in D which have not been said, so Alice can avoid saying an element of D until turn $2\lfloor \frac{N-|D_1|}{p} \rfloor + 1$ as the number of numbers remaining before Alice's turn would be $|D_1| + k - k_1 \geq |D_1| + a$. Hence, Alice can say a numbers not belonging to the set D_1.
2. Likewise, if we are in case 2, the argument proceeds symmetrically.

On the other hand, if no element of D is spoken by either player between turn $2r$ and turn $2\lfloor \frac{N-|D_1|}{p} \rfloor + 1$, then at turn $2\lfloor \frac{N-|D_1|}{p} \rfloor + 2$, Bob has to say b elements from $|D_1| + k - k_1 - a$ elements and hence, there would be intersection with elements of D (as $k - k_1 - a \leq p - 1 - a = b - 1$).

Let y_1 be the memory state of Bob when he first speaks an element of D in case 1, and define y_2 similarly. We claim that $y_1 = y_2$. Indeed, since Bob's strategy is deterministic and starts from x in both cases 1 and 2, Alice's strategy plays identically in both case 1 and case

2 until an element of D has been spoken. It follows that Bob must speak the same element of D in both cases. But if this element is in D_1, and they are in case 1, then this element has already been said before; similarly, if this element is in D_2, and they are in case 2, then this element has also been said before. Regardless of which element in D, Bob speaks at this state, there is some case where he loses, which contradicts the fact that Bob's strategy is successful. It follows that $|D_1| \bmod p \in \{k-a+1, ..., k-1, k, .., p-1\}$, as desired. ◀

▶ **Claim 24.** $|\mathcal{U}_x| \leq p\binom{N}{p-1}$.

Proof. We claim the sets in $\mathcal{U}_x$ form a $(p, \{1, 2, ..., p-1\})$-Modtown, from which this conclusion follows (Lemma 5). Each set in $\mathcal{U}_x$ has cardinality $pr \equiv 0 \mod p$. By Lemma 23, any pair of distinct sets $S_1, S_2 \in \mathcal{U}_x$ has $|S_1 \setminus S_2| \mod p \in \{k-a+1, ..., k-1, k, .., p-1\} \subseteq \{1, 2, .., p-1\}$ (as $k \geq a$). Note that since $|S_1 \cap S_2| = |S_1| - |S_1 \setminus S_2|$, it follows that $|S_1 \cap S_2| \mod p$ also belongs to $\{1, 2, .., p-1\}$. ◀

We established that for some r, $\mathcal{S}_r$ must have cardinality at least $2^{\Omega(N)}$. Since each element in $\mathcal{S}_r$ belongs to at least one collection $\mathcal{U}_x$, and since each collection $\mathcal{U}_x$ has cardinality at most pN^p (Claim 24), the number M of memory states x is at least $2^{\Omega(N)}/(pN^p)$. It follows that $m = \log M \geq \Omega(N)$, as desired. This proves Theorem 21. ◀

Proof of Theorem 22

Proof. In proof of Theorem 21, we use the fact that p is prime only to bound the size of $\mathcal{U}_x$ and thus, Lemma 23 ($a = 1$) hold even for composite m. As $\mathcal{S}_r$ must have cardinality at least $2^{\Omega(N)}$ and there are at most $2^{o(N)}$ memory states x, by pigeonhole principle, there exists x' with $|\mathcal{U}_{x'}| \geq 2^{\Omega(N)}$. Let's define the set of vectors U based on $\mathcal{U}_{x'}$. For every subset $S \in \mathcal{U}_{x'}$, we add a vector $v_S \in \{0,1\}^N$ to U where v_S is the characteristic vector of S ($v_S(l) = 1 \iff l \in S$). Clearly, $|U| = |\mathcal{U}_{x'}|$ and $U \subseteq \{0,1\}^N$. We claim the U and U form a Matching Vector family as $v_S^T v_S = |S| \mod m = 0$, for all S and $v_{S_1}^T v_{S_2} = |S_1 \cap S_2| \not\equiv 0 \mod m$, $\forall S_1 \neq S_2$ by Lemma 23 (when $N \equiv k \not\equiv 0 \mod m$). Hence, we have a exponential sized Matching Vector family (U, U) over $\mathbb{Z}_m^N$. ◀