

Polynomial Data Structure Lower Bounds in the Group Model

Alexander Golovnev
Harvard University
Cambridge, MA, USA
alexgolovnev@gmail.com

Gleb Posobin
Columbia University
New York, NY, USA
posobin@gmail.com

Oded Regev
Courant Institute
New York University
New York, NY, USA

Omri Weinstein
Columbia University
New York, NY, USA
omri@cs.columbia.edu

Abstract—Proving super-logarithmic data structure lower bounds in the static *group model* has been a fundamental challenge in computational geometry since the early 80’s. We prove a polynomial ($n^{\Omega(1)}$) lower bound for an explicit range counting problem of n^3 convex polygons in \mathbb{R}^2 (each with $n^{\tilde{O}(1)}$ facets/semialgebraic-complexity), against linear storage arithmetic data structures in the group model. Our construction and analysis are based on a combination of techniques in Diophantine approximation, pseudorandomness, and compressed sensing—in particular, on the existence and partial derandomization of optimal *binary* compressed sensing matrices in the polynomial sparsity regime ($k = n^{1-\delta}$). As a byproduct, this establishes a (logarithmic) separation between compressed sensing matrices and the stronger RIP property.

Keywords—compressed sensing; computational geometry; data structures; pseudorandomness.

I. INTRODUCTION

Understanding the tradeoff between the query time and storage space of data structures has been a long and active research endeavor for several decades. In the static model, which is the focus of our paper, the goal is to represent a database of n elements using the least amount of space $s \geq n$ (measured in memory words), so that queries $q \in \mathcal{R}$ on the input data can be answered quickly, in time t . The most compelling model for proving such tradeoffs is the *cell-probe* model [2], in which “query time” is measured only by the number t of memory accesses required to retrieve the answer to $q \in \mathcal{R}$ (as in query complexity), whereas all computations on “probed” memory cells are completely free of charge. Time-space lower bounds in this model are therefore a purely information-theoretic question, making *unconditional* lower bounds a viable possibility. Unfortunately, while a simple counting argument [3] shows that almost all static data structure problems with $|\mathcal{R}|$ queries require either near-trivial time $t \geq n^{0.99}$ ($t = n$ is easy by reading all the input database) or near-trivial space $s \geq |\mathcal{R}|^{0.99}$ ($s = |\mathcal{R}|$ is easy by storing all answers), the highest *explicit* cell-probe lower bound

known to date is merely

$$t \geq \Omega\left(\frac{\log(|\mathcal{R}|/n)}{\log(s/n)}\right). \quad (1)$$

In the interesting and realistic regime where $|\mathcal{R}| = \text{poly}(n)$ and the data structure is allowed to use *linear storage space* ($s = O(n)$), (1) yields a $t \geq \Omega(\log n)$ lower bound on the query time of several important data structure problems, such as evaluating hash functions, near-neighbor search, and 2D range counting to mention a few (see [4], [5], [6] and references therein). Proving an unconditional $\omega(\log n)$ lower bound for *any* explicit static problem (with polynomially many queries) against linear-space data structures, remains an outstanding open question in the field.

This gloomy state of affairs remains true even in restricted *arithmetic* models of data structures, which have been studied extensively over the past several decades, primarily in computational geometry and spatial databases [7], [8]. The most well-studied model in this line of work is the *group model* and its more restrictive *semigroup* variant [9], [10], where the canonical problem is that of geometric *range counting*. In the group model, the input database is a set of n points $x_1, \dots, x_n \in \mathbb{R}^d$ (where d is typically a small constant), each associated with a weight $w(x_i) \in G$ from an arbitrary commutative group (e.g., $G = \mathbb{Z}_m$), and the goal is to preprocess these points into small memory (by storing an arbitrary collection of s group elements), so that the weighted sum $\sum_{i: x_i \in r} w(x_i)$ of points in a given query “range” $r \in \mathcal{R}$ can be reported efficiently, where summation is over the underlying group. At query time, the data structure can only manipulate weights through the black-box *addition* and *subtraction* of weights stored in memory, and must work for any choice of the underlying group. The query time is the number of algebraic operations performed (additions/subtractions of memory elements). In particular, multiplication by scalars other than ± 1 is not allowed at query time. Any other computation, e.g., planning algebraic operations based on coordinates, is free of charge. The weaker

All proofs are deferred to the full version of the paper [1]

semigroup model is defined in the same way, with the crucial difference that only *addition* is allowed, but not subtraction.

Some of the most natural and well-studied examples of range searching problems are orthogonal range counting (where ranges \mathcal{R} are axis-parallel boxes), or counting with respect to more complex geometric objects such as balls, halfspaces or simplices. In its most general form, a range counting problem is defined by a family \mathcal{R} of subsets of \mathbb{R}^d , where a common measure of the “complexity” of the problem is the *semialgebraic complexity* of \mathcal{R} [11], i.e., the number of polynomial inequalities defining a range $r \in \mathcal{R}$ (Definition II.3).

Essentially all known range searching data structures can be implemented in the group model¹ (see [7], [8] and references therein). Two classic examples of such data structures for *orthogonal* range counting in d dimensions are *range trees* [12], which solve the problem using $s = O(n \log^d n)$ words of space and query time $t = O(\log^{d-1} n)$, and *Kd-trees* [13], which have linear space $s = O(n)$ at the price of polynomial $t \approx n^{1-1/d}$ query time. The latter upper bound can be achieved for *any* range counting problem where ranges have constant semialgebraic complexity [14].

For general polytopes or simplex range counting problems [8], such polynomial query times are believed to be inevitable unless near-trivial storage is used ($s \approx |\mathcal{R}|$), but this was only proved in rather weak arithmetic models (or pointer machines [15]). In the *semigroup* model, where only additions are allowed, the aforementioned data structures are known to be essentially optimal [10], [16]—in particular, this is the only arithmetic model where *polynomial* lower bounds are known on the query time of static range searching problems. By contrast, in the group model, where both additions and subtractions are allowed, the highest static lower bound (against linear storage) remains $\tilde{\Omega}(\log n)$, for 2D orthogonal range counting [17]. This challenge and disparity between the models was summarized by Pătraşcu [17] as follows:

*“Philosophically speaking, the difference in the type of reasoning behind **semigroup** lower bounds and **group** lower bounds is parallel to the difference between understanding geometry and understanding computation. Since we have been vastly more successful at the former, it should not come as a surprise that progress outside the semigroup model has been extremely slow.”*

¹ Up to $\text{poly}(\log \log n)$ factors, which can be shaved-off using standard word-level parallelism on the RAM.

Indeed, lower bounds in the semigroup model ultimately boil down to arguing that not all query ranges can be “covered” with a small number of subsets of input objects [8], [18]. Unfortunately, no such property holds for the group model, which makes proving lower bounds in the group model much harder.

Our main result is a polynomial lower bound for an explicit range counting problem in the static group model for data structures with linear storage:

Theorem I.1. *There is an explicit set \mathcal{P}_n of $m = n^3$ convex polygons in the plane \mathbb{R}^2 , and a prime p , such that any group-model data structure of size s and query time t for range counting with respect to \mathcal{P}_n (over \mathbb{Z}_p), must have*

$$t^s \geq n^{n/7}.$$

In particular, any linear storage ($s = O(n)$) data structure for \mathcal{P}_n (over \mathbb{Z}_p) must have $n^{\Omega(1)}$ query time. Moreover, the semialgebraic complexity of \mathcal{P}_n , i.e., number of facets of each convex polygon, is $n^{O((\log \log n)^2)}$.

We remark that Theorem I.1 holds even when preprocessing is allowed to depend on the specific group, i.e., for every group and input configuration, the data structure is allowed to store an arbitrary set of s group elements. In contrast, the standard model in range counting literature (as well as all known upper bounds) is *oblivious*: preprocessing uses the group as a black-box. We also note that n -facet polygons arise naturally in computational geometry, for example in the *planar point location* problem [19] (i.e., 2D nearest-neighbor search), where the input itself is a collection of $O(n)$ -facet *disjoint* polygons (Voronoi diagrams).

Lastly, we note that in the stronger *linear model* of data structures, where *multiplication* (by scalars) is allowed as well as addition and subtraction (hence the group is actually a ring), a recent line of results shows that proving super-logarithmic lower bounds for *any* explicit (range counting) problem would have dramatic implications to arithmetic circuit lower bounds and matrix rigidity [20], [21]. In that sense, the group model is the strongest arithmetic model in which polynomial lower bounds fall short of major circuit lower bounds, and Theorem I.1 meets precisely this frontier.

The conceptual message of this paper is that a range counting problem in \mathbb{R}^d is hard for the group model if (a subset of) its ranges \mathcal{R} forms a good pseudorandom generator against affine hyperplanes in $\mathbb{R}^{|\mathcal{R}|}$. This will be explained in the following section. While this may not be the *only* source of hardness of range counting in the group model, our work provides a new technique for analyzing such problems and the first substantial step

toward proving (exponentially higher) lower bounds on more natural range counting problems, i.e., with lower semialgebraic complexity: If such pseudorandom-generators (PRGs) can be realized as (indicators of) “simple” geometric ranges in reasonably low dimension, then any linear-storage arithmetic data structure should have high query time for the underlying range counting problem. In Section IV we formalize this by defining the notion of *Geometric PRGs*, which exploit the “geometry” of the d -dimensional input seed in a simple way. We prove the existence of such PRGs with nontrivial parameters, as a step toward polynomial lower bounds for *halfspace* range-counting in very high yet nontrivial dimension. We believe the concept of Geometric PRGs may have further applications in computational geometry and is of independent interest.

Techniques: Our proof introduces a new way to analyze arithmetic data structures (excluding multiplications), by combining ideas from Diophantine approximation, compressed sensing, and pseudorandomness (specifically, derandomization of anti-concentration theorems). Our results and analysis of “Geometric PRGs” (Theorem IV.7) rely on Fourier-analysis and hypercontractivity to establish pseudorandom properties of halfspace range counting. In the next two sections, we elaborate on these components and the role they play in our proof.

A. Binary compressed sensing and sparse recovery

A key ingredient in the proof of Theorem I.1 is the construction of *binary* matrices $M \in \{\pm 1\}^{m \times n}$ with $m = n^c$ and $c > 1$ where every k rows are linearly independent over \mathbb{R} (matrices with such property are commonly known as *compressed sensing*² matrices [22], [23]). If M is allowed to have $\text{poly}(n)$ -bit entries, the optimal value of $k = n$ is easily achievable, for example by taking an $n^c \times n$ Vandermonde matrix. However, with binary (or any constant-size) entries, which is crucially the case for our application, the answer is not clear [24]. One way to construct such explicit binary matrices over the finite field \mathbb{F}_2 (which is harder than over \mathbb{R}) is to take the parity-check matrix of rate-optimal binary error correcting codes (e.g., expander codes [25], [23], see the full version of the paper for more details), but unfortunately when the number of rows is polynomial ($m > n^{1+\delta}$) this approach only yields $k = O(n/\log n)$ -wise independence, and this is well known to be tight [26] (in fact the latter shows $k = O(n/\log_q n)$ for

²Indeed, it is not hard to see that if the rows of $M \in \mathbb{R}^{m \times n}$ are k -wise linearly independent, then M^\top is a compressed sensing matrix for $(k/2)$ -sparse m -dimensional vectors, with n “measurements”, hence the problem we study is equivalent.

any constant size field \mathbb{F}_q). Other binary constructions in the compressed-sensing and LDPC literature (e.g., [24], [27], [28]) achieve even worse parameters. As explained in Section I-C, for the proof of Theorem I.1 it is crucial to obtain the optimal value of $k = \Omega(n)$ (even settling for a much weaker $\text{poly} \log(n)$ query lower bound would still require $k = \omega(n/\log n)$).

This raises a more basic question: *Do optimal ($k = \Omega(n)$) binary compressed-sensing matrices over \mathbb{R} even exist?* Our first step is showing that a *random binary* $\text{poly}(n) \times n$ matrix is indeed $\Omega(n)$ -wise independent over \mathbb{R} (and in fact, with some more effort, also over any large enough finite field):

Theorem I.2 (Optimal Binary Compressed Sensing Matrices). *For every $c \geq 1$, and every large enough n , a uniformly random binary matrix $M \in \{\pm 1\}^{n^c \times n}$ is $\Omega(n/c)$ -wise independent over \mathbb{R} w.h.p. Moreover, such binary matrix M can be constructed using only $O(cn \log n (\log \log n)^2)$ random bits (instead of n^{c+1}).*

Note that a naïve union bound will not work here: The probability that a fixed set of k rows is linearly independent is, at best, 2^{-n} (the probability 2 rows are equal), which is not small enough to union bound over all $\binom{n^c}{k} = n^{\Theta(n)}$ choices of k -subsets for $k = \Omega(n)$. Hence, the proof requires a different analysis that somehow exploits the fact that we are working over fields of large characteristic, while overcoming the challenge that there is an unbounded number of linear combinations over \mathbb{R} .

A key observation is that our analysis for random matrices turns out to carry over even to certain *pseudo-random* binary matrices. This is one of the main insights of this paper, which leads to the partial derandomization ($O(cn \log n (\log \log n)^2)$ random bits) in Theorem I.2. While a fully explicit construction of optimal compressed sensing matrices with constant-size entries remains an intriguing question in the context of sparse recovery (see the next paragraph), the construction in Theorem I.2 turns out to be enough for obtaining a *fully explicit* data structure lower bound for range counting problems with reasonably low ($n^{\tilde{O}(1)}$) semialgebraic complexity. We elaborate on this step in Section I-C below.

Implications to Sparse Recovery and separation from RIP: An interesting consequence of Theorem I.2 is that it provides a (tight) separation between binary *compressed sensing* matrices in which any k rows are linearly independent, and the stronger *restricted isometry property* (RIP), which instead requires any k rows to be “nearly orthogonal”. Indeed, in the polynomial

sparsity regime $k = n^{1-\delta}$, it is well known that *any* RIP matrix must have $\Theta(k \log k)$ rows (and this is tight for random binary matrices, e.g. [28], [29]). By contrast, Theorem I.2 asserts that binary compressed sensing matrices only require $O(k)$ linear measurements. To the best of our knowledge, previous binary constructions, even non-explicit ones, were only guaranteed to work with $O(k \log k)$ measurements (see [24] and references therein).

B. Related Work in the Group Model

A sequence of papers initiated by Chazelle in the early 90's proved essentially tight lower bounds in the *semigroup* model, for orthogonal and halfspace range counting in any constant dimension [10], [30]. Similar bounds were shown in the *pointer machine* model [15]. In the *offline* group model, where no preprocessing is allowed, Chazelle [31] proved an $\Omega(n \log n)$ lower bound for halfspaces. No super-logarithmic data structure lower bounds were known in the static (i.e., online) setting.

In the *dynamic* group model [9], where the data structure needs to support insertions and deletions of points in \mathbb{R}^d , *polynomial* lower bounds ($\Omega(n^{1-1/d})$) have been proved only recently in a breakthrough result of Larsen [32] using combinatorial discrepancy arguments, but as the author notes himself, this technique does not apply to the static case.³ We remark that amortized dynamic lower bounds for “decomposable” data-structure problems⁴, such as range counting, also imply (up to a logarithmic loss) lower bounds for static data structures *with efficient preprocessing time* $p(n) \approx s$ [33]. Alas, in the group model this is a very severe restriction, since storing an arbitrary group element generally requires $O(|G|)$ group operations, which in the dynamic setting is prohibitive and trivialized the aforementioned reduction. For a broader overview of arithmetic lower bounds we refer the reader to [8, Section 3.4] and [32].

C. Technical Overview

We now provide a streamlined overview of our main result, Theorem I.1. Recall that a data structure in the group model must solve the given problem with respect to *any* underlying group. We will carefully design a range counting problem such that no (group model) data

³The discrepancy argument relies on the fact that dynamic arithmetic data structures induce a factorization of the query matrix into a product of two sparse matrices, which is not the case in the static setting.

⁴A data structure problem \mathcal{P} is *decomposable* if the answer to a query on the union of two input databases X, Y can be computed as a black-box from the marginal answers $\mathcal{P}(A, B) = f(\mathcal{P}(A), \mathcal{P}(B))$.

structure with space s and query time t can solve our problem even in the additive group of integers modulo p , \mathbb{Z}_p , for an explicit prime $p = \Theta(t)^s$.

Step one. Diophantine approximations: Let us fix the n input points $\mathbf{x} \in (\mathbb{R}^d)^n$ to a range counting problem, and the weights of these input points $\mathbf{w} \in \mathbb{Z}_p^n$. For this input, the data structure computes and stores $S \in \mathbb{Z}_p^s$, a set of s elements of the group. Later, for every query range $r \in \mathcal{R}$, the data structure will output a sum of at most t stored elements (or their negations). The first step in our lower bound uses a basic fact in Diophantine approximation (known as *Dirichlet's simultaneous approximation theorem*, which is an easy application of the pigeonhole principle), which implies that for $p := \Theta(t)^s$ and every set $|S| = s$ of integers from $\{0, \dots, p-1\}$, there is some $1 \leq q < p$ such that after multiplication by q , *all* those s numbers become smaller than $p/(4t)$ modulo p . This of course means that even t -sums of these numbers (and their negations) remain small modulo p (this is the only but crucial place where we exploit the restriction that only ± 1 coefficients are allowed in the group model at query time). Therefore, every arithmetic data structure with only s memory cells, on *every* set of input weights from \mathbb{Z}_p^n , must output a set which is small modulo p after multiplication by some $q = q(S) < p$.

The obvious next step is to construct an explicit range counting problem \mathcal{R} which does *not* have this property, i.e., for some configuration of the n inputs points in \mathbb{R}^d , there exists an assignment of weights $(w_1, \dots, w_n) \in \mathbb{Z}_p^n$, s.t. the output set *does not* become simultaneously too small after multiplication by *any* $1 \leq q < p$. We formalize this with the notion of *diversity*: A set of m elements $\mathbf{y} \in \mathbb{Z}_p^m$ (corresponding to m range counting queries) is *diverse*, if for every $1 \leq q < p$ there is some element (query answer) $i \in [m]$ such that $y_i \cdot q \in [\frac{p}{4}, \frac{3p}{4}] \pmod p$. Note that, since we aim for a polynomial ($t > n^\epsilon$) lower bound on query time, and trivially $s \geq n$, Dirichlet's theorem dictates that $p = \Omega(t)^s = 2^{\Omega(n \log n)}$. This fact will soon be important. We also remark that constructing an arbitrary diverse vector explicitly is actually quite easy⁵— the challenge is that this vector must correspond to an output of a *range counting* problem, and indeed this is our next goal.

⁵Let $p = 2^{0.5n \log n} - 1$, $m = 0.5n \log n$, and $\mathbf{y} \in \mathbb{Z}_p^m$ be a vector with $y_i = 2^i$ for $0 \leq i < m$. It is easy to see that \mathbf{y} is diverse (indeed, while multiplying q by all powers of 2, we must hit the interval $[p/4, 3p/4]$). Alas, \mathbf{y} is not an output of any *group model* problem, as generating it requires summing input weights with *multiplicities*.

Step two. Diversity from k -wise independence: For a range counting problem specified by a family of m ranges \mathcal{R} , and a set of input points $\mathbf{x} = (x_1, \dots, x_n) \in (\mathbb{R}^d)^n$, consider the *binary incidence matrix* $\mathcal{I}_{\mathcal{R}}(\mathbf{x}) \in \{0, 1\}^{m \times n}$ whose (i, j) th entry indicates whether point x_j lies in the i th range $r_i \in \mathcal{R}$. Note that for the weights $\mathbf{w} \in \mathbb{Z}_p^n$, the desired outputs to the m range queries are $\mathcal{I}_{\mathcal{R}}(\mathbf{x}) \cdot \mathbf{w} \in \mathbb{Z}_p^m$. Now our goal is to find a set of points \mathbf{x} , weights \mathbf{w} , and ranges \mathcal{R} for which $\mathcal{I}_{\mathcal{R}}(\mathbf{x}) \cdot \mathbf{w}$ is a diverse vector. Note that only the set of ranges \mathcal{R} needs to be explicit. The set of weights \mathbf{w} and the set of input points \mathbf{x} do not need to be explicit, because the data structure must solve the problem for any choice of input weights $\mathbf{w} \in \mathbb{Z}_p^n$ and points $\mathbf{x} \in (\mathbb{R}^d)^n$. We observe that *if for some input \mathbf{x}* , the rows of $\mathcal{I}_{\mathcal{R}}(\mathbf{x})$ are *k -wise independent* over the field \mathbb{F}_p for a k satisfying $(m/k)^{k/2} > p$, then there exist weights $\tilde{\mathbf{w}} \in \mathbb{Z}_p^n$ such that the answer vector $\mathcal{I}_{\mathcal{R}}(\mathbf{x}) \cdot \tilde{\mathbf{w}} \in \mathbb{Z}_p^m$ is diverse. This is an easy consequence of Chernoff bounds for k -wise independent random variables [34]. Hence, our problem boils down to designing a k -wise independent *binary matrix* over a large field \mathbb{F}_p , where k must satisfy $(m/k)^{k/2} > p = 2^{\Theta(n \log n)}$. With polynomially many queries ($m = |\mathcal{R}| = n^{O(1)}$), this condition requires that $k = \Omega(n)$ in order to get any useful lower bound⁶, which is as high as k can possibly be.

Step three. Binary compressed sensing matrices: As discussed in Section I-A, such optimal (“compressed sensing”) matrices with constant-size entries were not even known to *exist* in the polynomial sparsity regime $m = \text{poly}(k)$, and indeed the next step of our proof is showing that a *random* $n^c \times n$ binary matrix M is $\Omega(n)$ -wise independent with high probability (the first premise of Theorem I.2). The intuition for why random matrices work *over* \mathbb{R} is as follows: For any fixed nonzero linear combination of some fixed $k \approx n$ rows of M , the probability that each coordinate $i \in [n]$ of this linear combination is equal to 0, is $\approx 1/\sqrt{n}$, by the Littlewood-Offord Lemma [35]. (Note that this exploits the fact that we are working over a field of large characteristic.) Since columns are independent, the overall probability that this linear combination is the all-0 vector (and hence linearly dependent) is $\sim (\frac{1}{\sqrt{n}})^n = 2^{-\Theta(n \log n)}$, which is enough to union-bound over all $\binom{n^c}{k} = 2^{\Theta(n \log n)}$ subsets of k rows. Alas, there are two substantial flaws in this argument: Most importantly, it does not rule out *sparse* linear combi-

⁶Even settling for exponentially weaker $t > \text{poly} \log(n)$ lower bounds using this approach would still require $p > 2^{O(n \log \log n)}$, i.e., $k > \Omega(n \log \log n / \log n)$, which is already beyond known compressed sensing constructions with constant-size alphabets.

nations of rows—indeed, the probability that *two* rows are identical is 2^{-n} , which already dooms the entire claim. Fortunately, there are only $O(n^{2^c})$ such pairs, so intuitively a more careful union bound should work. Indeed, using a certain *chaining argument*, we can group the linear combinations by sparsity and handle them separately using some case analysis. The second flaw is that we “forgot” to union-bound over all possible linear combinations in \mathbb{R} as well, which appears daunting. We circumvent this hurdle using an approach (dating back to Komlós [36]) which exploits the observation that linear dependency of rows can be “charged” to the existence of a (minimal) square sub-matrix which is not full rank, making the union bound possible.

Step four. Derandomization: While this argument only shows that a random incidence matrix works (which implies a lower bound for a “trivial” range-counting problem⁷), the key observation of our proof is that the above analysis works even for certain *pseudo-random* matrices. Indeed, essentially all the above proof uses about M is that (i) columns are statistically independent, and (ii) that weighted sums of coordinates in each column are *anti-concentrated* for any choice of real-valued weights, i.e., each column ε -fools *affine hyperplanes* in \mathbb{R}^m (see Definition III.9), for $\varepsilon = n^{-0.1}$ (say). Fortunately, explicit pseudo-random generators (PRGs) against affine halfspaces (and therefore hyperplanes, see Proposition III.11) are known, with almost optimal seed-length $r = \tilde{O}(\log(m/\varepsilon)) = \tilde{O}(\log n)$ random bits [37]. As such, we can derandomize M by choosing each of the n columns independently using this PRG $g(s_i)$ with independent seeds, one per column. In other words, we can generate this \tilde{M} by sampling n columns from a *fully explicit* incidence matrix supported on only $2^r = n^{\tilde{O}(1)}$ columns. The “ranges” corresponding to this matrix are defined by its rows, i.e., $\mathcal{R} = \{g_i\}_{i=1}^m$, where g_i is the i th output bit of the PRG. Our theorem now guarantees that \tilde{M} contains an induced sub-matrix $\mathcal{I}_{\mathcal{R}} \subset \tilde{M}$ on n columns, which is $\Omega(n)$ -wise independent. Recall that these n columns need *not* be explicit, so long as they are sampled from a small enough superset (which is indeed the case for \tilde{M}).

Step five. Embedding into \mathbb{R}^2 : The final step of the proof is a simple geometric embedding of the ranges \mathcal{R} into the plane. This is done by laying out $2^r = n^{\tilde{O}(1)}$ points on a circle in \mathbb{R}^2 , each corresponding to a column of \tilde{M} , and observing that for any $g_i \in \mathcal{R}$, the subset of columns corresponding to inputs x for which $g_i(x) = 1$,

⁷I.e., when ranges have semi-algebraic complexity 2^m , which is the support-size of the columns of a random $m \times n$ matrix.




Figure 1. The family of $m = n^3$ convex polygons $\mathcal{R} = \{g_i\}_{i=1}^m$ is generated by mapping each of the $2^r = n^{O(\log \log n^2)}$ columns of the incidence matrix \tilde{M} to a point $p(j)$ on a circle in \mathbb{R}^2 . For each $i \in [m]$, g_i is the convex hull of all the points (columns) $p(j)$ on which g_i outputs 1: $\{p_j \in [2^r] : g_i(s_j) = +1\}$, where g_i is the i 'th output bit of the [37] PRG and $s_j \in \{0, 1\}^r$ is the input seed corresponding to the j th column of \tilde{M} . The figure illustrates this for $m = 3$ polygons. Black points on the circle denote the (unknown) n input points—the submatrix $\mathcal{I}_{\mathcal{R}}(\mathbf{x}) \subset \tilde{M}$ —for which the range counting problem is guaranteed to be hard.

can be trivially encoded as a convex polygon with at most 2^r facets, hence the semi-algebraic complexity of \mathcal{R} is $n^{\tilde{O}(1)}$, which completes the proof (see Figure 1).

Geometric PRGs: The proof of Theorem I.1 uses PRGs in a black-box manner – Since PRGs are agnostic to the representation of the input seed (i.e., treat it as a random bit-string with no underlying geometry), the semi-algebraic complexity of our range counting problem in Figure 1 (=number of facets of polygons) can only be as low as 2^r where r is the minimal seed length of the PRG, which for affine halfspaces in \mathbb{R}^m must be $r \geq \Omega(\log m)$ [38]. A natural question toward lower bounds for more natural range counting problems (i.e., with lower semi-algebraic complexity), is whether our lower bound approach can exploit the *dimensionality* of the input space, to decrease the complexity (number of polynomial equations) defining the ranges \mathcal{R} of the underlying range counting problem. Motivated by this question, we define the notion of *Geometric PRGs*: Here, the input “seed” is represented as a (random) point in a d -dimensional grid $[B]^d \subset \mathbb{R}^d$ as opposed to the standard bit-string representation, and each of the output bits g_i of the PRG is computed as a *low-degree d -variate polynomial threshold function (PTF)*. Such functions generalize (the incidence function of) geometric ranges such as d -dimensional *halfspaces*, which are among the most natural and well-studied range counting problems.

In contrast to standard PRGs, the mere existence of Geometric PRGs (in sublinear dimension $d \ll n$)

is a nontrivial question. We establish the existence of geometric PRGs in nontrivial ($d = n^\epsilon$) dimension and polynomial ($\sim n^\epsilon$) degree against affine halfspaces, and make a significant next step en-route to polynomial lower bounds in the group model for halfspace range counting (*degree-1* PTFs): A well known result of Diakonikolas et al. [39] asserts that *any* k -wise (statistically) independent distribution over $\{\pm 1\}^m$ $\tilde{\Omega}(1/\sqrt{k})$ -fools halfspaces in \mathbb{R}^m . Our main technical result is showing that, for a typical collection \mathcal{H} of $m = \text{poly}(d)$ hyperplanes in \mathbb{R}^d , the *signed inner-products* of \mathcal{H} with a random hyperplane $\mathbf{r} \in_{\mathbb{R}} \{\pm 1\}^d$ (the “seed”), forms a distribution on $\{\pm 1\}^m$ which is *almost k -wise independent* for any $k \leq d^{1/6}$, i.e. $\|\{\text{sign}\langle h_i, \mathbf{r} \rangle\}_{i \in [k]} - \mathcal{U}_d\|_1 = 1/\text{poly}(d)$ for any subset $\mathcal{H}' = \{h_1, \dots, h_k\} \subset \mathcal{H}$. This statement is essentially tight. We prove it by adapting a theorem due to Klartag and Regev [40] about random projections in *Gaussian* spaces to the setting of caps in the hypercube, using Fourier-analytic techniques. The main difficulty is that in Boolean space, the hypercontractivity argument of [40] is not enough to bound the desired measure of intersection, hence we use a convolution with the *majority* function and a “mirroring” trick to bound this measure (see Theorem IV.7). We leave open the question of finding an *explicit* construction of low-degree geometric PRGs, which could lead to group-model lower bounds for more natural range searching problems using our approach.

II. PRELIMINARIES

We identify the finite field \mathbb{F}_p with the set $\{0, \dots, p-1\}$. For a non-negative integer m , we denote by $[m]$ the set $\{1, \dots, m\}$. For a real number x , we denote by $\|x\|$ the distance from x to the nearest integer. All logarithms are to the base 2. We write $|b| < c \pmod p$ to mean that b is c -close to an integer divisible by p , i.e., that there exists an integer k such that $|b - kp| < c$, or in other words $\|b/p\| < c/p$. All logarithms in the paper are base 2 unless otherwise stated, and we use the shorthand $\log x = \log_2 x$.

Definition II.1 (*k -wise independent matrices*). A matrix $M \in \mathbb{F}^{m \times n}$ is k -wise independent if every set of k rows of M is linearly independent (over \mathbb{F}).

A. Range counting in the group model

Definition II.2 (Range counting in group model). Let $\mathcal{R} = \{\mathcal{R}_i \subset \mathbb{R}^d, i \in [m]\}$ be a range counting problem in \mathbb{R}^d . A data structure $\mathcal{D} = (P : (\mathbb{R}^d)^n \times G^n \rightarrow G^s, Q : G^s \rightarrow G^m)$ is said to be an (s, t) data structure in the group model for \mathcal{R} over an abelian group G , if:

- 1) For any configuration of n points x_1, \dots, x_n in \mathbb{R}^d , any tuple of n weights $\mathbf{w} = (w_1, \dots, w_n) \in G^n$, and any index $i \in [m]$, it holds that $Q_i \circ P((x_k)_{k=1}^n, \mathbf{w}) = \sum_{j: x_j \in \mathcal{R}_i} w_j$. That is, the answer to the i -th query is the sum of the weights of the points that fall into the i -th range.
- 2) Each query function Q_i computes a group sum of at most t memory elements or their negations, possibly with repetitions: $Q_i(g_1, \dots, g_s) = \sum_{j \leq t} \xi_j g_{i_j}$, with $\xi_j \in \{\pm 1\}$. Note that the function Q_i can be adaptive (i.e., a depth- t decision-tree over the s memory elements).

While this arithmetic model seems rather restrictive, all known data structures for range searching problems fall into this setting⁸. We remark that here we allow the preprocessing function P to depend on the specific group, i.e., for every group and every input $\mathbf{x} \in \mathbb{R}^d$, the data structure is allowed to compute and store an arbitrary set of s group elements. By contrast, in the standard “oblivious” definition of the group model [8], [10], as well as in all known upper bounds, the preprocessing function treats the group as a black-box: the s group elements are computed using a fixed set of group operations.

We use the common notion of semialgebraic sets to quantify the complexity of a range counting problem:

Definition II.3 (Semialgebraic Sets [14], [8]). A semialgebraic set is a subset of \mathbb{R}^d obtained from a finite number of sets of the form $\{\mathbf{x} \in \mathbb{R}^d \mid p(\mathbf{x}) \geq 0\}$, where p is a d -variate polynomial with integer coefficients, using arbitrary intersections, unions, and complementations. $\Gamma_{d, \Delta, s}$ denotes the collection of all the semialgebraic sets in \mathbb{R}^d obtained from at most s d -variate polynomials of total degree at most Δ each.

For brevity, when we say that a range counting problem $\mathcal{R} \subseteq \mathbb{R}^d$ has semialgebraic complexity K we mean that the number of linear inequalities defining each range $r \in \mathcal{R}$ is at most K . For example, a K -facet polytope in \mathbb{R}^d has semialgebraic complexity K under this convention.

The ranges we construct in Theorem I.1 are (convex) polygons in two dimensions with at most $n^{(\log \log n)^2}$ sides, so they will lie in $\Gamma_{2, 1, n^{(\log \log n)^2}}$.

B. Anti-concentration inequalities

We will use the following versions of the Littlewood-Offord Lemma.

⁸See Footnote 1.

Lemma II.4 ([35]). Suppose that $x_1, \dots, x_k \in \mathbb{R} \setminus \{0\}$, $y \in \mathbb{R}$. Then

$$\left| \left\{ \varepsilon \in \{-1, 1\}^k \mid \sum_{i=1}^k \varepsilon_i x_i = y \right\} \right| \leq \binom{k}{\lfloor k/2 \rfloor} \leq \frac{2^k}{\sqrt{k}}.$$

Lemma II.5 ([41, Theorem 6.3], [42, Theorem A.21]). Let \mathbb{F}_p be a prime field of order $p > 2$. Suppose that $x_1, \dots, x_k \in \mathbb{F}_p \setminus \{0\}$ and $y \in \mathbb{F}_p$. Then

$$\left| \left\{ \varepsilon \in \{-1, 1\}^k \mid \sum_{i=1}^k \varepsilon_i x_i = y \right\} \right| \leq 2^k \left(\frac{1}{p} + \sqrt{\frac{8}{k}} \right).$$

III. PROOF OF THEOREM I.1

A. Efficient Data Structures are Not Diverse

Let us fix the input points $(x_1, \dots, x_n) \in (\mathbb{R}^d)^n$, and vary their weights. Let us look at the possible answers that the composition $Q \circ P((x_1, \dots, x_n), \cdot)$ can produce. We will show that $Q \circ P$ always has structure (the resulting vectors are not “diverse”), while some range counting problems do not have this structure (are “diverse”), thus, cannot be computed by efficient data structures in the group model.

Definition III.1 (Diversity). We call a set $S \subset \mathbb{Z}_p$ diverse if for any $q \in \{1, \dots, p-1\}$, $q \cdot S \not\subseteq [-p/4, p/4] \pmod p$. Similarly, we say that a vector $\mathbf{w} \in \mathbb{Z}_p^m$ is diverse if the set of its coordinates is diverse.

Lemma III.2 (Outputs of an efficient DS are not diverse). For any $(4t)^s < p$, any (s, t) data structure $\mathcal{D} = (P : ([n]^d)^n \times \mathbb{Z}_p^n \rightarrow \mathbb{Z}_p^s, Q : \mathbb{Z}_p^s \rightarrow \mathbb{Z}_p^m)$ in the group model, any points $x_1, \dots, x_n \in [n]^d$, and any weights $w_1, \dots, w_n \in G$, the set of answers to queries

$$\{Q_j \circ P((x_i)_{i=1}^n, \mathbf{w})\}_{j=1}^m$$

is not diverse.

To compute each output coordinate, Q only adds and subtracts t of the memory cells, counting repetitions, so it is enough to show that there is a $1 \leq q \leq (4t)^s$ such that each of the memory cells is at most $p/(4t)$ modulo p after multiplication by q . The following lemma shows that such a q exists for any set of s integers.

Lemma III.3 (Dirichlet’s Simultaneous Approximation Theorem). For any set of s integer numbers x_1, \dots, x_s , a natural number $Q \geq 1$, and a natural number $p > 1$, there exists an integer $1 \leq q \leq Q^s$ such that

$$\forall i \in [s], |x_i \cdot q| < p/Q \pmod p.$$

Or, in greater generality: given s real numbers $\mathbf{r} = (r_1, \dots, r_s)$ and an integer Q , there is an integer $1 \leq$

$q \leq Q^s$ such that for every i , $\|q \cdot r_i\| < 1/Q$, i.e. all the given real numbers can be approximated by rational fractions k_i/q with some denominator $1 \leq q \leq Q^s$ up to an error strictly less than $1/(qQ)$.

B. k -wise Independence Implies Diversity

We have shown above that any set of outputs from an (s, t) data structure for range counting in the group model is not diverse if $(4s)^t < p$. Now we show that it is possible to construct a matrix $M \in \mathbb{Z}_p^{m \times n}$, such that if it is an incidence matrix of m shapes and n points, then there is a way to assign weights w to those points so that the answers Mw to the m queries form a diverse set.

The idea is that if w is uniformly random, then for a fixed $1 \leq q < p$ the probability of $q \cdot (Mw)_i$ being at most $p/4$ modulo p is $1/2$, and if the coordinates of Mw were independent, the probability of all of $q \cdot (Mw)_i$ being at most $p/4$ would be 2^{-m} . Of course, if $n < m$, some coordinates of Mw will be linearly, and thus statistically, dependent, but we can make do with a weaker requirement: that the coordinates are k -wise independent. Then the probability that all $q \cdot (Mw)_i$ are at most $p/4$ modulo p can be bounded by $1/p$ by a version of a Chernoff bound for k -wise independent variables. After applying union bound over all possible $q \in [1, p-1]$, this gives the existence of w that produces a diverse Mw . And the k -wise independence of the coordinates of Mw follows from the k -wise independence of the rows of the matrix M if it is viewed as a matrix in $\mathbb{F}_p^{m \times n}$, so then we are left with the problem of finding a binary matrix M with k -wise linearly independent rows over \mathbb{F}_p and a set of shapes and points for which M is an incidence matrix.

We will use the following form of the Chernoff bound for k -wise independent variables.

Theorem III.4 ([34]). *Let $k \geq 4$ be an even integer. Suppose that X_1, \dots, X_n are k -wise independent random variables taking values in $[0, 1]$. Let $X = X_1 + \dots + X_n$, and $\mu = E[X]$, and let $A > 0$. Then*

$$\Pr[|X - \mu| > A] < 2 \left(\frac{nk}{A^2} \right)^{k/2}.$$

Theorem III.5 (k -wise independence implies diverse weights). *Let $M \in \mathbb{F}_p^{m \times n}$ be k -wise independent and $p \leq 0.5 \left(\frac{m}{16k} \right)^{k/2}$, then there exists $w \in \mathbb{F}_p^n$ such that Mw is diverse.*

C. Existence of k -wise Independent Binary Matrices

In this section, we prove that almost all matrices from $\{\pm 1\}^{n^c \times n}$ are $\Omega(n)$ -wise independent (for every $c \geq 1$).

Theorem III.6. *For every $\varepsilon \in (0, 1/2]$, every large enough $n \leq m \leq n^{\frac{1}{2\varepsilon}}$, every prime field \mathbb{F}_p of order $p \geq \sqrt{n}$, a uniformly random matrix $M \in \{-1, 1\}^{m \times n}$ is $k := \varepsilon n$ -wise independent over \mathbb{F}_p with probability $1 - o(1)$.*

We remark that k -wise independence of M over \mathbb{F}_p implies k -wise independence over \mathbb{R} . Indeed, k -wise independence of M implies that any k rows of M contain a non-singular $k \times k$ submatrix over \mathbb{F}_p , this submatrix is also non-singular over \mathbb{R} , which in turn implies k -wise independence over \mathbb{R} .

In the proof of Theorem III.6 we will use the following claim.

Claim III.7. *Any k -dimensional subspace $L \subset \mathbb{F}_p^n$ contains at most 2^k vertices of the hypercube $\{\pm 1\}^n$.*

Now we are ready to prove the main result of this section.

Now we observe that the result of Theorem III.6 holds for $\{0, 1\}$ -matrices as well. While we cannot just take a $\{\pm 1\}$ matrix from Theorem III.6, increment each entry by one, and divide by two⁹, we can perform a similar operation while preserving k -wise independence.

Corollary III.8. *For every $\varepsilon \in (0, 1/2]$, every large enough $n \leq m \leq n^{\frac{1}{2\varepsilon}}$, every prime field \mathbb{F}_p of order $p \geq \sqrt{n}$, a uniformly random matrix $M \in \{0, 1\}^{(m-1) \times n}$ is $(k-1) := (\varepsilon n - 1)$ -wise independent over \mathbb{F}_p and \mathbb{R} with probability $1 - o(1)$.*

D. Derandomization

In this section we reduce the randomness used in the proof above from mn bits to $O(n \log m (\log \log m)^2)$ bits: in effect instead of sampling the columns from the set of all the possible columns in $\{\pm 1\}^m$, we show that we can limit ourselves to a much smaller subset and still get a very high probability of sampling a k -wise independent matrix. This will dramatically (exponentially) reduce the semi-algebraic complexity of the resulting range-counting problem in our final construction.

As discussed in the introduction, the main observation here is that the analysis of Theorem III.6 for random binary matrices, also carries over to the case where each column is chosen independently according to a distribution on $\{\pm 1\}^m$ which fools any affine hyperplane in \mathbb{R}^m . We now formalize this statement.

Definition III.9. *We say that a deterministic function $F : \{0, 1\}^r \rightarrow \{\pm 1\}^m$ ε -fools a family of sets $\mathcal{S} \subset$*

⁹for example, this could create an all zeros row, and the resulting matrix would not be k -wise independent

$2^{\{\pm 1\}^m}$, if for any $s \in \mathcal{S}$:

$$\left| \Pr_{\mathbf{x} \leftarrow \text{Unif}\{\pm 1\}^m}[\mathbf{x} \in s] - \Pr_{\mathbf{x} \leftarrow F}[\mathbf{x} \in s] \right| \leq \varepsilon,$$

where the second \mathbf{x} is the output of F when it is passed a uniformly random input string from $\{0, 1\}^r$ as an input. We also call such an F an ε -PRG against \mathcal{S} .

Our derandomization will use the following pseudo-random generator (PRG) of Gopalan, Kane and Meka [37], which ε -fools affine halfspaces with almost optimal seed-length.

Lemma III.10 ([37, Corollary 1.2]). *For every $\varepsilon > 0$ there exists a PRG $F_{\text{GKM}} : \{0, 1\}^r \rightarrow \{\pm 1\}^m$ that uses $r = O(\log(m/\varepsilon) \cdot (\log \log(m/\varepsilon))^2)$ uniformly random input bits and ε -fools affine halfspaces, i.e., the family of sets*

$$\{\mathbf{x} \in \{\pm 1\}^m \mid \langle \mathbf{x}, \alpha \rangle \leq \theta\}_{\alpha \in \mathbb{R}^m, \theta \in \mathbb{R}}.$$

Proposition III.11 (Fooling halfspaces \Rightarrow Fooling hyperplanes). *If F is an ε -PRG against affine halfspaces in \mathbb{R}^m , then it is a 2ε -PRG against hyperplanes in \mathbb{R}^m .*

Using this proposition, along with the GKM PRG for affine hyperplanes, we shall prove the following:

Theorem III.12. *For every $\varepsilon \in (0, 1/2]$, every large enough $n \leq m \leq n^{\frac{1}{2\varepsilon}}$, there is a polynomial-time algorithm that uses $O(n \log m (\log \log m)^2)$ random bits and outputs a matrix $W \in \{\pm 1\}^{m \times n}$, s.t. W is $k := \varepsilon n$ -wise independent over \mathbb{R} with probability $1 - o(1)$. As a consequence, W is also k -wise independent over \mathbb{F}_p for any prime $p > (\varepsilon n)!$.*

We observe that the matrix $W \in \{\pm 1\}^{m \times n}$ from the previous theorem can be also transformed into a $(k-1)$ -wise independent matrix from $W' \in \{0, 1\}^{m \times n}$.

Corollary III.13. *For every $\varepsilon \in (0, 1/2]$, every large enough $n \leq m \leq n^{\frac{1}{2\varepsilon}}$, there is a polynomial-time algorithm that uses $O(n \log m (\log \log m)^2)$ random bits and outputs a matrix $W' \in \{0, 1\}^{m \times n}$, s.t. W' is $(k-1) := (\varepsilon n - 1)$ -wise independent over \mathbb{R} (and over \mathbb{F}_p for any prime $p > (\varepsilon n)!$) with probability $1 - o(1)$.*

E. Putting it all together

Now we are ready to prove Theorem I.1.

Theorem I.1. *There is an explicit set \mathcal{P}_n of $m = n^3$ convex polygons in the plane \mathbb{R}^2 , and a prime p , such that any group-model data structure of size s and query time t for range counting with respect to \mathcal{P}_n (over \mathbb{Z}_p), must have*

$$t^s \geq n^{n/7}.$$

In particular, any linear storage ($s = O(n)$) data structure for \mathcal{P}_n (over \mathbb{Z}_p) must have $n^{\Omega(1)}$ query time. Moreover, the semialgebraic complexity of \mathcal{P}_n , i.e., number of facets of each convex polygon, is $n^{O((\log \log n)^2)}$.

We remark that a construction of a PRG against halfspaces with optimal seed length $O_\varepsilon(\log m)$ would improve the semialgebraic complexity of the polygons in the theorem statement: instead of polygons with $n^{O((\log \log n)^2)}$ facets, one could construct a set of n^3 polygons with $n^{O(1)}$ facets.

Corollary III.14. *If there exists an explicit PRG $G : \{0, 1\}^r \rightarrow \{0, 1\}^m$ which ε -fools halfspaces and uses optimal seed $r = O(\log(m/\varepsilon))$, then there exists an explicit set \mathcal{P}_n of $m = n^3$ convex polygons in the plane \mathbb{R}^2 with $\text{poly}(n)$ facets/semialgebraic complexity, such that any group-model data structure of size s and query time t for weighted range counting with respect to \mathcal{P}_n (over some \mathbb{Z}_p), must have $t^s \geq n^{\Omega(n)}$.*

IV. GEOMETRIC PRGS

The construction of the incidence matrix $\mathcal{I}_{\mathcal{R}}$ in the proof of Theorem I.1 uses affine PRGs in a *black-box* manner: the i th row of $\mathcal{I}_{\mathcal{R}}$ encodes the truth table of the i th output bit of the (GKM) PRG $g_i : \{\pm 1\}^r \mapsto \{\pm 1\}$, hence the complexity of the i th “range” $g_i \in \mathcal{R}$ (number of polytope facets) is trivially bounded by $2^r = n^{O(\log^2 \log n)}$. This is essentially inevitable if the PRG is applied in a black-box fashion (since $r > \log m$ is the best one could hope for [38]). However, we might hope to obtain lower bounds against more “natural” range counting problems (with lower semi-algebraic complexity) using our approach, by constructing PRGs that take advantage of the *underlying representation* of the input seed. Indeed, the input to a range counting problem is a point in a d -dimensional space, hence it is natural to ask if the PRG can exploit higher dimensionality in a “simple” way so as to reduce the semi-algebraic complexity of the underlying hard problem. This is the motivation of our next definition.

Definition IV.1 (Geometric PRGs). *A function $g : [B]^d \rightarrow \{\pm 1\}^m$ is said to be a (B, d, s, ε) -Geometric PRG (GPRG) against a family of sets $\mathcal{S} \subset 2^{\{\pm 1\}^m}$, if g ε -fools \mathcal{S} , and furthermore every output bit of g can be encoded as a low-degree polynomial threshold function (PTF), i.e.,*

$$\forall i \in [m], \quad g_i(r_1, \dots, r_d) = \text{sign}(P_i(r_1, \dots, r_d)),$$

where each P_i is a d -variate polynomial of (total) degree s .

Note that this definition generalizes the (incidence function of) halfspaces ($s = 1$), and their natural extension to ranges specified by low-degree PTFs. We now show that an explicit construction of Geometric PRGs would yield a data structure lower bound against a corresponding class of ranges.

To this end, note that in the proof of Theorem III.12, we showed that a PRG that $1/\sqrt{m}$ -fools halfspaces yields an εn -wise independent matrix (which, by Theorem I.1, implies a data structure lower bound). By inspection of the proof of Theorem III.12, we see that for any constant $\gamma > 0$, even a PRG that $1/m^\gamma$ -fools halfspaces still gives us an $\Omega_\gamma(n)$ -wise independent matrix, yielding a data structure lower bound of $t^s \geq n^{\Omega(n)}$ by Theorem I.1.

Using the preceding remark, a $(B, d, s, 1/m^\gamma)$ -Geometric PRG $g : [B]^d \rightarrow \{\pm 1\}^m$ gives an explicit set of m ranges, each of which defined by a single d -variate polynomial of degree s , such that range counting with respect to these ranges is hard for linear storage data structures:

Lemma IV.2. *Suppose $g : [B]^d \rightarrow \{\pm 1\}^m$ is a $(B, d, s, 1/m^\gamma)$ -Geometric PRG (GPRG) against affine halfspaces in \mathbb{R}^m for a constant $\gamma > 0$. Then there exists a range counting problem $\mathcal{P}_{\mathcal{R}}$ with $|\mathcal{R}| = m$ ranges, each of semi-algebraic complexity $\Gamma_{d,s,1}$, such that any S -space data structure for $\mathcal{P}_{\mathcal{R}}$ has query time $t \geq n^{\Omega(n/S)}$ in the static group model.*

A. Existence of high-degree Geometric PRGs

In contrast to the standard definition of PRGs (Definition III.9), whose mere existence with small seed length follows from a rather simple counting argument [38], the existence of GPRGs (with sublinear seedlength) against affine halfspaces, is a nontrivial question. The remainder of this section is devoted to this question.

Our first result shows that GPRGs against halfspaces exist in nontrivial dimension $d = m^\varepsilon$ (for any $\varepsilon > 0$), albeit with high degree $s \approx m^\varepsilon$ (which corresponds to ranges being high-degree PTFs). In fact, we will show an explicit construction of such GPRGs, by “adapting” BCH codes from \mathbb{F}_2 to \mathbb{R} .

An important ingredient of the proof is the following theorem due to Diakonikolas et al. [39], asserting that any k -wise (statistically) independent distribution over $\{\pm 1\}^m$ $\tilde{\Omega}(1/\sqrt{k})$ -fools halfspaces:

Theorem IV.3 (Bounded independence fools halfspaces [39]). *Let \mathcal{D} be a k -wise independent distribution on*

$\{\pm 1\}^m$. Then \mathcal{D} δ -fools halfspaces as long as

$$k \geq \frac{C}{\delta^2} \log^2 \left(\frac{1}{\delta} \right)$$

for some universal constant C .

By setting δ in Theorem IV.3 to $m^{-0.05}$, an immediate corollary is that any $k = (Cm^{0.1} \log^2 m)$ -wise independent distribution will fool halfspaces with small enough error for Lemma IV.2 to yields data structure lower bounds. Thus, it remains to construct a GPRG whose output is $Cm^{0.1} \log^2 m$ -wise statistically independent.

A k -wise independent distribution with uniform marginals is easy to construct from a parity check matrix M of a linear code of minimal distance $> k$ over \mathbb{F}_2^m : since any k of the M 's columns are linearly independent over \mathbb{F}_2 , taking the sum of a uniformly random subset of the rows of M gives a vector of length m whose entries are statistically k -wise independent. In fact, this is the optimal way to obtain such a k -wise independent distribution, as Alon, Babai, Itai showed:

Theorem IV.4 ([43]). *Suppose $(\xi_i)_{i=1}^m$ is a collection of k -wise independent random variables $\xi_i : \Omega \rightarrow \mathbb{R}$ each of which takes at least two distinct values with non-zero probability. Then the size of the sample space is at least the size of the Hamming ball of radius $\lfloor k/2 \rfloor$: $|\Omega| \geq \sum_{j=0}^{\lfloor k/2 \rfloor} \binom{m}{j}$.*

Moreover, when $m = 2^t - 1$ and $k = 2s + 1$, this bound is essentially achieved for $\xi_i = \langle H^i, x \rangle_{\mathbb{F}_2}$ where H^i is the i -th column of the parity check matrix $H \in \mathbb{F}_2^{(1+s)t \times m}$ of a binary BCH code and x is chosen uniformly at random from \mathbb{F}_2^{1+st} .

To make this construction “geometric”, treat H as a binary matrix over the reals and x as a binary vector over the reals. Then the entries of the vector $H^T x$ are integers that are k -wise independent modulo 2, and are between 0 and $k \log m/2$. Let $p(x) := \prod_{i=1}^{k \log m/2} (i - 1/2 - x)$, then $\text{sign}(p(x)) = (-1)^x$ for $x \in \{0, \dots, k \log m/2\}$, giving our desired theorem:

Theorem IV.5 (Explicit high-degree GPRGs). *Let $\mathcal{R}_i(x) = p(\langle H^i, x \rangle)$, where H is the parity check matrix defined above treated as a binary matrix over \mathbb{R} , and $x \in \{0, 1\}^r$ for $r = O(m^{0.1} \log^3 m)$. Then $\mathcal{R} : \{0, 1\}^r \rightarrow \{\pm 1\}^m$ is a $(2, O(m^{0.1} \log^3 m), O(m^{0.1} \log^3 m), 1/\text{poly}(m))$ -GPRG against affine halfspaces in \mathbb{R}^m .*

This theorem asserts a polynomial lower bound for range counting of $n^{O_\varepsilon(1)}$ explicit PTFs in $\mathbb{R}^{n^\varepsilon}$, each of degree $\tilde{O}(n^{0.1})$, against linear storage group-model

data structures. Is it possible to construct GPRGs using *constant degree* PTFs? This is the content of the next subsection.

B. Can Halfspaces Fool Halfspaces ?

The special case of degree-1 PTFs is particularly interesting in our context, as it corresponds to the very natural problem of *halfspace* range counting, which is among the most well-studied range counting problems in the literature [10], [8]. While Theorem IV.5 produces an *explicit* high-degree GPRG, for degree $s = 1$ it is not even clear whether GPRGs *exist* in nontrivial dimension, i.e., whether “halfspaces fool halfspaces” in sublinear dimension :

Problem IV.6. *Is there any set \mathcal{H} of $m = \text{poly}(d)$ halfspaces in \mathbb{R}^d for $d \ll m$, such that the function given by $g_i(r) := (\text{sign}\langle h_i, r \rangle)_{i \in [m]}$ for $r \in_R \{\pm 1\}^d$, $(1/\text{poly}(m))$ -fools affine halfspaces in \mathbb{R}^m ?*

We note that the (explicit) BCH-code construction in Theorem IV.5 shows these parameters are achievable over \mathbb{F}_2 , i.e., had we taken the *parity* of the inner-products $\langle h_i, r \rangle$ instead of their *signs*. Intuitively, threshold functions are much less sensitive than the parity operator (in particular, thresholds are monotone function and hence have low influence in sharp contrast to parities). Hence, achieving $n^{\Omega(1)}$ -wise independence (by Theorem IV.3) with signs of inner products seems to require a substantially different construction and analysis.

In the full version of the paper we show that it is possible to construct an *almost k -wise independent* distribution, i.e., a $(\delta/2^k, k)$ -independent distribution¹⁰ over $\{\pm 1\}^m$, with $\delta = 1/\text{poly}(d)$, from d -dimensional halfspaces over the Boolean hypercube, which implies that the same distribution is δ -far from $k = d^{\Omega(1)}$ -wise independence, i.e. all the k -marginals are δ -close to uniform in statistical distance:

Theorem IV.7 (Random halfspaces are almost k -wise independent). *For a fixed number p , pick $m = d^p$ uniformly random vectors v_1, \dots, v_m in $\{\pm 1\}^{2d+1}$ for large enough d , and let $\varepsilon < 1/6$. Then for a random vector $\mathbf{x} \in_R \{\pm 1\}^{2d+1}$, the distribution of the binary vector $f(\mathbf{x}) := \{\text{sign}\langle v_i, \mathbf{x} \rangle\}_{i=1}^m$ is $(d^{1/6}k/2^k, k)$ -independent with high probability (over the choice of the v_i 's), where $k = d^\varepsilon$.*

¹⁰This is a distribution where all marginals on k coordinates are $\delta/2^k$ -close in L_∞ norm to uniform, see [44] for details and its wide applications.

If Theorem IV.3 were “robust” to distributions *statistically-close* to being k -wise independent, then Theorem IV.7 would have established the existence of a degree-1 (halfspace) GPRG against halfspaces in n^ε dimensions. Unfortunately, the proof of Theorem IV.3 only applies to distributions which are *exponentially* close (in k) to k -wise independence (whereas the distance $\delta = 1/\text{poly}(k)$ obtained in Theorem IV.7 is essentially tight by considering the correlation between just 2 coordinates). Nonetheless, Theorem IV.3 is a necessary milestone toward the fooling question, and we believe the proof technique of Theorem IV.7 may be useful to directly settle the existence of GPRGs from halfspaces.

REFERENCES

- [1] A. Golovnev, G. Posobin, O. Regev, and O. Weinstein, “Polynomial data structure lower bounds in the group model,” *ECCC*, vol. 27, p. 57, 2020.
- [2] A. C. Yao, “Theory and applications of trapdoor functions (extended abstract),” in *FOCS 1982*, 1982, pp. 80–91.
- [3] P. B. Miltersen, “The bit probe complexity measure revisited,” in *STACS 1993*, 1993, pp. 662–671.
- [4] A. Siegel, “On universal classes of extremely random constant-time hash functions,” *SIAM J. Comput.*, vol. 33, no. 3, pp. 505–543, 2004.
- [5] M. Pătraşcu, “Unifying the landscape of cell-probe lower bounds,” in *FOCS 2008*, 2008, pp. 434–443.
- [6] K. G. Larsen, “Higher cell probe lower bounds for evaluating polynomials,” in *FOCS 2012*, 2012, pp. 293–301.
- [7] J. Matoušek, “Geometric range searching,” *ACM Comput. Surv.*, vol. 26, pp. 421–461, 1994.
- [8] P. K. Agarwal, “Simplex range searching and its variants: A review,” in *A Journey Through Discrete Mathematics*. Springer, 2017, pp. 1–30.
- [9] M. L. Fredman, “A lower bound on the complexity of orthogonal range queries,” *J. ACM*, vol. 28, no. 4, pp. 696–705, 1981.
- [10] B. Chazelle, “Lower bounds for orthogonal range searching: Part II. The arithmetic model,” *J. ACM*, vol. 37, no. 3, pp. 439–463, 1990.
- [11] P. K. Agarwal and J. Matoušek, “On range searching with semialgebraic sets,” *Discrete Comput. Geom.*, vol. 11, pp. 393–418, 1994.
- [12] G. S. Lueker, “A data structure for orthogonal range queries,” *SFCS 1978*, pp. 28–34, 1978.

- [13] J. L. Bentley, “Multidimensional binary search trees used for associative searching,” *Commun. ACM*, vol. 18, pp. 509–517, September 1975.
- [14] P. K. Agarwal, J. Matoušek, and M. Sharir, “On range searching with semialgebraic sets II,” in *FOCS 2012*. IEEE, 2012, pp. 420–429.
- [15] B. Chazelle and B. Rosenberg, “Simplex range reporting on a pointer machine,” *Comput. Geom.*, vol. 5, pp. 237–247, 1995.
- [16] P. Afshani, “A new lower bound for semigroup orthogonal range searching,” in *SoCG 2019*. LIPIcs, 2019, pp. 3:1–3:14.
- [17] M. Pătraşcu, “Lower bounds for 2-dimensional range counting,” in *STOC 2007*, 2007, pp. 40–46.
- [18] V. Koltun, “Sharp bounds for vertical decompositions of linear arrangements in four dimensions,” *Discrete Comput. Geom.*, vol. 31, p. 2004, 2004.
- [19] T. M. Chan and M. Patrascu, “Transdichotomous results in computational geometry, I: point location in sublogarithmic time,” *SIAM J. Comput.*, vol. 39, no. 2, pp. 703–729, 2009.
- [20] Z. Dvir, A. Golovnev, and O. Weinstein, “Static data structure lower bounds imply rigidity,” in *STOC 2019*. ACM, 2019, pp. 967–978.
- [21] S. N. Ramamoorthy and C. Rashtchian, “Equivalence of systematic linear data structures and matrix rigidity,” in *ITCS 2020*. LIPIcs, 2020, pp. 35:1–35:20.
- [22] E. J. Candès and T. Tao, “Near-optimal signal recovery from random projections: Universal encoding strategies?” *IEEE Trans. Inf. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [23] A. Dimakis, R. Smarandache, and P. Vontobel, “LDPC codes for compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 3093–3114, 5 2012.
- [24] P. Indyk, “Explicit constructions for compressed sensing of sparse signals,” in *SODA 2008*, 2008, pp. 30–33.
- [25] M. Sipser and D. A. Spielman, “Expander codes,” *IEEE Trans. Inf. Theory*, vol. 42, no. 6, pp. 1710–1722, 1996.
- [26] B. Chor, O. Goldreich, J. Hastad, J. Freidmann, S. Rudich, and R. Smolensky, “The bit extraction problem or t -resilient functions,” in *SFCS 1985*. IEEE, 1985, pp. 396–407.
- [27] W. Xu and B. Hassibi, “Efficient compressive sensing with deterministic guarantees using expander graphs,” in *ITW 2007*, Sep. 2007, pp. 414–419.
- [28] V. Guruswami, J. R. Lee, and A. Wigderson, “Euclidean sections of ℓ_1^N with sublinear randomness and error-correction over the reals,” in *RANDOM 2008*. Springer, 2008, pp. 444–454.
- [29] J. Błasiok, P. Lopatto, K. Luh, J. Marcinek, and S. Rao, “An improved lower bound for sparse reconstruction from subsampled Hadamard matrices,” in *FOCS 2019*. IEEE, 2019, pp. 1564–1567.
- [30] H. Brönnimann, B. Chazelle, and J. Pach, “How hard is half-space range searching,” *Discrete Comput. Geom.*, vol. 10, pp. 143–155, 1993.
- [31] B. Chazelle, “A spectral approach to lower bounds,” in *FOCS 1994*. IEEE, 1994, pp. 674–682.
- [32] K. G. Larsen, “On range searching in the group model and combinatorial discrepancy,” *SIAM J. Comput.*, vol. 43, no. 2, pp. 673–686, 2014.
- [33] M. H. Overmars and J. van Leeuwen, “Dynamization of decomposable searching problems yielding good worst-case bounds,” in *Theor. Comput. Sci.* Springer, 1981, pp. 224–233.
- [34] M. Bellare and J. Rompel, “Randomness-efficient oblivious sampling,” in *FOCS 1994*. IEEE, 1994, pp. 276–287.
- [35] P. Erdős, “On a lemma of Littlewood and Offord,” *Bull. Am. Math. Soc.*, vol. 51, no. 12, pp. 898–902, 1945.
- [36] J. Komlós, “On the determinant of $(0, 1)$ matrices,” *Studia Sci. Math. Hungar.*, vol. 2, pp. 7–21, 1967.
- [37] P. Gopalan, D. M. Kane, and R. Meka, “Pseudorandomness via the discrete Fourier transform,” *SIAM J. Comput.*, vol. 47, no. 6, pp. 2451–2487, 2018.
- [38] R. Meka and D. Zuckerman, “Pseudorandom generators for polynomial threshold functions,” *SIAM J. Comput.*, vol. 42, no. 3, pp. 1275–1301, 2013.
- [39] I. Diakonikolas, P. Gopalan, R. Jaiswal, R. A. Servedio, and E. Viola, “Bounded independence fools halfspaces,” *SIAM J. Comput.*, vol. 39, no. 8, pp. 3441–3462, 2010.
- [40] B. Klartag and O. Regev, “Quantum one-way communication can be exponentially stronger than classical communication,” in *STOC 2011*. ACM, 2011, pp. 31–40.
- [41] H. H. Nguyen and M. M. Wood, “Random integral matrices: universality of surjectivity and the cokernel,” *arXiv:1806.00596*, 2018.
- [42] H. H. Nguyen and E. Paquette, “Surjectivity of near-square random matrices,” *Comb. Probab. Comput.*, pp. 1–26, 2018.
- [43] N. Alon, L. Babai, and A. Itai, “A fast and simple randomized parallel algorithm for the maximal independent set problem,” *J. Algorithms*, vol. 7, no. 4, pp. 567–583, 1986.
- [44] N. Alon, O. Goldreich, J. Håstad, and R. Peralta, “Simple construction of almost k -wise independent random variables,” *Random Struct. Algorithms*, vol. 3, no. 3, pp. 289–304, 1992.