Electronic Colloquium on Computational Complexity, Report No. 1 (2026)

Multilinear Algebraic Branching Programs and the Min-Partition Rank Method

Théo Borém Fabris* Nutan Limaye† Srikanth Srinivasan‡ Amir Yehudayoff§

January 1, 2026

Abstract

It is a long-standing open problem in algebraic complexity to prove lower bounds against multilinear algebraic branching programs (mABPs). The best lower bounds in this setting are still quadratic (Alon, Kumar and Volk (*Combinatorica* 2020)). At the same time, it remains a possibility that the "min-partition rank" method introduced by Raz (*Theory Comput.* 2006), which is used to prove all known multilinear lower bounds, can also be used to prove superpolynomial lower bounds on the size of mABPs.

In this paper, we analyze the potential of the min-partition rank method to prove lower bounds on the size of mABPs and prove the following:

1. We relate this method to a purely combinatorial question regarding the minimum size of a set system whose chains satisfy a discrepancy condition. In the case of set-multilinear ABPs, this *characterizes* the best lower bound that can be achieved.
2. We prove a non-trivial upper bounds on the size of set systems with this combinatorial property. This recovers a superpolynomial separation between mABPs and multilinear formulas (Dvir, Malod, Perifel and Yehudayoff (*STOC 2012*)). Our proof is conceptually different, and may have wider consequences.
3. The property we study extends combinatorial notions of "balancing sets" considered in previous works, for which near-tight bounds are known via intervals. We show that intervals are very far from satisfying our property. This showcases how our methods capture combinatorial structure that evades previous techniques.

These results build a bridge between algebraic complexity theory and the behavior of random walks. The upper bound uses the fact (Csáki, Erdős, and Révész (*PTRF 1985*)) that a random walk of length n on the integers returns to its starting point once every $\leq n/\log n$ steps with noticeable probability. To show the latter result, we prove that two independent random walks are "far" from each other in discrete Fréchet distance.

*Department of Computer Science, University of Copenhagen, Denmark. Supported by the European Research Council (ERC) under grant agreement no. 101125652 (ALBA). Email: `thfa@di.ku.dk`

†IT University of Copenhagen, Denmark. Supported by Independent Research Fund Denmark (grant agreement No. 10.46540/3103-00116B) and Basic Algorithms Research Copenhagen (BARC), funded by VILLUM Foundation Grant 54451. Email: `nuli@itu.dk`

‡Department of Computer Science, University of Copenhagen, Denmark. Supported by the European Research Council (ERC) under grant agreement no. 101125652 (ALBA). Email: `srsr@di.ku.dk`

§Department of Computer Science, The University of Copenhagen, and The Technion-IIT. Supported by a DNRF Chair grant. Email: `amir.yehudayoff@gmail.com`

ISSN 1433-8092

Contents

1 Introduction

Basic background. Algebraic complexity theory is the study of algebraic problems in models such as algebraic formulas, branching programs and circuits. Analogous to the Boolean circuit classes $\mathbf{NC}^1 \subseteq \mathbf{NL} \subseteq \mathbf{NC}^2$, in algebraic complexity we study complexity classes $\mathbf{VF} \subseteq \mathbf{VBP} \subseteq \mathbf{VP}$, and a central goal is to prove lower bounds against all classes in this chain.

Notable success has been achieved in the *multilinear* setting where we restrict each of the algebraic models to compute only multilinear polynomials (in a syntactic[1] way). In particular, this gives us a multilinear hierarchy $\mathbf{mVF} \subseteq \mathbf{mVBP} \subseteq \mathbf{mVP}$. Well-known results of Raz [Raz09, Raz06] showed superpolynomial lower bounds on the size of multilinear formulas computing some explicit polynomials in $\mathbf{mVP}$, thus separating the first and third levels of this hierarchy. Subsequently, Dvir, Malod, Perifel and Yehudayoff [DMPY12] proved a stronger separation between the first and second levels of this hierarchy. There has been a large volume of follow-up work in algebraic complexity (e.g. [RSY08a, RY09, RY11, CLS19, AKV20, KS22, KS23, CKSS24]) and it also found applications in quantum computation [Aar04, RY11] and proof complexity [RT08, FSTW21]. Despite this interest, we have not yet been able to extend the lower bounds to higher classes in the multilinear hierarchy such as $\mathbf{mVBP}$. This is the question that motivates our work.

Multilinear algebraic branching programs (mABPs). The mABP computational model is an algebraic analogue of Boolean branching programs, which model space-eficient (Boolean) computation. An mABP is defined to be a directed acyclic graph G with a source a and sink b. The edges of G are labelled with affine linear polynomials (e.g. $1 - 2x_1 + x_2$). The polynomial computed is the sum, over all a-to-b paths, of the products of the edge-labels occurring along the path. To keep the polynomial multilinear, we impose the condition that the same variable cannot appear in more than one edge label along a path. Multilinear formulas and circuits are defined similarly using algebraic analogues of Boolean formulas and Boolean circuits, but we omit their formal definitions here as they are not crucial in what follows. The central open question we care about is the following:

Question 1. Do there exist explicit n-variate multilinear polynomials P_n that cannot be computed by mABPs of size $\mathrm{poly}(n)$?

As mentioned above, this question has been resolved positively for multilinear formulas by the work of Raz [Raz09, Raz06]. For mABPs and the stronger model of syntactially multilinear circuits, the best lower bounds are nearly quadratic [RSY08a, AKV20].

The lower bound method. Nearly all known lower bounds for multilinear models of computation (including all those mentioned above) can be proved using the "min-partition rank" technique developed by Raz [Raz09][2] and relies on ideas of Nisan [Nis91] and Nisan and Wigderson [NW97].

Given a multilinear polynomial $P(x_1, \ldots, x_n)$ with n even and a partition $\Pi = (Y, Z)$ of the underlying variable set into two parts, we define the *coefficient matrix* $M_\Pi(P)$ to be a $2^{|Y|} \times 2^{|Z|}$

[1] We can also define a *semantic* notion of multilinearity, which coincides with the syntactic notion for formulas in terms of computational power, but is potentially more powerful in the setting of branching programs and circuits. In this paper, the word 'multilinear' will only refer to the syntactic variants.

[2] It should be noted that *separations* between different multilinear models require us to sometimes use specializations of this technique that yield a more fine-grained understanding. But the basic technique suffices to just prove a lower bound.

matrix whose rows and columns are labelled by multilinear monomials m_Y and m_Z respectively. The (m_Y, m_Z)-th entry of $M_\Pi(P)$ is the coefficient of $m_Y m_Z$ in P. We define the *min-partition rank* of P by minimizing the rank of $M_\Pi(P)$ over all *equipartitions* (i.e. $|Y| = |Z| = n/2$):

$$\Gamma(P) := \min_{\substack{\Pi \\ \text{an equipartition}}} \text{rank}(M_\Pi(P)).$$

We call the polynomial P *full-rank* if $\Gamma(P) = 2^{n/2}$. In other words, M_Π has full rank for all Π.

Using this notion of the complexity of a polynomial, Raz established the following.

Theorem 1.1 (Raz [Raz09, Raz06]). Let $\mathbb{F}$ be any field. Any multilinear formula computing a full-rank multilinear polynomial $P(x_1, \ldots, x_n) \in \mathbb{F}[x_1, \ldots, x_n]$ must have size at least $n^{\Omega(\log n)}$.

Further, if $\mathbb{F}$ is infinite, then there exist multilinear circuits of size poly(n), and multilinear ABPs and formulas of size at most $n^{O(\log n)}$ computing full-rank polynomials.[3]

Using the above theorem, Raz proved the existence of explicit sequences of n-variate polynomials that have small multilinear circuits but no multilinear formulas of size $\leq n^{o(\log n)}$. A follow-up result of Dvir, Malod, Perifel and Yehudayoff [DMPY12] exhibited a similar separation between mABPs and multilinear formulas, but this was based on the fact that mABPs can compute a polynomial that is full-rank with respect to a suitable subfamily of all equipartitions.

Known lower bounds on mABP size. While Theorem 1.1 shows that full-rankness is not sufficient to prove a superpolynomial multilinear circuit lower bound, it could potentially prove such a lower bound on mABP size. However, the best mABP lower bound we have (which is also based on this method) is only $\Omega(n^2)$. We review this result [RSY08a, AKV20] below, as it naturally leads to our new approach to mABP lower bounds.

Assume that a full-rank polynomial $P(x_1, \ldots, x_n)$ has an mABP of size s and assume for simplicity that the mABP is 'layered' in the sense that vertices are divided into layers and edges go from one layer to the next, and for all vertices v at layer i, the set X_v of variables seen on the paths from a to v is a set of size i. To prove a lower bound for such an mABP, we note that for any $i \in [n]$, the polynomial P can be decomposed as

$$P = \sum_{v \text{ in layer } i} L_v(X_v) \cdot R_v(X \setminus X_v), \tag{1}$$

where L_v corresponds to the sum of all paths from a to v and R_v the sum of all paths from v to b. Proving a lower bound on the number of terms in any such decomposition yields a lower bound on the number of vertices at layer i.

To lower bound the number of terms in the decomposition, we turn to *discrepancy*. We show that if the number of summands on the right hand side of (1) is small, then we can find an equipartition $\Pi = (Y, Z)$ such that each summand has low-rank for this choice of Π. This can be guaranteed by ensuring that, for each v, the induced partition $\Pi' = (Y \cap X_v, Z \cap X_v)$ of X_v is sufficiently 'imbalanced'. Informally, this ensures that the coefficient matrices of L_v and R_v are both 'far from square' matrices and hence that $M_\Pi(f_v \cdot g_v)$ is far from full-rank.

[3] Raz's work actually proves this result for a suitable transcendental extension of any field but this can be circumvented over an infinite field by replacing the transcendental elements by random elements from (a large subset of) $\mathbb{F}$.

This leads to the combinatorial question of understanding the smallest set system where we have such a discrepancy bound [RSY08a]. More precisely, fix an $i \in [n]$. Call a set system $\mathcal{Y} \subseteq \binom{[n]}{i}$ *k-balancing* if for any balanced function $f : [n] \to \{\pm 1\}$ (representing the equipartition), there is a $Y \in \mathcal{Y}$ such that $|\sum_{y \in Y} f(y)| \leq k$. Then, the size of the smallest such k-balancing set system represents the best lower bound we can prove for the number of vertices in layer i of the mABP. Note that the answer to this is at most n, since the set of all intervals in $\{1, \dots, n\}$ of size i yields a set system with discrepancy bound 1. This simple example is nearly tight for $i \in [n/4, 3n/4]$ and small k: Alon, Kumar and Volk [AKV20] proved a tight lower bound of $\Omega(n/k)$ on the size of such 'balancing set systems' for all k , which leads to a near-linear bound on the number of vertices in each layer and hence a near-quadratic lower bound on mABP size.[4]

Our new technique. In order to prove a stronger lower bound, we suggest a modification of the decomposition in (1). We observe that such a decomposition of P can be obtained by looking at *any* vertex cut in the mABP separating the source a from the sink b. As long as we can find a partition Π and a suitable cut C such that for each vertex $v \in C$, the set X_v is imbalanced with respect to Π, we obtain a lower bound. This gives us potentially exponentially many possibilities to work with and hence a much better chance at proving a lower bound.

By graph-theoretic arguments, the non-existence of such a cut is equivalent to the existence of an a-b path such that all the vertices v along the path are balanced. These vertices correspond to a *chain* of subsets of the variable set, which naturally brings us the following definition.

Definition 1.2 (Balanced-chain set systems)**.** Let X be a finite set with n elements and let $\mathcal{X} \subseteq \mathcal{P}(X)$ be a family of subsets of X. Denote by $\mathcal{C}(\mathcal{X})$ the set of *maximal chains* of $\mathcal{P}(X)$ contained in $\mathcal{X}$; i.e., a chain $(C_0, \dots, C_l)$ is in $\mathcal{C}(\mathcal{X})$ iff $l = n$ and $|C_i| = i$ for each $i \in \{0, \dots, l\}$.

For a function $f \colon X \to \{-1, 1\}$ which we will call a *partition*, and $S \subseteq X$, let $f(S) := \sum_{x \in S} f(x)$. We say that f is *balanced* (or that f is a *balanced partition* of X) if $f(X) = 0$. We define the *chain-balance* of $\mathcal{X}$ with respect to f as

$$\mathrm{cbal}_{\mathcal{X}}(f) := \min_{(C_0, \dots, C_n) \in \mathcal{C}(\mathcal{X})} \max_{i \in [n]} |f(C_i)|.$$

We define the chain-balance of $\mathcal{X}$ to be

$$\mathrm{cbal}(\mathcal{X}) := \max_{\text{balanced } f \colon X \to \{-1,1\}} \mathrm{cbal}_{\mathcal{X}}(f).$$

We say that $\mathcal{X}$ is a *k-balanced-chain set system* if $\mathrm{cbal}(\mathcal{X}) \leq k$. ◀

The definition of chain-balance above resembles standard definitions in discrepancy theory. Think of X as $X = [n]$. The set $[n]$ has the canonical maximal chain $C^* = (\varnothing, [1], [2], \dots, [n])$. Every permutation π of $[n]$ defines a maximal chain C^π by applying π on each interval in C^* (that is, $C_i^\pi = \pi([i])$ for all i). The discrepancy of the permutations $\pi_1, \dots, \pi_m$ is defined as

$$\min_{f : X \to \{\pm 1\}} \max_{j \in [m]} \max_{i \in [n]} |f(\pi_j([i]))|.$$

In a nutshell, we care about the most balanced partition f of $[n]$. There are several known bounds on discrepancy of permutations (for example, the discrepancy is at most $O(m \log n)$), and it was algorithmically studied in the offline and online settings (see e.g. [NNN12] and references within).

[4] The general 'unlayered' case is more complicated.

Our definition is "dual" to the standard one in that the order of quantifiers is reversed ($\max_f \min_C$ versus $\min_f \max_C$). There are two additional differences between the standard formulation and ours. First, instead of the algorithmic task of finding f that is as balanced as possible, we care about "hardness" results that show that every f is not balanced. The second difference is more subtle. Instead of an explicit list of permutations, we care about an implicit list of permutations. The list of permutations we care about is described by all maximal chains in the given family of sets $\mathcal{X}$. Later on, we shall see that we mostly care about the case that $\mathcal{X}$ can be described via a small branching program. As we shall see below, there are many examples of exponentially large families of permutations that can be described by a small branching program. This makes proving lower bounds harder, but on the other hand, there are non-trivial constructions (upper bounds) that have implications in algebraic circuit complexity.

An example: intervals. Identify X with the set $[n]$. Consider the set system $\mathcal{I}$ of all intervals contained in $[n]$:

$$\mathcal{I} = \{[i,j] \mid 1 \leq i,j \leq n\}.$$

As noted above, this set system provides minimum-sized examples of 1-balanced set systems contained in $\binom{[n]}{i}$ for any i. It is therefore natural to understand if they yield balanced-*chain* set systems (the key new word is "chain"). This is a nice exercise, which we solve next.

The intervals set system, despite being balanced, is far from being a balanced-chain set system. Consider e.g. $f : [n] \to \{\pm 1\}$ given by

$$f(j) = \begin{cases} 1 & j \in [n/4] \cup [(3n/4)+1, n] \\ -1 & \text{otherwise} \end{cases}$$

Pictorially, the graph of $i \mapsto f([i])$ comprises a mountain of linear height and a valley of linear depth. It can be shown that any chain in $\mathcal{C}(\mathcal{I})$ contains a set S that is very much not balanced; i.e., $|f(S)| \geq \Omega(n)$. Note that unbalance $n/2$ is the worst possible.

1.1 Results

Balanced-chain set systems and lower bounds. Our first result relates balanced-chain set systems to the problem of proving mABP lower bounds. We show the following two-directional equivalence.

Theorem 1.3. Let $s = s(n)$ be any growing function of n with $s \geq n$. If any $(\log s)$-balanced-chain set system has size at least s, then any mABP computing a full-rank polynomial has size at least $s^{\Omega(1)}$. Conversely, if there is an $O(1)$-balanced-chain set system of size s, then there is an mABP computing a full-rank polynomial of size at most $s \cdot \mathrm{poly}(n)$.

The reader may find the forward direction fairly intuitive, given the discussion above. In the other direction, it is possible to turn this purely combinatorial construction back into a polynomial that satisfies the algebraic property of full-rankness.

Remark. The above theorem provides a near-tight characterization of the size of the smallest mABP in terms of the size of the smallest balanced-chain set system. The two directions differ only in the balance parameter: e.g., an mABP lower bound of the form $s \geq n^c$ requires us to show that every $\Omega(c \log n)$-balanced-chain set system needs size n^c, but an upper bound of $s \leq n^c$ can only be ensured by an $O(c)$-balanced-chain set system of size n^c. We can close this gap completely for the more restrictive setting of *set-multilinear* ABPs (defined below) where both directions use a balance parameter that is $\Theta(c)$. So at least in the set-multilinear setting, where the lower bound question is still open [CKSS24], we have a perfect characterization of the ABP complexity of full-rank polynomials in terms of the combinatorial question.

The next question is to understand how to analyze and prove lower bounds on the minimal size of a balanced-chain set system. A priori, this problem could be difficult as finding a partition $f : X \to \{\pm 1\}$ with respect to which $\mathrm{cbal}_{\mathcal{X}}(f)$ is small could be hard. However, we show via a simple covering argument that a *uniformly random* partition is as hard as a worst-case partition in the following sense: if there is a small set system that contains a balanced chain for a random partition with good probability, then there is also a small balanced-chain set system. We make this precise below.

Definition 1.4. For $\varepsilon \in [0, 1]$ and $k \in \mathbb{R}_{\geq 0}$, we call an mABP set system $\mathcal{X}$ an (ε, k)*-balanced-chain set system* if, for a uniformly random balanced $f : X \to \{\pm 1\}$, we have

$$\mathbb{P}_f \left[\mathrm{cbal}_{\mathcal{X}}(f) \leq k\right] \geq \varepsilon.$$

Note that a k-balanced-chain set system is also a $(1, k)$-balanced-chain set system. ◀

Lemma 1.5 (Worst-case to average-case reduction)**.** If there is an (ε, k)-balanced-chain set system of size s, then there is a k-balanced-chain set system of size at most $O(sn/\varepsilon)$.

We use both directions of Theorem 1.3 and Lemma 1.5 in what follows.

An improved upper bound on the size of balanced-chain set systems. We note that Theorem 1.3 in combination with the results of Raz (Theorem 1.1) immediately implies that there is an $O((\log n)^2)$-balanced-chain set system of size $n^{O(\log n)}$. In fact, it is not hard to give a direct proof of the slightly stronger fact that there is a 1-balanced-chain set system of size $n^{O(\log n)}$. This is proved using an inductive construction, using the fact (already mentioned above) that the family of intervals contained in $[n]$ are tight for the 'balancing set systems' problem. The above lemma, however, is still consistent with the fact that the mABP complexity of any full-rank polynomial could be $n^{\Omega(\log n)}$, which is no better than the upper bound in Theorem 1.1.

Using a more sophisticated inductive argument, we are able to improve the easy upper bound superpolynomially, marking the first such improvement.

Theorem 1.6. There is a 1-balanced-chain set system of size at most $n^{O(\log n/ \log\log n)}$.

Using Theorem 1.3, this immediately implies that there are full-rank mABPs of size $n^{O(\log n/ \log\log n)}$. An immediate consequence (using Theorem 1.1) is that mABPs are superpolynomially more powerful than multilinear formulas.

This recovers a superpolynomial separation between these models that was proved in [DMPY12]. Our result is weaker than this work both quantitatively (the superpolynomial lower bound) and qualitatively (the construction of the set system and hence the mABP is not uniform). However, we note that our proof is very different from that of [DMPY12]. While the proof of [DMPY12] reveals more about the structure of multilinear formulas by showing that they are not full-rank with respect to a much more restricted family of equipartitions (called arc-partitions), our proof yields a better understanding of the computational power of mABPs, which is what we need to prove lower bounds against the latter model.

From a broader computational complexity perspective, understanding the size of ABPs is closely related to studying the computational power of dynamic programming. The construction above shows that dynamic programs can be quite powerful. More concretely, consider the following non-deterministic read-once branching program model. The devices in this model are directed acyclic graphs with a designated source a and a designated sink b. The input is $x \in \{0,1\}^n$ for even n. Each edge in the graph is labeled by two variables (x_i, x_j). On every directed path every variable appears at most once. Assume that all a-to-b paths are of length exactly $n/2$. On input $x \in \{0,1\}^n$, an edge labeled by (x_i, x_j) is open iff $x_i + x_j = 1$ (in other words, it is balanced). The program accepts x iff there is an open path from a to b. We care about such programs that accept x iff the Hamming weight of x is exactly $n/2$. A program that achieves this functionality can actually "prove" its correctness. If it accepts an input x, then the open path that verifies acceptance shows that half of the entries of x are 1s and half are 0s.

What is the smallest program with this functionality? It is not hard to inductively construct such a program of size $n^{O(\log n)}$. The above theorem gives a non-trivial *superpolynomial* improvement. The exact complexity remains an open question.

Intervals and random partitions. As outlined above, set systems based on intervals have played an important role in previous investigations on multilinear lower bounds. Such families are optimal for constructions of balanced set systems and furthermore, results that separate multilinear circuits and ABPs from formulas [Raz06, DMPY12] are also based on combinatorial constructions that use intervals.

It is therefore natural to ask if intervals can be used to construct k-balanced-chain set systems for small k. We have already seen above that this family does not yield a good balanced-chain set system for a worst-case f. However, Lemma 1.5 shows that if this set system yields a balanced chain for a random partition, then that would allow us also to devise a somewhat larger set system that has balanced chains for all partitions. We are able to rule out this possibility with a strong negative result. This provides evidence that our technique may be able to prove superpolynomial lower bounds against general mABPs.

Theorem 1.7 (Intervals vs. random partitions)**.** There is a universal constant $c > 0$ such that for every large enough n, the set system $\mathcal{I}$ of intervals defined above is not an (ε, k)-balanced-chain set system for $\varepsilon \geq 2^{-n^c}$ and $k \leq n^c$.

In particular, the above is a considerable generalization of a recent result of Chatterjee, Kush, Saraf and Shpilka [CKSS24], whose results imply such a bound for the subset $\mathcal{I}_0$ of $\mathcal{I}$ that consists only of intervals of the form $[i]$ for $i \in \{0, \ldots, n\}$. The set system $\mathcal{I}_0$ contains just one maximal chain, while the set system $\mathcal{I}$ contains *exponentially* many chains, which makes the arguments from [CKSS24] inapplicable for $\mathcal{I}$. Using our results, we can prove the following lower bound for a generalization of the sum of ordered set-multilinear ABPs model introduced by [CKSS24].

Corollary 1.8 (Lower bounds for sum of interval-mABPs)**.** Let $\mathbb{F}$ be any field. Any Σ_πmABP program computing a full-rank polynomial $P(x_1, \ldots, x_n) \in \mathbb{F}[x_1, \ldots, x_n]$ must have size at least $2^{\Omega(n^{1/5})}$.

The main technical part of the proof of Theorem 1.7 analyzes the discrete Fréchet distance between two independent one-dimensional random walks. This is a basic problem of independent interest. We have two random walks, and we want to understand what is the best time parametrization that makes them as close as possible. Here is a slightly more formal description (but still not fully formal; for details see Section 4). A random walk defines a random function $f : \{0, 1, \dots, n\} \to \mathbb{Z}$ where $f(t)$ is the position of the walk at time t. The Fréchet distance between two functions f, g is

$$d_F(f, g) := \min_{\alpha,\beta} \max_t |f(\alpha(t)) - g(\beta(t))|$$

where α, β are time parameterizations (that may depend on both f and g). We have two independent (1-Lipschitz) random functions, we want to prove that we cannot hope to parametrize time to make them close (even when we know the future). The property that makes this problem challenging is the dependency on the future.

The set-multilinear setting. Two recent works have conducted investigations into the power of *set-multilinear* ABPs (smABPs). It is not hard to show that proving lower bounds against mABPs is at least as hard as the analogous problem for smABPs, and proving lower bounds against the latter is still an open question. The work of Bhargava, Dwivedi and Saxena [BDS25] showed that proving lower bounds on smABPs for low-degree polynomials implies general ABP lower bounds. A follow-up result of [CKSS24] showed restricted lower bounds against smABPs in both the low and high-degree settings. As mentioned above, this latter paper implicitly proves a result that is a special case of Theorem 1.7 above.

1.2 Outline of proofs

From ABPs to balanced-chain set systems and vice versa. We sketch the forward direction of Theorem 1.3 by showing if there is an mABP of size s computing a full-rank polynomial P, then there is an $O(\log s)$-balanced-chain set system of size at most s. Let X denote the set of n variables. We associate with each vertex v in the mABP a set $X_v \subseteq X$ of variables that appear on paths from the source vertex to v. The set system $\mathcal{X}$ we consider is the set of all such X_v (there are some hidden technicalities here). If there is a balanced $f : X \to \{\pm 1\}$ such that there is no $O(\log s)$-balanced-chain in $\mathcal{X}$, then we can find a cut C separating source from sink in the mABP containing only vertices v such that X_v is, for example, $10 \log s$-imbalanced with respect to f. Using the decomposition

$$P = \sum_{v \in C} L_v(X_v) \cdot R_v(X \setminus X_v),$$

we can show that P is not full-rank with respect to f, which is a contradiction.

In the setting of *set-multilinear* ABPs, we can get a tighter result by showing that there is an $O(\log_m s)$-balanced-chain set system of size at most s where m is the number of variables in each part of the variable-set.

For the other direction, we start with a k-balanced-chain set system $\mathcal{X}$ of size s and show that there is a combinatorial transformation that produces a 1-balanced-chain set system of size $s \cdot n^k$. In the 1-balanced-chain setting, we use an idea that was first used by Raz [Raz06] to construct full-rank polynomials. The observation is that, for a known equipartition $\Pi = (\{y_1, \dots, y_{n/2}\}, \{z_1, \dots, z_{n/2}\})$

of the n variables, the simple 'gadget' polynomial

$$Q_\Pi = \prod_{i=1}^{n/2} (y_i + z_i)$$

is full-rank with respect to Π. To handle an unknown partition, we construct a polynomial that is a sum of such polynomials for all possible equipartitions. A 1-balanced-chain set system gives us a recipe for constructing such a polynomial efficiently. Given any equipartition $\Pi = (Y, Z)$, we know that there is a chain $C = (C_0, \ldots, C_n)$ of sets in the set system such that each set has chain-balance either 0 or 1 with respect to Π. In particular, the even-sized sets $C_0, C_2, C_4, \ldots, C_n$ all are perfectly balanced and $C_{i+2} = C_i \cup \{y_{i/2+1}, z_{i/2+1}\}$ for each even $i < n$. A natural choice to construct a full-rank Q_Π is to pair up $y_{i/2+1}$ with $z_{i/2+1}$. It is possible to design an mABP that computes the sum of all such Q_Π for all chains in a given set system, as in [DMPY12]. To ensure that the resulting polynomial is full-rank with respect to all equipartitions, we need to take a suitable linear combination of all these Q_Π. Raz [Raz06] observed that this can be done over a transcendental extension of the base field $\mathbb{F}$ by using these additional transcendental field elements in the linear combination (the above gadget polynomial is actually from [RSY08b]). We note that if the field $\mathbb{F}$ is infinite, these transcendental elements can be replaced by random elements of $\mathbb{F}$, giving us a full-rank polynomial over the base field. A similar construction (with a different gadget polynomial) also works in the set-multilinear case.

Worst-case to average-case reduction. Assume that we have an (ε, k)-balanced-chain set system $\mathcal{X}$. The main observation is that, given a worst-case balanced partition $f : X \to \{\pm 1\}$, and a uniformly random permutation $\pi : X \to X$, the function $f \circ \pi$ is a uniformly random partition and hence, the set system $\mathcal{X}$ contains a k-balanced-chain with respect to $f \circ \pi$ with probability at least ε. This is the same as saying that the random set system $\mathcal{X}_\pi$ (obtained by permuting the elements of each set in $\mathcal{X}$ according to π) contains a k-balanced-chain with respect to f with probability at least ε. Choosing $O(n/\varepsilon)$ many such permutations π and taking the union of the corresponding set systems allows us to increase this probability to $1 - 2^{-n}$, at which point the probabilistic method guarantees the existence of a suitable set system that works for all partitions.

Constructing a 1-balanced-chain set system of size $n^{O(\log n/ \log\log n)}$. By the preceding argument, it suffices to construct a 1-balanced-chain set system of this size that contains a 1-balanced-chain for a uniformly random balanced partition f with good probability (e.g. $1/\operatorname{poly}(n)$). To do this, we identify the set X with $[n]$ and view $f : X \to \{\pm 1\}$ as a uniformly random walk of length n on the integers that starts and ends at the origin 0 (also known as a random bridge). We note that a maximal chain that is balanced with respect to f contains $n/2$ sets C that are balanced with respect to f. A random walk that returns to 0 many times would give us many *intervals* that satisfy this criterion. A result of Csáki, Erdős and Révész [CER85] tells us that a uniformly random walk f, with probability $\geq 1/\operatorname{poly}(n)$, returns to 0 at least once every $m \approx n/\log n$ steps. This implies that the set of all intervals, while not balanced, does contain a *non-maximal* chain of balanced intervals

$$\varnothing \subseteq [i_1] \subseteq [i_2] \subseteq \cdots \subseteq [i_r] = [n]$$

where each subsequent interval in the chain adds at most m elements to the previous one. To finish the argument, we only need to add sets to 'fill in' the missing sets in the chain to make it maximal.

This can be formulated as the same problem for the smaller set size m. Implementing a recursive construction leads to a set system of the given size.

Remark. While we only used the family of intervals in the above construction, it should be noted that the overall family constructed above consists of sets that are very different from intervals. This is because we need to carry out the aforementioned average-case to worst-case argument at each stage, which is done by randomly permuting the elements in the sets. This does not preserve the property of being an interval. It seems interesting to derandomize this "construction".

Intervals and random partitions. Let $\mathcal{I}$ denote the family of intervals of $[n]$ as defined above. Assume that $f : [n] \to \{\pm 1\}$ is chosen uniformly at random (for the sake of exposition, we drop for now the assumption that f is balanced). The claim is that the chance that there is a maximal chain $(C_0, \ldots, C_n)$ of intervals that make f balanced is exponentially small.

For the sake of exposition, assume that the chain starts at $C_1 = \{n/2\}$. The random walk f can be partitioned into two independent random walks: g that goes from $n/2$ to n and h that goes from $n/2$ to 1. The random walk f is chain-balanced only if there are time parameterizations under which g, h are close at all times (this is a discrete version of the Fréchet distance). We thus have two independent random walks (each of length $n/2$) and the claim is that it is very unlikely that there is a way to parametrize time so the distance between g, h is say $k < n^{1/3}$. The main idea is to first sample the random walk g and identify a pattern in it that h can not possibly follow. Fix the position of the random walk g at times that are multiples of some ℓ. Let p_t be the position of g at time $t\ell$ for $t \leq T := n/(2\ell)$. Typically, the distance between p_t and p_{t+1} is roughly $\sqrt{\ell}$. For concreteness, think of the case that $p_t = 0$ for even t and $p_t = \sqrt{\ell}$ for odd t. It remains to show that there is no sequence of (strictly increasing) time steps $\tau_1 < \tau_2 < \ldots < \tau_T$ such that $h(\tau_t)$ is close to p_t for all t.

Let $(C_0, \ldots, C_n)$ denote any maximal chain in $\mathcal{I}$ and assume that $C_1 = \{i\}$. We can decompose f into two random functions $g : [i] \to \{\pm 1\}$ and $h : [i+1, \ldots, n] \to \{\pm 1\}$. Each set C_j in the chain can be written as the disjoint union of two intervals $I'_j = [i - i' + 1, i]$ and $I''_j = [i + 1, i + j - i']$ contained in the domains of g and h respectively. The quantity $f(C_j)$ is $g(I'_j) + h(I''_j)$ and thus the balance of this chain with respect to f is equal to

$$k := \max_j |g(I'_j) + h(I''_j)| = \max_j |g(I'_j) - (-h(I''_j))|.$$

Thinking of g and $-h$ as random walks on the integers, the chain can be interpreted as an 'alignment' of these two random walks that ensures that they stay within distance k at all time steps, or equivalently, these two random walks are at distance at most k from each other in *discrete Fréchet distance* (defined formally below). The main technical part of the proof is to show that the probability that $k \leq n^{O(1)}$ is exponentially small.

This is done by finding some pattern in one of the two walks, say g, that does not appear in the other except with small probability. To simplify further, assume that g and h are random walks of the same length $n/2$. We divide the time steps of g into epochs of some suitable length ℓ and consider the positions of the random walk g at these various time steps. I.e.

$$p_1 = g([(n/2) - \ell, n/2]), p_2 = g([(n/2) - 2\ell, n/2]), \ldots, p_{n/2\ell} = g([n/2]).$$

As long as ℓ is large enough, there is enough randomness between any pair of time steps to ensure that most pairs of consecutive positions are quite far from one another.[5] Conditioning on these positions, we show that with high probability, there is no subsequence of (strictly increasing) timesteps when the positions of $-h$ are $p'_1, p'_2, \ldots, p'_{n/2\ell}$ such that $|p_t - p'_t| \le k$ for each t. This implies that the two walks cannot be aligned as described above.

To see that $-h$ cannot come close to the positions achieved by g, we think of the case when $p_1 = p_3 = p_5 = \cdots = 0$ and $p_2 = p_4 = p_6 = \cdots = \sqrt{\ell}$. Intuitively, this is a bad case as 0 is the most probable location of a random walk at any (even) time step and we expect a random walk to travel a distance of $\sqrt{\ell}$ every ℓ steps. Now, the expected number of times that a random walk of length $n/2$ returns to 0 is $\Theta(\sqrt{n})$ and we can also prove a concentration bound on the upper tail of this random variable. Thus, if we choose ℓ significantly smaller than $\sqrt{n}$, it is very unlikely that the random walk $-h$ can return to (a position close to) 0 at least $\Omega(n/\ell)$ many times.

1.3 Organization

We start with some preliminaries and then prove some combinatorial properties of the chain-balance of set systems in Section 2, including the worst-case to average-case reduction (Lemma 1.5). We will then show our main combinatorial theorems, starting with the construction of a 1-balanced-chain set system of size $n^{O(\log n/\log\log n)}$ (Theorem 1.6) in Section 3, followed by the analysis of intervals with respect to random partitions (Theorem 1.7) in Section 4. Finally, in Section 5, we prove the formal connection between the combinatorial constructions and the mABP complexity of full-rank polynomials (Theorem 1.3) and derive the consequences of our combinatorial theorems to algebraic complexity.

2 Properties of balanced-chain set systems and their chain-balance

In this section, we prove some properties of balanced-chain set systems related to bounds for the size of these systems and conversions between different notions of chain-balance.

2.1 Upper and lower bounds for the size of balanced-chain set systems

Now let us define a natural class of set systems.

Definition 2.1. For every $n, m \in \mathbb{N}$, we denote by $\mathcal{I}_m := \mathcal{I}_{n,m}$ the *m-interval set system* over $[n]$ defined as

$$\mathcal{I}_{n,m} := \{\, I_1 \cup \ldots \cup I_l \mid I_1, \ldots, I_l \text{ are intervals of } [n] \text{ and } l \le m\}.$$

◀

Assume that $m \le n/2$. Since any element of $\mathcal{I}_{n,m}$ is specified by the endpoints of the intervals, using the usual bounds on the binomial numbers, we get that the size of the m-interval set system is at most

$$|\mathcal{I}_m| \le (n/m)^{O(m)}.$$

We note below that using a simple inductive analysis, we can obtain a 1-balanced-chain set system from intervals.

[5] In the actual argument, we vary the lengths of these epochs to ensure that *every* consecutive pair of positions are indeed far from one another.

Theorem 2.2. For every even number $n \in \mathbb{N}$, the chain-balance of the $(2\lceil \lg n \rceil)$-interval set system is 1.

Proof. We prove this theorem by induction on n. First note that the statement is trivial for $n = 2$. Assume that $n := 2m$ for $m > 1$. Let $l := 2\lceil \lg n \rceil$ and f be a balanced partition of $[n]$. Our goal is to construct a maximal chain of $\mathcal{I}_{n,l}$ such that its elements have balance at most 1 with respect to f.

Informally, we start with the interval $[1, n]$, and look for a point $t \in [1, n]$ such that $f([1, t]) = 0$ (if no such a point exists, then we know that $f(1) \neq f(n)$; so we can pair them by adding the sets $\{1\}$ and $\{1, n\}$ to the chain and define 2 and $n-1$ as new endpoints, and repeat this argument until we find a point t at the same level as the endpoints). We then choose t_1 to be the point satisfying $f([1, t_1]) = 0$ and is the closest to the endpoints 1 or n, which implies that t_1 has distance at most $n/2$ to its closest endpoint t_1'. Now we consider f restricted to either the interval $I' := [t_1', t_1]$ or $I' := [t_1, t_1']$, and, as $|t_1' - t_1| \leq n/2$, we can obtain a maximal chain of I' balanced w.r.t. $f\restriction_{I'}$ using at most $2\lceil \lg n/2 \rceil$ intervals. We concatenate this chain after the chain that we already have, and obtain a f-balanced chain for I' using at most $2\lceil \lg n \rceil$ intervals. Now we repeat the process until we cover the whole interval $[1, n]$.

Formally, we consider the following procedure to inductively build such a maximal chain $(C_0, \ldots, C_n)$. We start with the $C_0 := \varnothing$, and set $l_0 := 1$, and $r_0 := n$. and consider the following cases.

- If there is no $t \in [l_0, r_0]$ such that $f([l_0, t]) = 0$, then we know that f is either positive or negative in $[l_0 + 1, r_0 - 1]$, and we find the smallest $h_1 > 0$ such that there is a $t \in [l_0, r_0]$ such that $|f([l_0, t])| = h_1$. We define, for every $j \in [h_1]$,

$$C_{2j-1} := C_{2j-2} \cup \{l_0 + j - 1\} \text{ and } C_{2j} := C_{2j-1} \cup \{r_0 - j + 1\}.$$

  We set $\ell_1 = \ell_0 + h_1$ and $r_1 = r_0 - h_1$. Note that these sets are unions of two intervals (one starting at 1 and other finishing at n) and have balance at most 1.

- Otherwise, we set $h_1 := 0$ and $l_1 := 1$, and $r_1 := n$.

After the analysis above, we obtain l_1 and r_1 such that $f([l_1, r_1]) = 0$ and there is a $t \in [l_1, r_1]$ such that $f([l_1, t]) = 0$, and we have a maximal balanced chain $(C_0, \ldots, C_{h_1})$ of $[l_0, l_1 - 1] \cup [r_1 + 1, r_0]$ that uses sets that are union of at most 2 intervals (in particular, C_{h_1} is the set $[l_0, l_1 - 1] \cup [r_1 + 1, r_0]$). Let t_1 be the element of $[l_1, r_1]$ closest to either l_1 or r_1 (breaking ties arbitrarily) such that $f([l_1, t_1]) = 0$. Now we consider two cases:

- If t_1 is closer to l_1, then

$$\lambda_1 := t_1 - l_1 + 1 \leq (r_1 - l_1 + 1)/2 \leq n/2,$$

  and $f_1 := f\restriction_{[l_1, t_1]}$ is a balanced partition of $[l_1, t_1]$. By applying induction on the interval $[l_1, t_1]$ and f_1, we obtain a maximal chain $(C_0', \ldots, C_{\lambda_1}')$ such that $C_0' = \varnothing$ and $C_{\lambda_1}' = [l_1, t_1]$, and, for each $j \in [\lambda_1]$, the set $C_j' \subseteq [l_1, t_1]$ is a union of at most $2\lceil \lg \lambda_1 \rceil$ intervals and

$$f\restriction_{[l_1, t_1]}(C_j') \leq 1.$$

- If t_1 is closer to r_1, then we perform similar operation to obtain

$$\lambda_1 := r_1 - t_1 + 1 \leq (r_1 - l_1 + 1)/2 \leq n/2$$

and a maximal chain $(C'_0, \ldots, C'_{\lambda_1})$ such that $C'_0 = \varnothing$ and $C'_{\lambda_1} = [t_1, r_1]$, and, for each $j \in [\lambda_1]$, the set $C'_j \subseteq [t_1, r_1]$ is a union of at most $2\lceil \lg \lambda_1 \rceil$ intervals and

$$f\restriction_{[t_1,r_1]}(C'_j) \leq 1.$$

In both cases, we now concatenate $(C_0, \ldots, C_{h_1})$ with $(C'_0, \ldots, C'_{\lambda_1})$; formally, we define, for every $j \in [\lambda_1]$,

$$C_{h_1+j} := C_{h_1} \cup C'_j.$$

Note that $(C_0, \ldots, C_{h_1+\lambda_1})$ is a chain with all elements in $[1, t_1] \cup [r_1, n]$ if t_1 is closer to l_1, and $[1, l_1] \cup [t_1, n]$ if t_1 is closer to r_1. Moreover, for every $j \in [\lambda_1]$, the set C_{h_1+j} is the union of at most $2 + 2\lceil \lg \lambda_1 \rceil \leq 2\lceil \lg n \rceil$ intervals and

$$f(C_{h_1+j}) = f(C_{h_1}) + f(C'_j) = f(C'_j) \leq 1.$$

Also note that $C_{h_1+\lambda_1}$ is again the union of two intervals $[l_0, l_2] \cup [r_2, r_0]$ for some l_2 and r_2. We now repeat the analysis above using l_2 instead of l_0 and r_2 instead of r_0, until we cover the whole interval $[1, n]$ with a maximal chain of $\mathcal{I}_{n,l}$ that is balanced w.r.t. f. ■

In Section 3, we use a probabilistic result about random balanced partitions, instead of the structural result used in our algorithm above (i.e., the existence of the point t_1), to obtain a 1-balanced-chain set system of size $n^{O(\lg n/\lg\lg n)}$.

Remark. We also note that using previous results about balancing systems, we can prove a $\Omega(n^2/k)$ lower bound for the size of any k-balanced-chain set system for $k \leq n/5$. More specifically, a result by Alon, Kumar, and Volk [AKV20] can be rephrased as follows: for every $k \in \mathbb{N}$ and sufficiently large $n \in \mathbb{N}$, if a collection $S_1, \ldots, S_m \subseteq [n]$ of sets satisfies that $2k \leq |S_i| \leq n - 2k$ for every $i \in [m]$ and, for every balanced partition f of n, there is an $i \in [m]$ such that $|f(S_i)| \leq 2k$, then

$$m \geq \Omega(n/k).$$

Hence, for any $2k$-balanced-chain set system $\mathcal{X}$ over $[n]$, we can consider, for every $l \in \{2k, \ldots, n - 2k\}$, the collection

$$\mathcal{X}_l := \{ S \in \mathcal{X} \mid |S| = l \}.$$

As $\mathrm{cbal}(\mathcal{X}) \leq 2k$, every balanced partition f has a maximal chain $(C_0, \ldots, C_n)$ of $\mathcal{X}$ such that, for every $i \in [n]$, the set C_i satisfies $|C_i| = i$ and

$$|f(C_i)| \leq 2k.$$

By the lower bound above, we obtain that $|\mathcal{X}_l| \geq \Omega(n/k)$, which implies that

$$|\mathcal{X}| \geq \Omega((n - 4k)n/k).$$

2.2 Conversion of (p,l)-balanced-chain systems into l-balanced-chain systems

As mentioned in the introduction, we can prove a worst-case to average-case reduction for balanced-chain set systems.

Lemma 2.3 (Lemma 1.5)**.** Let $n \in \mathbb{N}$ be an even number, and X be a set of n elements. Let $l \in [n]$ and $p \in (0,1]$. If $\mathcal{X}$ is a (p,l)-balanced-chain set system, over X, of size s, then there is an l-balanced-chain set system $\mathcal{Y}$, over X, of size $O(sn/p)$.

Proof. Let $r := \lceil n/(p \lg e) \rceil$. Let $\sigma := (\sigma_1, \dots, \sigma_r)$ be r independent random permutations of X. Let $\mathcal{Y}$ be the set system defined by the union of $\mathcal{X}$ with all the set systems obtained from applying each permutation σ_i to each element of $\mathcal{X}$, that is,

$$\mathcal{Y} := \mathcal{X} \cup (\cup_{i=1}^{r} \sigma_i \mathcal{X}),$$

where $\sigma_i \mathcal{X} := \{ \sigma_i(S) \mid S \in \mathcal{X} \}$. We claim that, with positive probability over σ, the set system $\mathcal{Y}$ has chain-balance $\mathrm{cbal}(\mathcal{Y})$ at most l. Therefore, there is a sequence of permutations such that the corresponding $\mathcal{Y}$ is an l-balanced-chain set system of size $(r+1)s$.

Let us now prove this claim. Let

$$B := \{ f \mid f \text{ is a balanced partition of } X \}$$

and

$$B_{\mathcal{X}} := \{ f \in B \mid \mathrm{cbal}_{\mathcal{X}}(f) \le l \}.$$

Let $g \in B$ be a fixed balanced partition of X, and let π be a random permutation of X. Note that, for $f \sim \mathrm{Unif}(B)$, we have

$$\begin{aligned}
\mathbb{P}_\pi[\forall h \in B_{\mathcal{X}}, g \neq h \circ \pi^{-1}] &= \mathbb{P}_\pi[\forall h \in B_{\mathcal{X}}, g \circ \pi \neq h] \\
&= \mathbb{P}_f[\forall h \in B_{\mathcal{X}}, f \neq h] \\
&= \mathbb{P}_f[f \notin B_{\mathcal{X}}] \\
&= 1 - \mathbb{P}_f[f \in B_{\mathcal{X}}] \\
&\le 1 - p < e^{-p}.
\end{aligned}$$

Thus, we obtain

$$\mathbb{P}_\sigma[\forall i \in [r],\ \forall h \in B_{\mathcal{X}}, g \neq h \circ \sigma_i^{-1}] = \prod_{i=1}^{r} \mathbb{P}_{\sigma_i}[\forall h \in B_{\mathcal{X}}, g \neq h \circ \sigma_i^{-1}] < e^{-rp},$$

and, as a consequence,

$$\mathbb{P}_\sigma[\forall g \in B,\ \exists i \in [r],\ \exists h \in B_{\mathcal{X}} \text{ s.t. } g = h \circ \sigma_i^{-1}] > 1 - 2^n e^{-rp} > 0.$$

Note that, for every σ in the event

$$E := \{\forall g \in B,\ \exists i \in [r],\ \exists h \in B_{\mathcal{X}} \text{ s.t. } g = h \circ \sigma_i^{-1}\},$$

we have that, for every $g \in B$, there are $i \in [r]$ and $h \in B_{\mathcal{X}}$ such that

$$h = g \circ \sigma_i.$$

As $h \in B_{\mathcal{X}}$, there is a chain $C := (C_0, \dots, C_n)$ of $\mathcal{X}$ such that, for every $j \in [n]$,

$$|h(C_j)| \le l.$$

By the definition of $\mathcal{Y}$ and $\sigma_i\mathcal{X}$, we know that $(\sigma_i(C_0), \dots, \sigma_i(C_n))$ is a maximal chain of $\mathcal{Y}$ such that, for every $j \in [n]$,

$$|g(\sigma_i(C_j))| = |(g \circ \sigma_i)(C_j)| = |h(C_j)| \le l.$$

Therefore, for every $\sigma \in E$, we have that, for every $g \in B$,

$$\mathrm{cbal}_{\mathcal{Y}}(g) \le l,$$

which implies that $\mathrm{cbal}(\mathcal{Y}) \le l$. ■

2.3 Conversion of l-balancedness to 1-balancedness

Let us now prove a combinatorial lemma that shows how to convert an l-balanced-chain set system into a 1-balanced-chain set system with some size blow-up.

Lemma 2.4. Let $n \in \mathbb{N}$ be an even number, and let X be a finite set with n elements. If $\mathcal{X}$ is an l-balanced-chain set system, over X, with size s, then there is a set system $\mathcal{Y}$ over X such that

- $\mathcal{Y}$ has chain-balance 1; and
- the size of $\mathcal{Y}$ is at most $s\binom{|X|}{\le l}$.

Proof. Let $n := |X|$ and let $m := n/2$. We define the set system $\mathcal{Y}$ as follows:

$$\mathcal{Y} := \left\{ S \setminus T \;\middle|\; S \in \mathcal{X}, T \in \binom{S}{\le l} \right\}.$$

First note that

$$|\mathcal{Y}| \le |\mathcal{X}| \binom{n}{\le l}.$$

Let us now prove that $\mathcal{Y}$ is 1-balanced. Let f be a balanced partition of X. As $\mathcal{X}$ is l-balanced, we know that there is a chain $C := (S_0, \dots, S_n)$ of $\mathcal{X}$ such that, for every $S \in C$,

$$|f(S)| \le l.$$

Let $\pi\colon [n] \to X$ be the bijection corresponding to the order in which elements of X are added in C. Consider the following process:

1. Let $T_0 := \varnothing$, and $M_0 := \varnothing$, and $l_0 := 0$.
2. For every $i \in [n]$:
   - if $T_{i-1} = \varnothing$, then we set
     $$T_i := \{\pi(i)\},\ l_i := l_{i-1}.$$
     Otherwise, let $x \in T_{i-1}$ be the element of T_{i-1} with maximum $\pi^{-1}(x)$.

- If $f(x) \neq f(\pi(i))$, then we set

$$T_i := T_{i-1} \setminus \{x\},\ l_i := l_{i-1} + 2,\ M_{l_i} := M_{l_{i-1}} \cup \{(x, \pi(i))\},$$

  and $u_{l_i} := \pi(i)$ and $v_{l_i} := x$, and we say that u_{l_i} and v_{l_i} were matched by this process.
- If $f(x) = f(\pi(i))$, then

$$T_i := T_{i-1} \cup \{\pi(i)\},\ l_i := l_{i-1}.$$

This process can be interpreted as an online algorithm that receives elements of $[n]$ in the order defined by C and defines a perfect matching of $[n]$ such that the elements of each matched pair are in different parts of the partition f. By induction on $i \in [n]$, we can prove that the following loop-invariant properties are true:

- We have that $T_i \subseteq S_i$, and $|f^{-1}(T_i)| \leq 1$ (i.e., all elements of T_i are in the same block of f), and $f(S_i \setminus T_i) = 0$, which implies that $f(S_i) = f(T_i)$ and

$$|T_i| = |f(T_i)| = |f(S_i)| \leq \min\{l, n - i\}$$

  as C is an l-balanced chain.

- The set M_{l_i} is a perfect matching of the set $S_i \setminus T_i$ (i.e., $V(M_{l_i}) = S_i \setminus T_i$) such that $|M_{l_i}| = l_i$ and every element $(u, v) \in M_{l_i}$ satisfies $f(u) \neq f(v)$.

- The matching $M_{l_{i-1}}$ is a submatching of M_{l_i}.

Moreover, we can show that the sets $\{\, l_i \mid i \in [n]\}$ and $\{\, 2i \mid i \in \{0, \ldots, m\}\}$ are the same. For every $i \in [m]$, we define

$$C_{2i-1} := V(M_{2i-2}) \cup \{u_{2i}\}, \text{ and } C_{2i} := V(M_{2i-2}) \cup \{u_{2i}, v_{2i}\}.$$

Let us now prove that $C' := (C_0, \ldots, C_n)$ is a 1-balanced chain of $\mathcal{Y}$. By the loop-invariants above, we know that, for every $i \in [m]$,

$$C_{2(i-1)} = V(M_{2i-2}) \subseteq V(M_{2i}) = C_{2i},$$

and that $|V(M_{2i-2})| = 2i - 2$ and $|V(M_{2i})| = 2i$. Moreover, there is a $j \in [n]$ such that

$$V(M_{2i}) = S_j \setminus T_j,$$

where $S_j \in \mathcal{X}$ and $T_j \subseteq S_j$ with $|T_j| \leq l$, so $C_{2i} = V(M_{2i}) \in \mathcal{Y}$. Using a similar argument, we can show that

$$C_{2i-1} \in \mathcal{Y}.$$

Hence, C' is a maximal chain of $\mathcal{Y}$. By the definition of M_{2i}, we have that, for every $i \in [m]$,

$$f(C_{2i}) = f(V(M_{2i})) = 0 \text{ and } |f(C_{2i-1})| = |f(C_{2i})| + 1 = 1.$$

Thus, C' is a 1-balanced chain of $\mathcal{Y}$ with respect to f. As f is an arbitrary balanced partition of X, we conclude that $\mathcal{Y}$ is a 1-balanced-chain set system. ■

3 Construction of a 1-balanced-chain set system of size $n^{O(\ln n/\ln\ln n)}$

In this section, our goal is to prove Theorem 1.6. Our proof uses concepts from one-dimensional random walk theory (Definition 3.1) and a result about the probability distribution of the length of the longest excursion in a random bridge (Lemma 3.2). In Section 5.3, the main result of this section is used to obtain a separation between multilinear formulas and mABPs.

Definition 3.1. Let $n \in \mathbb{N}$ and let $f\colon [n] \to \{-1, 1\}$. We associate a function $W_f\colon \{0, \dots, n\} \to \mathbb{Z}$ to f as follows: let $W_f(0) := 0$ and, for every $i \in [n]$,

$$W_f(i) := f([i]).$$

We call W_f the *walk defined from* f, and n the length of W_f. Note that the definition above is a bijection between f and W_f. Thus, for a given walk W, we denote by f_W the function such that $W = W_{f_W}$. Sometimes we also consider W_f to be the sequence $(W_f(0), \dots, W_f(n))$.

We say that W_f is a *bridge* if f satisfies $f([n]) = 0$. We say that W_f is a *excursion* if W_f is a bridge and, for every $i \in [n-1]$,

$$W_f(i) \neq 0.$$

For f sampled uniformly at random from the set of all functions from $[n]$ to $\{-1, 1\}$, we say $W := W_f$ is a *random walk* of length n, that is,

$$W \sim \mathrm{Unif}(\{\, W_f \mid f\colon [n] \to \{-1, 1\}\}.$$

Simarly, for even n, we define a *random bridge* B of length n as

$$B \sim \mathrm{Unif}(\{\, W_f \mid f\colon [n] \to \{-1, 1\}, f([n]) = 0\}).$$

Let B be a random bridge of length n and $f := f_B$. We say that a point $t \in \{0, \dots, n\}$ is a *zero* of B if

$$f([t]) = 0,$$

or, equivalently, $B(t) = 0$. Let $Z(B)$ be the (random) set of zeros of B. Note that $\{0, n\} \subseteq Z(B)$ by the definition of bridge. Let $\{z_0, \dots, z_l\} := Z(B)$ be the elements of $Z(B)$, where $z_0 < z_1 < \dots < z_l$. Let $\lambda(B)$ be the random variable defined as follows:

$$\lambda(B) := \max_{i \in [l]} z_i - z_{i-1}.$$

We call $\lambda(B)$ the *length of the longest excursion* in B. ◀

Lemma 3.2 (Lemma 10 from [CER85]). Let $n \in \mathbb{N}$ and let B be a random bridge of length $2n$. There are universal real constants $0 < C_1 \leq C_2$ and $\beta > 0$ such that, for every $a \geq n^{2/3}$, we have

$$\mathbb{P}_B[\lambda(B) \leq 2a] = C(n, a) \min\{(n+1)^{-1/2}, a^{-1/2}\} \exp(-\beta n/a),$$

for $C_1 \leq C(n, a) \leq C_2$.

Let us now prove our main theorem.

Theorem 3.3 (Theorem 1.6)**.** For every sufficiently large even number $n \in \mathbb{N}$, there is a 1-balanced-chain set system $\mathcal{X}_n$ over $[n]$ of size $n^{O(\ln n/\ln\ln n)}$.

Proof. For every $n \in \mathbb{N}$, let $N(n)$ be the minimum size of a 1-balanced-chain set system over $[n]$, and let A_n be a fixed 1-balanced-chain set system of size $N(n)$. Our goal is to obtain a recursive upper bound for $N(n)$.

For any $m \in [n]$, let B_m be the set system defined as follows:

$$B_m := \{\, [1, i] \odot A_t \mid t \in [m], 0 \le i \le n - t\},$$

where we define $[1, i] \odot A_t$ as

$$[1, i] \odot A_t := \{\, [1, i] \cup \{\, i + j \mid j \in R\} \mid R \in A_t\},$$

which can be interpreted as the translation of the set system A_t to the end of the interval $[1, i]$. Note that B_m is a set system over $[n]$ of size at most $n^2 N(m)$. Let f be a uniformly random balanced partition of $[n]$. We say that a point $t \in [n]$ is a *zero* of f if $\sum_{i \in [t]} f(i) = 0$, and we say that a set system A is *f-balanced* if $\mathrm{cbal}_A(f) \le 1$. Let E_m be the event that the maximum gap between two consecutive zeros of f is at most m, and let

$$\delta_m := \mathbb{P}_f[E_m].$$

Note that, by the definition of B_m, we get that if $f \in E_m$, then B_m is f-balanced. This follows from the fact that if $f \in E_m$, then the maximum gap between two consecutive zeros z_i and z_{i+1} of f is at most m, so we use the set system

$$[1, z_i] \odot A_{z_{i+1} - z_i}$$

to find a subchain of B_m that is $(f\restriction_{[z_i, z_{i+1}]})$-balanced, and then we just concatenate all these chain to find a f-balanced chain of B_m. By Lemma 3.2 and interpreting its result about random bridges B in the language of random balanced partitions f, there are positive universal constants $b, c \in \mathbb{R}$ such that, for every even $m \in [n]$ such that $(n/2)^{2/3} \le m/2$, we have

$$\delta_m = \mathbb{P}_B[\lambda(B) \le m] \ge c n^{-1/2} \exp(-bn/m).$$

Now fix the value of m to $m := 2\lceil n/(2 \ln n)\rceil$, which implies that

$$\delta := \delta_m \ge c/n^{b+1/2}.$$

Hence, for this choice of m, we have that B_m is a $(\delta, 1)$-balanced-chain set system over $[n]$. By Lemma 2.3, there is a 1-balanced-chain set system A, over $[n]$, of size

$$|A| \le O(|B_m| n/\delta) \le (n^2 N(n/\ln n)) n^{b+2} \le n^d N(n/\ln n),$$

for a positive constant d. By solving the recurrence relation $N(n) \le n^d N(n/\ln n)$, we obtain that

$$N(n) \le n^{O(\ln n/\ln\ln n)}.$$

∎

4 The chain-balance of interval set systems and the Fréchet distance of random walks

In this section, our main goal is to prove Theorem 1.7. In order to prove this theorem, we establish a connection between the chain-balance of interval set systems and the Fréchet distance between two random walks (Lemma 4.2), and prove that, with very high probability, the Fréchet distance of two random walks of length n is at least $n^{\Omega(1)}$ (Lemma 4.3). In Section 5.5, the main result of this section is used to obtain lower bounds on the size of sum of interval-mABPs computing full rank polynomials.

4.1 Fréchet distance of random walks

Let us start by defining the Fréchet distance of two discrete functions [EM94].

Definition 4.1. Let $n \in \mathbb{N}$, and let $l, r \in [n]$ such that $l + r = n$. Let $X\colon \{0, \ldots, l\} \to \mathbb{R}$ and $Y\colon \{0, \ldots, r\} \to \mathbb{R}$ be two functions. We define the *(discrete) Fréchet distance* $d_F(X, Y)$ between X and Y as

$$d_F(X, Y) := \min_{\alpha,\beta} \max_{t\in[n]} |X(\alpha(t)) - Y(\beta(t))|,$$

where the minimum ranges over all non-decreasing functions

$$\alpha\colon \{0, \ldots, n\} \to \{0, \ldots, l\} \text{ and } \beta\colon \{0, \ldots, n\} \to \{0, \ldots, r\}$$

satisfying the following properties:

- $\alpha(0) = 0 = \beta(0)$;
- $\alpha(n) = l$ and $\beta(n) = r$;
- For every $t \in [n]$, we have $\alpha(t) \le \alpha(t - 1) + 1$, and $\beta(t) \le \beta(t - 1) + 1$, and

$$\alpha(t) - \alpha(t - 1) + \beta(t) - \beta(t - 1) = 1.$$

◀

Now let us state the main technical lemmas of this section.

Lemma 4.2 (Discrepancy and Fréchet distance)**.** Let $n \in \mathbb{N}$ be an even number, and let $\mathcal{I} := \mathcal{I}_{n,1}$ be the 1-interval set system (Definition 2.1). Let $l, r \in [n]$ such that $l+r = n$, and let $W_{l,r} := (W_l, W_r)$ be a pair of two independent random walks W_l and W_r (Definition 3.1) of length l and r, respectively. For f being a uniformly random balanced partition, and $W := (X, Y) := W_{n,n}$, and $X_l := X{\restriction}_{\{0,\ldots,l\}}$ and $Y_r := Y{\restriction}_{\{0,\ldots,r\}}$, we get that, for every $k \in \mathbb{R}$,

$$\mathbb{P}_f[\mathrm{cbal}_{\mathcal{I}}(f) \le k] \le O(n^{5/2}) \max_{r\in[n]} \mathbb{P}_W[d_F(X_r, Y_{n-r}) \le k].$$

Proof. This lemma is proved in Section 4.2. ■

Lemma 4.3 (Fréchet distance of two random walks)**.** There is a positive constant $c_1 \in \mathbb{R}$ such that, for every sufficiently large $n \in \mathbb{N}$, and every $l := l(n) \in [n]$, and every $\varepsilon := \varepsilon(n) \in (0, 1/4)$ such that $n^{1/4-\varepsilon} \ge \frac{2}{3} \ln n$, we have, for $(X, Y) := W := W_{l,n-l}$,

$$\mathbb{P}_W[d_F(X, Y) < n^{1/4-\varepsilon}] \le \exp(-c_1 \cdot \min\{n^{1/4-\varepsilon}, n^{4\varepsilon}\}).$$

Proof. This lemma is proved in Section 4.4. ∎

Let us now prove Theorem 1.7.

Theorem 4.4 (Theorem 1.7)**.** There is a universal constant $c > 0$ such that, for every sufficiently large even number $n \in \mathbb{N}$, the 1-interval set system $\mathcal{I}_{n,1}$ is not an (ε, k)-balanced-chain set system for $\varepsilon > 2^{-cn^{1/5}}$ and $k < n^{1/5}$.

Proof. Suppose $\mathcal{I} := \mathcal{I}_{n,1}$ is an (ε, k)-balanced-chain set system for any $\varepsilon > 2^{-cn^{1/5}}$ and $k < n^{1/5}$ with $c := c_1/2$. By Lemma 4.2, we have that

$$\mathbb{P}_f[\mathrm{cbal}_{\mathcal{I}}(f) \le n^{1/5}] \le O(n^{5/2}) \max_{r \in [n]} \mathbb{P}_W[d_F(X_r, Y_{n-r}) \le n^{1/5}],$$

and, by Lemma 4.3,

$$\max_{r \in [n]} \mathbb{P}_W[d_F(X_r, Y_{n-r}) \le n^{1/5}] \le \exp(-c_1 n^{1/5}).$$

Hence,

$$\mathbb{P}_f[\mathrm{cbal}_{\mathcal{I}}(f) \le k] \le \mathbb{P}_f[\mathrm{cbal}_{\mathcal{I}}(f) \le n^{1/5}] \le O(n^{5/2}) \exp(-c_1 n^{1/5}) \le \exp(-c_1 n^{1/5}/2) = \exp(-cn^{1/5}),$$

a contradiction. ∎

4.2 Proof of Lemma 4.2

Let $\mathcal{I} := \mathcal{I}_{n,1}$ be the 1-interval set system over $[n]$. By the definition of chain-balance (Definition 1.2), we know that, for any balanced partition f of $[n]$ satisfying $\mathrm{cbal}_{\mathcal{I}}(f) \le k$, there is a maximal chain $C := C_f$ of $\mathcal{I}$ such that $|f(S)| \le k$ for every $S \in C$. Let s and e be the first and last elements of $[n]$ added to the chain C, and let $P := [s, e]$ be the cyclic interval[6] with extremities s and e such that $s + 1 \in P$, and $N := [n] \setminus P$ be the complement interval of P. For each element R of C with size $m := |R|$, let

$$S_+(m) := R \cap P, \text{ and } S_-(m) := R \cap N.$$

Note that for every nonempty $S_+(m)$ and $S_-(m)$, there are a_m and b_m in $[n]$ such that

$$S_+(m) = [s, a_m] \text{ and } S_-(m) = [b_m, s - 1], \tag{2}$$

where $[s, a_m]$ is the cyclic interval with extremities s and a_m that intersects P, and $[b_m, s-1]$ is the cyclic interval with extremities b_m and $s - 1$ that does not intersect P. Let

$$l := |P| \text{ and } r := |N|, \tag{3}$$

and, for every $m \in \{0, \ldots, n\}$, let

$$\alpha(m) := |S_+(m)| \text{ and } \beta(m) := |S_-(m)|. \tag{4}$$

Thus, the functions $\alpha\colon \{0, \ldots, n\} \to \{0, \ldots, l\}$ and $\beta\colon \{0, \ldots, n\} \to \{0, \ldots, r\}$ satisfy the following properties:

[6] We say that a set $I \subseteq [n]$ is a cyclic interval if there is a pair (s, e) such that: if $s \le e$, then $I = \{x \in [n] | s \le x \le e\}$; if $s > e$, then $I = \{x \in [n] \mid s \le x\} \cup \{x \in [n] \mid x \le e\}$.

- $\alpha(0) = 0 = \beta(0)$;
- $\alpha(n) = l$ and $\beta(n) = r$;
- For every $t \in [n]$, we have $\alpha(t) \leq \alpha(t-1)+1$, and $\beta(t) \leq \beta(t-1)+1$, and

$$\alpha(t) - \alpha(t-1) + \beta(t) - \beta(t-1) = 1.$$

Also note that, for every R in C with size $m := |R|$, the set $\{S_+(m), S_-(m)\}$ is a partition of R, so

$$f(R) = f(S_+(m)) + f(S_-(m)).$$

Furthermore, as C is a maximal chain of $\mathcal{I}$ avoiding the set of k-unbalanced elements, with respect to f, of $\mathcal{I}$, we get that

$$k \geq |f(R)| = |f(S_+(m)) + f(S_-(m))|,$$

Therefore, for every s and e in $[n]$, there is an injection from

the set $\mathcal{F}_{s,e}$ of balanced partitions f of $[n]$ such that $\text{cbal}_{\mathcal{I}}(f) \leq k$ and its chain C_f has s as first element and e as last element

to

the set of pairs (X, Y) of functions such that, for $l_{s,e} := l$ and $r_{s,e} := r$ defined from s and e as in Equation 3, X and Y are walks (Definition 3.1) of length l and r, respectively, and $d_F(X, Y) \leq k$.

This injection can be defined as follows. For each balanced partition f, we obtain the sequence $(a_m)_{m\in[n]}$ and $(b_m)_{m\in[n]}$ (as defined by Equation 2) and the function α and β (Equation 4). Then we map f to the pair (X, Y) of walks of length l and r defined as follows:

- $X(0) := 0 =: Y(0)$;
- For every $i \in [l]$ and $j \in [r]$, let x_i and y_j be respectively the i-th and j-th distinct values of the sequences $(a_m)_{m\in[n]}$ and $(b_m)_{m\in[n]}$ and let

$$X(i) := f([s, x_i]) \text{ and}$$
$$Y(j) := -f([y_j, s-1]).$$

Note that, for every $t \in [n]$,

$$X(\alpha(t)) = f([s, a_{\alpha(t)+\beta(t)}]) = f([s, a_t]) \text{ and}$$
$$Y(\beta(t)) = -f([b_{\alpha(t)+\beta(t)}, s-1]) = -f([b_t, s-1]).$$

Therefore, we get that X and Y are walks of length l and r, respectively, and, for every $t \in [n]$,

$$|X(\alpha(t)) - Y(\beta(t))| = |f([s, a_t]) - (-f([b_t, s-1]))| = |f(S_+(t)) + f(S_-(t))| \leq k,$$

which implies that $d_F(X, Y) \leq k$. Note that this mapping from f to (X, Y) is an injection, because we obtain X and Y from f by considering the sequences $(s+1, \ldots, e)$ and $(s-1, \ldots, e-1)$ (i.e., $(a_m)_m$ and $(b_m)_m$ without repetitions) and then defining $X(i) = f([s, s+i])$ and $Y = -f([s-i, s-1])$.

Finally, by the union bound over the first and last elements added to the chain C_f, we get that, for g being a uniform random partition of $[n]$ (i.e., $g \sim \text{Unif}\left(\{-1,1\}^{[n]}\right)$),

$$\begin{aligned}
\mathbb{P}_f[\text{cbal}_{\mathcal{I}}(f) \le k] &= \mathbb{P}_f[\exists s, e \in [n] \text{ s.t. } f \in \mathcal{F}_{s,e}] \le \sum_{s,e\in[n]} \mathbb{P}_f[f \in \mathcal{F}_{s,e}] \\
&= \sum_{s,e\in[n]} \mathbb{P}_g[g \in \mathcal{F}_{s,e} | g \text{ is balanced}] \\
&\le O(n^{1/2}) \sum_{s,e\in[n]} \mathbb{P}_g[g \in \mathcal{F}_{s,e}] \\
&\le O(n^{1/2}) \sum_{s,e\in[n]} \mathbb{P}_W[d(X_{l_{s,e}}, Y_{r_{s,e}}) \le k] \\
&\le O(n^{1/2}) \sum_{s,e\in[n]} \mathbb{P}_W[d(X_{|[s,e]|}, Y_{n-|[s,e]|}) \le k] \\
&\le O(n^{5/2}) \max_{l\in[n]} \mathbb{P}_W[d_F(X_l, Y_{n-l}) \le k].
\end{aligned}$$

4.3 Properties of random walks

In this section, we will prove some technical lemmas regarding properties of random walks that will be useful in the proof of Lemma 4.3. Our first lemma (Lemma 4.5) shows an asymptotic bound for the probability distribution of the first passage time of a random walk through a point.

Lemma 4.5. Let g be a random walk of length n. For any $\delta \in \mathbb{N}$, let F_δ be the random variable defined as follows:

$$F_\delta := \inf\{\, t > 0 \mid g(t) = \delta\}.$$

Then, for any $\delta := \delta(n) \in \mathbb{N}_{>0}$ and $z := z(n) \in \mathbb{N}$ such that $\Omega(\delta^2) \le z \le n/2$, we have

$$\mathbb{P}_g[F_\delta \ge z] = \Theta\left(\frac{\delta}{\sqrt{z}}\right).$$

Proof. For every $y \in \mathbb{N}$ such that $y + \delta$ is an even integer, we can use a standard fact about the distribution of first passage times for simple random walk (e.g. see Theorem 2 in Page 89 from [Fel68] or Equation 3.7 in page 29 from [BW21]) to obtain that

$$\mathbb{P}_g[F_\delta = y] = \frac{\delta}{y2^y}\binom{y}{\frac{y+\delta}{2}}.$$

If $y \to \infty$ and $y - \delta \to \infty$ as $n \to \infty$, we can apply Stirling's approximation and obtain

$$\begin{aligned}
\mathbb{P}_g[F_\delta = y] &= \Theta\left(\frac{\delta}{y2^y}\sqrt{\frac{1}{y(1-(\delta/y)^2)}}\frac{y^y}{(\frac{y+\delta}{2})^{\frac{y+\delta}{2}}(\frac{y-\delta}{2})^{\frac{y-\delta}{2}}}\right) \\
&= \Theta\left(\frac{\delta}{y^{3/2}}\frac{1}{(1+\delta/y)^{\frac{y+\delta}{2}}(1-\delta/y)^{\frac{y-\delta}{2}}}\right) \\
&= \frac{\delta}{y^{3/2}} \cdot \exp(-\Theta(\delta^2/y)).
\end{aligned}$$

Note that if $y \geq \Omega(\delta^2)$, then the estimate above become

$$\mathbb{P}_g[F_\delta = y] = \Theta\left(\frac{\delta}{y^{3/2}}\right),$$

and, as a consequence, for z satisfying $\Omega(\delta^2) \leq z \leq n/2$, we have

$$\mathbb{P}_g[F_\delta \geq z] = \sum_{y=z}^{n} \mathbb{P}_g[F_\delta = y] = \Theta\left(\frac{\delta}{\sqrt{z}}\right).$$

■

Lemma 4.6 below analyses the probability distribution of the sum of independent first passage random variables and shows a lower tail bound.

Lemma 4.6 (Lower tail bound for sums of first passage random variables)**.** There is a positive constant $c_4 \in \mathbb{R}$ such that, for every $r, n \in \mathbb{N}$, and every $\delta := \delta(n) \in \mathbb{N}$ and $k := k(n) \in \mathbb{N}$ such that

$$k\delta^2 \leq c_4 n \text{ and } (k\delta)^2 = \omega(n),$$

and for g being a random walk of length r, we have that if $F^{(1)}, \ldots, F^{(k)}$ are k independent copies of the random variable $F_\delta(g)$, then

$$\mathbb{P}_g\left[\sum_{i=1}^{k} F^{(i)} \leq n\right] \leq \exp(-\Omega((k\delta)^2/n)).$$

Proof. Let $t \in [n/2]$ be a parameter to be optimized later. For each $i \in [k]$, let $C_{t,i} := 1_{\{F^{(i)} \geq t\}}$ be a random variable. Note that $C_{t,1}, \ldots, C_{t,k}$ are i.i.d. Bernoulli random variables with success probability $p_t := \mathbb{P}_g[F_\delta \geq t]$. Let $C_t := \sum_{i=1}^{k} C_{t,i}$ be a random variable and let $\mu_t := \mathbb{E}[C_t]$ be its expectation. Note that, by Lemma 4.5, there is a positive constant c such that if $c\delta^2 \leq t \leq n/2$, then

$$\mu_t = kp_t \geq ck\delta/\sqrt{t},$$

and, by the Chernoff bound (see e.g. [DP09, Theorem 1.1]), we have

$$\mathbb{P}_g\left[C_t \leq \frac{\mu_t}{2}\right] \leq e^{-\Omega(\mu_t)}.$$

Moreover, if $\sum_{i=1}^{k} F^{(i)} \leq n$ holds, then $C_t \leq n/t$ holds. Hence, we obtain

$$\begin{aligned}
\mathbb{P}_g\left[\sum_{i=1}^{k} F^{(i)} \leq n\right] &\leq \min_{\substack{t \in [n/2] \\ t \geq c\delta^2 \\ n/t \leq ck\delta/(2\sqrt{t})}} \mathbb{P}_g[C_t \leq n/t] \\
&\leq \min_{\substack{t \in [n/2] \\ t \geq c\delta^2 \\ n/t \leq ck\delta/(2\sqrt{t})}} \mathbb{P}_g[C_t \leq \mu_t/2] \\
&\leq \min_{\substack{t \in [n/2] \\ t \geq c\delta^2 \\ n/t \leq ck\delta/(2\sqrt{t})}} e^{-\Omega(k\delta/\sqrt{t})}.
\end{aligned}$$

Note that, for

$$t := \left(\frac{2n}{ck\delta}\right)^2,$$

the following hold:

- $n/t \le ck\delta/(2\sqrt{t})$;
- if $(k\delta)^2 \ge \omega(n)$, then $t \le n/2$ for sufficiently large n;
- if $k\delta^2 \le c_4 n$ for $c_4 := 2/c^{3/2}$, then

$$t \ge \left(\frac{2n\delta}{cc_4 n}\right)^2 = \delta^2 \left(\frac{2}{cc_4}\right)^2 = c\delta^2.$$

Therefore, we obtain

$$\mathbb{P}_g\left[\sum_{i=1}^{k} F^{(i)} \le n\right] \le e^{-\Omega(k\delta/\sqrt{t})} = e^{-\Omega((k\delta)^2/n)}.$$

■

Lemma 4.7 shows that we can find, with high probability, a long spread sequence of points in the image of a long random walk g. In the next section, we will use such a sequence as a sequence of 'milestones' for g and show that an independent random walk h is very unlikely to hit the same milestones and this will allow us to bound the Fréchet distance of g and h.

Lemma 4.7. Let g be a random walk of length n and $\Delta := \Delta(n) \in \mathbb{N}$ with $\Delta \ge 2\log n$. We say that a sequence $\{x_i\}_{i\in[l]} \subseteq [n]$ is a (g, Δ)-*sequence* of length l if, for every $i \in [l]$, we have $|g(x_i) - g(x_{i-1})| \ge \Delta$, with $x_0 := 0$. Let L_Δ be the random variable equal to the maximum length of a (g, Δ)-sequence. Then, there is an universal positive constant c_3 such that, for $l(n, \Delta) := n/c_3\Delta^3$, we have

$$\mathbb{P}_g[L_\Delta < l(n, \Delta)] \le \exp(-\Omega(\Delta)).$$

Proof. Let $c_3 \in \mathbb{R}$ be a sufficiently large positive constant, whose value is defined later, and let $l := l(n, \Delta)$. Let D_Δ be the random variable defined as follows:

$$D_\Delta := \inf\{\, t > 0 \mid |g(t)| = \Delta\}.$$

Let us first obtain an upper bound for $\mathbb{P}_g[D_\Delta > c_3\Delta^3]$. Consider the first $T := c_3\Delta^3$ steps of g and call this random walk h. Split h further into Δ segments $(h_1, \ldots, h_\Delta)$, each of length T/Δ. In the event that $D_\Delta > T$, it must be the case that the random walk h is in the range $[-\Delta, \Delta]$ after each of the segments $(h_1, \ldots, h_i)$. In particular, this implies that the random walk h_i is contained in the range $[-2\Delta, 2\Delta]$ for each $i \in [\Delta]$. Note that for each $i \in [\Delta]$

$$\mathbb{P}_{h_i}\left[-2\Delta < h_i(j) < 2\Delta \;\forall j \in [c_3\Delta^2]\right] \le \mathbb{P}_{h_i}\left[F_{2\Delta} \ge c_3\Delta^2\right] \le O(2\Delta/(c_3\Delta^2)^{1/2}) < 1/2,$$

where the second to last inequality follows from Lemma 4.5 and the last inequality from choosing c_3 to be a sufficiently large constant. Hence, by considering all the segments $h_1, \ldots, h_\Delta$ of h, we get

$$\mathbb{P}_g[D_\Delta > c_3\Delta^3] \le 2^{-\Delta}.$$

Let $D^{(1)}, \dots, D^{(l)}$ be l independent copies of D_Δ (recall $l = n/c_3\Delta^3$). Note that if $\sum_{i=1}^{l} D^{(i)} \leq n$, then $L_\Delta \geq l$, by defining a (g, Δ)-sequence from the points obtained by each $D^{(i)}$. Therefore, we obtain

$$\begin{aligned}
\mathbb{P}_g[L_\Delta < l] \leq \mathbb{P}_g\left[\sum_{i=1}^{l} D^{(i)} > n\right] &\leq \mathbb{P}_g[\exists i \in [l] \text{ s.t. } D^{(i)} > n/l] \\
&\leq \sum_{i=1}^{l} \mathbb{P}_g[D^{(i)} > c_3\Delta^3] \leq l2^{-\Delta} \leq 2^{-\Delta+\log(n/c_3\Delta^3)} \\
&\leq e^{-\Omega(\Delta)},
\end{aligned}$$

where the last inequality follows from $\Delta \geq 2\log n$. ■

4.4 Proof of Lemma 4.3

Informal outline. We consider two random walks g and h of total length n. Assume w.l.o.g. that g has length at least $n/2$. Let Δ be a parameter that we will choose later. Using Lemma 4.7, we observe that w.h.p. there are approximately $k \approx n/\Delta^3$ points in g such that every consecutive pair of points is Δ-far from each other. We then analyze the probability that h can reach close to each of these points (which is required for the Fréchet distance between them to be small). The number of steps to go from one of these points to the next is captured by a first passage random variable F_δ with distance parameter $\delta = \Theta(\Delta)$. For the Fréchet distance to be small, the sum of k independent copies of these random variables must be at most n. However, we can bound the probability of this using Lemma 4.6. This probability is small as long as $k^2\delta^2 \gg n$ which is true for small enough Δ.

We now give the formal details. Assume without loss of generality that the random walk X has length l greater than the length $r := n - l$ of Y. Let $d := n^{1/4-\varepsilon}$, and $\Delta := 3d$, and $L := n/(2c_3\Delta^3)$. For the random walk Y and for every $z_1, \dots, z_k \in \mathbb{Z}$, define the following random variables:

$$T_{z_1} := \inf\{\, t \geq 0 \mid |Y(t) - z_1| \leq d\},$$

and, for every $i \in \{2, \dots, k\}$, let

$$T_{z_1,\dots,z_i} := \inf\{\, t > T_{z_1,\dots,z_{i-1}} \mid |Y(t) - z_i| \leq d\}. \tag{5}$$

Note that

$$\mathbb{P}_W[d_F(X,Y) < d] = \mathbb{P}_W[d_F(X,Y) < d \wedge L_\Delta(X) < L] + \mathbb{P}_W[d_F(X,Y) < d \wedge L_\Delta(X) \geq L]$$

and, by Lemma 4.7,

$$\mathbb{P}_W[d_F(X,Y) < d] \leq \exp(-\Omega(\Delta)) + \mathbb{P}_W[d_F(X,Y) < d \wedge L_\Delta(X) \geq L].$$

Condition on any X in the event $\{d_F(X,Y) < d \wedge L_\Delta(X) \geq L\}$, and let $\{x_i\}_{i \in L_\Delta(X)} \subseteq \mathbb{Z}$ be a fixed (X, Δ)-sequence. Let z_i denote $X(x_i)$ for each $i \in [L]$, and note that $|z_i - z_{i-1}| \geq \Delta$ for all $i \in [L]$ (for $z_0 := 0$).

By the definition of the Fréchet distance, if $d_F(X,Y) \le d$, then we know that, for every strictly increasing sequence $(a_1,\dots,a_m)$ of $\{0,\dots,l\}$, there is a strictly increasing sequence $(b_1,\dots,b_m)$ of $\{0,\dots,r\}$ such that

$$|X(a_i) - Y(b_i)| \le d$$

for every $i \in [m]$. This implies that $T_{X(x_1),\dots,X(x_L)}(Y) \le r$, thus

$$\mathbb{P}_W[d_F(X,Y) < d \wedge L_\Delta(X) \ge L \mid X] \le \mathbb{P}_W[T_{z_1,\dots,z_L}(Y) \le r].$$

By the definition of $T_{z_1,\dots,z_L}$ (Equation 5) and the translation invariance of random walks, we have

$$T_{z_1,\dots,z_L}(Y) = F_{|Y(T_{z_0,\dots,z_{L-1}}(Y)) - z_L| - d} + T_{z_1,\dots,z_{L-1}}(Y),$$

which implies that

$$T_{z_1,\dots,z_L}(Y) = \sum_{i=1}^{L} F_{|Y(T_{z_0,\dots,z_{i-1}}(Y)) - z_i| - d}.$$

As $|Y(T_{z_0,\dots,z_{i-1}}(g)) - z_{i-1}| \le d$ for all $i \in [L]$, we have that

$$\Delta \le |z_i - z_{i-1}| \le |z_i - Y(T_{z_0,\dots,z_{i-1}}(Y))| + |Y(T_{z_0,\dots,z_{i-1}}(Y)) - z_{i-1}| \le |z_i - Y(T_{z_0,\dots,z_{i-1}}(Y))| + d,$$

and, consequently, $F_{\Delta-2d} \le F_{|z_i - Y(T_{z_0,\dots,z_{i-1}}(Y))| - d}$. Thus, we obtain

$$\mathbb{P}_Y[T_{z_1,\dots,z_L}(Y) \le r] \le \mathbb{P}_Y\left[\sum_{i=1}^{L} F_{\Delta-2d}^{(i)} \le r\right],$$

for $F_{\Delta-2d}^{(1)},\dots,F_{\Delta-2d}^{(L)}$ being L independent copies of $F_{\Delta-2d}$. Let $k := L$ and $\delta := \Delta - 2d = d$ so that

$$k\delta^2 = \frac{n}{2c_3\Delta^3} \cdot \frac{\Delta^2}{3^2} \le c_4 n/2, \text{ and } k\delta = \frac{n}{2c_3\Delta^3} \cdot \frac{\Delta}{3} \ge \Omega\left(\frac{n}{n^{1/2-2\varepsilon}}\right) \ge \omega(\sqrt{n}) \tag{6}$$

for sufficiently large n. Thus, by Lemma 4.6 using k and δ defined above, which satisfies the lemma's hypothesis by Equation 6, we get

$$\begin{aligned}
\mathbb{P}_Y[T_{z_1,\dots,z_L}(Y) \le r] &\le \mathbb{P}_Y\left[\sum_{i=1}^{L} F_{\Delta-2d}^{(i)} \le r\right] \\
&\le \mathbb{P}_Y\left[\sum_{i=1}^{k} F_{\delta}^{(i)} \le n/2\right] \\
&\le \exp(-\Omega((k\delta)^2/n)) \\
&\le \exp\left(-\Omega\left(\left(\frac{n}{2c_3\Delta^3} d\right)^2 /n\right)\right) \\
&\le \exp(-\Omega(n/d^4)).
\end{aligned}$$

Therefore, we obtain

$$\begin{aligned}
\mathbb{P}_W[d_F(X,Y) < d] &\le \exp(-\Omega(\Delta)) + \exp(-\Omega(n/d^4)) \\
&\le \exp(-\Omega(\min\{d, n/d^4\})) \\
&\le \exp(-c_1 \cdot \min\{n^{1/4-\epsilon}, n^{4\epsilon}\}),
\end{aligned}$$

for some absolute constant c_1.

5 Full rank multilinear ABPs and balanced-chain set systems

In this section, our goals are to show a relationship between constructions of mABPs computing full rank polynomials and balanced-chain set systems, and use this relationship together with our results from Sections 3 and 4 to, respectively, obtain a separation between the power of multilinear formulas and mABPs to compute full rank polynomials, and prove lower bounds for the size of sums of *interval-mABPs* computing full rank polynomials. We first recall the definition of coefficient matrix used by the min-partition rank criterion and the main properties of its rank (Section 5.1.2), and then we proceed to prove the conversion from balanced-chain set systems to full-rank mABPs (Section 5.2) and its consequences (Section 5.3), and the conversion from full-rank mABPs to balanced-chain set systems (Section 5.4) and its consequences (Section 5.5).

5.1 Preliminaries

5.1.1 Multilinear and Set-Multilinear polynomials

Given a set X of variables, a multivariate polynomial $P(X)$ is *multilinear* if it is a linear combination of multilinear monomials, i.e. monomials in which the individual degree of each variable is at most 1.

Given a partition $\mathcal{P} := \{X_1, \dots, X_n\}$ of X, we say that a multivariate polynomial $\tilde{P}(X)$ is *set-multilinear w.r.t.* $\mathcal{P}$ (or just set-multilinear if $\mathcal{P}$ is clear from context) if $\tilde{P}$ is a linear combination of set-multilinear monomials, which are monomials that are a product of exactly one variable from each X_i. In what follows, we will mostly be interested in the case when $\mathcal{P}$ is an *equipartition*, i.e. $|X_1| = |X_2| = \cdots = |X_n| = N$.

A set-multilinear ABP w.r.t. $\mathcal{P}$ is defined similarly to an mABP except that every edge is labelled with a homogeneous linear polynomial in one of the variable sets in the partition $\mathcal{P}$ and no two edges in a path can be labelled with polynomials over the same variable set. Note that the paths from any vertex u to any vertex v in such an ABP always compute a set-multilinear polynomial w.r.t. a partition $\mathcal{P}' \subseteq \mathcal{P}$.

5.1.2 Coefficient matrices

We defined already the notion of a coefficient matrix above but we repeat the definition here and also widen it to include set-multilinear polynomials.

Definition 5.1 (Coefficient matrices for multilinear and set-multilinear polynomials)**.** Let X be a set of variables and $P(X)$ be a multilinear polynomial. Recall that we define by a *partition* as a function $f : X \to \{\pm 1\}$.

For a partition f, we define the corresponding coefficient matrix $M_f(P)$ to be the following $2^r \times 2^c$ matrix where $r := |f^{-1}(1)|$ and $c := |f^{-1}(-1)|$. The rows of $M_f(P)$ are labelled by multilinear monomials in $f^{-1}(1)$ and the columns by multilinear monomials in $f^{-1}(-1)$. Given multilinear monomials m_b over $f^{-1}(b)$ for each $b \in \{\pm 1\}$, we define the (m_1, m_{-1})th entry of $M_f(P)$ to be the coefficient of their product $m_1 m_{-1}$ in the polynomial P.

Now we consider the case of a set-multilinear polynomial $\tilde{P}(X)$ w.r.t. a partition $\mathcal{P} := \{X_1, \dots, X_n\}$ of X. In this case a partition is defined by $f : \mathcal{P} \to \{-1, 1\}$, and the corresponding coefficient matrix $M_f(\tilde{P})$ has rows and columns labelled by *set-multilinear* monomials w.r.t. the partitions $f^{-1}(1)$ and $f^{-1}(-1)$ respectively. The entries of the coefficient matrix are defined in the same way. ◀

The following simple properties of the above notion will be useful.

Fact 5.2. The coefficient matrix defined above satisfies the following properties:

1. **Subadditivity:** Assume that a multilinear polynomial $P(X)$ is the sum of multilinear polynomials $Q_1(X), \ldots, Q_s(X)$. Then, for any partition f,

$$\mathrm{rank}(M_f(P)) \leq \sum_{i=1}^{s} \mathrm{rank}(M_f(Q_i)).$$

   This is also true if $P, Q_1, \ldots, Q_s$ are all set-multilinear w.r.t. the same partition.

2. **Disjoint multiplicativity:** Assume that a multilinear polynomial $P(X)$ is the product of two multilinear polynomials $Q(X')$ and $R(X \setminus X')$ on disjoint sets of variables. Then

$$\mathrm{rank}(M_f(P)) = \mathrm{rank}(M_{f|_{X'}}(Q)) \cdot \mathrm{rank}(M_{f|_{X \setminus X'}}(R))$$

   where $f|_{X'}$ and $f|_{X \setminus X'}$ represent the restriction of f to X' and $X \setminus X'$ respectively.

   A similar fact is also true in the case that $P(X)$ is a set-multilinear polyomial w.r.t. $\mathcal{P} := \{X_1, \ldots, X_n\}$ and Q, R are set-multilinear polynomials w.r.t. $\mathcal{P}' \subseteq \mathcal{P}$ and $\mathcal{P} \setminus \mathcal{P}'$ respectively.

3. **Trivial bound:** The rank of any matrix is bounded by the number of rows and the number of columns of the matrix. In the multilinear case, this implies

$$\mathrm{rank}(M_f(P)) \leq 2^{\min\{|f^{-1}(1)|, |f^{-1}(-1)|\}}$$

   for any $f : X \to \{\pm 1\}$.

   In the case of a set-multilinear polynomial P w.r.t. an equipartition $\mathcal{P} := \{X_1, \ldots, X_n\}$ and a partition $f : \mathcal{P} \to \{\pm 1\}$, we have similarly

$$\mathrm{rank}(M_f(P)) \leq N^{\min\{|f^{-1}(1)|, |f^{-1}(-1)|\}},$$

   where $N := |X_1| = \cdots = |X_n|$.

We now define the notion of a full-rank polynomial, which is the main hardness criterion that is used to prove lower bounds in the multilinear setting [Raz09, Raz06].

Definition 5.3 (Full-rank polynomials)**.** We start with the multilinear case. Let X be a set of n variables with n even. A multilinear polynomial $P(X)$ is said to be *full-rank* if its rank w.r.t. every balanced partition $f : X \to \{\pm 1\}$ is as large as possible, i.e. $2^{n/2}$. More generally for $p, \varepsilon \in [0, 1]$ the polynomial P is said to be (p, ε)-*almost full-rank* if for a uniformly random balanced partition f

$$\mathbb{P}_f \left[\mathrm{rank}(M_f(P)) \geq \varepsilon 2^{n/2} \right] \geq p.$$

We have identical definitions for set-multilinear polynomials w.r.t. an equipartition into n sets of size N each, except that the maximum rank is $N^{n/2}$. ◀

5.2 Construction of full rank multilinear ABPs from balanced-chain set systems

In this section, our goals are to prove our main construction of full rank mABPs from 1-balanced-chain set systems (Theorem 5.4) and generalize it to the case of l-balanced-chain set systems (Theorem 5.5). In order to obtain our generalization, we use a combinatorial lemma (Lemma 2.4) that shows how to convert an l-balanced-chain set system into a 1-balanced-chain set system with some size blow-up.

We start with the 1-balanced case. As is standard [Raz06], we will start by proving it by enlarging the field by adding transcendental elements. Later on, we will show how to remove the transcendental elements at the expense of making the polynomial less 'explicit' (by replacing these elements by random elements from the field).

Theorem 5.4. Let $n \in \mathbb{N}$ be an even number, and let X be a set of n elements. Let $W_X := \{\, w_t | t := (i,u) \in [n] \times X\}$ and $V_X := \{\, x_i \mid i \in X\}$ be two disjoint sets of variables with, respectively, n^2 and n elements. If $\mathcal{X}$ is a 1-balanced-chain set system over X, then there is a polynomial $P_{\mathcal{X}}$ over the variables $V_X \cup W_X$ such that the following hold:

- All coefficients of $P_{\mathcal{X}}$ are 0 or 1;
- For every field $\mathbb{F}$, the polynomial $P_{\mathcal{X}}$ is full rank over $\mathbb{F}(W_X)[V_X]$ and can be computed by an mABP of size at most $|\mathcal{X}|$ over $\mathbb{F}(W_X)$.

Proof. Let $l := n/2$ and $m := |\mathcal{X}|$. For any maximal chain $C \in \mathcal{C}(\mathcal{X})$, let $\pi_C \colon [n] \to X$ be the bijection defined by the order in which the elements of X are added in the chain C. We define the 0/1 polynomials Q_C over the variables V_X and P_C over the variables $V_X \cup W_X$ as follows:

$$Q_C := \prod_{i=1}^{l} (x_{\pi_C(2i-1)} + x_{\pi_C(2i)}), \text{ and}$$

$$P_C := Q_C \cdot \prod_{i=1}^{2l} w_{i,\pi_C(i)}.$$

Then let $P_{\mathcal{X}}$ be the following 0/1 polynomial over the variables $V_X \cup W_X$:

$$P := P_{\mathcal{X}} := \sum_{C \in \mathcal{C}(\mathcal{X})} P_C.$$

For any field $\mathbb{F}$, let us first construct an mABP M of size at most m computing P over $\mathbb{F}(W_X)$. We add all the elements of the set system $\mathcal{X}$ with even cardinality to the nodes of M, and choose the source and sink of M to be the nodes corresponding to the empty set and $[n]$ respectively. For every chain (R, S, T) in $\mathcal{X}$ between two sets R and T with even cardinality and

$$|T| = |S| + 1 = |R| + 2,$$

let u, v be the elements of X such that $\{u\} = S \setminus R$ and $\{v\} = T \setminus S$. We then add the edge (R, T) to the edges of M and label it with the following linear form of $\mathbb{F}(W_X)[V_X]$:

$$l(R,T) := w_{|S|,u}(x_u + x_v)w_{|T|,v}. \tag{7}$$

By the definition of P, we can show that M computes P, as there is a bijective correspondence between the source-to-sink paths and their labels in M and the maximal chains $C \in \mathcal{C}(X)$ of $\mathcal{X}$ and their corresponding terms

$$\prod_{i=1}^{l} w_{2i-1,\pi_C(2i-1)}(x_{\pi_C(2i-1)} + x_{\pi_C(2i)})w_{2i,\pi_C(2i)}$$

in the polynomial P. It is also not hard to show that M is syntactically multilinear, as the variables x_u's appear at most once in each computational path of M (again using the correspondence between the source-to-sink paths of M and the maximal chains of $\mathcal{X}$). Note that we consider M an mABP over the field $\mathbb{F}(W_X)$, so the labels are linear forms.

Now let us prove that P is a full rank polynomial over $\mathbb{F}(W_X)$. Let f be a balanced partition of X. By the definition of chain-balance (Definition 1.2) and the fact that $\mathcal{X}$ is 1-balanced, we know that there is a maximal chain $C := C_f$ of $\mathcal{X}$ such that every element S of C satisfies $|f(S)| \leq 1$. Note that if we also interpret the polynomial P as a polynomial in $\mathbb{F}[V_X \cup W_X]$, then we get that

$$Q_C = P\restriction_{\{w_{i,u}=[\pi_C(i)=u]\ \forall i\in[n],\ u\in X\}}, \tag{8}$$

which is a polynomial in $\mathbb{F}[V_X]$. By standard inequalities between rank of polynomial matrices after substitution by zero and one values, we get

$$\text{rank}_{\mathbb{F}(W_X)}(M_f(P)) \geq \text{rank}_{\mathbb{F}}(M_f(Q_C)),$$

since any rank-1 decomposition of $M_F(P)$ over $\mathbb{F}(W_X)$ can be converted into a rank-1 decomposition of $M_f(Q_C)$ over $\mathbb{F}$ by using the substitution in Equation 8. As the i-th term of the product in Q_C is a polynomial with respect to $\{x_{\pi(2i-1)}, x_{\pi(2i)}\}$ and $f(\pi(2i-1)) \neq f(\pi(2i))$ (by the fact that C is a 1-balanced chain w.r.t. f), we get that

$$\begin{aligned}\text{rank}_{\mathbb{F}}(M_f(Q_C)) &= \prod_{i=1}^{l} \text{rank}_{\mathbb{F}}(M_f(x_{\pi(2i-1)} + x_{\pi(2i)}))\\ &= \prod_{i=1}^{l} 2 = 2^l,\end{aligned}$$

where the first equality follows from the multiplicative property of the coefficient matrix with respect to a partition of the variables (Fact 5.2). Therefore, for every balanced f,

$$\text{rank}_{\mathbb{F}(W_X)}(M_f(P)) \geq \text{rank}_{\mathbb{F}}(M_f(Q_C)) \geq 2^l,$$

thus P is a full rank polynomial over $\mathbb{F}(W_X)$. ■

More generally, we have the following stronger version of Theorem 5.4 for l-balanced-chain set systems.

Theorem 5.5. Let $n \in \mathbb{N}$ be an even number, let $l \in [n]$ and X finite with n elements. Let $W_X := \{\, w_t \mid t := (i, u) \in [n] \times X\}$ and $V_X := \{\, x_i \mid i \in X\}$ be two disjoint sets of variables with, respectively, n^2 and n elements. If $\mathcal{X}$ is a l-balanced-chain set system over X, then there is a polynomial $P_{\mathcal{X}}$ over the variables $V_X \cup W_X$ such that

- All coefficients of $P_{\mathcal{X}}$ are 0 or 1;
- For every field $\mathbb{F}$, the polynomial $P_{\mathcal{X}}$ is full rank over $\mathbb{F}(W_X)[V_X]$ and can be computed by an mABP of size at most $|\mathcal{X}|\binom{n}{\leq l}$ over $\mathbb{F}(W_X)$.

Proof. By applying Lemma 2.4 to the l-balanced-chain set system $\mathcal{X}$ given as input for Theorem 5.5, we obtain a 1-balanced-chain set system $\mathcal{Y}$ over X with size at most $|\mathcal{X}|\binom{n}{\leq l}$. Now we apply Theorem 5.4 to $\mathcal{Y}$ and obtain the polynomial $P_{\mathcal{Y}}$ over the variables $V_X \cup W_X$ such that

- All coefficients of $P_{\mathcal{Y}}$ are 0 or 1;
- For every field $\mathbb{F}$, the polynomial $P_{\mathcal{Y}}$ is full rank over $\mathbb{F}(W_X)[V_X]$ and can be computed by an mABP of size at most $|\mathcal{Y}|\binom{n}{\leq l}$ over $\mathbb{F}(W_X)$.

Therefore, $P_{\mathcal{Y}}$ is a polynomial satisfying the properties required by Theorem 5.5. ■

5.3 A separation between multilinear formulas and mABPs

As mentioned in the introduction, we can use Theorem 1.3 together with Theorem 3.3 to obtain a superpolynomial separation between multilinear formulas and mABPs. We show this separation (Corollary 5.7) in this section, after we prove the following formal version of the second part of Theorem 1.3.

Theorem 5.6 (Second part of Theorem 1.3)**.** Let $n \in \mathbb{N}$ be an even number, let $l \in [n]$ and X finite with n elements. If $\mathcal{X}$ is a l-balanced-chain set system over X, then, for every infinite field $\mathbb{F}$, there is an n-variate full rank polynomial $P_{\mathcal{X}}$ over $\mathbb{F}$ that can be computed by an mABP over $\mathbb{F}$ with size at most $|\mathcal{X}|\binom{n}{\leq l}$.

Proof. Let $m := n/2$. By applying Theorem 5.5 with the l-balanced-chain set system $\mathcal{X}$, we obtain a polynomial P over the variables $V_X \cup W_X$ (as defined in Theorem 5.5) such that, for every field $\mathbb{F}$, the polynomial P is full rank over $\mathbb{F}(W_X)[V_X]$ and can be computed by an mABP of size at most $|\mathcal{X}|\binom{n}{\leq l}$ over $\mathbb{F}(W_X)$. By the definition of full-rank polynomial (Definition 5.3), we have that, for every balanced partition f of X, there is a 2^m-square submatrix $M_f := M_f(P)[R_f, C_f]$, for some sets R_f and C_f, of $M_f(P)$ with nonzero determinant (as a polynomial over $\mathbb{F}(W_X)$). By Definition 5.1 and the fact that the degree of $P_{\mathcal{X}}$ over the variables in W_X is at most n, we know that all the entries of $M_f(P)$ are polynomials of $\mathbb{F}[W_X]$ of degree at most n, so the determinant of M_f is a non-zero polynomial of degree at most $\Delta := n2^m$. We denote by D_f the determinant of M_f:

$$D_f := \det(M_f).$$

Let B_n be the set of all balanced partitions of $[n]$, and let $S \subseteq \mathbb{F}$ be a fixed subset of $\mathbb{F}$ of cardinality $|B_n|\Delta|W_X|$. By Ore-DeMillo-Lipton-Schwartz-Zippel lemma [Ore22, DL78, Sch80, Zip79], we know that, for w being a uniformly random element of S^{W_X},

$$\mathbb{P}_w[D_f(w) = 0] \leq \frac{\deg(D_f)}{|S|} \leq \frac{\Delta}{|S|},$$

and, consequently,

$$\mathbb{P}_w[\exists f \in B_n \text{ s.t. } D_f(w) = 0] \leq \frac{|B_n|\Delta}{|S|} = \frac{1}{|W_X|}.$$

Hence, there is a $w \in S^{W_X}$ such that, for every $f \in B_n$, we have

$$D_f(w) \neq 0.$$

Let $Q \in \mathbb{F}[V_X]$ be the polynomial obtained by substituting all the variables in W_X of P by the corresponding value in w. Note that, for every $f \in B_n$, all the entries of $M_f(Q)$ are equal to the value of the corresponding entry of $M_f(P)$ evaluated at w. Thus, we obtain

$$\det(M_f(Q)[R_f, C_f]) = \det(M_f(P)[R_f, C_f])(w) = D_f(w) \neq 0,$$

which implies that $M_f(Q)$ has rank at least 2^m. Therefore, $Q \in \mathbb{F}[V_W]$ is a full rank polynomial over $\mathbb{F}$.

Finally, we just need to prove that Q can be computed by an mABP over $\mathbb{F}$ with size at most $|\mathcal{X}|\binom{n}{\leq l}$. This follows from the fact that we can change all linear forms on the labels of any mABP computing P over $\mathbb{F}(W_X)$ to the evaluation of this label on w (which is a linear form over $\mathbb{F}$) in order to obtain a mABP of the same size computing Q. ■

Now let us use Theorem 5.6 and Theorem 3.3 to prove a separation between of multilinear formulas and mABPs.

Corollary 5.7. Let $n \in \mathbb{N}$ be a sufficiently large even number, and $\mathbb{F}$ be an infinite field. Then there is a n-variate polynomial P_n over $\mathbb{F}$ computed by a mABP of size $n^{O(\ln n/\ln\ln n)}$, but any multilinear formula computing P_n has size at least $n^{\Omega(\ln n)}$.

Proof. Let $\mathcal{X} := \mathcal{X}_n$ be the 1-balanced-chain set system over $[n]$ of size $n^{O(\ln n/\ln\ln n)}$ obtained from Theorem 3.3. Let $P := P_n$ be the n-variate polynomial obtained from Theorem 5.6 using $\mathcal{X}$ as its 1-balanced-chain set system. Hence, P is a full rank polynomial that can be computed by an mABP of size at most

$$|\mathcal{X}|n = n^{O(\ln n/\ln\ln n)},$$

and, by Theorem 1.1, any multilinear formula over $\mathbb{F}$ computing P has size at least $n^{\Omega(\ln n)}$. ■

5.4 Construction of balanced-chain set systems from full-rank multilinear ABPs

In this section, our goal is to prove our formal version of the first part of Theorem 1.3. In particular, we are going to prove a more general theorem (Theorem 5.13), which allows to relax the hypothesis from the first part of Theorem 1.3 to average-case almost-full-rank guarantees.

Notation. Given an mABP $\mathcal{A}$ with source a and sink b, and a vertex v of $\mathcal{A}$, we use $X_v(\mathcal{A})$ to denote the set of all variables that appear in the label of an edge on a a-v path. Similarly, we use $\overline{X}_v(\mathcal{A})$ to denote the set of all variables that appear in the label of an edge on a v-b path. If the mABP $\mathcal{A}$ is clear from context, we simplify the notation to $X_v, \overline{X}_v$.

We will now define a structured kind of mABP for which some arguments are simpler. The following notion may be thought of as an analogue of 'homogeneous ABPs' in the multilinear setting.

Definition 5.8 (Valid labelings and layered mABPs)**.** Fix an mABP $\mathcal{A}$ computing a polynomial P in the variable set X. A *valid labeling* of $\mathcal{A}$ is an assignment to each vertex v a subset $Y_v \subseteq X$

such that $X_v \subseteq Y_v \subseteq X \setminus \overline{X}_v$; further, if (u, v) is an edge of $\mathcal{A}$, it must be the case that $Y_u \subseteq Y_v$. Note that the assignment $Y_v := X_v(\mathcal{A})$ in particular meets the conditions of this definition.

A *layered mABP* is an mABP along with a valid labeling such that the following holds. Partition the vertex set V of the mABP into subsets $V_0, \ldots, V_n$ where $n = |X|$, where the i-th layer contains exactly those vertices v such that their vertex label Y_v has size i. Then

All edges go from vertices in layer i, for some i, to layer $i+1$. Further, the label of an edge (u, v) is a linear polynomial in the unique variable in $Y_v \setminus Y_u$, or a field constant.

◀

The following two statements showcase how we will use valid labelings.

Observation 5.9. If $\mathcal{A}$ is an mABP, we note that $X_v(\mathcal{A}) \cap \overline{X}_v(\mathcal{A}) = \varnothing$ for any vertex v of $\mathcal{A}$.

Further, if we consider the sub-mABPs $\mathcal{A}'_v$ and $\mathcal{A}''_v$ obtained by redefining v as the sink or source respectively, we note that the polynomial L_v computed by $\mathcal{A}'_v$ is a multilinear polynomial in the variables $X_v(\mathcal{A})$ and the polynomial R_v computed by $\mathcal{A}''_v$ is a multilinear polynomial in the variables $\overline{X}_v(\mathcal{A})$.

If $Y \subseteq X$ is such that $X_v(\mathcal{A}) \subseteq Y \subseteq \overline{X}_v(\mathcal{A})$, we can also consider L_v as a multilinear polynomial in Y and R_v a multilinear polynomial in $X \setminus Y$.

Lemma 5.10 (mABP decomposition based on any vertex cut)**.** Let $\mathcal{A}$ be any mABP over a set of variables X with a valid labeling. Let C be any cut in $\mathcal{A}$ separating the source a from the sink b. Then we can decompose the polynomial P computed by $\mathcal{A}$ as

$$P = \sum_{v \in C} L_v(Y_v) \cdot R_v(X \setminus Y_v)$$

where for each v, the set $Y_v \subseteq X$ is the label of v, and the polynomials L_v and R_v are multilinear polynomials on the specified subsets of X.

Proof. We prove the lemma by induction on the size of the cut C.

If $|C| = 0$, then there are no a-b paths in the ABP and the polynomial computed by the ABP is 0. Thus the lemma follows trivially in this case.

Otherwise, let $v \in C$ be any fixed vertex in C. Removing v from $\mathcal{A}$ yields an mABP $\mathcal{A}'$ and a cut $C' := C \setminus \{v\}$ of smaller cardinality than C. We note that the same labeling is still a valid labelling for $\mathcal{A}'$. By the induction hypothesis, we have

$$P' = \sum_{v' \in C'} L_{v'}(Y_{v'}) \cdot R_{v'}(X \setminus Y_{v'}).$$

Finally, we observe that the difference $P - P'$ is the sum, over all paths π in $\mathcal{A}$ *containing* v, of the product of the labels of edges in π. This sum can be factored into the corresponding sum of all a-v paths and the sum of all v-b paths. Using Observation 5.9, we are led to the factorization

$$P - P' = L_v(Y_v) \cdot R_v(X \setminus Y_v)$$

where L_v corresponds to the sum of all a-v paths and R_v to the sum of all v-b paths. This implies the lemma statement. ■

With the above lemma in hand, we can prove the main result for layered mABPs.

Lemma 5.11 (Balanced-chain set systems from layered mABPs)**.** Let $n \in \mathbb{N}$ be an even number, and X be a set of n variables. Let $p := p(n) \in [0,1], \varepsilon := \varepsilon(n) \in (0,1]$. Fix any *layered mABP* $\mathcal{A}$ of size s computing a (p,ε)-almost full-rank polynomial P over the variables X, and let Y_v denote the label of each vertex $v \in V(\mathcal{A})$ under the valid labeling guaranteed by the definition of layered mABPs. Then

$$\mathcal{X} := \{ Y_v \mid v \in V(\mathcal{A})\}$$

is a $(p, \lg(s/\varepsilon))$-balanced-chain set system.

Proof. Since $\mathcal{A}$ is (p,ε)-almost full-rank, we know that for a uniformly random balanced partition $f : X \to \{\pm 1\}$, we have

$$\mathbb{P}_f \left[\operatorname{rank}(M_f(P)) \geq \varepsilon 2^{n/2}\right] \geq p. \tag{9}$$

We will show that $\mathcal{X}$ as defined in the statement of the lemma is $(1+\lg(s/\varepsilon))$-balanced w.r.t. any f that satisfies the event in (9) above, which will prove the lemma. For the rest of the lemma, fix such an f.

Define $k := \lg(s/\varepsilon)$ and let U denote the set of vertices v whose labels that are k-imbalanced w.r.t. f, i.e.,

$$U := \{v \in V(\mathcal{A}) \mid |f(Y_v)| > k\}.$$

If there is a path from the source a of $\mathcal{A}$ to its sink b that avoids all the vertices in U, then we are done because the labels of the vertices along the path yield a maximal chain of subsets in $\mathcal{X}$ that is k-balanced w.r.t. f.

On the other hand, if there is no such path, then the set of vertices U form an a-b cut. In this case, we will derive a contraditicon to the fact that $M_f(P)$ has high rank. By Lemma 5.10 above, we can write

$$P = \sum_{v \in U} L_v(Y_v) R_v(X \setminus Y_v)$$

for multilinear polynomials L_v and R_v in the given sets of variables. By the subadditivity and multiplicativity (Fact 5.2) of the rank of the coefficient matrix, we see that

$$\begin{aligned}
\operatorname{rank}(M_f(P)) &\leq \sum_{v\in U} \operatorname{rank}(M_f(L_v \cdot R_v)) \\
&\leq s \cdot \max_{v \in U} \operatorname{rank}(M_f(L_v \cdot R_v)) \\
&= s \cdot \max_{v \in U} \operatorname{rank}(M_{f|_{Y_v}}(L_v)) \cdot \operatorname{rank}(M_{f|_{X\setminus Y_v}}(R_v))
\end{aligned}$$

where $f|_Y$ represents the restriction of f to Y. Fix any $v \in U$ and note that $M_{f|_{Y_v}}(L_v)$ is a matrix with dimensions $2^{r_v} \times 2^{c_v}$ where

$$|r_v - c_v| = |f(Y_v)| > k.$$

Thus we have

$$\operatorname{rank}(M_{f|_{Y_v}}(L_v)) \leq 2^{\min\{r_v, c_v\}} \leq 2^{(r_v+c_v)/2 - |r_v - c_v|/2} < 2^{(r_v+c_v)/2-k/2}.$$

In a similar way, we also see that

$$M_{f|_{X\setminus Y_v}}(R_v) \leq 2^{\min\{r'_v, c'_v\}} \leq 2^{(r'_v+c'_v)/2 - |r'_v - c'_v|/2} < 2^{(r'_v+c'_v)/2-k/2}.$$

where $M_{f|_{X\setminus Y_v}}(R_v)$ is a $2^{r'_v} \times 2^{c'_v}$ matrix, and

$$|r'_v - c'_v| = |n - r_v - (n - c_v)| = |r_v - c_v| > k.$$

Plugging these rank bounds into the computation above, we get

$$\varepsilon \cdot 2^{n/2} \leq \text{rank}(M_f(P_i)) < s \cdot \max_{v \in U} 2^{(r_v+c_v)/2+(r'_v+c'_v)/2-k} = s \cdot 2^{n/2-k} \tag{10}$$

which is a contradiction since $k = \lg(s/\varepsilon)$. ■

To finish the proof of the main theorem of this section, we show how to convert any mABP to a layered mABP with a not too large blowup.

Lemma 5.12 (General mABPs to layered mABPs)**.** Let $\mathcal{A}$ be any mABP of size s computing a multilinear polynomial P is a set X of n variables. Then, there is a layered mABP $\mathcal{A}'$ of size $O(s^2 n)$ computing P.

Proof. We construct $\mathcal{A}'$ by starting with the mABP $\mathcal{A}$ and adding more vertices. We also start with the valid vertex labeling that assigns to each vertex v the set $Y_v := X_v(\mathcal{A})$. This partitions the vertex set into $n+1$ layers $V_0, \ldots, V_n$ depending on the size of Y_v.

To fulfill the condition that edges go from one layer to the next (see the item in Definition 5.8), we proceed as follows. Call an edge $e := (u, v)$ *bad* if it joins u in layer i with v in layer $j \neq i+1$. Since we have a valid labeling, it must be the case that $Y_u \subseteq Y_v$ and hence that $j \geq i$. Hence for a bad edge, we have either $j = i$ (both endpoints are within the same layer) or $j > i+1$ (endpoints in different layers).

We remove the bad edges in two steps (at all times maintaining a valid labeling).

Bad edges within the same layer. Arrange the vertices of the mABP in topological order $u_1 := a, \ldots, u_s := b$. As long as there are bad edges within the same layer, we repeat the following. We start with a bad edge $e := (u_i, u_j)$ with *length* $\text{len}(e) := |i - j|$ as small as possible. The label of e must be a field constant α as u_i and u_j are in the same layer (and hence $Y_{u_i} = Y_{u_j}$). We remove the edge e and then do the following.

- If $i \neq 1$, we add in the edge $e' := (u_k, u_j)$ for each edge (u_k, u_i) in the mABP. The label of e' is set to α times the label of (u_k, u_i).
- If $i = 1$, then note that $j \neq s$ (otherwise the polynomial computed by the mABP is a constant and the lemma is trivial). In this case, we add the edge $e' := (u_i, u_k)$ for each edge (u_j, u_k) in the mABP. The label of e' is again set to α times the label of (u_j, u_k).

In both cases, each path containing e is replaced by a different path computing the same polynomial. Hence, the polynomial computed by the new mABP is the same. Further, the same labeling remains valid. Finally, the above process can create parallel edges, which can be handled by replacing them with a single edge whose label is the sum of the labels of the parallel edges.

This way, we reduce the number of bad edges with the smallest possible length and only create bad edges of a greater length. Repeating this process removes all bad edges between the same layer. Note that this process maintains the size of the mABP.

From now onwards, we will assume that there are no more bad edges within the same layer.

Bad edges between different layers. Fix any bad edge $e := (u, v)$ where $u \in V_i$ and $v \in V_j$ for $j > i + 1$. The label of e is a linear polynomial which is w.l.o.g.

$$\alpha_0 + \sum_{i=1}^{r} \alpha_i x_i.$$

Note that $\{x_1, \ldots, x_r\} \subseteq Y_v \setminus Y_u$. By adding some more summands to the label of e with coefficient 0, we can assume w.l.o.g. that in fact $Y_v = Y_u \cup \{x_1, \ldots, x_r\}$ and hence $r = j - i > 1$. We now remove e, and add a simple layered gadget with $O(n)$ vertices between vertices u and v that will simulate the presence of e. The gadget is given in Figure 1. Note that the gadget computes the polynomial $\alpha_0 + \sum_{i=1}^{r} \alpha_i x_i$. It is also naturally layered in $r + 1$ layers as shown. Assigning every vertex w in layer i, for $i \in \{2, \ldots, r\}$, the label

$$Y_w := Y_u \cup \{x_1, \ldots, x_{i-1}\}$$

leads to a valid labeling for the mABP.

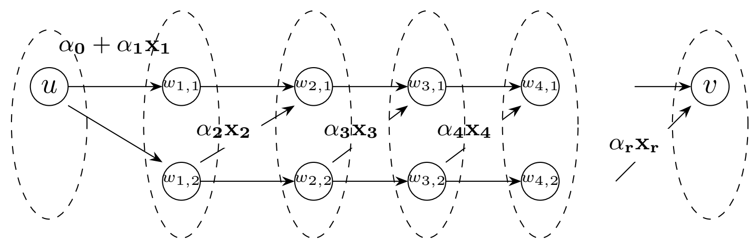

Figure 1: The gagdet construction used in Lemma 5.12. All unlabeled edges are assumed to have the label 1. The dashed ellipses show the natural layering of the gadget into $t + 1$ layers $0, \ldots, t$.

Repeating the above process for all the bad edges leads to a layered mABP $\mathcal{A}'$ computing the same polynomial. Since $\mathcal{A}$ has size s and we add $O(n)$ edges for each edge of $\mathcal{A}$, the mABP $\mathcal{A}'$ has size $O(s^2 n)$. This concludes the proof of the lemma. ■

The above results directly imply a stronger version of the first part of Theorem 1.3 from the introduction.

Theorem 5.13 (Balanced-chain set systems from layered mABPs (first part of Theorem 1.3))**.** Let $n \in \mathbb{N}$ be an even number, and X be a set of n variables. Let $p := p(n) \in [0, 1], \varepsilon := \varepsilon(n) \in (0, 1]$. Fix any mABP $\mathcal{A}$ of size $s \geq n$ computing a (p, ε)-almost full-rank polynomial P over the variables X, Then there exists a $(p, O(\log(sn/\varepsilon))$-balanced-chain set-system $\mathcal{X}$ of size at most $O(s^2 n)$.

Proof. Follows directly from the statements of Lemma 5.12 and Lemma 5.11. ■

5.5 Lower bounds against interval-mABPs

We show in this section how to use the combinatorial results in Section 4 and the connection between multilinear ABPs and balanced-chain set systems to obtain a stronger version of a result of [CKSS24]. While we state the result for multilinear ABPs, the same idea also works in the set-multilinear setting.

Let X be a set of n variables. Given a bijection $\pi : X \to [n]$, we denote by $\mathcal{I}_{X,\pi}$ the set system consisting of intervals with respect to π contained in the set X. Formally, we have

$$\mathcal{I}_{X,\pi} := \{\{x_{\pi(i)}, x_{\pi(i+1)}, \ldots, x_{\pi(j)}\} \mid 1 \leq i, j \leq n\}.$$

Interval-mABPs. A π-interval mABP is defined to be an ABP whose vertices are divided into $n+1$ layers. Each vertex v is labelled with a π-*interval* $Y_v \in \mathcal{I}_{X,\pi}$, and all vertices from the i-th layer are labelled with sets of size i. All edges go from vertices in layer i, for some i, to vertices in layer $i+1$. An edge (u, v) from vertex u in layer i to vertex v in layer $i+1$ is labelled by either a linear polynomial in the unique variable $x \in Y_v \setminus Y_u$, or a field constant. We say that an ABP $\mathcal{A}$ is an *interval mABP* if there is a bijection π such that $\mathcal{A}$ is a π-interval mABP.

We make a simple observation regarding interval mABPs that follows easily from their definition.

Remark. Fix any vertex v in a π-interval mABP. Then all paths from the source a to v are only labelled by variables in Y_v and paths from v to the sink b are only labelled by variables in $X \setminus Y_v$. In particular, for any π, a π-interval mABP $\mathcal{A}$ is indeed an mABP, and the labelling $\{Y_v\}_{v \in V(\mathcal{A})}$ is a valid witness for the fact that $\mathcal{A}$ is a layered mABP (Definition 5.8).

We denote by Σ_πmABP the computational model that is a sum of interval mABPs (possibly with respect to different orderings π).[7] The size of an instance is the sum of all the sizes of the individual mABPs that make up the instance.

Comparison with [CKSS24]. This model is a generalization of the model considered in [CKSS24] who prove a lower bound for the sum of *ordered* (set-)multilinear ABPs. The difference between the ordered and interval settings is that in the former, all the X_v are intervals of the form $[i]$ for some $i \in \{0, \ldots, n\}$. Note that the underlying set system has only one maximal chain in the ordered case, but exponentially many chains in the interval setting. This makes our result below more challenging.

An exponential lower bound. We show the following strong lower bound against Σ_πmABP.

Theorem 5.14 (Corollary 1.8)**.** Let $n \in \mathbb{N}$ be sufficiently large even number, and X be a set of n variables, and $\mathbb{F}$ be any field. Let P be a full rank polynomial in $\mathbb{F}[X]$. If $\mathcal{A}$ is a Σ_πmABP of size s computing P, then

$$s \geq 2^{\Omega(n^{1/5})}.$$

Proof. The proof is an easy consequence of our previous results. As $\mathcal{A}$ is a Σ_πmABP that computes P, it can be specified as an explicit sum of interval mABPs

$$\mathcal{A} = \mathcal{A}_1 + \cdots + \mathcal{A}_t$$

where each $\mathcal{A}_i$ is an interval mABP with respect to an ordering $\pi_i : X \to [n]$ computing a polynomial P_i.

Let $f : X \to \{-1, 1\}$ be a uniformly random balanced partition of X. It follows from the subadditivity of matrix rank that

$$2^{n/2} = \text{rank}(M_f(P)) \leq \sum_{i \in [t]} \text{rank}(M_f(P_i)),$$

and hence by averaging, there is an i such that P_i is $(1/t, 1/t)$-almost full-rank. Since $\mathcal{A}_i$ is layered with valid labelling comming from the set-system $\mathcal{I}_{X,\pi_i}$, Lemma 5.11 implies that $\mathcal{I}_{X,\pi_i}$ is a

[7] The notation here is from [CKSS24] but we add an additional subscript 'π' to recall that the mABPs are all interval mABPs.

$(1/t, \lg(st))$-balanced-chain set-system. Note that $t \leq s$ and hence $\mathcal{I}_{X,\pi_i}$ is a $(1/s, 2\lg(s))$-balanced-chain set-system.

On the other hand, it follows from Theorem 4.4 that $\mathcal{I}_{X,\pi_i}$ is *not* an (ε, k)-balanced-chain set-system, for $k < n^{1/5}$ and some $\varepsilon = \exp(-\Omega(n^{1/5}))$. This implies that $s \geq \exp(\Omega(n^{1/5}))$, proving the theorem. ■

5.6 The set-multilinear case

Both parts of Theorem 1.3 work also in the set-multilinear setting with nearly the same proofs. We sketch the ideas here, noting only the points of departure in the proofs. We assume that we are working with set-multilinear polynomials w.r.t. an equipartition $\mathcal{P} := \{X_1, \ldots, X_n\}$ into sets of size N.

For the second part of Theorem 1.3, we note that the proofs of Theorem 5.4, Theorem 5.5 and Theorem 5.6 above work almost without change in the set-multilinear case. The only difference in the construction above is that in (7), we replace the polynomial $x_u + x_v$ by a set-multilinear ABP of size $O(n)$ computing the inner product between the corresponding *variable sets* X_u and X_v.[8]

For the first part of Theorem 1.3, we note that we can get something *stronger* in the set-multilinear setting by showing that if a set-multilinear ABP of size s computes a (p, ε)-almost full-rank polynomial, then there is a $(p, O(\log_N(sn/\varepsilon))$-balanced-chain set-system of size $O(s^2 n)$. The improvement is in the chain-balance parameter in the set-system where the base of the logarithm is now N instead of 2. The source of this improvement can be traced to the 'layered' case Lemma 5.11 (defined similarly in the set-multilinear case), where the base of the exponent is N instead of 2 in all the rank bounds.

In the interesting case that $N = n^{\Theta(1)}$, the above paragraphs show that we get a tight connection between the smallest set-multilinear ABP computing a full-rank polynomial. For any c, there is a set-multilinear ABP of size $n^{O(c)}$ computing a full-rank polynomial if and only if there is an $O(c)$-balanced-chain set-system of size $n^{O(c)}$.

Finally, these results can be used to argue the lower bound for sums of interval-set-multilinear ABPs just as for the multilinear case.

6 Further questions

The main questions related to our work are the following:

- Can we close the gap between the lower bound $\Omega(n^2/k)$ and upper bound $n^{O(\ln n/\ln\ln n)}$ for the size of k-balanced-chain set systems?
- Can we obtain a uniform construction of full rank multilinear ABPs of size $n^{o(\ln n)}$? A possible approach for this question is a derandomization of our construction. In particular, it is sufficient to derandomize our application of the worst-case to average-case reduction, which uses random permutations of a given set system (Lemma 2.3).

[8]There is nothing special about Inner Product: any non-singular bilinear polynomial in $X_u \cup X_v$ would also work just as well, assuming that it has a set-multilinear ABP of linear size.

References

[Aar04] Scott Aaronson. Multilinear formulas and skepticism of quantum computing. In László Babai, editor, *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 118–127. ACM, 2004. doi:10.1145/1007352.1007378.

[AKV20] Noga Alon, Mrinal Kumar, and Ben Lee Volk. Unbalancing sets and an almost quadratic lower bound for syntactically multilinear arithmetic circuits. *Comb.*, 40(2):149–178, 2020. URL: https://doi.org/10.1007/s00493-019-4009-0, doi:10.1007/S00493-019-4009-0.

[BDS25] C. S. Bhargav, Prateek Dwivedi, and Nitin Saxena. Lower bounds for the sum of small-size algebraic branching programs. *Theor. Comput. Sci.*, 1041:115214, 2025. URL: https://doi.org/10.1016/j.tcs.2025.115214, doi:10.1016/J.TCS.2025.115214.

[BW21] Rabi Bhattacharya and Edward C. Waymire. *Random walk, Brownian motion, and martingales*, volume 292 of *Grad. Texts Math.* Cham: Springer, 2021. doi:10.1007/978-3-030-78939-8.

[CER85] E Csáki, P Erdős, and P Révész. On the length of the longest excursion. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 68(3):365–382, 1985.

[CKSS24] Prerona Chatterjee, Deepanshu Kush, Shubhangi Saraf, and Amir Shpilka. Lower bounds for set-multilinear branching programs. In Rahul Santhanam, editor, *39th Computational Complexity Conference, CCC 2024, July 22-25, 2024, Ann Arbor, MI, USA*, volume 300 of *LIPIcs*, pages 20:1–20:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. URL: https://doi.org/10.4230/LIPIcs.CCC.2024.20, doi:10.4230/LIPICS.CCC.2024.20.

[CLS19] Suryajith Chillara, Nutan Limaye, and Srikanth Srinivasan. Small-depth multilinear formula lower bounds for iterated matrix multiplication with applications. *SIAM J. Comput.*, 48(1):70–92, 2019. doi:10.1137/18M1191567.

[DL78] Richard A. Demillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978. URL: https://www.sciencedirect.com/science/article/pii/0020019078900674, doi:10.1016/0020-0190(78)90067-4.

[DMPY12] Zeev Dvir, Guillaume Malod, Sylvain Perifel, and Amir Yehudayoff. Separating multilinear branching programs and formulas. In Howard J. Karloff and Toniann Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 615–624. ACM, 2012. doi:10.1145/2213977.2214034.

[DP09] Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009. URL: http://www.cambridge.org/gb/knowledge/isbn/item2327542/.

[EM94] Thomas Eiter and Heikki Mannila. Computing discrete fréchet distance. techreport cd-tr 94/64, 1994.

[Fel68] W. Feller. An introduction to probability theory and its applications. I. New York-London-Sydney: John Wiley and Sons, Inc. XVIII, 509 p. (1968)., 1968.

[FSTW21] Michael A. Forbes, Amir Shpilka, Iddo Tzameret, and Avi Wigderson. Proof complexity lower bounds from algebraic circuit complexity. *Theory Comput.*, 17:1–88, 2021. URL: https://doi.org/10.4086/toc.2021.v017a010, doi:10.4086/TOC.2021.V017A010.

[KS22] Deepanshu Kush and Shubhangi Saraf. Improved low-depth set-multilinear circuit lower bounds. In Shachar Lovett, editor, *37th Computational Complexity Conference, CCC 2022, July 20-23, 2022, Philadelphia, PA, USA*, volume 234 of *LIPIcs*, pages 38:1–38:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. URL: https://doi.org/10.4230/LIPIcs.CCC.2022.38, doi:10.4230/LIPICS.CCC.2022.38.

[KS23] Deepanshu Kush and Shubhangi Saraf. Near-optimal set-multilinear formula lower bounds. In Amnon Ta-Shma, editor, *38th Computational Complexity Conference, CCC 2023, July 17-20, 2023, Warwick, UK*, volume 264 of *LIPIcs*, pages 15:1–15:33. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023. URL: https://doi.org/10.4230/LIPIcs.CCC.2023.15, doi:10.4230/LIPICS.CCC.2023.15.

[Nis91] Noam Nisan. Lower bounds for non-commutative computation (extended abstract). In Cris Koutsougeras and Jeffrey Scott Vitter, editors, *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing, May 5-8, 1991, New Orleans, Louisiana, USA*, pages 410–418. ACM, 1991. doi:10.1145/103418.103462.

[NNN12] Alantha Newman, Ofer Neiman, and Aleksandar Nikolov. Beck's three permutations conjecture: A counterexample and some consequences. In *53rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2012, New Brunswick, NJ, USA, October 20-23, 2012*, pages 253–262. IEEE Computer Society, 2012. doi:10.1109/FOCS.2012.84.

[NW97] Noam Nisan and Avi Wigderson. Lower bounds on arithmetic circuits via partial derivatives. *Comput. Complex.*, 6(3):217–234, 1997. doi:10.1007/BF01294256.

[Ore22] O. Ore. Über höhere Kongruenzen. Norsk matem. Forenings Skrifter 1, Nr. 7, 15 S. (1922)., 1922.

[Raz06] Ran Raz. Separation of multilinear circuit and formula size. *Theory Comput.*, 2(6):121–135, 2006. URL: https://doi.org/10.4086/toc.2006.v002a006, doi:10.4086/TOC.2006.V002A006.

[Raz09] Ran Raz. Multi-linear formulas for permanent and determinant are of super-polynomial size. *J. ACM*, 56(2):8:1–8:17, 2009. doi:10.1145/1502793.1502797.

[RSY08a] Ran Raz, Amir Shpilka, and Amir Yehudayoff. A lower bound for the size of syntactically multilinear arithmetic circuits. *SIAM J. Comput.*, 38(4):1624–1647, 2008. doi:10.1137/070707932.

[RSY08b] Ran Raz, Amir Shpilka, and Amir Yehudayoff. A lower bound for the size of syntactically multilinear arithmetic circuits. *SIAM Journal on Computing*, 38(4):1624–1647, 2008.

[RT08] Ran Raz and Iddo Tzameret. The strength of multilinear proofs. *Comput. Complex.*, 17(3):407–457, 2008. URL: https://doi.org/10.1007/s00037-008-0246-0, doi:10.1007/S00037-008-0246-0.

[RY09] Ran Raz and Amir Yehudayoff. Lower bounds and separations for constant depth multilinear circuits. *Comput. Complex.*, 18(2):171–207, 2009. URL: https://doi.org/10.1007/s00037-009-0270-8, doi:10.1007/S00037-009-0270-8.

[RY11] Ran Raz and Amir Yehudayoff. Multilinear formulas, maximal-partition discrepancy and mixed-sources extractors. *J. Comput. Syst. Sci.*, 77(1):167–190, 2011. doi:10.1016/j.jcss.2010.06.013.

[Sch80] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, October 1980. doi:10.1145/322217.322225.

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation*, pages 216–226, Berlin, Heidelberg, 1979. Springer Berlin Heidelberg.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il