Electronic Colloquium on Computational Complexity, Report No. 43 (2026)

An Unconditional Barrier for Proving Multilinear Algebraic Branching Program Lower Bounds

Deepanshu Kush*
University of Cambridge

April 1, 2026

Abstract

Since the breakthrough superpolynomial multilinear formula lower bounds of Raz (Theory of Computing 2006), proving such lower bounds against multilinear algebraic branching programs (mABPs) has been a longstanding open problem in algebraic complexity theory. All known multilinear lower bounds rely on the *min-partition rank method*, and the best bounds against mABPs have remained quadratic (Alon, Kumar, and Volk, Combinatorica 2020).

We show that the min-partition rank method *cannot* prove superpolynomial mABP lower bounds: there exists a full-rank multilinear polynomial computable by a polynomial-size mABP. This is an *unconditional barrier*: new techniques are needed to separate mVBP from higher classes in the multilinear hierarchy.

Our proof resolves an open problem of Fabris, Limaye, Srinivasan, and Yehudayoff (ECCC 2026), who showed that the power of this method is governed by the minimum size $N(n)$ of a combinatorial object called a 1-*balanced-chain set system*, and proved $N(n) \leq n^{O(\log n/\log\log n)}$. We prove $N(n) = n^{O(1)}$ by giving the chain-builder a binary choice at each step, biasing what was a symmetric random walk into one where the imbalance increases with probability at most $1/4$; a supermartingale argument combined with a multi-scale recursion yields the polynomial bound.

*This research is supported by Tom Gur's ERC Starting Grant 101163189 and UKRI Future Leaders Fellowship MR/X023583/1.

ISSN 1433-8092

Contents

1 Introduction

1.1 Algebraic complexity and multilinear models

Algebraic complexity theory studies the minimum number of arithmetic operations needed to evaluate multivariate polynomials. The foundational work of Valiant [Val79] introduced algebraic analogues of the central objects in Boolean complexity. An *algebraic circuit* computes a polynomial via a directed acyclic graph of addition and multiplication gates; an *algebraic formula* is the special case where the underlying graph is a tree (every intermediate result is used exactly once); and an *algebraic branching program* (ABP) is a layered directed acyclic graph with a designated source and sink, whose edges are labelled by linear forms, and whose output polynomial is the sum over all source-to-sink paths of the product of edge labels along the path. These three models give rise to the complexity classes $\mathsf{VF} \subseteq \mathsf{VBP} \subseteq \mathsf{VP}$[1], the algebraic analogues of $\mathsf{NC}^1 \subseteq \mathsf{NL} \subseteq \mathsf{P}$. Proving strict separations between these classes is a central goal of the field, and has remained stubbornly open despite decades of effort.

A polynomial is *multilinear* if every variable has individual degree at most 1 in each monomial. A circuit is *multilinear* if it is restricted to maintain multilinearity at every gate, not just at the output. A *multilinear algebraic branching program* (mABP) is an ABP in which the same variable does not appear in more than one edge label along any source-to-sink path; the polynomial it computes is therefore multilinear by construction. Multilinear formulas and multilinear circuits are defined analogously, giving a multilinear hierarchy

$$\mathsf{mVF} \subseteq \mathsf{mVBP} \subseteq \mathsf{mVP}.$$

The systematic study of multilinear computation can be traced back at least to the work of Nisan [Nis91], who proved exponential lower bounds for non-commutative ABPs[2] using a rank argument on coefficient matrices. Nisan and Wigderson [NW97] extended these ideas to the commutative setting and introduced the method of partial derivatives, which has since become a cornerstone of algebraic lower bound techniques. In the years since, multilinear models (as well as their more restricted cousins, *set*-multilinear models) have become one of the most natural and intensively studied settings for proving lower bounds [Raz06, Jan08, RSY08, Raz09, RY09, HY11, RY11, DMPY12, FLMS15, AR16, MST16, CELS18, CLS19, RR19, KNS20, RR20, AKV20, GR21, KS22, LST22, TLS22, KS23, CKSS24, BDS25, LST25, FLSY26]. This is natural: arguably the most important polynomials studied in algebraic complexity theory (the elementary symmetric polynomials, the permanent, the determinant, iterated matrix multiplication, etc.) are themselves multilinear, and any circuit lower bound for these polynomials is in particular a multilinear lower bound, so progress in the multilinear setting is a necessary first step toward the general case.

1.2 Raz's breakthrough and the mABP challenge

The most celebrated result in multilinear complexity is the work of Raz [Raz06, Raz09], who proved the first superpolynomial lower bounds on the size of multilinear formulas. Specifically, Raz exhibited explicit polynomials in mVP that require multilinear formulas of size $n^{\Omega(\log n)}$, thereby separating mVF from mVP. These results sparked a large body of follow-up work, including extensions to constant-depth multilinear circuits [RY09, CLS19], connections to quantum computation [Aar04, RY11], and applications to proof complexity [RT08, FSTW21]. This was subsequently

[1] We refer the reader to the surveys of Shpilka and Yehudayoff [SY10] and Saptharishi [Sap16] for a comprehensive treatment of algebraic complexity theory.

[2] These are equivalent to *read-once oblivious* ABPs (ROABPs) and *ordered set-multilinear* ABPs (see, e.g., [CKSS24]), a restricted form of multilinear ABPs.

strengthened by Dvir, Malod, Perifel, and Yehudayoff [DMPY12], who exhibited a polynomial computable by a polynomial-size mABP but requiring multilinear formulas of size $n^{\Omega(\log n)}$, separating mVF from mVBP, and thereby clarifying the multilinear hierarchy:

$$\mathsf{mVF} \subsetneq \mathsf{mVBP} \subseteq \mathsf{mVP}.$$

These lower bounds for multilinear formulas were proved over two decades ago, and the natural next challenge has been clear ever since: extend the lower bounds from multilinear *formulas* to multilinear *ABPs*. That is, prove superpolynomial lower bounds on the size of mABPs computing an explicit polynomial—or better yet, separate mVBP from mVP. Despite sustained effort, progress on this front has been remarkably limited. Since the multilinear *formula* lower bound breakthrough, the best lower bounds against mABPs have only been $\tilde{\Omega}(n^{4/3})$ by Raz, Shpilka, and Yehudayoff [RSY08], improved to $\tilde{\Omega}(n^{3/2})$ by Jansen [Jan08], and to $\tilde{\Omega}(n^2)$ by Alon, Kumar, and Volk [AKV20][3]. Jumping ahead, our work shows that this apparent barrier against substantial improvements on these lower bounds is not merely a technical obstacle: it reflects an inherent limitation of the underlying method.

1.3 The set-multilinear setting

A natural refinement of multilinearity arises when the variable set is partitioned into blocks $X_1, \ldots, X_d$, and each monomial is required to contain exactly one variable from each block. Polynomials and models satisfying this property are called *set-multilinear*. Many important polynomials—most notably the iterated matrix multiplication polynomial $\mathrm{IMM}_{n,d}$, which is complete for VBP—are set-multilinear with respect to their natural variable partition. Since set-multilinearity is a further restriction on top of multilinearity, lower bounds in this setting are a natural stepping stone toward general multilinear lower bounds.

Set-multilinear models have become especially prominent due to a phenomenon known as *hardness escalation*[4]. Analogous to the multilinear hierarchy, we have a set-multilinear hierarchy

$$\mathsf{smVF} \subseteq \mathsf{smVBP} \subseteq \mathsf{smVP}.$$

Lower bounds for set-multilinear formulas already follow from Raz's work [Raz09], but set-multilinear formulas have been studied in much greater detail [KS22, KS23, TLS22], mainly owing to the hardness escalation phenomenon and the pursuit of general class separations in algebraic complexity. The separation $\mathsf{smVF} \subsetneq \mathsf{smVBP}$ follows from the work of Kush and Saraf [KS23][5], who proved near-optimal set-multilinear formula lower bounds for a polynomial in smVBP, thereby implying:

$$\mathsf{smVF} \subsetneq \mathsf{smVBP} \subseteq \mathsf{smVP}.$$

Proving superpolynomial lower bounds against set-multilinear ABPs—or better yet, separating smVBP from smVP—is the next frontier. Since set-multilinear ABPs are a restriction of mABPs, one might hope that stronger lower bounds than the $\tilde{\Omega}(n^2)$ of [AKV20] could be achievable. However, as we show, our barrier applies in the set-multilinear setting as well[6].

[3] The lower bounds of [RSY08] and [AKV20] are stated for syntactically multilinear *circuits*, a model at least as powerful as mABPs. Quadratic lower bounds for general (i.e., possibly non-multilinear) ABPs were proved by Chatterjee, Kumar, She, and Volk [CKSV22].

[4] Lower bounds for set-multilinear models at small degree can be "lifted" to lower bounds for unrestricted models; see Section 9 for a detailed discussion.

[5] They build on the aforementioned work of [DMPY12] but require additional ideas to make the hard polynomial *set*-multilinear.

[6] That said, rank-based methods beyond the min-partition rank framework may still have more traction in the set-multilinear setting than in the general multilinear one; see Section 9.

1.4 The min-partition rank method and full-rank polynomials

All known lower bounds for multilinear models rely on a single technique: the *min-partition rank method*, introduced by Raz [Raz09] and building on Nisan [Nis91] and Nisan–Wigderson [NW97].

The idea is as follows. Given a multilinear polynomial $P(x_1, \ldots, x_n)$ with n even, and an *equipartition* of the variables into two sets Y and Z of size $n/2$ each, one forms the *coefficient matrix* $M_{Y,Z}(P)$: a $2^{n/2} \times 2^{n/2}$ matrix whose rows are indexed by multilinear monomials in Y, columns by multilinear monomials in Z, and entries by the corresponding coefficients of P. The key algebraic fact is that if P admits a decomposition $P = \sum_{i=1}^{s} L_i \cdot R_i$ where each L_i depends only on Y and each R_i only on Z, then $\mathrm{rank}(M_{Y,Z}(P)) \leq s$. For a multilinear formula of size s, the tree structure guarantees the existence of a "balanced" vertex whose removal yields such a decomposition for some equipartition, with s summands. Thus a lower bound on $\min_{(Y,Z)} \mathrm{rank}(M_{Y,Z}(P))$—the *min-partition rank* of P—gives a lower bound on formula size. The same ideas, applied to vertex cuts in an mABP rather than balanced vertices in a formula, yield lower bounds on mABP size as well, though the resulting bounds are substantially weaker quantitatively [RSY08, AKV20].

(In the set-multilinear case, the definition is analogous: given a set-multilinear polynomial with respect to a partition $\mathcal{P} = \{X_1, \ldots, X_n\}$ into blocks of size N each, one partitions the *blocks* into two equal groups and forms the coefficient matrix indexed by set-multilinear monomials from each group. The matrix is $N^{n/2} \times N^{n/2}$, and the method works similarly.)

A multilinear polynomial P is called *full-rank* if $M_{Y,Z}(P)$ has maximum possible rank ($2^{n/2}$) for every equipartition (Y, Z). (The analogous statement holds in the set-multilinear case, with $N^{n/2}$ replacing $2^{n/2}$.) Full-rank polynomials are the "hardest instances" for the min-partition rank method: any lower bound the method can prove applies most strongly to them. Raz [Raz09] proved:

Any multilinear formula computing a full-rank polynomial requires size $n^{\Omega(\log n)}$.

Raz [Raz06] showed that there are explicit full-rank polynomials computable by multilinear circuits of size $\mathrm{poly}(n)$, and by multilinear formulas and mABPs of size $n^{O(\log n)}$. The formula lower bound of $n^{\Omega(\log n)}$ combined with the circuit upper bound of $\mathrm{poly}(n)$ immediately gives the separation $\mathsf{mVF} \neq \mathsf{mVP}$. The separation $\mathsf{mVF} \neq \mathsf{mVBP}$ required additional work: in an ingenious construction, Dvir, Malod, Perifel, and Yehudayoff [DMPY12] showed that mABPs can compute a polynomial that is full-rank with respect to a suitable *subfamily* of equipartitions (called arc-partitions), and that this restricted notion of full-rankness already suffices to force multilinear formulas to have size $n^{\Omega(\log n)}$.

This leads to the natural question: But what about mABPs? Can a full-rank polynomial be computed by an mABP of size $\mathrm{poly}(n)$, or does every full-rank polynomial require mABPs of superpolynomial size? If the latter, the min-partition rank method would prove $\mathsf{mVBP} \neq \mathsf{mVP}$. If the former, then full-rankness is insufficient to witness mABP hardness, and the min-partition rank method—the only technique known to prove lower bounds for multilinear models—hits a fundamental barrier. This is the question that motivates our work.

1.5 The combinatorial characterization of Fabris et al.

The recent work of Fabris, Limaye, Srinivasan, and Yehudayoff [FLSY26] converts this algebraic question into a purely combinatorial one. They introduce the notion of a *balanced-chain set system* (Definition 2.2 below) and prove the following two-directional connection between these objects and the mABP complexity of full-rank polynomials (Theorem 1.3 of [FLSY26]):

- If every $(c \log n)$-balanced-chain set system over $[n]$ has size $n^{\Omega(c)}$, then every mABP computing a full-rank polynomial has size $n^{\Omega(c)}$.

- Conversely, if there is an $O(1)$-balanced-chain set system of size s, then there is an mABP of size $s \cdot \mathrm{poly}(n)$ computing a full-rank polynomial.

Note the slight asymmetry in the balance parameter ($c \log n$ versus $O(1)$) between the two directions. In the set-multilinear setting, this gap closes completely: both directions use balance $\Theta(c)$, giving a tight characterization of the set-multilinear ABP complexity of full-rank polynomials in terms of balanced-chain set system size [FLSY26, Section 5.6].

Write $N(n)$ for the minimum size of a 1-balanced-chain set system over $[n]$. The lower bound $N(n) \geq \Omega(n^2)$ follows from earlier work on balancing set systems by Alon, Kumar, and Volk [AKV20]. For the upper bound, $N(n) \leq n^{O(\log n)}$ already follows from Raz's constructions [Raz06] combined with the equivalence above; Fabris et al. improve this to $N(n) \leq n^{O(\log n/\log\log n)}$ via an elegant recursive argument based on the return times of random walks (see Section 2.2 for an overview). The determination of $N(n)$ is posed as an open problem in [FLSY26].

The asymptotic behaviour of $N(n)$ has a direct complexity-theoretic interpretation:

- If $N(n) = n^{\omega(1)}$, then $\mathsf{smVBP} \neq \mathsf{smVP}$: every full-rank polynomial would require superpolynomial-size set-multilinear ABPs, confirming the power of the min-partition rank method.[7]
- If $N(n) = n^{O(1)}$, then the min-partition rank method *cannot* prove $\mathsf{mVBP} \neq \mathsf{mVP}$: some full-rank polynomial would have a polynomial-size mABP, and full-rankness—the only known witness for multilinear hardness—would be insufficient to force large mABPs.

1.6 Our results and implications for algebraic complexity

We resolve this question.

Theorem 1.1. $N(n) = n^{O(1)}$. *That is, for every sufficiently large n, there exists a 1-balanced-chain set system over $[n]$ of size $n^{O(1)}$.*

The proof appears in Sections 3–8 and is the main technical contribution of this paper. We now spell out the consequences for algebraic complexity.

The implication from Theorem 1.1 to the next corollary uses the easier of the two directions of the Fabris et al. equivalence: given a 1-balanced-chain set system of size s, one can construct a full-rank polynomial computable by an mABP of size $s \cdot \mathrm{poly}(n)$ over any infinite field. The construction combines a standard gadget polynomial (originating in [Raz06]) with a Schwartz–Zippel derandomization; we include a self-contained proof in Appendix A for completeness. Combined with Theorem 1.1, this gives:

Corollary 1.2. *For every infinite field $\mathbb{F}$ and all sufficiently large n, there exists an n-variate full-rank multilinear polynomial over $\mathbb{F}$ computable by an mABP of size $n^{O(1)}$.*

The min-partition rank method proves lower bounds by exhibiting a full-rank polynomial and arguing that circuit classes cannot compute it efficiently. The maximum lower bound it can achieve is therefore the minimum complexity of *any* full-rank polynomial. Corollary 1.2 shows this minimum is polynomial for mABPs:

Corollary 1.3 (Main result)**.** *The min-partition rank method cannot prove superpolynomial lower bounds on mABP size.*

[7] For mABPs, the asymmetry in the balance parameter means that $N(n) = n^{\omega(1)}$ alone does not immediately imply $\mathsf{mVBP} \neq \mathsf{mVP}$; one would additionally need to show that every $(c \log n)$-balanced-chain set system has size $n^{\Omega(c)}$. In the set-multilinear case, the tight characterization means $N(n) = n^{\omega(1)}$ suffices directly.

This is an *unconditional barrier* result, and it applies to the only technique known to prove lower bounds for multilinear models. The practical implication is that proving superpolynomial mABP lower bounds—in particular, separating mVBP from mVP—requires fundamentally new ideas, going beyond the min-partition rank framework.

Combining Corollary 1.2 with Raz's formula lower bound [Raz09] gives:

Corollary 1.4. *There exist n-variate multilinear polynomials computable by mABPs of size $n^{O(1)}$ but requiring multilinear formulas of size $n^{\Omega(\log n)}$.*

This recovers the qualitative separation $\mathsf{mVF} \subsetneq \mathsf{mVBP}$ first proved by Dvir, Malod, Perifel, and Yehudayoff [DMPY12], via a completely different route. (Our construction also matches the quantitative parameters of [DMPY12]: the mABP has polynomial size while the formula lower bound is $n^{\Omega(\log n)}$.)

Finally, by Theorem 1.1 combined with the set-multilinear construction in Appendix A (Remark A.1), all three results above extend to set-multilinear ABPs: there exist polynomial-size set-multilinear ABPs computing full-rank polynomials, the min-partition rank method cannot prove superpolynomial lower bounds against set-multilinear ABPs either, and we recover the result $\mathsf{smVF} \subsetneq \mathsf{smVBP}$ of [KS23].

Remark 1.5. Our result does *not* necessarily imply $\mathsf{mVBP} = \mathsf{mVP}$ (or $\mathsf{smVBP} = \mathsf{smVP}$). The class mVP contains polynomials that are not full-rank, and our result says nothing about those. What it shows is that the property of full-rankness—the *only* property currently known to witness multilinear hardness—is insufficient to force large mABPs.

2 Definitions and Proof Overview

We begin by defining the combinatorial objects central to this paper (Section 2.1), then recall the approach of Fabris et al. and explain why it yields a quasipolynomial rather than polynomial upper bound (Section 2.2). We then describe our two-block steering idea (Section 2.3) and give a detailed outline of the proof (Section 2.4).

2.1 Definitions

Let n be a positive even integer and $[n] = \{1, \ldots, n\}$.

Definition 2.1 (Set systems and maximal chains)**.** A *set system* over $[n]$ is a family $\mathcal{X} \subseteq \mathcal{P}([n])$. A *maximal chain* in $\mathcal{X}$ is a sequence $(C_0, C_1, \ldots, C_n)$ satisfying:

- $C_i \in \mathcal{X}$ for all $i \in \{0, 1, \ldots, n\}$;
- $|C_i| = i$ (so $C_0 = \emptyset$ and $C_n = [n]$);
- $C_i \subset C_{i+1}$ for all i.

Think of a maximal chain as building $[n]$ one element at a time, with every intermediate set belonging to $\mathcal{X}$.

Definition 2.2 (Balanced colorings and chain-balance)**.** A *balanced coloring* is a function $f\colon [n] \to \{\pm 1\}$ with $f([n]) := \sum_{x\in[n]} f(x) = 0$ (equal numbers of $+1$'s and -1's). For $S \subseteq [n]$, the *imbalance* of S is $f(S) := \sum_{x\in S} f(x)$.

The *chain-balance* of a set system $\mathcal{X}$ is defined by a two-player game. An adversary picks a balanced coloring f; the builder responds with a maximal chain in $\mathcal{X}$; the cost is the worst imbalance along the chain. Formally:

$$\mathrm{cbal}(\mathcal{X}) = \max_{\text{balanced } f} \min_{\substack{(C_0,\dots,C_n) \\ \text{maximal chain in } \mathcal{X}}} \max_{0 \le i \le n} |f(C_i)| .$$

$\mathcal{X}$ is *k-balanced-chain* if $\mathrm{cbal}(\mathcal{X}) \le k$. Write $N(n)$ for the minimum size of a 1-balanced-chain set system over $[n]$.

Definition 2.3 (Average-case version)**.** $\mathcal{X}$ is (ε, k)-balanced-chain if for uniformly random balanced f,

$$\Pr_f[\mathcal{X} \text{ contains a chain with all } |f(C_i)| \le k] \ge \varepsilon.$$

The following result of [FLSY26] allows us to pass from the average case to the worst case.

Theorem 2.4 (Worst-case to average-case, [FLSY26, Lemma 2.3])**.** *If there is an (ε, k)-balanced-chain set system of size s, then there is a k-balanced-chain set system of size $O(sn/\varepsilon)$.*

The proof is a probabilistic argument: take $O(n/\varepsilon)$ random permutations of $[n]$, apply each to every set in $\mathcal{X}$, and take the union. A counting argument shows that the resulting set system works for all balanced colorings.

2.2 The approach of Fabris et al. and its limitations

To motivate our construction, we first recall the approach of [FLSY26] to prove their improved upper bound on $N(n)$. By Theorem 2.4, it suffices to construct a set system that contains a 1-balanced chain for a "noticeable fraction" of random balanced colorings; the resulting set system can then be converted to one that works for *all* balanced colorings with only a polynomial blowup. The starting point is the family of all intervals $\{[i,j] : 1 \le i \le j \le n\}$, which has $O(n^2)$ sets. A maximal chain in this family corresponds to growing an interval from both endpoints. For a 1-balanced-chain set system, we need every balanced coloring f to admit a chain along which $|f(C_i)| \le 1$ at every step. But for a random balanced coloring, the imbalance of the prefix $[1,t]$ is a random walk $W(t) = \sum_{i=1}^{t} f(i)$, whose maximum absolute value $\max_t |W(t)|$ is easily seen to be $\Theta(\sqrt{n})$ with high probability (see, e.g., [MU17, Chapter 13]). The interval family has many other chains (growing from both ends), but Fabris et al. show (Theorem 1.7 of [FLSY26]) that even the best among exponentially many chains has imbalance $n^{\Omega(1)}$ with overwhelming probability, via an analysis of the discrete Fréchet distance between two independent random walks. In short, no chain in the interval family comes close to achieving balance 1. Despite this negative result, the interval family is useful in a *recursive* way. A classical result of Csáki, Erdős, and Révész [CER85] implies that the random walk $W(t)$ returns to zero at least once every $\approx n/\log n$ steps with noticeable probability. This gives a partial chain of balanced intervals $\emptyset \subset [1, i_1] \subset [1, i_2] \subset \cdots \subset [1, n]$ with gaps of size $\approx n/\log n$ between consecutive elements. To fill each gap, one recurses on a segment of size $\approx n/\log n$. This leads to a set system of size $n^{O(\log n/\log\log n)}$: the recursion has depth $\Theta(\log n/\log\log n)$, and each level contributes a polynomial factor.

The bottleneck is clear: the gap between consecutive zeros of a random walk is $\Theta(n/\log n)$, so each recursion level reduces the problem size by only a $\log n$ factor. This is fundamentally a consequence of the symmetric nature of the random walk: with a single ordering, the builder has no choice at each step, and the walk moves up or down with equal probability.

2.3 Our idea: two-block steering

Our key idea is to give the builder a *binary choice* at each step by partitioning $[n]$ into two blocks (i.e., a left half and right half) and considering sets formed by taking a prefix of each block. At each step, the builder chooses which block to extend, always picking the option that reduces the imbalance (or a fair coin if both are equally bad).

This converts the imbalance process from a *symmetric* random walk (with a single ordering) into a *negatively biased* one. The mechanism is simple: if the current imbalance is $h > 0$ (too many $+1$'s have been added), the builder wants to add an element with f-value -1. With two blocks, she is forced to increase h only if *both* blocks' next elements have f-value $+1$. Since $f([n]) = 0$ forces the remaining elements to sum to $-h$, the $+1$ elements are the *minority* in the pool, and the probability that both draws land in the minority is at most $1/4$.

With this bias, the imbalance process stays within $O(\log n)$ of the origin (by a supermartingale argument), and returns to $h \leq 1$ every $O(\log n)$ steps (by an exponential tail bound on excursion lengths). Gaps are now $O(\log n)$ rather than $O(n/\log n)$, and filling each gap requires enumerating all $2^{O(\log n)} = n^{O(1)}$ subsets—a polynomial cost.

There is one complication: the two-block structure eventually breaks down when one block is exhausted. At that point (after about $n - \tilde{O}(\sqrt{n})$ steps), the remaining elements form a contiguous residual of size $\tilde{O}(\sqrt{n})$, and we must recurse. A logarithmic number of recursions reduces the problem to constant size. And because the gap parameter at each level is $O(\log m)$ where m is the *current* segment size (not $\log n$), the total cost remains polynomial.

2.4 Proof outline

We describe a randomized strategy that, given a random balanced coloring f, builds a 1-balanced chain with high probability. The set system $\mathcal{S}$ is then defined to include every subset that could appear as a chain set during any execution of this strategy. We describe the four components below.

1. Two-block steering (Section 3). We partition a contiguous segment I of size m into two halves I^L and I^R. A set formed by taking a prefix of length a from I^L and a prefix of length b from I^R corresponds to a grid point (a, b). A maximal chain in this family is a monotone lattice path from $(0, 0)$ to $(m/2, m/2)$. The *steered path* (Definition 3.1) is the specific path chosen by the builder: at each step, she extends whichever block reduces the absolute imbalance, breaking ties by a fair coin flip. See Figure 1 on page 11 for a worked example.

2. Forced probability and supermartingale (Sections 4–5). The key observation is that the builder's strategy makes the imbalance process negatively biased at every step. Writing C_t for the chain set at step t and $H(t)$ for its imbalance, the pool of unassigned elements is $[n] \setminus C_t$, which sums to $-H(t)$ because $f([n]) = 0$. In particular, when the imbalance $|H(t)| \geq 1$, the pool elements sharing the sign of $H(t)$ are the strict minority: there are $P = (R - |H|)/2 < R/2$ of them out of R total. By Lemma 3.2, $|H|$ increases only if both candidate next elements (one from each block) have f-value $\operatorname{sgn}(H)$. Since both are drawn from the pool, this probability is $P(P-1)/(R(R-1)) < (1/2)^2 = 1/4$.

This bound—$p_t \leq 1/4$—holds unconditionally at every step, with no conditions on the pool size or the current imbalance. We exploit it via the potential function $\varphi(h) = 3^h$, which is a martingale for the $(1/4, 3/4)$ birth-death chain (since $\frac{1}{4} \cdot 3^{h+1} + \frac{3}{4} \cdot 3^{h-1} = 3^h$). Since our chain has upward probability $p_t \leq 1/4$ at every step, $3^{|H(t)|}$ is a supermartingale. Doob's maximal inequality (Theorem 5.3) then gives $\max_t |H(t)| \leq |H(0)| + K \ln m$ with high probability ($\geq 1 - O(1/m^2)$), where $K = 4$ (Theorem 5.6).

Furthermore, by computing the moment generating function of the first-passage time from state 1 to state 0 in the $(1/4, 3/4)$ birth-death chain (Section 5.4), we obtain an exponential tail bound: the probability that a single excursion lasts more than t steps decays as $(\sqrt{3}/2)^t$. This implies that the imbalance returns to $|H| \le 1$ every $\le C_1 \ln m$ steps (Corollary 5.8, with $C_1 = 28$), and that descents from height $O(\log m)$ to height 1 take $O(\log m)$ steps (Lemma 5.9). See Figure 2 on page 14 for an illustration.

3. Gap filling (Section 6). Between consecutive times the steered path visits $\{|H| \le 1\}$, the chain must traverse a "gap" of $\le C_1 \ln m$ elements during which $|H|$ may exceed 1. We cannot control the imbalance element-by-element during these gaps, but we do not need to: a deterministic greedy lemma (Lemma 6.1) shows that whenever $|f(S)| \le 1$ and $|f(S \cup I)| \le 1$, the elements of I can be ordered so that $|f|$ never exceeds 1 along the way. The proof is simple: the constraint $f(S \cup I) - f(S) = f(I)$ guarantees that correcting-sign elements are always available. Since we do not know which ordering will be used, the set system must include *all* subsets of each gap region; each gap has size $\le C_1 \ln m$, contributing $2^{C_1 \ln m} = m^{O(1)}$ intermediate sets, so the *local set system* at each scale has polynomial size.

4. Multi-scale recursion (Sections 7–8). Eventually one of the two half-blocks is exhausted. By a block-deviation bound (Lemma 3.3, via Azuma's inequality), the two blocks advance at roughly the same rate, so the residual—the portion of the unexhausted block that remains—is a contiguous segment $I_1 \subseteq I_0$ of size $m_1 \lesssim \sqrt{m_0}$. At this point, the imbalance may be as large as $O(\log m_0)$, but a descent (in at most $O(\log m_0)$ steps, by the exponential tail bound above) brings it back to ≤ 1.

The builder now recurses: she applies two-block steering to the residual I_1, producing a new residual $I_2 \subseteq I_1$ of size $m_2 \lesssim \sqrt{m_1}$, and so on. After $J = O(\log \log n)$ scales, the residual has constant size and the chain is completed by brute force.

The set system $\mathcal{S}$ is defined as the collection of all sets that could arise from any such multi-scale execution (Definition 7.4). Each set in $\mathcal{S}$ is a *composite*: a union of local patterns, one from each active scale, where each local pattern belongs to the local set system $\mathcal{L}(I_j)$ of the corresponding segment. The size of $\mathcal{S}$ is controlled by the key telescoping estimate $\sum_{j=0}^{J} \log m_j \le 3 \log n$ (since $m_{j+1} \lesssim \sqrt{m_j}$), which ensures that the product of the local set system sizes across all scales remains polynomial (Lemma 7.5).

By Theorem 2.4 (the worst-case to average-case reduction of [FLSY26]), the fact that $\mathcal{S}$ contains a 1-balanced chain for 90% of balanced colorings implies $N(n) \le |\mathcal{S}| \cdot O(n) = n^{O(1)}$.

3 The Two-Block Construction

3.1 The two-interval grid

Partition $[n]$ into two equal blocks $B_1 = [1, n/2]$ and $B_2 = [n/2 + 1, n]$. A *two-interval set* is any set of the form $[1, a] \cup [n/2 + 1, n/2 + b]$ for integers $0 \le a, b \le n/2$. We identify it with the *grid point* $(a, b) \in \{0, \dots, n/2\}^2$; its *level* (cardinality) is $a + b$.

A maximal chain in the family of all two-interval sets corresponds to a *monotone lattice path* from $(0, 0)$ to $(n/2, n/2)$: a sequence of n steps, each incrementing either the first or second coordinate by one. At each step, the chain-builder chooses which block to extend.

The same construction applies to any contiguous segment $I = [s, s+m-1]$ of even size m: split I at its midpoint into halves I^L and I^R, and consider prefix pairs on this segment's grid. Between two grid points (a, b) and (a', b') with $a \le a'$ and $b \le b'$, the *gap region* consists of elements at positions $a + 1, \dots, a'$ in I^L and positions $b + 1, \dots, b'$ in I^R. See Figure 1 for an illustration.

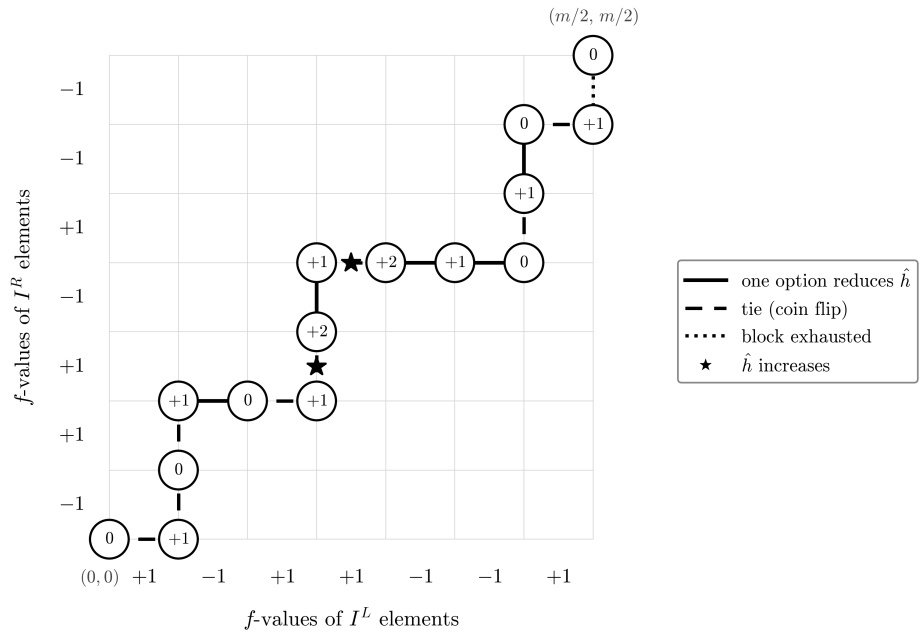

Figure 1: An example of a steered path on the two-interval grid with $m/2 = 7$. The ± 1 labels along each axis show the f-values of the elements of I^L (horizontal) and I^R (vertical); the numbers inside each node show the running imbalance $H(t)$; let $\hat{h}(t) = |H(t)|$ denote the absolute imbalance. The path begins at $(0, 0)$ with $H = 0$. At step 1, both options give $\hat{h} = 1$, so a coin flip decides (dashed); the coin selects right, giving $H = +1$. At step 4, from $(1, 2)$ with $H = +1$: extending I^L adds $f = -1$ (reducing $\hat{h}$ to 0), while extending I^R adds $f = +1$ (increasing $\hat{h}$ to 2); the builder picks the reducing option (solid). At step 6 (★), from $(3, 2)$ with $H = +1$: both candidates have $f = +1$, so $\hat{h}$ must increase to 2 regardless of the builder's choice—this is a "bad" event whose probability is bounded in Section 4. At the final step, I^L is exhausted and only I^R remains (dotted).

3.2 The steered path

For a given coloring $f: I \to \{\pm 1\}$ (since I can be any segment of $[n]$, we do not assume f to be balanced here), the steered path on I is determined by the following rule.

Definition 3.1 (Steered path)**.** The *steered path* on a segment I (split into halves I^L, I^R) with respect to f proceeds as follows. At each step with current grid position (a, b): if both blocks are active ($a < |I^L|$ and $b < |I^R|$), compute the two candidate overall chain imbalances (one for incrementing a, one for incrementing b) and choose the one with smaller absolute value. Break ties by a fair coin flip. If one block is exhausted, extend the other.

At each step, the chain set C_t grows by one element, and the overall chain imbalance is $H(t) := f(C_t)$. Write $\hat{h}(t) = |H(t)|$ for the absolute imbalance.

Lemma 3.2 (Steering transitions)**.** *Suppose $\hat{h}(t) \geq 1$ and both blocks are active. Then $\hat{h}(t+1) = \hat{h}(t) + 1$ if and only if both candidate next elements have f-value $\operatorname{sgn}(H(t))$; otherwise $\hat{h}(t+1) = \hat{h}(t) - 1$. When $H(t) = 0$: both candidates give $\hat{h}(t+1) = 1$.*

Proof. If $H > 0$: adding an element with f-value -1 gives $H' = H - 1$, so $|H'| = \hat{h} - 1$. The builder selects this option (or one of two such options). The only way $\hat{h}$ increases is if both candidates have $f = +1 = \operatorname{sgn}(H)$. The case $H < 0$ is symmetric. □

3.3 Block deviation

The steered path's two block indices a_t and b_t need not advance at the same rate. The following lemma controls the imbalance between blocks.

Lemma 3.3 (Block deviation)**.** *Let f be a uniformly random balanced coloring of $[n]$. Along the steered path on a segment of size m, write $D(t) = a_t - b_t$ for the difference between the two block indices. With probability at least $1 - 1/m$ (over f and the tie-breaking coins),*

$$\max_{t \le m} |D(t)| \le 4\sqrt{m \ln m}.$$

Proof. At each step, exactly one of a_t or b_t increases by 1, so D changes by $+1$ (if I^L is extended) or -1 (if I^R is extended). We claim $D(t)$ is a martingale (Definition 5.2): conditional on all information through step $t-1$ (the f-values revealed and coin outcomes so far), $\mathbb{E}[D(t) - D(t-1)] = 0$.

To see this, let x and y denote the next unassigned elements of I^L and I^R, respectively. Their f-values have not yet been revealed. Since f is a uniformly random balanced coloring of $[n]$, the conditional distribution of the unrevealed f-values (given the history) is exchangeable: in particular, $\Pr[f(x) = +1,\ f(y) = -1] = \Pr[f(x) = -1,\ f(y) = +1]$. Now consider the two cases. If $f(x) \ne f(y)$: the builder must extend whichever block holds the imbalance-reducing element, and by exchangeability this is equally likely to be I^L or I^R. If $f(x) = f(y)$: both options give the same absolute imbalance, and the fair coin decides. In both cases, D goes up or down with equal probability.

Since $D(t)$ is a martingale with increments bounded by 1, the Azuma–Hoeffding inequality (see, e.g., [MU17, Chapter 13]) gives

$$\Pr\left[\max_{t \le m} |D(t)| \ge \lambda\right] \le 2\exp\left(-\frac{\lambda^2}{2m}\right).$$

Setting $\lambda = 4\sqrt{m \ln m}$: $2\exp(-8\ln m) = 2/m^8 \le 1/m$ for $m \ge 2$. □

At block exhaustion (say $a_T = m/2$): the residual consists of the remaining $m/2 - b_T = D(T)$ elements of I^R, a contiguous segment. If instead I^R exhausts first ($b_T = m/2$), the residual has size $|D(T)|$. In either case, the residual has size $|D(T)| \le 4\sqrt{m \ln m} \le m^{2/3}$ (for $m \ge 100$).

4 The Forced Probability

This section contains the central observation of the paper. The argument uses exactly one fact about the coloring: $f([n]) = 0$.

At any point during the construction (at any scale of the multi-scale recursion described later), the chain set C_t is some subset of $[n]$, and the *pool* is $\mathrm{Pool}_t := [n] \setminus C_t$ (the elements not yet assigned). Since $C_t \cup \mathrm{Pool}_t = [n]$ and $f([n]) = 0$:

$$f(\mathrm{Pool}_t) = -f(C_t) = -H(t). \tag{1}$$

The pool has $R = |\mathrm{Pool}_t|$ elements. Writing $h = \hat{h}(t) = |H(t)|$, the number of pool elements with f-value $\mathrm{sgn}(H(t))$ is

$$P = \frac{R-h}{2}, \qquad Q := \frac{R+h}{2} \text{ (the count with the opposite sign).} \tag{2}$$

Since P and Q are counts of actual elements, both are non-negative; in particular $R \ge h$. When $h = 0$: $P = Q = R/2$ (the pool is evenly split). When $h \ge 1$: $P < R/2 < Q$, i.e., the elements sharing the sign of H are the *minority*.

Lemma 4.1 (Forced probability)**.** *At any step of the construction, if $\hat{h}(t) \geq 1$ and both blocks of the current segment are active (so $R \geq 2$), then*

$$p_t := \Pr[\hat{h}(t+1) = \hat{h}(t) + 1 \mid \mathcal{F}_t] \leq \frac{1}{4},$$

where $\mathcal{F}_t$ denotes the σ-algebra generated by all revealed f-values and coin-flip outcomes through step t (formally defined in Definition 5.1 below). The probability is over the remaining randomness in f: the f-values of elements not yet assigned.

Proof. By Lemma 3.2, $\hat{h}$ increases iff both candidate next elements have f-value $\operatorname{sgn}(H)$. Both are drawn from the pool.

If $P \leq 1$: it is impossible to draw two elements from at most one, so $p_t = 0 \leq 1/4$.

If $P \geq 2$: the probability is

$$p_t = \frac{P}{R} \cdot \frac{P-1}{R-1}.$$

Since $h \geq 1$: $P = (R-h)/2 \leq (R-1)/2$, so $P/R \leq (R-1)/(2R) < 1/2$. Also $P - 1 < P$ and $P \leq (R-1)/2$ imply $(P-1)/(R-1) < P/R < 1/2$. Therefore $p_t < (1/2)(1/2) = 1/4$.

(In fact $p_t = 1/4$ is never achieved for finite R with $h \geq 1$. But this distinction plays no role; $p_t \leq 1/4$ suffices for the supermartingale argument.) □

Remark 4.2 (Why no conditions are needed)**.** The bound $p_t \leq 1/4$ holds at every step, at every scale, during every transition and descent, with no conditions on R beyond $R \geq 2$. The reason is purely global: $f([n]) = 0$ forces the pool to sum to $-H(t)$, making $\operatorname{sgn}(H)$-elements the minority. This unconditional bound is what allows the supermartingale argument (Section 5) to apply uniformly across all scales.

5 The Supermartingale Argument

We use Lemma 4.1 to show that $\hat{h}(t)$ stays logarithmic with high probability. The tool is the *supermartingale method*, which we review for readers unfamiliar with martingale theory. Figure 2 illustrates the key quantities defined in this section—the height (imbalance) bound, excursions, balanced visits, and gaps—on a sample trajectory.

5.1 Background on supermartingales

We collect the probabilistic definitions and tools needed; standard references include Williams [Wil91]. Readers familiar with basic martingale theory may skip to Section 5.3.

Definition 5.1 (Filtration)**.** A *filtration* $(\mathcal{F}_t)_{t\geq 0}$ is an increasing sequence of σ-algebras on the same probability space: $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \cdots$. Think of $\mathcal{F}_t$ as the information available at time t. In our setting, $\mathcal{F}_t$ encodes the f-values of all elements assigned by step t, together with all coin-flip outcomes up to step t.

Definition 5.2 (Supermartingale)**.** A sequence $(Y_t)_{t\geq 0}$ adapted to $(\mathcal{F}_t)$ is a *supermartingale* if $\mathbb{E}[Y_{t+1} \mid \mathcal{F}_t] \leq Y_t$ for all t. It is a *martingale* if equality holds. Informally: the expected future value is at most (respectively exactly, for a martingale) the current value.

The key property of supermartingales is that they are unlikely to become much larger than their starting value.

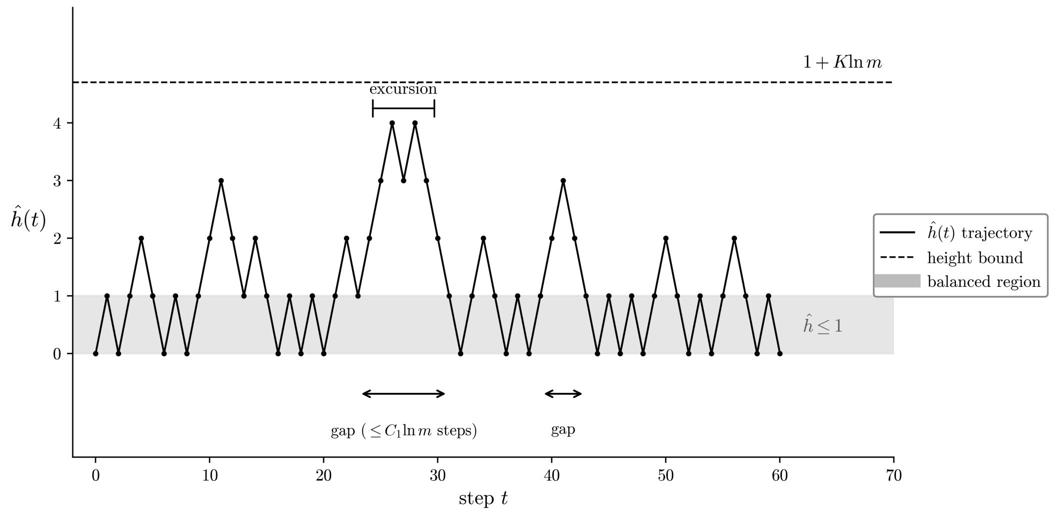

Figure 2: A schematic trajectory of the absolute imbalance $\hat{h}(t) = |H(t)|$ along the steered path within a single scale of the construction. The shaded band marks the *balanced region* $\hat{h} \leq 1$: when the trajectory is in this region, the chain set C_t has $|f(C_t)| \leq 1$. The dashed line marks the height bound $1 + K \ln m$ (Theorem 5.6): the trajectory stays below this bound with high probability. Each departure from and return to the balanced region constitutes an *excursion*; a *gap* is the interval between two consecutive balanced visits. By Corollary 5.8, every gap has length at most $C_1 \ln m$. During each gap, the gap-filling lemma (Lemma 6.1) provides a 1-balanced ordering of the gap elements, and the set system includes all subsets of each gap region to accommodate every possible such ordering. The key qualitative features are that the trajectory returns to the balanced region frequently and never strays far from it—both consequences of the negative bias ($p \leq 1/4$). Compare this with the approach of Fabris et al. (Section 2.2), where the imbalance follows a *symmetric* random walk: there, the maximum height is $\Theta(\sqrt{m})$ and the gaps between returns are $\Theta(m/\log m)$, necessitating a deep recursion.

Theorem 5.3 (Doob's maximal inequality [Wil91, Chapter 14])**.** *If $(Y_t)_{0 \leq t \leq T}$ is a non-negative supermartingale and $\lambda > 0$:*

$$\Pr\left[\max_{0 \leq t \leq T} Y_t \geq \lambda\right] \leq \frac{\mathbb{E}[Y_0]}{\lambda}.$$

Definition 5.4 (Birth-death chain)**.** A *birth-death chain* on $\{0, 1, 2, \ldots\}$ with parameters (p, q) (where $p + q = 1$) is a discrete-time Markov chain where, from state $h \geq 1$, the chain moves to $h+1$ with probability p ("birth") and to $h - 1$ with probability q ("death"). When $q > p$: the chain has a net drift toward 0.

Definition 5.5 (Geometric random variable)**.** A *geometric random variable* with success probability $\theta \in (0, 1]$, written $\mathrm{Geom}(\theta)$, is a random variable taking values in $\{1, 2, 3, \ldots\}$ with $\Pr[\mathrm{Geom}(\theta) = k] = \theta(1 - \theta)^{k-1}$. It models the number of independent trials (each succeeding with probability θ) needed until the first success. Its mean is $1/\theta$ and $\Pr[\mathrm{Geom}(\theta) > k] = (1 - \theta)^k$.

5.2 Why 3^h is the right potential function

For a $(1/4, 3/4)$ birth-death chain (Definition 5.4), consider the function $\phi(h) = 3^h$. A direct calculation shows it is a martingale:

$$p \cdot 3^{h+1} + q \cdot 3^{h-1} = 3^{h-1}\left(\tfrac{9}{4} + \tfrac{3}{4}\right) = 3^{h-1} \cdot 3 = 3^h.$$

This is a special case of the "gambler's ruin" martingale: for any (p, q) birth-death chain, α^h is a martingale when $\alpha = q/p$.

In our setting, $p_t \leq 1/4$ at every step (Lemma 4.1), so $3^{\hat{h}(t)}$ is a supermartingale:

$$\mathbb{E}[3^{\hat{h}(t+1)} \mid \mathcal{F}_t] = 3^{\hat{h}-1}(8p_t + 1) \leq 3^{\hat{h}-1} \cdot 3 = 3^{\hat{h}(t)}.$$

Doob's inequality (Theorem 5.3) then bounds the maximum of $3^{\hat{h}}$, and hence of $\hat{h}$ itself.

5.3 The imbalance bound

We now apply the supermartingale method to bound the maximum imbalance along the steered path. In the multi-scale construction (Section 8), the invariant maintained across scales ensures that the absolute imbalance is at most 1 at the start of each scale. We state the bound under this assumption.

Fix $K = 4$.

Theorem 5.6 (Imbalance bound)**.** *Consider the steered path on any segment with m' unassigned elements. If $\hat{h} \leq 1$ at the start, then*

$$\Pr\Big[\max_t \hat{h}(t) \geq 1 + K \ln m'\Big] \leq \frac{3}{(m')^2},$$

where the maximum is over all steps with both blocks active.

Proof. Decompose the trajectory of $\hat{h}$ into *excursions* above level h_0: each excursion starts when $\hat{h}$ first exceeds h_0 (reaching h_0+1) and ends when $\hat{h}$ returns to h_0. There are at most $m'/2$ excursions (each takes at least 2 steps).

Within an excursion, $\hat{h}(t) \geq 1$ and both blocks are active, so Lemma 4.1 gives $p_t \leq 1/4$ at every step. Define $Y_t = 3^{\hat{h}(t\wedge\tau)}$, where $t \wedge \tau := \min(t, \tau)$ denotes the stopped time and τ is the excursion's end. For $t < \tau$:

$$\mathbb{E}[Y_{t+1} \mid \mathcal{F}_t] = p_t \cdot 3^{\hat{h}+1} + (1-p_t) \cdot 3^{\hat{h}-1} = 3^{\hat{h}-1}(9p_t + 1 - p_t) = 3^{\hat{h}-1}(8p_t + 1) \leq 3^{\hat{h}} = Y_t.$$

So $(Y_{t\wedge\tau})$ is a non-negative supermartingale.

By Doob's inequality (Theorem 5.3):

$$\Pr[\text{excursion reaches height } L] = \Pr\big[\max_t Y_{t\wedge\tau} \geq 3^L\big] \leq \frac{Y_0}{3^L} = \frac{3^{h_0+1}}{3^L} = 3^{1-K\ln m'}.$$

Since $K = 4$: $3^{K \ln m'} = (m')^{K \ln 3}$. As $\ln 3 \approx 1.099 > 1$: $(m')^{K\ln 3} > (m')^4$. So $\Pr[\text{excursion reaches } L] < 3/(m')^4$.

Union over at most $m'/2$ excursions:

$$\Pr[\max \hat{h} \geq L] \leq \frac{m'}{2} \cdot \frac{3}{(m')^4} = \frac{3}{2(m')^3} \leq \frac{3}{(m')^2}. \qquad \square$$

5.4 The moment generating function of the first-passage time

The next two subsections bound, respectively, the length of a single excursion (Section 5.5) and the time to descend from height $O(\log m)$ to height 1 (Section 5.6). Both reduce to bounding the first-passage time τ_0 from state 1 to state 0 in the $(1/4, 3/4)$ birth-death chain: the excursion bound needs an exponential tail on τ_0, and the descent bound needs to control a sum of independent copies of τ_0. The natural tool for both is the moment generating function $M := \mathbb{E}[e^{\gamma\tau_0}]$, evaluated at a suitable rate parameter $\gamma > 0$.

Specifically: an exponential tail $\Pr[\tau_0 > t] \le O(\rho^t)$ (for some constant $0 < \rho < 1$) follows immediately from Markov's inequality applied to $e^{\gamma\tau_0}$, provided M is finite. For the descent, the time is a sum $S = \tau_0^{(1)} + \cdots + \tau_0^{(h_0-1)}$ of independent copies, and $\mathbb{E}[e^{\gamma S}] = M^{h_0-1}$; a Chernoff-type bound then requires knowing the *exact value* of M to verify that the base $\alpha = Me^{-\gamma C_3}$ of the exponential decay (computed later in Section 5.6) is less than 1.

We now compute M via a closed-form expression for the probability generating function of τ_0.

The *probability generating function* of τ_0 is $G(s) := \mathbb{E}[s^{\tau_0}] = \sum_{t=1}^{\infty} s^t \Pr[\tau_0 = t]$, defined for $|s| \le 1$ and potentially for some $|s| > 1$. A first-step analysis gives a recursion. From state 1, one step is taken (contributing a factor of s), after which the chain is at state 0 with probability $q = 3/4$, or at state 2 with probability $p = 1/4$. If the chain reaches state 0, the passage is complete, contributing s in total. If the chain reaches state 2, it must return to 0. Since the chain is nearest-neighbor, it must first pass through state 1 (taking time distributed as τ_0 by translation invariance), and then from state 1 reach state 0 again (an independent copy of τ_0, by the Markov property). The probability generating function of a sum of independent random variables is the product of their generating functions, so the contribution from this case is $s \cdot G(s)^2$. Combining:

$$G(s) = qs + ps \cdot G(s)^2.$$

Rearranging: $ps \cdot G(s)^2 - G(s) + qs = 0$, a quadratic in $G(s)$. The quadratic formula gives two roots:

$$G(s) = \frac{1 \pm \sqrt{1 - 4pqs^2}}{2ps}.$$

To select the correct root, note that τ_0 is almost surely finite (since $q > p$), so $G(1) = 1$. At $s = 1$: $\sqrt{1 - 4pq} = \sqrt{1/4} = 1/2$, giving roots $(1 - 1/2)/(2p) = 1$ and $(1 + 1/2)/(2p) = 3$. Since $G(1) = 1$, we take the minus sign:

$$G(s) = \frac{1 - \sqrt{1 - 4pqs^2}}{2ps}.$$

Define $\gamma := \ln(2/\sqrt{3})(\approx 0.144)$ and $\rho := e^{-\gamma} = \sqrt{3}/2(\approx 0.866)$. At $s = e^{\gamma} = 2/\sqrt{3}$: $4pqs^2 = 4 \cdot \frac{1}{4} \cdot \frac{3}{4} \cdot \frac{4}{3} = 1$, so the discriminant vanishes and

$$M := \mathbb{E}[e^{\gamma\tau_0}] = G(e^{\gamma}) = \frac{1}{2p \cdot e^{\gamma}} = \frac{1}{2 \cdot \frac{1}{4} \cdot \frac{2}{\sqrt{3}}} = \sqrt{3}. \tag{3}$$

An immediate consequence is an exponential tail bound. By Markov's inequality applied to the non-negative random variable $e^{\gamma\tau_0}$: for any $t \ge 1$,

$$\Pr[\tau_0 > t] = \Pr[\tau_0 \ge t + 1] \le \frac{M}{e^{\gamma(t+1)}} = \sqrt{3} \cdot \rho^{t+1} = \frac{3}{2}\,\rho^t. \tag{4}$$

5.5 The excursion length bound

Theorem 5.6 bounds the *maximum* imbalance (the dashed line in Figure 2). For the construction, we also need the imbalance to *return* to small values frequently—specifically, every $O(\log m')$ steps. We establish this by bounding the length of each excursion (a trip from $\hat{h} = 1$ up to some peak and back to $\hat{h} = 0$).

Since $p_t \le 1/4$ at every step (Lemma 4.1), the excursion lengths in our chain are stochastically dominated by those of a $(1/4, 3/4)$ birth-death chain (Definition 5.4).

Fix $C_1 = \lceil 4/\gamma \rceil = 28$.

Lemma 5.7 (Excursion bound)**.** *Each excursion from $\hat{h} = 1$ to $\hat{h} = 0$ along the steered path takes at most $C_1 \ln m'$ steps with probability $\ge 1 - 2/(m')^3$.*

Proof. It is easy to see that the excursion length is *stochastically dominated* by the first-passage time τ_0 from state 1 to state 0 in the $(1/4, 3/4)$ birth-death chain (since $p_t \le 1/4$ at every step). However, for the sake of completeness, let us include a formal proof.

We can use a coupling argument. Consider a $(1/4, 3/4)$ birth-death chain (Z_t) starting at $Z_0 = 1$, run alongside the excursion $(\hat{h}(t))$ (which also starts at $\hat{h}(0) = 1$). At each step, we couple the two chains using shared randomness as follows: draw $U \sim \mathrm{Unif}[0, 1]$; the reference chain steps up iff $U \le 1/4$; our chain steps up iff $U \le p_t$. Since $p_t \le 1/4$ (Lemma 4.1): whenever our chain steps up, so does the reference chain. Equivalently, whenever the reference chain steps down, our chain also steps down. Under this coupling, $\hat{h}(t) \le Z_t$ for all t (by induction: both start at 1, and at each step the reference chain is at least as likely to increase). In particular, if $\tau_0 := \inf\{t \ge 1 : Z_t = 0\}$ denotes the first-passage time of the reference chain to state 0, then our excursion ends by time τ_0.

It remains to bound $\Pr[\tau_0 > C_1 \ln m']$. By the tail bound (4):

$$\Pr[\tau_0 > C_1 \ln m'] \le \frac{3}{2} \cdot \rho^{C_1 \ln m'} = \frac{3}{2} \cdot (m')^{-C_1 \gamma}.$$

Since $C_1 \gamma = \lceil 4/\gamma \rceil \cdot \gamma \ge 4$: this is at most $\frac{3}{2}/(m')^4 \le 2/(m')^3$. □

Corollary 5.8 (Gap bound)**.** *With probability $\ge 1 - 1/m'$: all excursions along the steered path on a segment of m' unassigned elements have length $\le C_1 \ln m'$.*

Proof. Union over $\le m'/2$ excursions: $(m'/2) \cdot 2/(m')^3 = 1/(m')^2 \le 1/m'$. □

In other words, the trajectory returns to the balanced region (the shaded band in Figure 2) every $\le C_1 \ln m'$ steps with high probability.

5.6 The descent bound

During a scale transition (Section 8), the overall imbalance may need to *descend* from some height $\hat{h} = h_0$ to $\hat{h} \le 1$. Since $p_t \le 1/4$ at every step, the descent is stochastically dominated by the corresponding process in a $(1/4, 3/4)$ birth-death chain.

The descent from h_0 to 1 can be decomposed into $h_0 - 1$ successive "drops" (from k to $k - 1$, for $k = h_0, h_0 - 1, \ldots, 2$). Each drop is a first-passage time from state k to state $k - 1$. In the $(1/4, 3/4)$ birth-death chain, the transition probabilities are the same at every state, so by translation invariance the first-passage time from k to $k - 1$ has the same distribution as from 1 to 0. The total descent time is dominated by a sum of $h_0 - 1$ independent copies of τ_0.

Fix $C_3 = 8$.

Lemma 5.9 (Descent bound)**.** *The descent from* $\hat{h} = h_0$ *to* $\hat{h} \leq 1$ *takes at most* $C_3 h_0$ *steps with probability* $\geq 1 - 1/(m')^2$*, provided* $h_0 \leq 1 + K \ln m'$*.*

Proof. Let τ_0 denote the first-passage time from state 1 to state 0 in the $(1/4, 3/4)$ birth-death chain. Write $S = \tau_0^{(1)} + \cdots + \tau_0^{(h_0-1)}$ where $\tau_0^{(1)}, \ldots, \tau_0^{(h_0-1)}$ are independent copies of τ_0.

Since the $\tau_0^{(i)}$ are independent, $\mathbb{E}[e^{\gamma S}] = M^{h_0-1}$ where $M = \sqrt{3}$ is the moment generating function computed in (3). By Markov's inequality applied to the non-negative random variable $e^{\gamma S}$:

$$\Pr[S > C_3 h_0] = \Pr[e^{\gamma S} > e^{\gamma C_3 h_0}] \leq \frac{\mathbb{E}[e^{\gamma S}]}{e^{\gamma C_3 h_0}} = \frac{M^{h_0-1}}{e^{\gamma C_3 h_0}} = \frac{(M e^{-\gamma C_3})^{h_0}}{M}.$$

With $C_3 = 8$:

$$\alpha := M e^{-8\gamma} = \sqrt{3} \cdot (e^{-\gamma})^8 = \sqrt{3} \cdot \left(\frac{\sqrt{3}}{2} \right)^8 = \sqrt{3} \cdot \frac{3^4}{2^8} = \frac{81\sqrt{3}}{256} \approx 0.548 < 1.$$

Since $h_0 \leq 1 + K \ln m'$ with $K = 4$:

$$\frac{\alpha^{h_0}}{M} \leq \frac{\alpha \cdot \alpha^{K \ln m'}}{M} = \frac{\alpha}{M} \cdot (m')^{K \ln \alpha} = \frac{\alpha}{M} \cdot (m')^{-K \ln(1/\alpha)}.$$

Now $K \ln(1/\alpha) = 4 \ln(256/(81\sqrt{3})) \approx 4 \cdot 0.60 = 2.40 > 2$, and $\alpha/M = 81/256 \approx 0.316 < 1$. Therefore $\Pr[S > C_3 h_0] \leq 1/(m')^2$. □

6 Gap Filling

The steered path visits $\{\hat{h} \leq 1\}$ frequently, but between visits the imbalance may exceed 1. The following deterministic lemma shows that the gap elements can always be reordered to maintain $|f| \leq 1$.

Lemma 6.1 (Greedy gap filling)**.** *Let* $f \colon [n] \to \{\pm 1\}$ *be any balanced coloring. Let* $S \subseteq [n]$ *with* $|f(S)| \leq 1$*, and let* $I \subseteq [n] \setminus S$ *with* $|f(S \cup I)| \leq 1$*. Then there exists an ordering* $x_1, x_2, \ldots, x_{|I|}$ *of* I *such that*

$$|f(S \cup \{x_1, \ldots, x_k\})| \leq 1 \qquad \text{for all } k \in \{0, 1, \ldots, |I|\}.$$

The idea is simple: at each step, greedily add an element of the "correcting" sign. The key insight: such an element always exists, because the starting and ending imbalances are both ≤ 1.

Proof. Construct the ordering greedily. Write $h_k = f(S \cup \{x_1, \ldots, x_k\})$ and maintain $|h_k| \leq 1$. Let $R = |I| - k$ (remaining elements) and $\sigma = f(I \setminus \{x_1, \ldots, x_k\})$ (their sum). The key identity is $h_k + \sigma = f(S \cup I)$, so $\sigma = f(S \cup I) - h_k$.

Case $h_k = 0$**.** Any element gives $|h_{k+1}| = 1 \leq 1$.

Case $h_k = 1$**.** We need $f(x_{k+1}) = -1$. Since $|f(S \cup I)| \leq 1$ and $h_k = 1$: $\sigma = f(S \cup I) - 1 \leq 0$. The number of remaining -1 elements is $(R - \sigma)/2 \geq R/2 \geq 1$.

Case $h_k = -1$**.** Symmetric: $\sigma = f(S \cup I) + 1 \geq 0$, giving $(R + \sigma)/2 \geq 1$ elements with $f = +1$. □

7 The Set System

We now define the set system $\mathcal{S}$ and prove it has polynomial size. The set system must contain every set that could appear as a chain set during any execution of the multi-scale builder's strategy (Section 8). Since the strategy operates at multiple scales—steering on $I_0 = [n]$, then on a residual $I_1 \subseteq I_0$, then on $I_2 \subseteq I_1$, and so on—each chain set is a union of contributions from several scales. The main challenge is to show that the total number of such unions is polynomial in n, despite the recursion having $J_{\max} = O(\log \log n)$ levels.

The key insight is that the *logarithms* of the local set system sizes form a geometric series. At scale j, the segment has size m_j and the local set system has $m_j^{O(1)}$ elements, contributing $O(\log m_j)$ to the logarithm of the total count. Since $m_{j+1} \leq m_j^{2/3}$, we have $\log m_j \leq (2/3)^j \log n$, and the sum $\sum_j \log m_j$ converges to $O(\log n)$ regardless of the number of levels. This is why the recursion depth $J_{\max} = O(\log \log n)$ causes no difficulty.

7.1 Named constants

All constants below are universal (independent of n).

- $K = 4$: imbalance bound coefficient (Theorem 5.6).
- $C_1 = 28$: excursion length constant (Lemma 5.7).
- $C_3 = 8$: descent time constant (Lemma 5.9).
- $\Gamma = C_1 + KC_3 = 60$: maximum transition size per unit of $\ln m$. This accounts for: tail from gap bound ($\leq C_1 \ln m$) plus descent ($\leq C_3 + KC_3 \ln m$); total $\leq \Gamma \ln m + C_3$.
- $M_0 = 700$: base-case threshold.
- $C_2 = \Gamma + 4 = 64$: local set system size exponent.

7.2 Local set systems

Recall (Section 3.1) that on a segment I of size m, the steered path traces a monotone lattice path on the grid of prefix pairs (a, b) with $0 \leq a, b \leq m/2$. Between two consecutive balanced visits at grid points (a, b) and (a', b'), the elements added to the chain form a *gap region*: positions $a+1, \ldots, a'$ of I^L and $b+1, \ldots, b'$ of I^R. The gap-filling lemma (Lemma 6.1) allows us to order these elements arbitrarily (subject to the 1-balanced constraint), so the local set system must include *all subsets* of each gap region.

Definition 7.1 (Local set system)**.** For a contiguous segment I of size m, the *local set system* $\mathcal{L}(I)$ consists of all subsets $A \subseteq I$ of the form $A = P_L \cup P_R \cup T$, where:

- (a, b) and (a', b') are grid points with $a \leq a' \leq m/2$, $b \leq b' \leq m/2$, and $(a' - a) + (b' - b) \leq \Gamma \ln m + M_0$;
- P_L is the prefix of I^L of length a and P_R is the prefix of I^R of length b;
- $T \subseteq G$ where G is the gap region between (a, b) and (a', b').

The budget $\Gamma \ln m + M_0$ is chosen to accommodate three scenarios that arise in the main proof (Section 8): within-scale gaps between consecutive balanced visits ($\leq C_1 \ln m$ elements), scale transitions including the descent ($\leq \Gamma \ln m + C_3$ elements), and the base case where a small residual is absorbed ($\leq C_1 \ln m + M_0$ elements). All three fit within $\Gamma \ln m + M_0$.

Claim 7.2. *$|\mathcal{L}(I)| \leq m^{C_2}$ for m above an absolute constant.*

Proof. We count the choices that specify an element of $\mathcal{L}(I)$.

Grid pairs. The starting grid point (a, b) has at most $(m/2 + 1)^2 \leq m^2$ choices. Given (a, b), the ending grid point (a', b') satisfies $(a' - a) + (b' - b) \leq \Gamma \ln m + M_0$, so there are at most $(\Gamma \ln m + M_0 + 1)^2$ choices for $(a' - a, b' - b)$. For m above a constant, $(\Gamma \ln m + M_0 + 1)^2 \leq m$, giving at most $m^2 \cdot m = m^3$ grid pairs total.

Gap subsets. The gap region G has $|G| = (a' - a) + (b' - b) \leq \Gamma \ln m + M_0$ elements, so $T \subseteq G$ has at most $2^{\Gamma \ln m + M_0}$ choices. Now $2^{\Gamma \ln m} = m^{\Gamma \ln 2} \leq m^{\Gamma}$ (since $\ln 2 < 1$), and 2^{M_0} is a (large) absolute constant. So the number of gap subsets is at most $2^{M_0} \cdot m^{\Gamma}$.

Total. $|\mathcal{L}(I)| \leq m^3 \cdot 2^{M_0} \cdot m^{\Gamma} = 2^{M_0} \cdot m^{\Gamma+3}$. For $m \geq 2^{M_0}$: $2^{M_0} \leq m$, giving $|\mathcal{L}(I)| \leq m^{\Gamma+4} = m^{C_2}$. For $m < 2^{M_0}$: $|\mathcal{L}(I)| \leq 2^{2M_0}$, a constant, which is at most m^{C_2} for $m \geq 2$. □

7.3 Multi-scale decompositions

Definition 7.3 (Multi-scale decomposition)**.** A decomposition of depth J consists of:

(i) *Nested segments:* contiguous segments $[n] = I_0 \supseteq I_1 \supseteq \cdots \supseteq I_J$ with $|I_{j+1}| \leq |I_j|^{2/3}$ for all $j < J$.

(ii) *Local patterns:* for each $j < J$, a set $L_j \in \mathcal{L}(I_j)$ with $L_j \subseteq I_j \setminus I_{j+1}$.

(iii) *Terminal pattern:* a set $L_J \in \mathcal{L}(I_J)$.

The *composite set* is $S = L_0 \cup \cdots \cup L_J$. The sets $L_0, \ldots, L_J$ are pairwise disjoint: for $j < J$, $L_j \subseteq I_j \setminus I_{j+1}$, and $L_J \subseteq I_J$.

Definition 7.4 (The set system $\mathcal{S}$)**.** $\mathcal{S}$ is the collection of all composite sets from all multi-scale decompositions of all depths $0 \leq J \leq J_{\max}$, where $J_{\max}$ is the smallest integer with $n^{(2/3)^{J_{\max}}} \leq M_0$.

Since $(2/3)^{J_{\max}} \leq \ln M_0 / \ln n$, we have $J_{\max} \leq \ln(\ln n / \ln M_0)/\ln(3/2) \leq 3 \ln \ln n$ for large n. One might be concerned that summing over this many depths causes a super-polynomial blowup. The next lemma shows this is not the case.

7.4 Size bound

Lemma 7.5. $|\mathcal{S}| \leq n^{3C_2+7}$.

Proof. Write $m_j = |I_j|$. The crux of the argument is the following telescoping estimate.

Key estimate. The shrinkage condition $m_{j+1} \leq m_j^{2/3}$ gives $\ln m_{j+1} \leq \frac{2}{3} \ln m_j$, so $\ln m_j \leq (2/3)^j \ln n$. Summing:

$$\sum_{j=0}^{J} \ln m_j \leq \ln n \cdot \sum_{j=0}^{\infty} (2/3)^j = 3 \ln n. \tag{5}$$

This is the reason the recursion depth does not matter: each additional scale contributes a geometrically smaller term, and the total is bounded by $3 \ln n$ regardless of J.

Counting decompositions of a fixed depth J. A decomposition is specified by two choices: the nesting (the segments $I_1, \dots, I_J$), and the patterns $(L_0, \dots, L_J)$.

Nesting. $I_0 = [n]$ is fixed. For $j \geq 1$, the segment I_{j+1} is a contiguous sub-segment of I_j, specified by a start position and a length: at most m_j^2 choices. The total number of nestings is at most $\prod_{j=0}^{J-1} m_j^2$. Taking logarithms: $\ln(\prod_{j<J} m_j^2) = 2 \sum_{j<J} \ln m_j \leq 6 \ln n$.

Patterns. For each $j \leq J$: $|L_j| \leq |\mathcal{L}(I_j)| \leq m_j^{C_2}$ choices (Claim 7.2). The total number of pattern tuples is at most $\prod_{j=0}^{J} m_j^{C_2}$. Taking logarithms: $\ln(\prod_{j \leq J} m_j^{C_2}) = C_2 \sum_{j \leq J} \ln m_j \leq 3C_2 \ln n$.

Combining. The number of decompositions of depth J is at most

$$\exp\big((6 + 3C_2) \ln n\big) = n^{6+3C_2}.$$

Note that this bound is *independent of* J: the geometric convergence of $\sum_j \ln m_j$ absorbs the depth.

Summing over depths. There are at most $J_{\max} + 1 \leq 3 \ln \ln n + 2$ possible depths. Therefore:

$$|\mathcal{S}| \leq (3 \ln \ln n + 2) \cdot n^{6+3C_2} \leq n^{3C_2+7}$$

for large enough n. □

Remark 7.6 (Explicit construction)**.** The set system $\mathcal{S}$ is fully explicit: given 1^n, a description of $\mathcal{S}$ can be computed in poly(n) time by enumerating all valid nestings and all local patterns within each nesting. No randomness is involved in the definition of $\mathcal{S}$; the randomness in our proof appears only in the *analysis* in Section 8 (showing that $\mathcal{S}$ is $(9/10, 1)$-balanced-chain) and in the worst-case to average-case reduction (Theorem 2.4; see also Section 9).

8 Proof of the Main Theorem

We now prove Theorem 1.1: $N(n) = n^{O(1)}$. The proof combines the set system $\mathcal{S}$ (Section 7) with a multi-scale builder's strategy that produces a 1-balanced chain in $\mathcal{S}$ for at least 90% of balanced colorings, and then invokes the worst-case to average-case reduction (Theorem 2.4).

8.1 Hypotheses on M_0

The multi-scale recursion requires M_0 to be large enough that certain arithmetic conditions hold at every scale. We set $M_0 = 700$ and verify two conditions for all $m_{j+1} \geq M_0$:

(H1) **Descent fits.** $m_{j+1} - C_3 - KC_3 \cdot \frac{3}{2} \ln m_{j+1} \geq M_0/2$.

This ensures that after a descent consumes at most $C_3(1 + K \ln m'_j) \leq C_3 + KC_3 \cdot \frac{3}{2} \ln m_{j+1}$ elements from I_{j+1} (using $\ln m'_j \leq \frac{3}{2} \ln m_{j+1} + 1$), at least $M_0/2$ elements remain.

Check at $m_{j+1} = 700$: $700 - 8 - 48 \ln 700 \approx 700 - 8 - 314 = 378 \geq 350$.

(H2) **Gap fits.** For all $m' \geq M_0/2 = 350$: $m' - (m')^{2/3} \geq C_1 \ln m'$.

This ensures that the steered path runs long enough (at least $C_1 \ln m'$ steps before block exhaustion) to guarantee a balanced visit.

Check at $m' = 350$: $350 - 350^{2/3} \approx 350 - 50 = 300 \geq 28 \ln 350 \approx 164$.

8.2 Probability space

The set system $\mathcal{S}$ is *fixed*: it depends only on n and the universal constants, not on the coloring f. The builder's strategy, however, is randomized: it depends on f (which determines the steered path) and on auxiliary tie-breaking coins (from the steered path's coin flips when both options give the same $|H|$). We will show that

$$\Pr_{f,\text{coins}}[\text{strategy produces a 1-balanced chain in } \mathcal{S}] \geq 9/10.$$

Since every chain the strategy produces lies in $\mathcal{S}$ regardless of the coin outcomes (because $\mathcal{S}$ includes all possible gap-filler subsets for all possible nestings), it follows that

$$\Pr_f[\mathcal{S} \text{ contains a 1-balanced chain}] \geq \Pr_{f,\text{coins}}[\text{strategy succeeds}] \geq 9/10.$$

8.3 The builder's strategy

Theorem 8.1. $\Pr_f[\mathcal{S}$ *contains a* 1*-balanced maximal chain*$] \geq 9/10$.

Proof. The strategy processes $[n]$ in scales $j = 0, 1, 2, \ldots$. At each scale, the builder steers a two-block path through a segment, gap-fills between balanced visits, and eventually exhausts one block, producing a smaller residual segment for the next scale. The following invariant is maintained at the start of each scale.

Invariant (j)**.**

(i) The chain has been constructed through some level ℓ_j, with $|f(C)| \leq 1$ at every level $\leq \ell_j$. The chain set $C^{(j)}$ at level ℓ_j satisfies $\left|f(C^{(j)})\right| \leq 1$.

(ii) There is a contiguous segment $I_j \subseteq [n]$ of size m_j and a grid position (a_0, b_0) in the local coordinates of I_j (meaning a_0 elements of I_j^L and b_0 of I_j^R have already been consumed during the transition from scale $j-1$; at scale 0, $(a_0, b_0) = (0, 0)$).

(iii) The pool $[n] \setminus C^{(j)}$ consists of the $m'_j := m_j - a_0 - b_0$ unassigned elements of I_j, with $m'_j \geq M_0/2$.

At scale 0: $I_0 = [n]$, $(a_0, b_0) = (0, 0)$, $C^{(0)} = \emptyset$, and $m'_0 = n$. All three conditions hold.

Processing scale j (while $m'_j \geq M_0/2$).

The steered path on I_j continues from grid position (a_0, b_0). At each step, the overall chain imbalance is $H(t) = f(C_t)$ and $\hat{h}(t) = |H(t)|$. By (i), $\hat{h}(0) = \left|f(C^{(j)})\right| \leq 1$.

Step 1: Forced probability. By Lemma 4.1, $p_t \leq 1/4$ at every step with $\hat{h} \geq 1$ and both blocks active. As discussed in Remark 4.2, this requires no conditions beyond $f([n]) = 0$, so no scale-specific verification is needed.

Step 2: Imbalance and gap bounds. By Theorem 5.6 (applied with $m' = m'_j$), the imbalance stays bounded:

$$\Pr\Big[\max_t \hat{h}(t) \geq 1 + K \ln m'_j\Big] \leq \frac{3}{(m'_j)^2}.$$

By Corollary 5.8, the steered path visits $\{\hat{h} \leq 1\}$ every $\leq C_1 \ln m'_j$ steps, with probability $\geq 1 - 1/m'_j$. To ensure that at least one balanced visit occurs before block exhaustion, we need the path to run for at least $C_1 \ln m'_j$ steps in the "active phase" (when both blocks are still available). Since block exhaustion leaves a residual of size $\leq (m'_j)^{2/3}$, the active phase lasts at least $m'_j - (m'_j)^{2/3}$

steps. By (H2), this exceeds $C_1 \ln m'_j$ (using $m'_j \geq M_0/2 = 350$), so at least one balanced visit is guaranteed.

Step 3: Gap filling within scale j. Let $\tau_0 < \tau_1 < \cdots < \tau_r$ be the times at which the steered path visits $\{\hat{h} \leq 1\}$. Between consecutive visits τ_i and τ_{i+1}, both chain sets C_{τ_i} and $C_{\tau_{i+1}}$ have $|f| \leq 1$, and the elements $C_{\tau_{i+1}} \setminus C_{\tau_i}$ form a subset of a gap region on I_j's grid of size at most $C_1 \ln m'_j$. By Lemma 6.1, there exists a 1-balanced ordering of these elements. Each intermediate chain set is a prefix pair plus a subset of the gap region, so it lies in $\mathcal{L}(I_j)$ (the gap size $C_1 \ln m'_j \leq \Gamma \ln m_j + M_0$ fits within the budget of Definition 7.1).

Step 4: Block exhaustion. By Lemma 3.3 (applied with m'_j remaining steps), with probability $\geq 1 - 1/m'_j$, one block is exhausted at some time T_j, leaving a contiguous residual $I_{j+1} \subseteq I_j$ of size

$$m_{j+1} \leq 4\sqrt{m'_j \ln m'_j} \leq (m'_j)^{2/3} \leq m_j^{2/3},$$

where the last two inequalities hold for $m'_j \geq M_0/2$.

Step 5: Transition or base case.
Let $S = C_{\tau_r}$ denote the chain set at the last balanced visit of scale j, so $|f(S)| \leq 1$.

Case A: $m_{j+1} \geq M_0$ (continue recursion).
At block exhaustion, the overall imbalance $\hat{h}(T_j)$ may be as large as $1 + K \ln m'_j$. To restore $\hat{h} \leq 1$, we run the steered path on I_{j+1} starting from grid position $(0,0)$. Since Lemma 4.1 applies at every step (the global pool argument requires only $f([n]) = 0$), Lemma 5.9 guarantees that the descent from $\hat{h}(T_j)$ to $\hat{h} \leq 1$ takes at most $C_3(1 + K \ln m'_j)$ steps, with probability $\geq 1 - 1/(m'_j)^2$. Let (a', b') denote the grid position on I_{j+1} at the moment the first balanced visit occurs, so $a' + b' \leq C_3 + KC_3 \ln m'_j$.
Let S' be the chain set at this first balanced visit on I_{j+1}. By construction, $|f(S')| \leq 1$. The transition elements $S' \setminus S$ consist of two parts:

(a) the tail from scale j (elements consumed between the last balanced visit τ_r and block exhaustion T_j): at most $C_1 \ln m'_j$ elements;

(b) descent elements from I_{j+1}: at most $C_3 + KC_3 \ln m'_j$ elements.

The total is at most $C_1 \ln m'_j + C_3 + KC_3 \ln m'_j = \Gamma \ln m'_j + C_3 \leq \Gamma \ln m_j + M_0$ (using $m'_j \leq m_j$ and $C_3 \leq M_0$). Since $|f(S)| \leq 1$ and $|f(S')| \leq 1$, Lemma 6.1 provides a 1-balanced ordering of these transition elements. The intermediate sets lie in $\mathcal{L}(I_j)$, as the transition size fits within the gap budget.
We now set $C^{(j+1)} = S'$ and $(a_0^{(j+1)}, b_0^{(j+1)}) = (a', b')$, and verify the invariant for scale $j+1$:

(i) The chain has $|f| \leq 1$ at every level through ℓ_{j+1} (the level of S'), and $\left|f(C^{(j+1)})\right| = |f(S')| \leq 1$.

(ii) I_{j+1} is a contiguous segment with valid grid position (a', b').

(iii) The effective size is $m'_{j+1} = m_{j+1} - a' - b' \geq m_{j+1} - C_3 - KC_3 \ln m'_j$. By (H1) (using $\ln m'_j \leq \frac{3}{2} \ln m_{j+1} + 1$, which follows from $m_{j+1} \geq (m'_j)^{2/3}/2$), this is at least $M_0/2$.

Case B: $m_{j+1} < M_0$ (base case).
When the residual is smaller than M_0, there is no need to recurse. The remaining elements consist of the tail from scale j (at most $C_1 \ln m'_j$ elements) plus the entirety of I_{j+1} (fewer than

M_0 elements), for a total of at most $C_1 \ln m'_j + M_0 \le \Gamma \ln m_j + M_0$. Since $|f(S)| \le 1$ and $f(S \cup \text{remaining}) = f([n]) = 0$, Lemma 6.1 provides a 1-balanced ordering of the remaining elements. The intermediate sets lie in $\mathcal{L}(I_j)$ (the total size fits within the gap budget). The chain is now complete.

Chain validity. It remains to verify that every chain set C_ℓ produced by the strategy lies in $\mathcal{S}$. At level ℓ, suppose scales $0, \dots, j-1$ have been completed and scale j is currently active. Then $C_\ell = L_0 \cup \dots \cup L_{j-1} \cup L'_j$ where, for each $i < j$, the set $L_i \in \mathcal{L}(I_i)$ is scale i's completed local pattern with $L_i \subseteq I_i \setminus I_{i+1}$, and $L'_j \in \mathcal{L}(I_j)$ is the partial pattern from scale j (a prefix pair plus a gap-region subset). This is a depth-j decomposition (Definition 7.3) with terminal $L'_j \in \mathcal{L}(I_j)$.

During a transition from scale j to scale $j+1$: the transition elements lie in I_j (covered by $\mathcal{L}(I_j)$), and the contribution from I_{j+1} is a prefix pair (in $\mathcal{L}(I_{j+1})$), giving a depth-$(j+1)$ decomposition. In the base case, the terminal pattern $L_j \in \mathcal{L}(I_j)$ covers the tail and the residual. In all cases, $C_\ell \in \mathcal{S}$.

Failure probability. At each scale j, the strategy can fail due to: the imbalance exceeding $1 + K \ln m'_j$ (probability $\le 3/(m'_j)^2$ from Theorem 5.6), some excursion exceeding $C_1 \ln m'_j$ steps (probability $\le 1/m'_j$ from Corollary 5.8), the block deviation being too large (probability $\le 1/m'_j$ from Lemma 3.3), or the descent taking too long (probability $\le 1/(m'_j)^2$ from Lemma 5.9). The total failure probability at scale j is at most

$$\frac{3}{(m'_j)^2} + \frac{1}{m'_j} + \frac{1}{m'_j} + \frac{1}{(m'_j)^2} \le \frac{7}{m'_j}.$$

To sum over all scales, we read backward from the last active scale. Since $m_j \ge n^{(2/3)^j}$ (by the shrinkage condition) and $m'_j \ge m_j/2$, the contribution from scale $J_{\max} - k$ is at most $2/M_0^{(3/2)^k}$. The sum forms a rapidly convergent series:

$$\sum_{j=0}^{J_{\max}} \frac{7}{m'_j} \le 14 \sum_{k=0}^{\infty} \frac{1}{M_0^{(3/2)^k}} \le \frac{14}{M_0}\left(1 + \frac{1}{M_0^{1/2}} + \frac{1}{M_0^{5/4}} + \cdots\right) < \frac{28}{M_0} = \frac{28}{700} < \frac{1}{10},$$

where the parenthetical series converges to less than 2 since each term is smaller than the previous by a factor of at least $M_0^{1/2} \ge 26$. □

Proof of Theorem 1.1. By Theorem 8.1, $\mathcal{S}$ is a $(9/10,\ 1)$-balanced-chain set system. By Lemma 7.5, $|\mathcal{S}| \le n^{3C_2+7}$. Applying Theorem 2.4:

$$N(n) \le O\left(\frac{|\mathcal{S}| \cdot n}{9/10}\right) = O(|\mathcal{S}| \cdot n) \le n^{3C_2+8}.$$

With $C_2 = 64$: $N(n) \le n^{200}$. □

9 Conclusion and open questions

Comparison with Fabris et al. The construction of [FLSY26] uses a single ordering of $[n]$ (prefix sets $[1,t]$) and relies on the zeros of the random walk $W(t) = f([1,t])$. By the classical result of Csáki, Erdős, and Révész [CER85], these zeros are spaced $\Theta(n/\log n)$ apart, leading to a gap-filling recursion of depth $\Theta(\log n/\log\log n)$ and the upper bound $n^{O(\log n/\log\log n)}$.

Our construction uses two orderings (one per block) with active steering. With a single ordering, the imbalance is a symmetric random walk ($p = q = 1/2$), and controlling it requires the entire recursion. With two orderings, the (absolute) imbalance $\hat{h}$ is negatively biased ($p \le 1/4$, $q \ge 3/4$), and the supermartingale $3^{\hat{h}}$ controls it directly. Gaps shrink from $\Theta(n/\log n)$ to $O(\log n)$, making the gap-filler cost polynomial ($2^{O(\log n)} = n^{O(1)}$) rather than requiring deep recursion.

The optimal exponent. Our proof gives $N(n) \le n^{200}$. The best lower bound is $N(n) \ge \Omega(n^2)$, which follows from the work of Alon, Kumar, and Volk [AKV20] on balancing set systems as reinterpreted by [FLSY26]. What is the true growth rate of $N(n)$? Is $N(n) = \Theta(n^2)$? Using $d > 2$ blocks in the steering construction would reduce the forced probability from $\approx 1/4$ to $\approx 1/2^d$, potentially improving the exponent, though we have not pursued this.

Uniformity. Fabris et al. [FLSY26] ask whether there is a *uniform* construction of a full-rank polynomial computable by a polynomial-size mABP. As noted in Remark 7.6, our set system $\mathcal{S}$ is fully explicit. The sole source of non-uniformity is the worst-case to average-case reduction (Theorem 2.4), which converts $\mathcal{S}$ into a 1-balanced-chain set system by applying $O(n)$ random permutations and arguing by the probabilistic method that the result works for all balanced colorings. Once a worst-case 1-balanced-chain set system is in hand, the construction of [FLSY26, Theorem 5.4] (see also Appendix A) deterministically produces a full-rank polynomial and a polynomial-size mABP over the transcendental extension $\mathbb{F}(W)$. To obtain a construction over the base field $\mathbb{F}$ itself (when $\mathbb{F}$ is infinite), one additionally needs to project the transcendental parameters W to elements of $\mathbb{F}$; a random restriction works with high probability (by Schwartz–Zippel), but a deterministic choice would require further derandomization. We note that neither source of non-uniformity affects the barrier (Corollary 1.3), which is a statement about existence, not constructivity.

The set-multilinear setting and hardness escalation. As mentioned in Section 1.3, set-multilinear models have become prominent due to a phenomenon known as *hardness escalation*: lower bounds for set-multilinear models at small degree can be lifted to lower bounds for general models. Raz [Raz13] showed that superpolynomial lower bounds for set-multilinear formulas at degree $r \le O(\log n/\log\log n)$ would imply superpolynomial lower bounds for *general* algebraic formulas. This was a key ingredient in the breakthrough of Limaye, Srinivasan, and Tavenas [LST25], who proved the first superpolynomial lower bounds against constant-depth algebraic circuits[8]: their proof proceeds by first establishing lower bounds against constant-depth set-multilinear circuits computing the IMM polynomial, and then lifting via set-multilinearization. On the ABP side, Bhargav, Dwivedi, and Saxena [BDS25] established an analogous escalation: superpolynomial lower bounds for $\Sigma_\pi$smABPs[9] at degree $r \le O(\log n/\log\log n)$ would imply superpolynomial lower bounds for general ABPs. Chatterjee, Kush, Saraf, and Shpilka [CKSS24] made progress toward this by proving superpolynomial $\Sigma_\pi$smABP lower bounds for a set-multilinear polynomial of degree $\omega(\log n)$. While this appears tantalizingly close to the escalation threshold, they in fact prove their lower bound for a polynomial in smVBP—for which escalation to general ABP lower bounds is, by definition, impossible. New techniques are therefore needed to pursue general ABP lower bounds via escalation through the $\Sigma_\pi$smABP model.

Our barrier shows that the min-partition rank method cannot separate smVBP from smVP. However, we believe there may be more room for progress in the set-multilinear setting than in the general multilinear one. The reason is that recent set-multilinear lower bounds—specifically the breakthrough result of Limaye, Srinivasan, and Tavenas and its follow-ups [LST25, TLS22,

[8] Over fields of characteristic 0; this was subsequently extended to all fields by Forbes [For24].

[9] A $\Sigma_\pi$smABP is a sum of ordered set-multilinear ABPs, each with a possibly different variable ordering. See [CKSS24] for a precise definition.

LST22]—rely on a technique that seems to be appreciably different from the min-partition rank method: the *lopsided partial derivative method*. The standard min-partition rank method partitions the variable blocks $X_1, \dots, X_d$ into two groups and analyzes the rank of the resulting coefficient matrix, where all blocks have the same size N. The lopsided method, by contrast, first restricts or projects the variable sets so that different blocks X_i have *different sizes*, creating an asymmetry that the standard method does not exploit. It is this asymmetry that enabled the constant-depth circuit lower bounds of [LST25] and the unbounded-depth set-multilinear formula lower bounds of [TLS22] for the IMM polynomial.

Crucially, the characterization of Fabris et al. [FLSY26]—and hence our barrier—applies specifically to the min-partition rank method, which uses equipartitions (all blocks of equal size). It does not immediately rule out the lopsided partial derivative method as a tool for proving set-multilinear ABP lower bounds. We therefore propose the following directions:

1. *Set-multilinear ABP lower bounds via lopsided methods.* Can the lopsided partial derivative technique, or extensions of it, prove superpolynomial lower bounds for smVBP? Our barrier does not apply to such methods, and the success of the lopsided approach for set-multilinear formulas [LST25, TLS22] suggests this is a natural avenue.

2. *Understanding the gap between $\Sigma_\pi$smABP and smVBP.* The results of [CKSS24] show that the $\Sigma_\pi$smABP model is strictly weaker than smVBP (superpolynomially so, even for polynomials with superlogarithmic degree in smVBP). The results of [BDS25] show that proving $\Sigma_\pi$smABP lower bounds at slightly smaller degree would have dramatic consequences. Understanding the precise boundary—and whether the lopsided method or other techniques can reach the escalation threshold—is a compelling open problem.

3. *Characterizing the lopsided method.* Is there an analogue of the Fabris et al. characterization for the lopsided partial derivative method? Such a characterization would clarify the power of this technique for ABPs, and either provide a new path to lower bounds or reveal further barriers.

Acknowledgments

This project began as an earnest attempt to prove $N(n) = n^{\omega(1)}$ (and thereby $\mathsf{smVBP} \neq \mathsf{smVP}$) via expander-based methods; I am thankful to Gaia Carenini for those early discussions. I am also grateful to Robert Andrews, Tom Gur, Shubhangi Saraf, and Ben Lee Volk for detailed comments that helped improve the presentation of this paper.

A From balanced-chain set systems to full-rank mABPs

We prove Corollary 1.2: a 1-balanced-chain set system of polynomial size yields a full-rank polynomial computable by a polynomial-size mABP (or smABP; see Remark A.1). The argument follows [FLSY26, Theorems 5.4 and 5.6]; we include it here for completeness and because the 1-balanced case admits a clean presentation.

Let $X = [n]$ with n even, and let $\mathcal{X}$ be a 1-balanced-chain set system over X of size s. Write $\mathcal{C}(\mathcal{X})$ for the set of maximal chains in $\mathcal{X}$. Let $V = \{x_i : i \in X\}$ be a set of n variables and $W = \{w_{t,u} : t \in [n],\, u \in X\}$ a set of n^2 formal indeterminates (which will later be restricted to field elements).

Step 1: The polynomial over a transcendental extension. For each maximal chain $C = (C_0, C_1, \dots, C_n) \in \mathcal{C}(\mathcal{X})$, let $\pi_C\colon [n] \to X$ be the bijection defined by $\pi_C(i) = C_i \setminus C_{i-1}$ (the element added at step i). Define the *gadget polynomial*

$$Q_C \;=\; \prod_{i=1}^{n/2} \bigl(x_{\pi_C(2i-1)} + x_{\pi_C(2i)}\bigr) \;\in\; \mathbb{F}[V],$$

and the *weighted polynomial*

$$P_C \;=\; Q_C \cdot \prod_{i=1}^{n} w_{i,\pi_C(i)} \;\in\; \mathbb{F}[V \cup W].$$

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

Step 2: The mABP. We construct an mABP M over $\mathbb{F}(W)$ of size at most s computing P. The vertices of M are the sets in $\mathcal{X}$ of even cardinality, with source $\emptyset$ and sink X. For every chain (R, S, T) in $\mathcal{X}$ with $|R|$ even, $|S| = |R| + 1$, and $|T| = |R| + 2$, let $u = S \setminus R$ and $v = T \setminus S$, and add an edge from R to T with label

$$w_{|S|,u} \cdot (x_u + x_v) \cdot w_{|T|,v} \;\in\; \mathbb{F}(W)[V].$$

Each source-to-sink path in M corresponds to a maximal chain C and computes the monomial $\prod_{i=1}^{n/2} w_{2i-1,\pi_C(2i-1)}(x_{\pi_C(2i-1)} + x_{\pi_C(2i)}) w_{2i,\pi_C(2i)}$, so M computes P. Since each variable x_u appears on at most one edge per path, M is syntactically multilinear. The number of vertices is at most $|\mathcal{X}| = s$.

Step 3: Full-rankness over $\mathbb{F}(W)$. Fix a balanced partition $f\colon X \to \{\pm 1\}$. Since $\mathcal{X}$ is 1-balanced-chain, there exists a chain $C_f \in \mathcal{C}(\mathcal{X})$ with $|f(C_i)| \le 1$ for all i. In particular, for each $i \in [n/2]$, the two elements $\pi_{C_f}(2i-1)$ and $\pi_{C_f}(2i)$ satisfy $f(\pi_{C_f}(2i-1)) \neq f(\pi_{C_f}(2i))$ (since consecutive even-indexed sets have balance 0, each pair must cross the partition).

Consider the substitution $w_{t,u} \mapsto [\pi_{C_f}(t) = u]$ (i.e., $w_{t,u} = 1$ if $\pi_{C_f}(t) = u$ and 0 otherwise). Under this substitution, P projects to Q_{C_f}. Since the coefficient matrix of a product of polynomials on disjoint variable sets has rank equal to the product of the individual ranks (see, e.g., [FLSY26, Fact 5.2]), and each factor $(x_{\pi(2i-1)} + x_{\pi(2i)})$ has rank 2 with respect to f (as the two variables lie on opposite sides of the partition), we get

$$\operatorname{rank}_{\mathbb{F}}\bigl(M_f(Q_{C_f})\bigr) = \prod_{i=1}^{n/2} 2 = 2^{n/2}.$$

Since restricting the W-variables to 0/1 values can only decrease rank: $\operatorname{rank}_{\mathbb{F}(W)}(M_f(P)) \ge \operatorname{rank}_{\mathbb{F}}(M_f(Q_{C_f})) = 2^{n/2}$. Thus P is full-rank over $\mathbb{F}(W)$.

Step 4: Derandomization over $\mathbb{F}$. It remains to replace the transcendental elements W by elements of $\mathbb{F}$. For each balanced partition f, the full-rankness above means there exists a $2^{n/2} \times 2^{n/2}$ submatrix of $M_f(P)$ with nonzero determinant $D_f \in \mathbb{F}[W]$. Since each entry of $M_f(P)$ has degree at most n in W, the polynomial D_f has degree at most $\Delta := n \cdot 2^{n/2}$.

Let $S \subseteq \mathbb{F}$ be any subset of size $|S| > 2^n \cdot \Delta$ (which exists since $\mathbb{F}$ is infinite). By the Schwartz–Zippel lemma [Ore22, DL78, Sch80, Zip79], for a uniformly random $w \in S^W$:

$$\Pr_w[D_f(w) = 0] \;\leq\; \frac{\Delta}{|S|}.$$

By a union bound over all $\binom{n}{n/2} \leq 2^n$ balanced partitions:

$$\Pr_w[\exists \text{ balanced } f \text{ with } D_f(w) = 0] \;\leq\; \frac{2^n \cdot \Delta}{|S|} \;<\; 1.$$

Hence there exists $w^* \in S^W$ such that $D_f(w^*) \neq 0$ for every balanced f. Define $Q^* \in \mathbb{F}[V]$ by substituting w^* into P. Then $M_f(Q^*)$ has rank $2^{n/2}$ for every balanced partition f, so Q^* is full-rank over $\mathbb{F}$. The mABP for Q^* is obtained from M by the same substitution: its size is unchanged and it computes Q^* over $\mathbb{F}$.

Applying the above with $\mathcal{X}$ being the set system from Theorem 1.1 (of size $s = n^{O(1)}$) yields Corollary 1.2.

Remark A.1 (Set-multilinear case)**.** The same argument extends to set-multilinear polynomials with a single modification. Given a partition $\mathcal{P} = \{X_1, \ldots, X_n\}$ of the variable set into blocks of size N each, replace the gadget factor $(x_{\pi(2i-1)} + x_{\pi(2i)})$ by the inner product

$$\mathrm{IP}(X_{\pi(2i-1)}, X_{\pi(2i)}) \;=\; \sum_{j=1}^{N} x_{\pi(2i-1),j} \cdot x_{\pi(2i),j},$$

which is a set-multilinear polynomial in $X_{\pi(2i-1)} \cup X_{\pi(2i)}$ computable by a set-multilinear ABP of size $O(N)$. Since the inner product has full rank ($= N$) with respect to any partition that separates $X_{\pi(2i-1)}$ from $X_{\pi(2i)}$, the multiplicativity of coefficient matrix rank gives $\mathrm{rank}(M_f(Q_{C_f})) = N^{n/2}$ for every balanced partition f of the blocks. The Schwartz–Zippel derandomization proceeds identically (with $N^{n/2}$ replacing $2^{n/2}$). The resulting set-multilinear ABP has size $s \cdot O(nN)$, which is polynomial when $s = n^{O(1)}$ and $N = n^{O(1)}$.

References

[Aar04] Scott Aaronson. Multilinear formulas and skepticism of quantum computing. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 118–127. ACM, 2004.

[AKV20] Noga Alon, Mrinal Kumar, and Ben Lee Volk. Unbalancing sets and an almost quadratic lower bound for syntactically multilinear arithmetic circuits. *Combinatorica*, 40(2):149–178, 2020.

[AR16] Vikraman Arvind and S. Raja. Some lower bound results for set-multilinear arithmetic computations. *Chic. J. Theor. Comput. Sci.*, 2016, 2016.

[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.

[CELS18] Suryajith Chillara, Christian Engels, Nutan Limaye, and Srikanth Srinivasan. A near-optimal depth-hierarchy theorem for small-depth multilinear circuits. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 934–945. IEEE Computer Society, 2018.

[CER85] Endre Csáki, Paul Erdős, and Pál 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 *39th Computational Complexity Conference (CCC 2024)*, volume 300 of *LIPIcs*, pages 20:1–20:20, 2024.

[CKSV22] Prerona Chatterjee, Mrinal Kumar, Adrian She, and Ben Lee Volk. Quadratic lower bounds for algebraic branching programs and formulas. *Computational Complexity*, 31(2):8, 2022.

[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.

[DL78] Richard A. DeMillo and Richard J. Lipton. A probabilistic remark on algebraic program testing. *Information Processing Letters*, 7(4):193–195, 1978.

[DMPY12] Zeev Dvir, Guillaume Malod, Sylvain Perifel, and Amir Yehudayoff. Separating multilinear branching programs and formulas. In *Proceedings of the 44th STOC*, pages 615–624. ACM, 2012.

[FLMS15] Hervé Fournier, Nutan Limaye, Guillaume Malod, and Srikanth Srinivasan. Lower bounds for depth-4 formulas computing iterated matrix multiplication. *SIAM Journal on Computing*, 44(5):1173–1201, 2015.

[FLSY26] Théo Borém Fabris, Nutan Limaye, Srikanth Srinivasan, and Amir Yehudayoff. Multilinear algebraic branching programs and the min-partition rank method. *Electron. Colloquium Comput. Complex.*, TR26-001, 2026.

[For24] Michael A. Forbes. Low-depth algebraic circuit lower bounds over any field. In Rahul Santhanam, editor, *39th Computational Complexity Conference, CCC 2024, Ann Arbor, MI, USA, July 22-25, 2024*, LIPIcs, pages 31:1–31:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024.

[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.

[GR21] Purnata Ghosal and B. V. Raghavendra Rao. Limitations of sums of bounded read formulas and ABPs. In *Computer Science – Theory and Applications (CSR 2021)*, volume 12730 of *Lecture Notes in Computer Science*, pages 143–171. Springer, 2021.

[HY11] Pavel Hrubeš and Amir Yehudayoff. Homogeneous formulas and symmetric polynomials. *Comput. Complex.*, 20(3):559–578, 2011.

[Jan08] Maurice J. Jansen. Lower bounds for syntactically multilinear algebraic branching programs. In *Proceedings of the 33rd International Symposium on Mathematical Foundations of Computer Science*, MFCS '08, page 407–418, Berlin, Heidelberg, 2008. Springer-Verlag.

[KNS20] Neeraj Kayal, Vineet Nair, and Chandan Saha. Separation between read-once oblivious algebraic branching programs (roabps) and multilinear depth-three circuits. *ACM Trans. Comput. Theory*, 12(1):2:1–2:27, 2020.

[KS22] Deepanshu Kush and Shubhangi Saraf. Improved low-depth set-multilinear circuit lower bounds. In *37th Computational Complexity Conference (CCC 2022)*, volume 234 of *LIPIcs*, pages 38:1–38:16, 2022.

[KS23] Deepanshu Kush and Shubhangi Saraf. Near-optimal set-multilinear formula lower bounds. In *38th Computational Complexity Conference (CCC 2023)*, volume 264 of *LIPIcs*, pages 15:1–15:33, 2023.

[LST22] Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. On the partial derivative method applied to lopsided set-multilinear polynomials. In Shachar Lovett, editor, *37th Computational Complexity Conference, CCC 2022, Philadelphia, PA, USA, July 20-23, 2022*, volume 234 of *LIPIcs*, pages 32:1–32:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[LST25] Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Superpolynomial lower bounds against low-depth algebraic circuits. *J. ACM*, 72(4):26:1–26:35, 2025.

[MST16] Meena Mahajan, Nitin Saurabh, and Sébastien Tavenas. Vnp=vp in the multilinear world. *Inf. Process. Lett.*, 116(2):179–182, 2016.

[MU17] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*. Cambridge University Press, 2nd edition, 2017.

[Nis91] Noam Nisan. Lower bounds for non-commutative computation (extended abstract). In *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 410–418. ACM, 1991.

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

[Ore22] Oystein Ore. Über höhere kongruenzen. *Norsk Matematisk Forenings Skrifter*, 1(7):15, 1922.

[Raz06] Ran Raz. Separation of multilinear circuit and formula size. *Theory of Computing*, 2(6):121–135, 2006.

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

[Raz13] Ran Raz. Tensor-rank and lower bounds for arithmetic formulas. *J. ACM*, 60(6):40:1–40:15, 2013.

[RR19] C. Ramya and B. V. Raghavendra Rao. Lower bounds for multilinear order-restricted ABPs. In *44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019)*, volume 138 of *LIPIcs*, pages 52:1–52:14, 2019.

[RR20] C. Ramya and B. V. Raghavendra Rao. Lower bounds for special cases of syntactic multilinear ABPs. *Theor. Comput. Sci.*, 809:1–20, 2020.

[RSY08] 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.

[RT08] Ran Raz and Iddo Tzameret. The strength of multilinear proofs. *Comput. Complex.*, 17(3):407–457, 2008.

[RY09] Ran Raz and Amir Yehudayoff. Lower bounds and separations for constant depth multilinear circuits. *Comput. Complex.*, 18(2):171–207, 2009.

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

[Sap16] Ramprasad Saptharishi. A survey of lower bounds in arithmetic circuit complexity, 2016. Available at `https://github.com/dasarpmar/lowerbounds-survey`.

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

[SY10] Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3–4):207–388, 2010.

[TLS22] Sébastien Tavenas, Nutan Limaye, and Srikanth Srinivasan. Set-multilinear and non-commutative formula lower bounds for iterated matrix multiplication. In Stefano Leonardi and Anupam Gupta, editors, *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, pages 416–425. ACM, 2022.

[Val79] Leslie G. Valiant. Completeness classes in algebra. In *Proceedings of the 11th Annual ACM Symposium on Theory of Computing (STOC)*, pages 249–261. ACM, 1979.

[Wil91] David Williams. *Probability with Martingales.* Cambridge University Press, 1991.

[Zip79] Richard Zippel. Probabilistic algorithms for sparse polynomials. In *Symbolic and Algebraic Computation (EUROSAM '79)*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer, 1979.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il