

Blackbox Polynomial Identity Testing for Depth 3 Circuits

Neeraj Kayal* Shubhangi Saraf†

April 16, 2009

Abstract

We study $\Sigma\Pi\Sigma(k)$ circuits, i.e., depth three arithmetic circuits with top fanin k . We give the first *deterministic* polynomial time *blackbox* identity test for $\Sigma\Pi\Sigma(k)$ circuits over the field \mathbb{Q} of rational numbers, thus resolving a question posed by Klivans and Spielman (STOC 2001).

Our main technical result is a structural theorem for $\Sigma\Pi\Sigma(k)$ circuits that compute the zero polynomial. In particular we show that if a $\Sigma\Pi\Sigma(k)$ circuit $C = \sum_{i \in [k]} A_i = \sum_{i \in [k]} \prod_{j \in [d]} \ell_{ij}$ computing the zero polynomial, where each A_i is a product of linear forms with coefficients in \mathbb{R} , is *simple* ($\gcd\{A_i \mid i \in [k]\} = 1$) and *minimal* (for all proper nonempty subsets $S \subset [k]$, $\sum_{i \in S} A_i \neq 0$), then the *rank* (dimension of the span of the linear forms $\{\ell_{ij} \mid i \in [k], j \in [d]\}$) of C can be upper bounded by a function only of k . This proves a weak form of a conjecture of Dvir and Shpilka (STOC 2005) on the structure of identically zero depth three arithmetic circuits. Our blackbox identity test follows from this structural theorem by combining it with a construction of Karnin and Shpilka (CCC 2008).

Our proof of the structure theorem exploits the geometry of finite point sets in \mathbb{R}^n . We identify the linear forms appearing in the circuit C with points in \mathbb{R}^n . We then show how to apply high dimensional versions of the Sylvester–Gallai Theorem, a theorem from incidence-geometry, to identify a special linear form appearing in C , such that on the subspace where the linear form vanishes, C restricts to a simpler circuit computing the zero polynomial. This allows us to build an inductive argument bounding the rank of our circuit. While the utility of such theorems from incidence geometry for identity testing has been hinted at before, our proof is the first to develop the connection fully and utilize it effectively.

*Microsoft Research India. neeraka@microsoft.com.

†MIT CSAIL. shibs@mit.edu.

1 Introduction

Identity testing is the following problem: given an arithmetic circuit¹ computing a multivariate polynomial $f(X_1, \dots, X_n)$ over a field \mathbb{F} , determine if the polynomial is identically zero. Algorithms for primality testing [AB03], perfect matching [MVV87] and some fundamental structural results in complexity such as the PCP Theorem and IP=PSPACE involve testing if a particular polynomial is zero.

Schwartz [Sch80] and Zippel [Zip90] observed that by evaluating a polynomial at randomly chosen points from a sufficiently large domain, we can determine if the polynomial is nonzero with high probability. The correctness of their algorithm follows from the simple observation that any polynomial of total degree d cannot have many roots over a field whose size is much larger than d . The Schwartz–Zippel Lemma combined with a standard counting argument implies that for every integer s , there is a *poly*(s)-sized set of points P such that for every circuit C of size s , C computes the zero polynomial if and only if $C(\mathbf{a}) = 0$ for every $\mathbf{a} \in P$. Blackbox Identity testing is the problem of giving an *explicit*² construction of such a test set P . Any explicit construction of such a set of points immediately gives, via interpolation, an explicit polynomial f which cannot be computed by circuits of size s [Agr05].

A more surprising connection between identity testing and the task of proving arithmetic circuit lower bounds was discovered by Impagliazzo and Kabanets [IK03] who showed that *any* polynomial-time algorithm for identity testing (not necessarily a blackbox identity test³) would also imply certain arithmetic circuit lower bounds. More specifically, they showed that if identity testing has an efficient deterministic polynomial time algorithm then (almost) NEXP does not have polynomial size arithmetic circuits. For the pessimist, this indicates that derandomizing identity testing is a hopeless problem. For the optimist, this means on the contrary that to obtain an arithmetic circuit lower bound, we “simply” have to prove a good upper bound on identity testing.

Because of the difficulty of the general problem, research has focussed on bounded depth arithmetic circuits. Grigoriev and Karpinski [GK98] have shown that any depth three arithmetic circuit over a finite field computing the permanent or the determinant requires exponential size⁴. But progress in this direction stalled and very recently, an “explanation” for this was discovered by Agrawal and Vinay [AV08] who showed that there is chasm at depth four - proving exponential lower bounds for depth four arithmetic circuits already implies exponential lower bounds for arbitrary depth arithmetic circuits. They also showed that a complete blackbox derandomization of Identity Testing problem for depth four circuits with multiplication gates of small fanin implies a nearly complete derandomization of general Identity Testing. As most of these questions are fairly easy

¹Arithmetic circuits are circuits with two types of internal nodes/gates: a \times gate computes the product of its inputs whereas a $+$ gate is allowed to compute an arbitrary linear combination of its inputs, and the wires carry elements of a field \mathbb{F} .

²Over infinite fields such as the rationals, we require more so that the evaluation can be carried out efficiently - the bit-length of the coordinates of the points in P need to be polynomially bounded. Furthermore, if the degree of the polynomial computed is huge then we also require the construction to give a prime of small bit-length so that the computation can be carried out modulo p .

³A non-blackbox algorithm is given a full description of the circuit and it needs to decide whether it is identically 0.

⁴The size of the field is held constant and a lower bound is obtained on the size of the circuit as a function of the dimension of the matrix.

for depth two circuits, we see that depth three circuits stand between the relatively easy (depth-two) and the difficult, fairly general case (depth-four). Hence it is a worthwhile goal to get a good understanding of depth-three circuits.

Another important direction of research, pursued in the works of Chen and Kao [CK00], Lewin and Vadhan [LV98], Klivans and Spielman [KS01] and Agrawal and Biswas [AB03] on the identity testing problem has been the effort to take advantage of the structure of a polynomial to reduce the number of random bits needed for identity testing. In this process, Klivans and Spielman gave a blackbox identity testing algorithm for depth two arithmetic circuits and posed as a challenge the problem of devising blackbox identity testers for depth three circuits with bounded top fanin. Recall that a depth three arithmetic circuit, also called a $\Sigma\Pi\Sigma$ -circuit, has an addition gate at the top (output) layer, followed by multiplication gates at the middle layer, followed by addition gates at the bottom layer, the gates being of arbitrary fanin⁵. In other words, a $\Sigma\Pi\Sigma$ circuit is a sum of terms, each of which is a product of a linear function of the input variables. We denote the set of n input, depth three circuits, where the top addition gate has fanin k , and the middle multiplication gates have fanin at most d , by $\Sigma\Pi\Sigma(k, d, n)$. The challenge posed by Klivans and Spielman was taken up by Dvir and Shpilka [DS05] and then by Kayal and Saxena [KS06] and a non-blackbox deterministic polynomial-time algorithm was devised (see also [AM07]). Recently Karnin and Shpilka [KS08a] obtained a quasi-polynomial time *blackbox* identity test for $\Sigma\Pi\Sigma(k, d, n)$ circuits. Despite the progress made on this question, a deterministic polynomial-time blackbox test had remained elusive.

In this paper we fully resolve the Klivans–Spielman challenge for arithmetic circuits with rational coefficients by giving the first deterministic blackbox identity test for $\Sigma\Pi\Sigma(k, d, n)$ circuits whose running time, for every fixed value of k , is *polynomial* in d and n . Our main technical contribution towards the proof of this result is a structural theorem for such circuits that answers a weak form of a conjecture by Dvir and Shpilka. In particular we prove that the “rank” of bounded top-fanin $\Sigma\Pi\Sigma$ circuits is a constant depending only on the size of the top fanin. Combined with a result from [KS08a] which says that a good rank bound suffices for black-box identity testing, we obtain the full result. The proof of our structure theorem uses results from the incidence geometry of \mathbb{R}^n . In particular we invoke a high dimensional version of the Sylvester–Gallai Theorem that enables us to identify certain configurations of linear forms appearing in any high rank circuit that prevent it from being identically zero. The survey by Borwein and Moser [BM90] contains a good introduction to the Sylvester–Gallai Theorem - its history, its proofs and its many generalizations. Before we state the conjecture and our main result, we introduce some terminology.

2 Definitions and statement of results

$[k]$ denotes the set $\{1, 2, \dots, k\}$. \mathbb{Q} denotes the field of rational numbers and \mathbb{R} the field of real numbers.

Depth Three Arithmetic Circuits. We consider arithmetic circuits with coefficients in a field \mathbb{F} (in this paper, \mathbb{F} will always be either \mathbb{Q} or \mathbb{R}) . A $\Sigma\Pi\Sigma$ circuit C is a formal expression of the

⁵If the depth three circuit has a multiplication gate at the top then problems pertaining to identity testing and lower bounds boil down to the relatively easy case of depth two circuits.

form

$$\sum_{i \in [k]} A_i = \sum_{i \in [k]} \prod_{j \in [d]} \ell_{ij},$$

where the ℓ_{ij} are linear forms of the type $\ell_{ij} = \sum_{k \in [n]} a_k \cdot X_k = \mathbf{a} \cdot \mathbf{X}$, where $\mathbf{a} = (a_1, \dots, a_n)$ is a fixed vector in \mathbb{Q}^n , and $\mathbf{X} = (X_1, \dots, X_n)$ is the tuple of indeterminates. k is the fanin of the top gate of the circuit and d is the fanin of each multiplication gate A_i . The A_i 's are referred to as the terms or the constituent $\Pi\Sigma$ subcircuits of C and the ℓ_{ij} 's as the set of linear forms that belong to the circuit. Recall that we denote the class of such circuits by $\Sigma\Pi\Sigma(k, d, n)$.

Remark Note that the above definition only captured homogeneous circuits. For the purpose of identity testing, we can assume this without loss of generality. Indeed, notice that a polynomial $C(X_1, \dots, X_n)$ of degree less than or equal to d is zero if and only if the corresponding homogeneous polynomial $Z^d \cdot C(\frac{X_1}{Z}, \dots, \frac{X_n}{Z})$ is zero. This observation can be used to homogenize the input circuit C . Notice also that after the homogenization all multiplication gates $\{A_i\}$ have the same fanin d . In the rest of the paper we will assume that the input circuit is homogeneous, i.e. the ℓ_{ij} 's as linear forms (zero constant term) and all the multiplication gates have the same fanin d .

When the context requires it, we will drop some of the parameters in talking about this class of circuits. $\Sigma\Pi\Sigma(k)$ will denote the set of $\Sigma\Pi\Sigma$ circuits with top fanin k and $\Sigma\Pi\Sigma(k, d)$ will be the set of $\Sigma\Pi\Sigma$ circuits with top fanin k and middle (multiplication gate) fanin d . Given such a circuit, we can naturally associate it with the polynomial computed by it. We say that $C \equiv 0$ if the polynomial computed by C is identically zero. We are now ready to state our main result which shows that for every fixed k , there is a deterministic blackbox identity tester for the class of $\Sigma\Pi\Sigma(k, d, n)$ circuits over the field \mathbb{Q} that runs in time polynomial in n and d .

Theorem 2.1 [Blackbox PIT for $\Sigma\Pi\Sigma(k)$ circuits] *There is a deterministic algorithm that takes as input a triple (k, d, n) of natural numbers and in time $\text{poly}(n) \cdot d^{2^{O(k \cdot \log k)}}$, outputs a set $P \subset \mathbb{Z}^n$ with the following properties:*

1. Any $\Sigma\Pi\Sigma(k, d, n)$ circuit C with rational coefficients computes the zero polynomial if and only if $C(\mathbf{a}) = 0$ for every $\mathbf{a} \in P$.
2. The number of points in P is $\text{poly}(n) \cdot d^{2^{O(k \cdot \log k)}}$.
3. For every $(a_1, \dots, a_n) \in P$ and every $i \in [n]$: $|a_i| \leq \text{poly}(2^{n^2} \cdot d) \cdot 2^{2^{O(k \cdot \log k)}}$. In particular, the bit-length of each point in P is $2^{O(k \cdot \log k)} \cdot O(n^3 \cdot \log d)$.

Remark.

1. Notice that in the theorem above, the number of points in P and the bit-lengths of these points are both independent of the bit-lengths of the constants from \mathbb{Q} used in the circuit. Hence we can allow arbitrary constants from \mathbb{Q} to be used on the edges coming into addition gates in the circuit. We get this feature because the two main components of the proof, the structure theorem (Theorem 2.2) as well as the result from [KS08a] (Lemma 2.3) are independent of the bit-lengths of the constants from \mathbb{Q} used in the circuit.

2. For every fixed value of k , the algorithm for the construction of the set P alluded to in the above theorem can in fact be implemented in TC^0 . Combined with the observation that a given depth three circuit can be evaluated at a given point in TC^0 , we get a deterministic P-uniform TC^0 -algorithm for identity testing of $\Sigma\Pi\Sigma(k, d, n)$ circuits. Previously no efficient deterministic algorithm, not even a non-blackbox one, for identity testing of $\Sigma\Pi\Sigma(k)$ was known which can be implemented in TC^0 . We do not stress the constant depth computability because it is not the main point of our result. But the ability to do identity testing using small depth uniform circuits can potentially be useful at other places, such as in the context of the question [MVV87]: Is $\text{BipartiteMatching} \in \text{NC}$?
3. For concreteness, we only state our results over \mathbb{Q} but our theorem is valid over any field that can be embedded into the real numbers, in particular for any totally real extension of \mathbb{Q} . Over such fields, the same set of points as constructed above suffices for identity testing.
4. As noted in the introduction, a blackbox identity testing algorithm for any class of circuits can in general be used to construct explicit polynomials that are hard to compute by the corresponding class of circuits. For the class of $\Sigma\Pi\Sigma(k)$ arithmetic circuits however, such polynomials were already known. For example, Shpilka and Wigderson [SW01] have shown that the higher degree elementary symmetric polynomials cannot be computed by such circuits.

After introducing the requisite terminology, we state the two main ingredients leading to this theorem - our proof of a conjecture by Dvir and Shpilka and a construction of rank preserving subspaces by Karnin and Shpilka.

2.1 Notions related to a $\Sigma\Pi\Sigma$ circuit

Let $C = \sum_{i \in [k]} A_i = \sum_{i \in [k]} \prod_{j \in [d]} \ell_{ij}$ be a $\Sigma\Pi\Sigma$ circuit. We give the following definitions:

minimal: We say that C is *minimal* if no strict nonempty subset of its constituent $\Pi\Sigma$ polynomials $\{A_1, \dots, A_k\}$ sums to zero.

simple: We say that C is *simple* if the gcd of its constituent $\Pi\Sigma$ polynomials, $\gcd(A_1, A_2, \dots, A_k)$ equals one. The simplification of a $\Sigma\Pi\Sigma$ -circuit C , denoted $\text{Sim}(C)$, is the $\Sigma\Pi\Sigma$ circuit obtained by dividing each term by the gcd of all the terms. i.e.,

$$\text{Sim}(C) \stackrel{\text{def}}{=} \sum_{i \in [k]} \frac{A_i(\mathbf{X})}{g(\mathbf{X})}, \quad \text{where } g(\mathbf{X}) = \gcd(A_1, \dots, A_k)$$

rank: Identifying each linear form $\ell = \sum_{i \in [k]} a_i \cdot X_i$ with the vector $(a_1, \dots, a_n) \in R^n$, we define the *rank* of C to be the dimension of the vector space spanned by the set $\{\ell_{ij} \mid i \in [k], j \in [d]\}$.

pairwise-rank: For a $\Sigma\Pi\Sigma(k)$ circuit $C = \sum_{i \in [k]} A_i$, we define the *pairwise-rank* of C to be $\min_{1 \leq i < j \leq k} \{\text{rank}(\text{Sim}(A_i + A_j))\}$, where $A_i + A_j$ is the subcircuit of C containing just the two multiplication gates A_i and A_j . If C has only one multiplication gate, we say the pairwise-rank of C is ∞ .

2.2 The Dvir–Shpilka conjecture and the main Structural Result

Very roughly, the conjecture of Dvir and Shpilka [DS05] asserts that for every k , there is a constant $c(k)$ such that if a depth three circuit C with top fanin k computes the zero polynomial then the rank of C is at most $c(k)$ (independent of the degree d of the intermediate polynomials computed at the different multiplication gates). As a step towards the conjecture, a $\text{poly}(2^{k^2} \cdot \log d)$ upper bound on the rank was obtained by Dvir and Shpilka [DS05]. This was subsequently improved by Saxena and Seshadri [SS08] to $(\text{poly}(k) \cdot \log d)$. Over *finite fields*, this conjecture was disproved by Kayal and Saxena [KS06] but the situation over fields of characteristic zero remained unclear. The conjecture soon revealed its fundamental nature - the weaker polylogarithmic upper bound was used by Karnin and Shpilka [KS08a] to give a quasipolynomial time deterministic blackbox identity test for $\Sigma\Pi\Sigma$ circuits with bounded top fanin. It was also used by Shpilka [Shp07] and by Karnin and Shpilka [KS08b] to give a quasipolynomial time algorithm for reconstruction of $\Sigma\Pi\Sigma$ circuits. In this paper, we prove the conjecture of Dvir and Shpilka over the field \mathbb{R} of real numbers, and therefore also over all subfields of \mathbb{R} such as \mathbb{Q} , the field of rational numbers. We then combine this result with ideas from [KS08a] to get an efficient algorithm for blackbox identity testing of $\Sigma\Pi\Sigma$ -circuits with bounded top-fanin.

Theorem 2.2 [Structure Theorem: Rank bound for $\Sigma\Pi\Sigma(k)$ circuits] *For every k , there exists a constant $c(k)$ (where $c(k) \leq 3^k((k+1)!)^2 = 2^{O(k \cdot \log k)}$) such that every $\Sigma\Pi\Sigma(k)$ circuit C with coefficients in \mathbb{R} that is simple, minimal, and computes the zero polynomial has $\text{rank}(C) \leq c(k)$.*

Remark.

1. Dvir and Shpilka conjectured that $c(k)$ is in fact a polynomially increasing function of k . We are able to only prove the weaker $\text{poly}(2^{k \cdot \log k})$ upper bound on $c(k)$. The best previously known bound was $(\text{poly}(k) \cdot \log d)$ [SS08] (note the dependence on d).
2. Our proof techniques also enable us to prove the structure theorem above (and hence blackbox identity testing) for the case $k = 3$ over complex numbers and over prime fields of very large characteristic. For more discussion about these results and why our proof does not go through for larger values of k over these fields, see Appendix A.

2.3 From the Rank Bound to Identity Testing

We give below the construction of Karnin and Shpilka [KS08a] which used ideas from an earlier work of Gabizon and Raz [GR05] to show how the rank bound of Theorem 2.2 translates into the blackbox identity testing algorithm of Theorem 2.1.

Lemma 2.3 [Translating rank bounds into a blackbox identity test.] *[KS08a] Let \mathbb{F} be a field and $R(k, d)$ be an integer such that every minimal and simple $\Sigma\Pi\Sigma(k, d, n)$ circuit over \mathbb{F} computing the zero polynomial has rank at most $R(k, d)$. For $\alpha \in \mathbb{F}$ let A_α be the $n \times R(k, d)$ matrix for which $(A_\alpha)_{i,j} = \alpha^{i(j+1)}$. Let $b_\alpha \stackrel{\text{def}}{=} (\alpha, \alpha^2, \dots, \alpha^n)$. Let S, T be subsets of \mathbb{F} such that $|S| = n \cdot \binom{kd}{2} + 2^k$ and $|T| = d + 1$. Let $P \subset \mathbb{F}^n$ be the following set of points.*

$$P \stackrel{\text{def}}{=} \left\{ A_\alpha \cdot \mathbf{x} + b_\alpha : \alpha \in S \text{ and } \mathbf{x} \in T^{R(k,d)} \right\}.$$

Then for every $\Sigma\Pi\Sigma(k, d, n)$ circuit C , C is identically zero if and only if $C(\mathbf{a}) = 0$ for all $\mathbf{a} \in P$.

This lemma can be applied to our situation as follows.

Proof of Theorem 2.1 We set $\mathbb{F} = \mathbb{Q}$ and using Theorem 2.2, we get that $R(k, d) = c(k) = 2^{O(k \cdot \log k)}$ is independent of d . We choose S to be $\{1, 2, \dots, m\}$ where

$$m = n \cdot \left(\binom{kd}{2} + 2^k \right) \cdot \binom{c(k) + 2}{2} + 1$$

and T to be $\{1, 2, \dots, d + 1\}$. We apply the above lemma to these choices of $\mathbb{F}, R(k, d), S$ and T . We thus get a set P which satisfies property (1). The number of points in P is $|S| \cdot |T|$ so that $|P| = \text{poly}(n) \cdot d^{2^{O(k \cdot \log k)}}$. Thus P satisfies property (2). Every coordinate of a point in P is the dot product of two vectors of length $R(k, d) = c(k)$ whose entries have bit-length $2^{O(k \cdot \log k)} \cdot O(n^2 \cdot \log d)$. This means that the bit-length of each point in P is $2^{O(k \cdot \log k)} \cdot O(n^3 \cdot \log d)$. This proves property (3). Clearly, this set P is very explicit - in fact it is so explicit that for every fixed k , P can be computed in the complexity class P-uniform TC^0 . This completes the proof of the theorem. □

The rest of this article is devoted to a proof of Theorem 2.2.

3 Organization

The rest of this paper is organized as follows. In Section 4, we give an overview of the techniques that we use in the proof of Theorem 2.2. In Section 5, we give the proof of Theorem 2.2 while deferring the proof of a key lemma used in the proof to Section 7. We give a sketch of the proof of this lemma and its connection to the Sylvester–Gallai Theorem and a related hyperplane decomposition lemma in Section 6. In Section 7 we prove the main technical (key) lemma, which we call the fanin reduction lemma, and in Section 8 we prove the hyperplane decomposition lemma. We conclude with a discussion of open problems in Section 9. Finally, in Appendix A, we discuss some of the conjectures formulated in the conclusion.

4 Overview of Proof of Rank Bound

In this section we give an overview of the proof of the structure theorem (Theorem 2.2). The proof proceeds by induction on the number of multiplication gates in the circuit. As induction hypothesis, we assume that any simple, minimal $\Sigma\Pi\Sigma$ circuit with fewer than k multiplication gates that is identically 0 cannot have high rank. Now if possible, let $C = \sum_{i=1}^k A_i = \sum_{i=1}^k \prod \ell_{ij}$ be a simple and minimal circuit in n variables that has high rank, and such that $C \equiv 0$. We will obtain a contradiction. For the sake of simplicity, we assume that each linear form ℓ_{ij} that appears in a gate of the circuit C , appears there with multiplicity one only. In Section 7, when we give the full argument, we remove this assumption.

Looking at the circuit modulo a linear form. We will be looking at the circuit modulo an appropriately chosen linear form. If $\ell = a_1 \cdot X_1 + \dots + a_n \cdot X_n$ is a linear form with $a_1 \neq 0$, then the image of a circuit C modulo ℓ is defined to be circuit obtained by replacing X_1 by $-\frac{1}{a_1} \cdot (a_2 \cdot X_2 + \dots + a_n \cdot X_n)$ in C . i.e.

$$C'(X_2, X_3, \dots, X_n) = C \pmod{\ell} \stackrel{\text{def}}{=} C\left(-\frac{1}{a_1} \cdot (a_2 \cdot X_2 + \dots + a_n \cdot X_n), X_2, \dots, X_n\right).$$

(see Section 7 for a more accurate definition that avoids the degenerate case when $a_1 = 0$.) Observe that if A_i is a $\Pi\Sigma$ polynomial of rank r and ℓ is a linear form, then either A_i equals zero modulo ℓ (i.e. ℓ divides A_i) or the rank of A_i drops by at most one to $r - 1$. Now if we pick a linear form ℓ which occurs in one of the constituent $\Pi\Sigma$ polynomials, say in A_1 , then A_1 equals zero modulo ℓ so that the resulting circuit $C' = C \pmod{\ell}$ would have at most $k - 1$ multiplication gates, each surviving gate having rank at most one less than what it had previously. Notice that if C computes the zero polynomial then so does the circuit C' . If this circuit C' was both simple and minimal we would be immediately done by the induction hypothesis.

However, in general it may not be possible to ensure C' is simple and minimal, and hence we use an intermediate notion, pairwise-rank, that very effectively captures and deals with the issues of simplicity and minimality. We first show that (1) any simple and minimal circuit computing the zero polynomial that has high rank must also have high pairwise rank. We then show that (2) no circuit with high pairwise rank can compute the zero polynomial.

Step (1) is the easier of the two steps. We show that if the circuit C has low pairwise-rank, then by setting some of the variables of the circuit to random values, we can obtain a new circuit that is still simple, minimal, has high rank, computes the zero polynomial, but has fewer multiplication gates (see Lemma 5.3). This contradicts the induction hypothesis.

Step (2) again uses the induction hypothesis. One of the key lemmas used here, which we refer to as the fanin reduction lemma, roughly asserts that if C is a simple circuit with high pairwise rank, then there exists a linear form ℓ in C such that if we go modulo ℓ , we get a circuit C' that still has high pairwise rank, but with fewer multiplication gates. Also, if $C \equiv 0$, then $C' \equiv 0$. From C' , we then show how to extract a subcircuit that is simple, minimal, computes the zero polynomial and has high rank. This will contradict the induction hypothesis. The bulk of the work goes into proving the fanin reduction lemma, Lemma 5.5. The vital ingredient in the proof of Lemma 5.5 is a theorem from incidence geometry called the high-dimensional Sylvester–Gallai Theorem. Before we state the Sylvester–Gallai Theorem, we first translate our problem into geometrical language.

4.1 A correspondence between $\Sigma\Pi\Sigma(k, d, n)$ circuits and k -colored points in \mathbb{R}^n

We identify the linear forms appearing in C with colored points in \mathbb{R}^n . A linear form $\ell = a_1 \cdot X_1 + \dots + a_n \cdot X_n$ corresponds to the point $P_\ell = (1, \frac{a_2}{a_1}, \dots, \frac{a_n}{a_1}) \in \mathbb{R}^n$ (see Section 7 for a more accurate definition of this correspondence which avoids the degenerate case when $a_1 = 0$). If the linear form $\ell \in A_i$ then we assign the color i to the point P_ℓ . Since a linear form could appear in multiple gates, in general a point could have many colors (see Section 7 for details). Our choice of the mapping of linear forms to points satisfies the property that 3 linear forms are linearly dependent iff they map

to collinear points. For two points $P \neq Q \in \mathbb{R}^n$ we will denote by $\lambda(P, Q)$ the line joining P and Q . For a point P and a color i we will denote by \mathcal{L}_i^P the pencil of lines $\{\lambda(P, Q) : Q \text{ has color } i\}$.⁶

Translating the search for a suitable linear form into the search for a suitable point.

Let the set of all points in the image of the set of linear forms in C be S . For a color $i \in [k]$, we will denote by S_i the set of points of color i . Now fix a linear form ℓ that occurs in C and consider two multiplication gates, say A_1 and A_2 occurring in C which do not contain ℓ . Consider the set $S_1 \Delta S_2$ which is the symmetric difference of the two sets of points of color 1 and 2 respectively (see Section 7 for a more accurate definition which takes care of degenerate cases when a linear form occurs in a gate with a higher multiplicity). Then the dimension of the space spanned by points in $S_1 \Delta S_2$ corresponds to the rank of the simplification of the circuit $A_1 + A_2$. Now consider the two pencils of lines $\mathcal{L}_1^{P_\ell}$ and $\mathcal{L}_2^{P_\ell}$. Notice that $\mathcal{L}_1^{P_\ell} \cap \mathcal{L}_2^{P_\ell}$ is again a pencil of lines through P_ℓ . Now $\gcd(A_1 \pmod{\ell}, A_2 \pmod{\ell})$ is nontrivial ($\neq 1$) if and only if there exists a line common to these two pencils, i.e. $\mathcal{L}_1^{P_\ell} \cap \mathcal{L}_2^{P_\ell} \neq \emptyset$. In fact the degree of $\gcd(A_1 \pmod{\ell}, A_2 \pmod{\ell})$ is exactly the number of lines in the pencil $\mathcal{L}_1^{P_\ell} \cap \mathcal{L}_2^{P_\ell}$. Now let us consider the symmetric difference $\mathcal{L}_1^{P_\ell} \Delta \mathcal{L}_2^{P_\ell}$ which is again another pencil of lines through P_ℓ . The requirement that C modulo ℓ has high pairwise rank, i.e. for all pairs of gates A_1, A_2 that do not contain ℓ , the simplification of $A_1 \pmod{\ell} + A_2 \pmod{\ell}$ should have high rank, then exactly translates into the requirement that the lines in this pencil $\mathcal{L}_1^{P_\ell} \Delta \mathcal{L}_2^{P_\ell}$ should span a high dimensional space.

Applying the Sylvester–Gallai Theorem. At this point it is not a priori clear as to why there should exist even a single line in the pencil $\mathcal{L}_1^{P_\ell} \Delta \mathcal{L}_2^{P_\ell}$. In fact if we fix the point P_ℓ then such an assertion is easily seen to be false. We show that there indeed exists a linear form ℓ , and the corresponding point P_ℓ such that for every pair of colors i and j (P_ℓ has color neither i nor j), the lines in the pencil $\mathcal{L}_i^{P_\ell} \Delta \mathcal{L}_j^{P_\ell}$ span a space of large enough dimension. The proof of this fact forms the main substance of the proof of our fanin reduction lemma, Lemma 5.5. In order to prove this result, we crucially use the Sylvester–Gallai Theorem, a result from incidence geometry. The basic Sylvester–Gallai Theorem roughly states that if S is a finite set of points in \mathbb{R}^n that are not all collinear, then there exists a line passing through exactly 2 points of S . This kind of statement is already in the spirit of what we want to show. We use a high dimensional version of the Sylvester–Gallai Theorem along with some colorful combinatorics to obtain our final result⁷.

5 The Rank Bound

In this section we the circuit rank bound, Theorem 2.2. The main technical result that we use is Lemma 5.5 (which we call the fanin reduction lemma). A sketch of the proof of the fanin reduction lemma is given in Section 6, and the full proof given in Section 7. We first state some lemmas and definitions related to circuit transformations that will be useful in the proof of Theorem 2.2. We then state the fanin reduction lemma (Lemma 5.5) and show how to combine the results on circuit transformations and the fanin reduction lemma to give a proof of Theorem 2.2.

⁶A pencil of lines is just a set of lines through a common point.

⁷Dvir and Shpilka [DS05] even observed that a certain colorful analog of the Sylvester–Gallai Theorem would imply the rank bound for the special case of $k = 3$. Such a result had in fact been proved much earlier by Edelman and Kelly [EK66]. Unfortunately, such a direct approach does not generalize for higher values of k . For more discussion about these results, see Appendix A.

Remark For the rest of this paper, all $\Sigma\Pi\Sigma$ circuits will have coefficients in \mathbb{R} .

5.1 Circuit Transformations

In this section we discuss some operations on circuits that will be useful in the proof of the rank bound (Theorem 2.2).

Lemma 5.1 [Invariance of circuit properties under invertible linear transformations of the variables.] *Let $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be an invertible linear transformation. Let $C = \sum_{i \in [k]} A_i = \sum_{i \in [k]} \prod_{j \in d} \ell_{ij}$ be a $\Sigma\Pi\Sigma$ circuit, and let $\pi(C)$ be the circuit $\sum_{i \in [k]} \pi(A_i) = \sum_{i \in [k]} \prod_{j \in d} \pi(\ell_{ij})$, where for a linear form $\ell = \mathbf{a} \cdot \mathbf{X}$, $\pi(\ell) = \pi(\mathbf{a}) \cdot \mathbf{X}$. Then, $\pi(C)$ is simple $\iff C$ is simple, $\pi(C)$ is minimal $\iff C$ is minimal, $\pi(C) \equiv 0 \iff C \equiv 0$, and $\text{rank}(\pi(C)) = \text{rank}(C)$.*

The proof of this lemma is immediate from definitions, and we omit it. We say that two circuits C and C' are equivalent, denoted by $C \sim C'$, if there exists an invertible linear transformation $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that $C = \pi(C')$.

Lemma 5.2 [Schwartz–Zippel Lemma] *Let $f(X_1, X_2, \dots, X_n)$ be a nonzero n variate polynomial of degree d over \mathbb{R} . Then for (a_1, a_2, \dots, a_n) chosen uniformly at random in $[0, 1]^n$, the probability that $f(a_1, a_2, \dots, a_n) = 0$ is zero.*

Lemma 5.3 [Setting linear forms to random values.] *Let $C \equiv 0$ be a simple and minimal $\Sigma\Pi\Sigma$ circuit in the n indeterminates X_1, X_2, \dots, X_n . Let $\text{rank}(C) = r$. Let $t \in [n]$, and let $\alpha_1, \alpha_2, \dots, \alpha_t$ be real numbers picked independently and uniformly from $[0, 1]$. Let Z be an indeterminate, and consider the new circuit C' formed by replacing X_i by $\alpha_i Z$ for all $i \in [t]$. Then with probability 1, C' is minimal and $\text{rank}(\text{Sim}(C')) \geq r - t$.*

Proof Let $C = \sum_{i \in [k]} A_i = \sum_{i \in [k]} \prod_{j \in d_j} \ell_{ij}$. and after replacing X_i by $\alpha_i Z$ for all $i \in [t]$, let $C' = \sum_{i \in [k]} A'_i = \sum_{i \in [k]} \prod_{j \in d_j} \ell'_{ij}$.

With probability 1 C' is minimal: We will show that with probability 1, for all $S \subseteq [k]$, $\sum_{i \in S} A'_i \not\equiv 0$. Let $S \subset [k]$, and let $P_S = P_S(X_1, X_2, \dots, X_n)$ be the polynomial computed by $\sum_{i \in S} A_i$. Since C is minimal, P_S is nonzero. Let the polynomial computed by $\sum_{i \in S} A'_i$ be $P'_S = P_S(\alpha_1 Z, \alpha_2 Z, \dots, \alpha_t Z, X_{t+1}, \dots, X_n)$. If $P'_S \equiv 0$, then this must even hold for the setting of $Z = 1$, i.e. $P_S(\alpha_1, \alpha_2, \dots, \alpha_t, X_{t+1}, \dots, X_n) \equiv 0$. Hence this must also be true for any settings of the variables X_{t+1}, \dots, X_n to values in $[0, 1]$, and in particular to random values. By the Schwartz–Zippel Lemma (Lemma 5.2), this can happen only with probability 0. Hence, for any $S \subset [k]$, we get that with probability 1, $\sum_{i \in S} A'_i \not\equiv 0$. Hence, by the union bound, with probability 1, for all $S \subset [k]$, $\sum_{i \in S} A'_i \not\equiv 0$, i.e. C' is minimal.

With probability 1 $\text{rank}(\text{Sim}(C')) \geq r - t$: Observe that the Schwartz–Zippel Lemma (Lemma 5.2), implies that with probability 1, no linear form appearing in C gets mapped to zero under the random substitutions. Also, with probability 1, just changing X_1 to $\alpha_1 Z$ will not change the rank of the circuit at all. With probability 1, the definition of rank implies that the remaining $t - 1$ substitutions could have made the rank of C drop by at most $t - 1$. Hence with probability 1 we get that $\text{rank}(C') \geq \text{rank}(C) - (t - 1) = r - t + 1$. Also, since⁸ $\text{rank}(C') \leq \text{rank}(\text{Sim}(C')) + \text{rank}(\text{gcd}(C'))$,

⁸where for $\text{gcd}(C') = \prod \ell_i$, we have $\text{rank}(\text{gcd}(C')) = \dim(\text{span}\{\ell : \ell \mid \text{gcd}(C')\})$

it suffices to show that with probability 1, $\text{rank}(\text{gcd}(C')) \leq 1$. In fact, we will show that with probability 1, for all linear forms ℓ such that $\ell \mid \text{gcd}(C')$ (i.e. ℓ divides $\text{gcd}(C')$), it also holds that $\ell \mid Z$, and this will imply the result.

Let $\ell_a = a_1X_1 + a_2X_2 + \dots + a_nX_n$ and $\ell_b = b_1X_1 + b_2X_2 + \dots + b_nX_n$ be any two nonzero linear forms in C such that $\text{gcd}(\ell_a, \ell_b) = 1$. Hence, for all $\gamma \in \mathbb{R}$, $\ell_a - \gamma\ell_b \neq 0$. Let $\ell'_a = a_1\alpha_1Z + a_2\alpha_2Z + \dots + a_t\alpha_tZ + a_{t+1}X_{t+1} \dots + a_nX_n$, and $\ell'_b = b_1\alpha_1Z + b_2\alpha_2Z + \dots + b_t\alpha_tZ + b_{t+1}X_{t+1} \dots + b_nX_n$. We will show that with probability 1, $\text{gcd}(\ell'_a, \ell'_b) \mid Z$, and the result will follow easily from this.

Let $\ell_a^* = a_{t+1}X_{t+1} \dots + a_nX_n$, and let $\ell_b^* = b_{t+1}X_{t+1} \dots + b_nX_n$.

- **Case 1:** Suppose that at least one of ℓ_a^*, ℓ_b^* is 0. Hence at least one of ℓ'_a, ℓ'_b will divide Z . The Schwartz–Zippel Lemma (Lemma 5.2) implies that with probability 1, both ℓ'_a and ℓ'_b are nonzero and hence $\text{gcd}(\ell'_a, \ell'_b) \mid Z$.
- **Case 2:** Suppose that both $\ell_a^*, \ell_b^* \neq 0$, and $\text{gcd}(\ell_a^*, \ell_b^*) = 1$. In this case clearly $\text{gcd}(\ell'_a, \ell'_b) = 1$, and hence $\text{gcd}(\ell'_a, \ell'_b) \mid Z$.
- **Case 3:** Suppose that both $\ell_a^*, \ell_b^* \neq 0$, and there exists nonzero $\gamma^* \in \mathbb{R}$ such that $\ell_a^* = \gamma^*\ell_b^*$. In this case, if $\text{gcd}(\ell'_a, \ell'_b) \nmid Z$, it must be that $\ell'_a = \gamma^*\ell'_b$. Also, we know that $\ell_a - \gamma^*\ell_b \neq 0$. If $\ell'_a - \gamma^*\ell'_b \equiv 0$, then this must even hold for the setting of $Z = 1$, i.e. $a_1\alpha_1 + a_2\alpha_2 + \dots + a_t\alpha_t + a_{t+1}X_{t+1} + \dots + a_nX_n - \gamma^*(b_1\alpha_1 + b_2\alpha_2 + \dots + b_t\alpha_t + b_{t+1}X_{t+1} + \dots + b_nX_n) \equiv 0$. Hence this must also be true for any settings of the variables X_{t+1}, \dots, X_n to values in $[0, 1]$, and in particular to random values. By the Schwartz–Zippel Lemma (Lemma 5.2), this can happen only with probability 0.

Hence in all three cases we get that with probability 1, $\text{gcd}(\ell'_a, \ell'_b) \mid Z$. Since any linear form in $\text{gcd}(C')$ must be of the form $\text{gcd}(\ell'_a, \ell'_b)$, where the original two linear forms ℓ_a, ℓ_b were initially such that $\text{gcd}(\ell_a, \ell_b) = 1$, the above case analysis implies that after taking the union bound over all such pair of linear forms appearing in C , with probability 1, $\text{gcd}(\ell'_a, \ell'_b) \mid Z$. This completes the proof.

■

Definition 5.4 [Setting a linear form to 0: $C|_{\ell=0}$] Let $C = \sum_{i \in [k]} A_i = \sum_{i \in [k]} \prod_{j \in d} \ell_{ij}$ be a $\Sigma\Pi\Sigma$ circuit. Let ℓ be a linear form appearing in C . Let $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is any linear map of rank $n - 1$ such that $\text{kernel}(\pi) = \text{span}(\ell)$ ⁹. We let $C|_{\ell=0}$ denote the class of circuits obtained by applying such a transformation π to C , to get a circuit $\pi(C)$, where $\pi(C)$ is the circuit $\sum_{i \in [k]} \prod_{j \in d} \pi(\ell_{ij})$, where for a linear form $\ell = \mathbf{a} \cdot \mathbf{X}$, $\pi(\ell) = \pi(\mathbf{a}) \cdot \mathbf{X}$. Under such a transformation, all the constituent $\Pi\Sigma$ polynomials that contain ℓ get set to 0, and we remove all such gates from the circuit.

It is easy to see that if C_1 and C_2 are two circuits in $C|_{\ell=0}$, then $C_1 \sim C_2$, and we omit the proof. We abuse notation by using $C' = C|_{\ell=0}$ to refer to any circuit C' in the class $C|_{\ell=0}$. Note that if $C \equiv 0$, then for any circuit C' in the class $C|_{\ell=0}$, we have $C' \equiv 0$.

⁹Where for $\ell = \mathbf{a} \cdot \mathbf{X}$, $\text{span}(\ell)$ denotes the one dimensional vector space spanned by the vector \mathbf{a}

5.2 The Fanin Reduction Lemma

Lemma 5.5 is the main technical result of our paper that allows us to apply an induction argument to reason about the rank of $\Sigma\Pi\Sigma$ circuits computing the zero polynomial. We show that if a simple circuit C has high pairwise-rank, then by “setting a linear form to 0”, we can transform it to a new circuit C' with fewer multiplication gates that still has high pairwise-rank. Also, if $C \equiv 0$, then $C' \equiv 0$.

Lemma 5.5 [Fanin Reduction Lemma] *Let $k, A > 0$ be integers. Let $B = 3(A + 1)k^2$. Let C be a simple $\Sigma\Pi\Sigma(k)$ circuit such that $\text{pairwise-rank}(C) \geq A$, and $\text{rank}(C) \geq B$. Then there exists a linear form ℓ in the circuit C such that for $C' = C|_{\ell=0}$, $\text{pairwise-rank}(C') \geq A$.*

A sketch of the proof of this lemma is given in Section 6 and full proof given in Section 7. With these tools in hand, we are now ready to prove the main theorem of this paper.

5.3 Proof of Theorem 2.2: The Rank Bound

Theorem 2.2 [Rank bound for $\Sigma\Pi\Sigma(k)$ circuits]: *Let $c(k) = 3^k((k + 1)!)^2$. Let C be a simple and minimal $\Sigma\Pi\Sigma(k)$ circuit that computes the zero polynomial. Then $\text{rank}(C) \leq c(k)$.*

The proof proceeds by an induction on k , the number of multiplication gates in C . We first show that C must have high pairwise-rank. If it does not have high pairwise-rank, then we can use Lemma 5.3 to obtain a new circuit that is still simple and minimal and has high rank, but with fewer multiplication gates. This would contradict the induction hypothesis. We then use Lemma 5.5 to find a linear form ℓ such that the circuit $C' = C|_{\ell=0}$ also has high pairwise-rank, and is such that $C' \equiv 0$, but has fewer multiplication gates. Any minimal subset of the multiplication gates of C' that sums to 0 will give a circuit C_{\min} that also has high pairwise-rank, is minimal, and still computes the zero polynomial. The simplification of C_{\min} will then be simple, minimal, have high rank, and will have fewer than k multiplication gates, contradicting the induction hypothesis.

Proof We will prove the above theorem by induction on k .

For $k = 1, 2$, the result is vacuously true. Let $k \geq 3$ and assume the theorem is true for $\Sigma\Pi\Sigma(m)$ circuits for all $m \leq k - 1$.

If possible let

$$C = \sum_{i=1}^k A_i = \sum_{i=1}^k \prod_{j=1}^d \ell_{ij}$$

be simple and minimal such that $C \equiv 0$, and $\text{rank}(C) > c(k)$. Let the indeterminates appearing in C be X_1, X_2, \dots, X_n .

Case 1: $\text{pairwise-rank}(C) < c(k) - c(k - 1)$. Hence there exist $i, j \in [k]$ such that $i \neq j$ and $\text{rank}(\text{Sim}(A_i + A_j)) < c(k) - c(k - 1)$. Let $A_i + A_j = \text{gcd}(A_i, A_j) \cdot \text{Sim}(A_i + A_j)$. Without loss of generality, by Lemma 5.1 (equivalence up to linear transformations), let $\text{span}((\text{Sim}(A_i + A_j)))^{10}$ be spanned by X_1, X_2, \dots, X_t , where $t = \text{rank}(\text{Sim}(A_i + A_j))$. For each $i \in [t]$, let α_i be a uniformly

¹⁰For $C = \sum_{i \in [k]} \prod_{j \in [d]} \ell_{ij}$, we let $\text{span}(C) = \text{span}(\{\ell_{ij} \mid i \in [k], j \in [d]\})$

random real number in $[0, 1]$. For $i \in [t]$, set $X_i = \alpha_i Z$. By Lemma 5.3, with probability 1 we get a (homogeneous) circuit C' such that $C' \equiv 0$, it has at most $k - 1$ gates (since after the random substitution, both A_i and A_j will have the same set of linear forms up to scalar multiples, and they can be merged into a single gate), is still minimal, and its ged has rank at most 1. Also, since $t < c(k) - c(k - 1)$ we get that $\text{rank}(\text{Sim}(C'))$ is strictly greater than $c(k - 1)$. Hence $\text{Sim}(C')$ is simple, minimal, computes the zero polynomial, has at most $k - 1$ multiplication gates, and has rank strictly greater than $c(k - 1)$, contradicting the induction hypothesis.

Case 2: $\text{pairwise-rank}(C) \geq c(k) - c(k - 1)$. Hence for all $i, j \in [k]$ such that $i \neq j$, $\text{rank}(\text{Sim}(A_i + A_j)) \geq c(k) - c(k - 1) > c(k)/2 > c(k - 1) + 1$. Notice that by choice of the function c , $c(k) \geq 3k^2((c(k - 1) + 1) + 1)$. By Lemma 5.5 (With $A = c(k - 1) + 1$ and $B = c(k)$), there exists a linear form ℓ in C such that for $C' = C|_{\ell=0}$, $\text{pairwise-rank}(C') \geq c(k - 1) + 1$.

Now, since the gates containing ℓ got set to 0, the number of gates in C' is at most $k - 1$. Also, $\text{pairwise-rank}(C') \geq c(k - 1) + 1$ implies that for all subsets $S \subseteq [k]$ such that S indexes at least two nonzero gates of C , $\text{rank}(\text{Sim}(C'|_S)) \geq c(k - 1) + 1$ (where $C'|_S$ is the subcircuit of C' obtained by restricting to only those multiplication gates of C' that are indexed by S). We know that $\sum_{i \in [k]} A'_i = 0$ (where at least one of the A'_i is set to 0). Now take the smallest nonempty such set S for which $\sum_{i \in S} A'_i = 0$. Then, $\sum_{i \in S} A'_i$ is a minimal circuit such that its simple part has rank at least $c(k - 1) + 1$. This contradicts the induction hypothesis.

Thus we conclude that $\text{rank}(C) \leq c(k)$.

■

6 The Sylvester–Gallai Theorem and the Fanin Reduction Lemma

In this section we will sketch a proof of the fanin reduction lemma (Lemma 5.5) and highlight the main ingredients in the proof. The full proof is given in Section 7. Our proof of the fanin reduction lemma first translates the problem from a question about circuits to a question purely about the incidence properties of colored points in \mathbb{R}^n . The main tools that we use in analyzing the points is the Sylvester–Gallai Theorem, and a related hyperplane decomposition theorem. Before we state these results, we first introduce some terminology that we will use.

Affine spaces and hyperplanes. We say that $H \subseteq \mathbb{R}^n$ is an *affine space* if it is a translation of a linear space. In other words, there exists a linear vector space $H' \subseteq \mathbb{R}^n$ and a vector $v \in \mathbb{R}^n$ such that $H = v + H' = \{v + u \mid u \in H'\}$. The dimension $\dim(H)$ of the affine space is the dimension of the corresponding linear space $\dim(H')$. We will be using the term *hyperplane* interchangeably with affine space. In this terminology, a point is a hyperplane of dimension 0, a line is a hyperplane of dimension 1 etc. For a set $S \subseteq \mathbb{R}^n$ of points, the *affine span* of S , denoted $\text{affine-span}(S)$, is the intersection of all the affine spaces containing S . Note that the affine span of a set is also an affine space. Also note the difference between this notion of affine-span and the notion of vector space span ¹¹.

The Sylvester–Gallai Theorem (see the survey by Borwein and Moser [BM90] for details) asserts the following:

¹¹Informally, as sets, the vector space span of a set of points/vectors S would equal the affine span of $S \cup \{\mathbf{0}\}$, where $\{\mathbf{0}\}$ denotes the origin (or zero vector).

Theorem 6.1 [Sylvester–Gallai] *Let S be a finite set of points spanning an affine space $V \subseteq \mathbb{R}^n$ such that $\dim(V) \geq 2$. Then there exists a line $L \subseteq V$ such that $|L \cap S| = 2$.*

We state below the high dimensional Sylvester–Gallai Theorem. It was first proved in a slightly different form by Hansen [Han65]. The version below is a slightly refined version of Hansen’s result, and was obtained by Bonnice and Edelstein [BE67, Theorem 2.1].

Theorem 6.2 [Generalized Sylvester–Gallai for high dimensions] *([Han65], [BE67]) Let S be a finite set of points spanning an affine space $V \subseteq \mathbb{R}^n$ such that $\dim(V) \geq 2t$. Then, there exists a t dimensional hyperplane H such that $|H \cap S| = t + 1$, and such that H is spanned by the points of S . i.e $\text{affine-span}(H \cap S) = H$.*

Using the above result, we obtain the following ‘decomposition’ theorem. A similar decomposition procedure was carried out by Edelstein and Kelly [EK66] (to obtain a Sylvester–Gallai kind of theorem for colored points), and by Bonnice and Edelstein [BE67]. We defer the proof of the hyperplane decomposition lemma (Lemma 6.3) to Section 8.

Lemma 6.3 [Hyperplane decomposition] *Let V be an m dimensional affine space over \mathbb{R} . Let $S \subset V$ be a finite set, such that $\text{affine-span}(S) = V$. Let $S_{\text{core}} \subseteq S$, and let $H_{\text{core}} = \text{affine-span}(S_{\text{core}})$ be an affine space of dimension m_{core} . Then for some $r \geq \frac{m - m_{\text{core}}}{2}$, there exist hyperplanes $H_1, H_2, \dots, H_r \subset V$, such that letting $H = \text{affine-span}(\{H_i \mid i \in [r]\})$, we have the following properties.*

1. For all $i \in [r]$, $H_{\text{core}} \subseteq H_i$ and $\dim(H_i) = m_{\text{core}} + 1$.
2. $\dim(H) = m_{\text{core}} + r$. In particular, if $R \subseteq [r]$ is such that for each $i \in R$, P_i is a point in $H_i \setminus H_{\text{core}}$, then $\dim(\text{affine-span}(\{P_i \mid i \in R\})) = |R| - 1$.
3. For all $i \in [r]$, $(H_i \setminus H_{\text{core}}) \cap S \neq \phi$.
4. For every point $P \in S \cap H$, there exists $i \in [r]$ such that $P \in H_i$. Note that it is not necessary that every point in S lies on one of the H_i ’s but every point of S inside H certainly does lie on at least one of the H_i ’s.

We present below a very informal outline of the proof of the fanin reduction lemma (Lemma 5.5) to demonstrate how the hyperplane decomposition lemma is used in its proof. For the full details of the proof, see Section 7.

Outline of proof of the fanin reduction lemma: Lemma 5.5: In Section 4.1 we saw how to map the linear forms appearing in the circuit C to colored points in \mathbb{R}^n . We will assume the terminology used in Section 4.1. Recall that we want to show that there exists a linear form ℓ in C and the corresponding point P_ℓ such that for all pairs of colors i and j such that ℓ does not occur in A_i and A_j , the set of lines in the pencil $\mathcal{L}_i^{P_\ell} \Delta \mathcal{L}_j^{P_\ell}$ span a high dimensional space. We will use the hyperplane decomposition lemma to accomplish this. Let the set of points S corresponding to linear forms in C span the affine space V . We choose a (relatively low dimensional) subspace $H_{\text{core}} \subseteq V$ such that for every pair of colors i and j , the symmetric difference of the points

of those colors, $S_i \Delta S_j$, contained within H_{core} spans a high dimensional subspace. We apply the hyperplane decomposition theorem to V and H_{core} to get a large collection of hyperplanes H_1, H_2, \dots, H_r each containing H_{core} and satisfying the properties listed in Lemma 6.3. Let $H = \text{affine-span}(H_1, H_2, \dots, H_r)$. Observe that property (2) implies that if P_i and P_j are two points of S in $H_i \setminus H_{\text{core}}$ and $H_j \setminus H_{\text{core}}$ respectively, then the line through them does not contain any other point of S . This will be a very useful property. For a pair of colors i, j , let $(S_i \Delta S_j)^H$ denote the set of points in $S_i \Delta S_j$ that lie in $H \setminus H_{\text{core}}$. We say a pair of colors (i, j) is over-split if a large subset of the hyperplanes $\{H_i\}$ contain an element of $(S_i \Delta S_j)^H$. Otherwise we say the pair is under-split. Since for each pair of under-split colors (i, j) the set $(S_i \Delta S_j)^H$ occurs in few hyperplanes, and since the total number of hyperplanes is large, the pigeon-hole principle implies that there exists a hyperplane H^* that does not contain any member of $(S_i \Delta S_j)^H$ for any under-split pair of colors $\{i, j\}$. By property (3), there exists a point P_ℓ contained in $H^* \setminus H_{\text{core}}$. Let the corresponding linear form be ℓ . We will show that for all pairs of colors i and j such that ℓ does not occur in A_i and A_j , the set of lines in the pencil $\mathcal{L}_i^{P_\ell} \Delta \mathcal{L}_j^{P_\ell}$ span a high dimensional space. From now on we will only mention pairs of colors corresponding to multiplication gates in the circuit that do not contain ℓ . Now for any pair of colors $\{i, j\}$, since a line through P_ℓ and any point in $(S_i \Delta S_j)^H \setminus H^*$ does not contain any other point of S (by property (3)), the set of such lines is contained in the pencil $\mathcal{L}_i^{P_\ell} \Delta \mathcal{L}_j^{P_\ell}$. If the pair of colors is over-split, then this pencil will span a high dimensional space, and hence this pair of colors will not create any worry. If the pair of colors (i, j) is under-split, then recall that $H^* \setminus H_{\text{core}}$ does not contain any element of $(S_i \Delta S_j)^H$. Now consider the intersection of $(S_i \Delta S_j)$ with H_{core} and call it $(S_i \Delta S_j)^{\text{core}}$. Recall that by the choice of H_{core} , $(S_i \Delta S_j)^{\text{core}}$ spans a high dimensional space. Also, any line through P_ℓ and a point of $(S_i \Delta S_j)^{\text{core}}$ lies entirely within H^* and does not contain any other point of H_{core} . Since there are no points of $S_i \Delta S_j$ in $H^* \setminus H_{\text{core}}$, hence any such line is a line in the pencil $\mathcal{L}_i^{P_\ell} \Delta \mathcal{L}_j^{P_\ell}$. Since $(S_i \Delta S_j)^{\text{core}}$ spans a high dimensional space, so does the pencil $\mathcal{L}_i^{P_\ell} \Delta \mathcal{L}_j^{P_\ell}$. Hence under-split pairs of colors do not create a problem either, and we are done.

7 Proof of the Fanin Reduction Lemma

In this section we give a full proof of the fanin reduction lemma. Recall the fanin reduction lemma from Section 5.

Lemma 5.5[Fanin Reduction Lemma] *Let $k, A > 0$ be integers. Let $B = 3(A + 1)k^2$. Let C be a simple $\Sigma\Pi\Sigma(k)$ circuit such that $\text{pairwise-rank}(C) \geq A$, and $\text{rank}(C) \geq B$. Then there exists a linear form ℓ in the circuit C such that for $C' = C|_{\ell=0}$, $\text{pairwise-rank}(C') \geq A$.*

Proof Let $C = \sum_{i \in [k]} A_i = \sum_{i \in [k]} \prod_{j \in [d]} \ell_{ij}$. Let the indeterminates appearing in C be (X_1, X_2, \dots, X_n) . Let $L_C = \{\ell_{ij} \mid i \in [k], j \in [d]\}$ be the multiset of linear forms in C , and for each $i \in [k]$, let $L_{A_i} = \{\ell_{ij} \mid j \in [d]\}$ be the multiset of linear forms in A_i . Hence $|L_{A_i}| = d$.

From linear forms to points in \mathbb{R}^n . For a linear form $\ell = \sum_{i \in [n]} a_i \cdot X_i = \mathbf{a} \cdot \mathbf{X}$, let $S(\ell) = \{\lambda \mathbf{a} \mid \lambda \in \mathbb{R}\}$ be the line in \mathbb{R}^n corresponding to the span of ℓ . Let $V \subseteq \mathbb{R}^n$ be an $n - 1$ dimension hyperplane such that (i) V does not contain the origin $\mathbf{0}$, and (ii) for all linear forms $\ell \in L_C$, $V \cap S(\ell) \neq \emptyset$. (A random hyperplane satisfies these conditions with probability 1.)

We map each linear form in L_C to a point in $V \subseteq \mathbb{R}^n$ by the following map $g : L_C \rightarrow \mathbb{R}^n$, where $g(\mathbf{a} \cdot \mathbf{X}) = \lambda_{\mathbf{a}} \mathbf{a}$, where $\lambda_{\mathbf{a}}$ is the unique nonzero constant such that $\lambda_{\mathbf{a}} \mathbf{a} \in V$. In other words, g

maps the linear form $\ell = \sum_{i \in [n]} a_i \cdot X_i$ to the unique point of intersection of the hyperplane V with the line $S(\ell) = \{\lambda \mathbf{a} \mid \lambda \in \mathbb{R}\}$ ¹². Observe that in doing this, two linear forms get mapped to the same point if and only if they are scalar multiples of each other. Let the set of points in the image of g be $S = \{P_1, P_2, \dots, P_N\}$.

Coloring the points. We assign a *multiset of colors* $K(P)$ to each point P in the image of g . For $i \in [k]$, we let the gate A_i represent the color i . Each linear form ℓ_{ij} in A_i is given the color i , and if ℓ_{ij} maps to the point P under g , then i will be one of the colors in the multiset of colors $K(P)$ assigned to P . In particular, letting $\text{mult}(P, i)$ denote the multiplicity of color i in the set $K(P)$, we have that for all $i \in [k]$, $\text{mult}(P, i) = |g^{-1}(P) \cap L_{A_i}|$. Let C_P be the *set* of colors assigned to the point P (i.e. the set of colors in the multiset $K(P)$, but taken without multiplicity).

Differentiating between colors. We say that a point $P \in S$ differentiates between the colors i and j , if $\text{mult}(P, i) \neq \text{mult}(P, j)$. For a color i , we let S_i denote the multiset of all points of color i , each point $P \in S$ having multiplicity $\text{mult}(P, i)$. For our purposes, we extend the notion of symmetric difference of two sets to a symmetric difference of two *multisets* as follows. For $i, j \in [k]$, we define $S_i \Delta S_j$ to be the *set* $\{P \in S \mid \text{mult}(P, i) \neq \text{mult}(P, j)\}$

Observe that for a linear form $\ell = \mathbf{a} \cdot \mathbf{X}$, the function g maps it to a point of the form $\lambda \mathbf{a}$. Hence by definition, linear forms that are linearly independent (over the vector space \mathbb{R}^n) get mapped to points that are independent¹³. In particular, for linear forms $\ell_1, \dots, \ell_r \in L_C$, $\dim(\text{span}(\{\ell_i \mid i \in [r]\})) = d \iff \dim(\text{affine-span}(\{g(\ell_i) \mid i \in [r]\})) = d - 1$. Thus the condition “ $\text{rank}(C) \geq B$ ” exactly translates to the condition $\dim(\text{affine-span}(S)) \geq B - 1$. Similarly the condition “ $\text{pairwise-rank}(C) \geq A$ ” translates to the condition that for all $i, j \in [k]$ such that $i \neq j$, $\dim(\text{affine-span}(S_i \Delta S_j)) \geq A - 1$.

Choosing a core hyperplane. We now describe how to pick a set $S_{\text{core}} \subseteq S$ satisfying the following conditions (i) $\dim(\text{affine-span}(S_{\text{core}})) \leq k^2 A$ and (ii) for all $i, j \in [k]$ such that $i \neq j$, $\dim(\text{affine-span}(S_{\text{core}} \cap (S_i \Delta S_j))) \geq A - 1$.

For $i, j \in [k]$ such that $i \neq j$, since $\dim(\text{affine-span}(S_i \Delta S_j)) \geq A - 1$, there exists a set $S_{ij} \subseteq (S_i \Delta S_j)$ such that $|S_{ij}| = A$, and $\dim(\text{affine-span}(S_{ij})) = A - 1$. Pick such a set S_{ij} for every $i, j \in [k]$ such that $i \neq j$, and let $S_{\text{core}} = \cup_{i, j \in [k], i \neq j} S_{ij}$. Then $|S_{\text{core}}| \leq \binom{k}{2} A$, and it satisfies the required conditions. Let $H_{\text{core}} = \text{affine-span}(S_{\text{core}})$, and let $\dim(H_{\text{core}}) = m_{\text{core}}$.

Applying the Hyperplane decomposition lemma. By Lemma 6.3 applied the sets S and S_{core} , we get that for some $r \geq \frac{m - m_{\text{core}}}{2}$, there exist hyperplanes $H_1, H_2, \dots, H_r \subset V$, such that letting $H = \text{affine-span}(\{H_i \mid i \in [r]\})$, we have the following properties.

1. For all $i \in [r]$, $H_{\text{core}} \subseteq H_i$ and $\dim(H_i) = m_{\text{core}} + 1$.
2. $\dim(H) = m_{\text{core}} + r$. In particular, if $R \subseteq [r]$ is such that for each $i \in R$, \tilde{P}_i is a point in $H_i \setminus H_{\text{core}}$, then $\dim(\text{affine-span}(\{\tilde{P}_i \mid i \in R\})) = |R| - 1$.

¹²Alternatively we could view the line $S(\ell)$ as a point in the real projective space \mathbb{RP}^{n-1} , and we could map the linear forms in L_C to the corresponding points in \mathbb{RP}^{n-1} . The rest of the discussion would follow almost identically, with the correct analog of Lemma 6.3 for \mathbb{RP}^{n-1} . In some ways it might be more natural to work in projective space, since the projective embedding precisely captures the equivalence of linear forms that differ only by a nonzero scalar factor. By taking the projection of the lines onto the hyperplane V , we essentially get the same features as those of the projective embedding.

¹³We say that r points in \mathbb{R}^n are independent if they span an $r - 1$ dimensional affine space.

3. For all $i \in [r]$, $(H_i \setminus H_{\text{core}}) \cap S \neq \phi$.
4. For every point $P \in S \cap H$, there exists $i \in [r]$ such that $P \in H_i$. Note that it is not necessary that every point in S lies on one of the H_i 's but every point of S inside H certainly does lie on at least one of the H_i 's.

Observe that property (2) implies that if \tilde{P}_i and \tilde{P}_j are two points of S in $H_i \setminus H_{\text{core}}$ and $H_j \setminus H_{\text{core}}$ respectively, then the line through them does not contain any other point of S . This will be a very useful property later on. For a pair of colors $i, j \in [k]$, let $(S_i \Delta S_j)^H$ denote the set of points in $S_i \Delta S_j$ that lie in $H \setminus H_{\text{core}}$. Also, let $(S_i \Delta S_j)^{\text{core}}$ denote the set of points in $S_i \Delta S_j$ that lie in H_{core} .

Over-split and under-split pairs of colors. We say that a hyperplane $\tilde{H} \in \{H_i \mid i \in [r]\}$ *splits* a pair of colors $\{i, j\}$ if there exists a point $Q \in \tilde{H} \cap (S_i \Delta S_j)^H$. We say that a pair of colors $\{i, j\}$ is *over-split* if there exist at least $A + 1$ distinct hyperplanes contained in $\{H_i \mid i \in [r]\}$ that each splits $\{i, j\}$. Else we say that $\{i, j\}$ is *under-split*.

By the hyperplane decomposition theorem, $r \geq \frac{B - k^2 A}{2} \geq (A + 1)k^2$. Since each under-split pair of colors is split by at most A hyperplanes, and there are at most k^2 pairs of under-split colors, by the pigeon-hole principle there must exist a hyperplane $H^* \in \{H_i \mid i \in [r]\}$ such that for every under-split pair of colors $\{i, j\}$, H^* does not contain any member of $(S_i \Delta S_j)^H$. In other words, for all under-split pairs of colors $\{i, j\}$, and for all points $Q \in H^* \setminus H_{\text{core}}$, $\text{mult}(Q, i) = \text{mult}(Q, j)$.

Setting a linear form to zero. By property 3 in Lemma 6.3, there exists a point $P_\ell \in (H^* \setminus H_{\text{core}}) \cap S$. Let $\ell = g^{-1}(P_\ell)$ ¹⁴ be the linear form corresponding to P_ℓ . We will set ℓ to zero, and show that for $C' = C|_{\ell=0}$, we have $\text{pairwise-rank}(C') \geq A$. Note that since the circuit C is simple, the set of colors C_{P_ℓ} assigned to P_ℓ cannot have all k colors (since otherwise the linear form corresponding to the point P_ℓ would be present in all gates of C), i.e. $|C_{P_\ell}| < k$. Also, if $|C_{P_\ell}| = k - 1$, then the linear form ℓ would be present in all but one of the gates of C , and hence $C|_{\ell=0}$ would have only one multiplication gate. In this case the theorem is vacuously true. From now on, we assume $|C_{P_\ell}| \leq k - 2$.

Translating the problem to colored points. To show that for $C' = C|_{\ell=0}$ that $\text{pairwise-rank}(C') \geq A$, we need to show that for all nonzero gates A'_i and $A'_j \in C'$, $\text{rank}(\text{Sim}(A'_i + A'_j)) \geq A$. Let $\ell = \mathbf{a} \cdot \mathbf{X}$, and let $\pi : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a linear map such that π has nullity 1, and $\text{kernel}(\pi) = \text{span}(\mathbf{a})$. π is simply a linear transformation that maps points in \mathbb{R}^n to the subspace on which ℓ vanishes. Hence $P_\ell \equiv \lambda_a \mathbf{a}$ for some nonzero $\lambda_a \in \mathbb{R}$, and $\pi(P_\ell) = \mathbf{0}$.

For all $i \in C_{P_\ell}$ (i.e. the color i occurs in the set of colors of P_ℓ), we get that ℓ occurs in the gate A_i . Since we 'set' $\ell = 0$, hence $A'_i \equiv 0$, and A'_i is eliminated from C' . For all other colors, the corresponding multiplication gates stay nonzero, and the linear forms get transformed by π . Recall that for a linear form $\tilde{\ell} = \mathbf{b} \cdot \mathbf{X}$, there exists nonzero $\lambda_b \in \mathbb{R}$ such that $g(\tilde{\ell}) = \lambda_b \mathbf{b}$. Hence $g(\tilde{\ell})$ recovers $\tilde{\ell}$ up to multiplication by a scalar. Also observe that if we apply π to $g(\tilde{\ell})$, we would recover $\pi(\tilde{\ell})$ up to multiplication by a scalar. Recall that the multisets of colors associated with the points of S indicate the multiplicities of occurrence of the corresponding linear forms in the gates of C . We can similarly associate colors with the points obtained by applying π to the points in S to indicate multiplicities of occurrence of linear forms in the gates of C' .

¹⁴Note that many linear forms in the multiset L_C might map to P , however they will all be identical (up to multiplication by a constant in \mathbb{R}).

Let $S' = \{\pi(Q) \mid Q \in S\}$. For $Q' \in S'$ and $i \in [k]$ such that $i \notin C_{P_\ell}$, let $\text{mult}'(Q', i) = \sum_{Q \in \pi^{-1}(Q')} \text{mult}(Q, i)$. Notice that if $Q' = \pi(g(\tilde{\ell}))$, then $\text{mult}'(Q', i)$ represents the number of times the linear form $\pi(\tilde{\ell})$ appears in the gate $A'_i \in C'$.

For $i, j \in [k]$, let $(S_i \Delta S_j)' = \{Q' \in S' \mid \text{mult}'(Q', i) \neq \text{mult}'(Q', j)\}$ ¹⁵. Then showing $\text{pairwise-rank}(C') \geq A$ exactly translates to showing that for all $i, j \in [k]$ such that $i \neq j$, and $i, j \notin C_{P_\ell}$, $\dim(\text{affine-span}((S_i \Delta S_j)')) \geq A - 1$.

Finishing up the proof. Let i, j be any two colors not in C_{P_ℓ} . We consider two cases depending on whether $\{i, j\}$ is over-split or under-split.

Case 1: Suppose that $\{i, j\}$ is over-split. Hence by the definition of being over-split, there exist at least A hyperplanes distinct from H^* that split $\{i, j\}$. Without loss of generality (by reindexing the hyperplanes if necessary) let these hyperplanes be H_1, H_2, \dots, H_A , and for $1 \leq t \leq A$, let $Q_t \in H_t$ be such that $\text{mult}(Q_t, i) \neq \text{mult}(Q_t, j)$. Let $S_c = \{Q_t \mid 1 \leq t \leq A\}$. By property 2 of Lemma 6.3, we have $\dim(\text{affine-span}(\{P_\ell\} \cup S_c)) = A$. Since $\text{affine-span}(\{P_\ell\} \cup S_c) \subseteq V$, and since $\mathbf{0} \notin V$, hence $\dim(\text{affine-span}(\{\mathbf{0}, P_\ell\} \cup S_c)) = A + 1$. Let $\pi(Q_t) = Q'_t$. Let $S'_c = \{Q'_t \mid 1 \leq t \leq A\}$. Also, $\pi(\mathbf{0}) = \mathbf{0}$, and $\pi(P_\ell) = 0$. Since π has nullity 1 we get that $\dim(\text{affine-span}(\{\mathbf{0}\} \cup S'_c)) \geq A$. Hence $\dim(\text{affine-span}(S'_c)) \geq A - 1$.

Hence, to show that $\dim(\text{affine-span}((S_i \Delta S_j)')) \geq A - 1$, it suffices to show that $S'_c \subseteq (S_i \Delta S_j)'$. Let $Q'_t \in S'_c$. We will show that $\text{mult}'(Q'_t, i) \neq \text{mult}'(Q'_t, j)$, and this will imply that $Q'_t \in (S_i \Delta S_j)'$. Let $\pi(Q_t) = Q'_t$, where Q_t is a point in S_c . By definition of S_c we know that $\text{mult}(Q_t, i) \neq \text{mult}(Q_t, j)$. Also Q_t is the only point of S that maps to Q'_t under π , since if there was another point $Q^* \in S$ that maps to Q'_t , it would mean that P_ℓ, Q_t and Q^* are collinear. However there is no such point (by property 2 and 4 of Lemma 6.3). This implies that $\text{mult}'(Q'_t, i) \neq \text{mult}'(Q'_t, j)$, and hence $Q'_t \in (S_i \Delta S_j)'$. Since Q'_t was an arbitrary point in S'_c , this implies that $S'_c \subseteq (S_i \Delta S_j)'$, thus completing the proof for Case 1.

Case 2: Suppose that $\{i, j\}$ is under-split. Hence, by choice of H^* , for all points $Q \in (H^* \setminus H_{\text{core}}) \cap S$, we have $\text{mult}(Q, i) = \text{mult}(Q, j)$. Let $(S_i \Delta S_j)^{\text{core}} = H_{\text{core}} \cap (S_i \Delta S_j) = \{Q_1, Q_2, \dots, Q_s\}$. By choice of H_{core} , we have that $\dim(\text{affine-span}((S_i \Delta S_j)^{\text{core}})) \geq A - 1$. Since $\text{affine-span}((S_i \Delta S_j)^{\text{core}}) \subseteq H_{\text{core}}$, and $P_\ell \notin H_{\text{core}}$, we get $\dim(\text{affine-span}(\{P_\ell\} \cup (S_i \Delta S_j)^{\text{core}})) = A$. Now, letting $(S_i \Delta S_j)^{\text{core}' } = \{\pi(Q) \mid Q \in (S_i \Delta S_j)^{\text{core}}\}$, exactly as in the previous case we get that $\dim(\text{affine-span}((S_i \Delta S_j)^{\text{core}'})) \geq A - 1$.

To show that $\dim(\text{affine-span}((S_i \Delta S_j)')) \geq A - 1$, we will show that $(S_i \Delta S_j)^{\text{core}' } \subseteq (S_i \Delta S_j)'$. For $t \in [s]$, let $\pi(Q_t) = Q'_t$. Hence $Q'_t \in (S_i \Delta S_j)^{\text{core}' }$. By definition of $(S_i \Delta S_j)^{\text{core}}$, we know that $\text{mult}(Q_t, i) = \text{mult}(Q_t, j)$. Note that the line joining P_ℓ (which lies in $H^* \setminus H_{\text{core}}$) and Q_t (which lies in H_{core}) is contained in H^* . Also, since $\dim(H^*) = \dim(H_{\text{core}}) + 1$, any line through P_ℓ that intersects H_{core} will do so in a unique point. Also, any point in $\pi^{-1}(Q'_t)$ other than Q_t must lie on the line through P_ℓ and Q_t , and hence must lie in $H^* \setminus H_{\text{core}}$. All such points have equal multiplicity of i and j (because H^* does not split i and j). Now when we compute the multiplicity of colors i and j at the point Q'_t , we take the sum of multiplicities over all preimages of Q'_t under π . Q_t has different multiplicities for i and j , but all other points in the preimage have same multiplicities. Thus we get $\text{mult}'(Q'_t, i) \neq \text{mult}'(Q'_t, j)$, and hence $Q'_t \in (S_i \Delta S_j)'$. Since Q'_t was an arbitrary point in

¹⁵Each point in $(S_i \Delta S_j)'$ corresponds to a line in the pencil $\mathcal{L}_i^{P_\ell} \Delta \mathcal{L}_j^{P_\ell}$ as given in the outline of the proof of the lemma.

$(S_i \Delta S_j)^{\text{core}'}$, this implies that $(S_i \Delta S_j)^{\text{core}'} \subseteq (S_i \Delta S_j)'$. Hence $\dim(\text{affine-span}((S_i \Delta S_j)')) \geq A - 1$, thus completing the proof for Case 2.

■

8 The Hyperplane Decomposition Lemma

In this section we give a proof of the hyperplane decomposition lemma (Lemma 6.3) which was critically used in the proof of the fanin reduction lemma to identify structure in the linear forms of any high rank $\Sigma\Pi\Sigma$ circuit. The hyperplane decomposition lemma follows from an application of the high dimensional Sylvester–Gallai Theorem (Theorem 6.2), and we now give its proof.

Lemma 6.3 [Hyperplane decomposition] *Let V be an m dimensional affine space over \mathbb{R} . Let $S \subset V$ be a finite set, such that $\text{affine-span}(S) = V$. Let $S_{\text{core}} \subseteq S$, and let $H_{\text{core}} = \text{affine-span}(S_{\text{core}})$ be an affine space of dimension m_{core} . Then for some $r \geq \frac{m - m_{\text{core}}}{2}$, there exist hyperplanes $H_1, H_2, \dots, H_r \subset V$, such that letting $H = \text{affine-span}(\{H_i \mid i \in [r]\})$, we have the following properties.*

1. For all $i \in [r]$, $H_{\text{core}} \subseteq H_i$ and $\dim(H_i) = m_{\text{core}} + 1$.
2. $\dim(H) = m_{\text{core}} + r$. In particular, if $R \subseteq [r]$ is such that for each $i \in R$, P_i is a point in $H_i \setminus H_{\text{core}}$, then $\dim(\text{affine-span}(\{P_i \mid i \in R\})) = |R| - 1$.
3. For all $i \in [r]$, $(H_i \setminus H_{\text{core}}) \cap S \neq \phi$.
4. For every point $P \in S \cap H$, there exists $i \in [r]$ such that $P \in H_i$. Note that it is not necessary that every point in S lies on one of the H_i 's but every point of S inside H certainly does lie on at least one of the H_i 's.

Proof

For every point $P \in S$ such that $P \notin H_{\text{core}}$, let $H_P = \text{affine-span}(P, H_{\text{core}})$. Let \mathcal{H} be the system (pencil) of hyperplanes H_P thus obtained.

Let H' be a hyperplane in V of $m - m_{\text{core}}$ dimensions which intersects H_{core} in a single point Q , and intersects each $H_P \in \mathcal{H}$ in a single line through Q , which we denote by L_P (a random hyperplane in H of $m - m_{\text{core}}$ dimensions would have these properties with high probability). Thus we get a pencil of lines \mathcal{L} through Q that are contained in H' . Now let H'' be a hyperplane of $m - m_{\text{core}} - 1$ dimensions that is contained in H' , and intersects each line $L_P \in \mathcal{L}$ in a single point, such that these points span H'' , and $H'' \cap H_{\text{core}} = \phi$. Moreover no line contained in H'' is parallel to any line contained in H_{core} (again a random hyperplane would work with high probability). Let $S_{\mathcal{L}}$ be the set of points of intersection of H'' with the lines in \mathcal{L} .

By Lemma 6.2, there exists an $r - 1$ dimensional elementary hyperplane H_{elem} (for some $r \geq \frac{m - m_{\text{core}}}{2}$) in H'' that is spanned by and containing exactly r points (of $S_{\mathcal{L}}$), say Q_1, \dots, Q_r . Let the lines (in H') corresponding to these points be L_1, L_2, \dots, L_r . Observe that since $H'' \cap H_{\text{core}} = \phi$, hence $H_{\text{elem}} \cap H_{\text{core}} = \phi$. Also, no line contained in H_{elem} is parallel to any line contained in H_{core} . Clearly $\text{affine-span}(\{L_i \mid i \in [r]\}) = \text{affine-span}(Q, H_{\text{elem}})$.

For every $i \in [r]$, let $H_i = \text{affine-span}(H_{\text{core}}, L_i)$. Observe that $H_i \in \mathcal{H}$, where \mathcal{H} was the pencil of hyperplanes we originally considered. Let $H = \text{affine-span}(\{H_i \mid i \in [r]\}) = \text{affine-span}(H_{\text{core}} \cup \{L_i \mid i \in [r]\})$. We now show that the hyperplanes H_i satisfy the requisite properties.

Property 1: Property 1 holds immediately by the fact that $H_i = \text{affine-span}(H_{\text{core}}, L_i)$, where L_i intersects H_{core} in the single point Q .

Property 2: We need to show that $\dim(H) = m_{\text{core}} + r$. Observe that $H_{\text{elem}} \subset \text{affine-span}(\{L_i \mid i \in [r]\}) \subset H$. Also, $\dim(H_{\text{elem}}) = r - 1$, $H_{\text{elem}} \cap H_{\text{core}} = \emptyset$ and no line contained in H_{elem} is parallel to any line contained in H_{core} . These properties imply that $\dim(\text{affine-span}(H_{\text{elem}}, H_{\text{core}})) = m_{\text{core}} + r$. Now, for $i \in [r]$, $H_i = \text{affine-span}(Q_i, H_{\text{core}})$, where $Q_i \in H_{\text{elem}}$. Hence $H = \text{affine-span}(H_{\text{elem}}, H_{\text{core}})$, and $\dim(H) = m_{\text{core}} + r$.

Similarly, if $R \subseteq [r]$ and if for $i \in R$, P_i is a point in $H_i \setminus H_{\text{core}}$, then since $H_i = \text{affine-span}(P_i, H_{\text{core}})$, hence $\text{affine-span}(\{H_i \mid i \in R\}) = \text{affine-span}(H_{\text{core}}, \{P_i \mid i \in R\})$. Also, since $\dim(H) = m_{\text{core}} + r$, it must be that $\dim(\text{affine-span}(\{H_i \mid i \in R\})) = m_{\text{core}} + |R|$ (since each additional hyperplane H_i can increase the dimension of the span by at most 1). As $\dim(\text{affine-span}(H_{\text{core}})) = m_{\text{core}}$, it must be that $\dim(\text{affine-span}(\{P_i \mid i \in R\})) \geq |R| - 1$. Since the span of $|R|$ points can have dimension at most $|R| - 1$, it must be that $\dim(\text{affine-span}(\{P_i \mid i \in R\})) = |R| - 1$.

Property 3: Property 3 holds since $H_i \in \mathcal{H}$, which was the pencil of hyperplanes originally considered, and by definition of \mathcal{H} we have that every hyperplane in \mathcal{H} contains a point $P \in S$ such that $P \notin H_{\text{core}}$.

Property 4: If possible let $P \in S \cap H$ be such that $P \notin H_i$ for all $i \in [r]$. Then the hyperplane $H_P = \text{affine-span}(H_{\text{core}}, P) \in \mathcal{H}$ is such that $H_P \subset H$, and H_P is distinct from all H_i , $i \in [r]$. Let H' intersect H_P in the line L_P . Then L_P is not contained in H_i for all $i \in [r]$, since if it was, then H_P would equal H_i . Hence L_P is a line through Q such that $L_P \subset H$, and $L_P \not\subset H_{\text{core}}$. This we conclude that $L_P \subset \text{affine-span}(\{L_i \mid i \in [r]\})$. Let $H'' \cap L_P = Q_P$. Then observe that $Q_P \subset \text{affine-span}(\{Q_i \mid i \in [r]\})$, contradicting the fact that Q_1, \dots, Q_r span an elementary hyperplane. Hence property 4 must hold.

This completes the proof of the hyperplane decomposition lemma. ■

9 Conclusion

Our paper invites further work in several directions.

1. **Proving high-dimensional Sylvester–Gallai Theorem over the field of complex numbers.** Such a theorem would extend our results on the rank bound and identity testing to the complex numbers. We conjecture the following analogue of Lemma 6.2 over the field \mathbb{C} of complex numbers: *There exists a constant $c \geq 2$ such that if S is a finite set of points spanning an affine space $V \subseteq \mathbb{C}^n$ with $\dim(V) \geq c \cdot t$ then there exists a t dimensional hyperplane $H \subseteq V$ such that $|H \cap S| = t + 1$ and H is spanned by the points of S . i.e $\text{affine-span}(H \cap S) = H$.*

2. **Conjecture for Sylvester–Gallai over finite fields.** Such a theorem would extend our results on the rank bound and identity testing to large finite fields. We conjecture that the following version of Sylvester–Gallai is true over finite fields: *There exists a constant $c \geq 2$ such that if p is a prime and S is a subset of points in the 3-dimensional projective space $\mathbb{P}^3(\mathbb{F}_p)$ and $p > |S|^c$, then there exists a pair of points in S such that the line through this pair of points contains no other point from S .*
3. **Devising a blackbox identity testing algorithm over finite fields.** As shown in [KS06], even the weaker form of the conjecture of Dvir and Shpilka is false over finite fields. Perhaps there is some other neat classification of the structure of depth three arithmetic circuits computing the zero polynomial over finite fields. We challenge the interested reader to devise a blackbox identity testing algorithm for $\Sigma\Pi\Sigma(k)$ circuits over finite fields.
4. **Resolving the stronger Dvir–Shpilka conjecture over fields of characteristic zero.** Prove or disprove that the rank $c(k)$ in Theorem 2.2 can be improved to a polynomially growing function of the top fanin k .

See Appendix A for a discussion on how the first two conjectures will affect our understanding of identity testing.

Acknowledgements

We wish to thank Amit Deshpande, Madhu Sudan, Partha Mukhopadhyay, Satya Lokam and Vijay Patankar for lots of encouragement and helpful discussions. We also want to thank Madhu Sudan for commenting on an earlier draft of this paper and improving its presentation. Many thanks to Swastik Kopparty for very patiently listening to many versions of the proof as it evolved. This work was initiated while the second author was visiting Microsoft Research Lab India and we wish to thank Microsoft Research Lab for their warm hospitality.

References

- [AB03] Manindra Agrawal and Somenath Biswas. Primality and identity testing via chinese remaindering. *Journal of the ACM*, 50(4):429–443, 2003.
- [Agr05] Manindra Agrawal. Proving lower bounds via pseudo-random generators. In R. Ramanujan and Sandeep Sen, editors, *FSTTCS 2005: Foundations of Software Technology and Theoretical Computer Science, 25th International Conference, Hyderabad, India, December 15-18, 2005, Proceedings*, volume 3821 of *Lecture Notes in Computer Science*, pages 92–105. Springer, 2005.
- [AM07] Vikraman Arvind and Partha Mukhopadhyay. The ideal membership problem and polynomial identity testing. *Electronic Colloquium on Computational Complexity (ECCC)*, 14(095), 2007.
- [AV08] Manindra Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *FOCS*, pages 67–75, 2008.
- [BE67] W. Bonnice and M. Edelstein. Flats associated with finite sets in \mathbf{P}^d . *Nieuw. Arch. Wisk.*, 15:11–14, 1967.
- [BM90] P. Borwein and W. O. J. Moser. A survey of sylvester’s problem and its generalizations. *Aequationes Mathematicae*, 40(1):111–135, 1990.
- [CK00] Zhi-Zhong Chen and Ming-Yang Kao. Reducing randomness via irrational numbers. *SIAM J. Comput.*, 29(4):1247–1256, 2000.
- [DS05] Zeev Dvir and Amir Shpilka. Locally decodable codes with 2 queries and polynomial identity testing for depth 3 circuits. In *STOC*, pages 592–601, 2005.
- [EK66] M. Edelstein and L. M. Kelly. Bisecants of finite collections of sets in linear spaces. *Canadian Journal of Mathematics*, 18:375–380, 1966.
- [EPS06] Noam D. Elkies, Lourens M. Pretorius, and Conrad Johann Swanepoel. Sylvester-gallai theorems for complex numbers and quaternions. *Discrete & Computational Geometry*, 35(3):361–373, 2006.
- [GK98] Dima Grigoriev and Marek Karpinski. An exponential lower bound for depth 3 arithmetic circuits. In *STOC*, pages 577–582, 1998.
- [GR05] Ariel Gabizon and Ran Raz. Deterministic extractors for affine sources over large fields. In *FOCS*, pages 407–418, 2005.
- [Han65] Sten Hansen. A generalization of a theorem of sylvester on the lines determined by a finite point set. *Mathematica Scandinavica*, 16:175–180, 1965.
- [IK03] Impagliazzo and Kabanets. Derandomizing polynomial identity tests means proving circuit lower bounds. In *STOC: ACM Symposium on Theory of Computing (STOC)*, 2003.

- [Kel86] L. M. Kelly. A resolution of the sylvester - gallai problem of j. -p. serre. *Discrete & Computational Geometry*, 1:101–104, 1986.
- [KS01] Adam Klivans and Daniel A. Spielman. Randomness efficient identity testing of multivariate polynomials. In *STOC*, pages 216–223, 2001.
- [KS06] Neeraj Kayal and Nitin Saxena. Polynomial identity testing for depth 3 circuits. In *Proceedings of the twenty-first Annual IEEE Conference on Computational Complexity (CCC)*, 2006.
- [KS08a] Zohar Shay Karnin and Amir Shpilka. Black box polynomial identity testing of generalized depth-3 arithmetic circuits with bounded top fan-in. In *IEEE Conference on Computational Complexity*, pages 280–291, 2008.
- [KS08b] Zohar Shay Karnin and Amir Shpilka. Reconstruction of generalized depth-3 arithmetic circuits with bounded top fan-in, 2008. Manuscript.
- [LV98] Daniel Lewin and Salil P. Vadhan. Checking polynomial identities over any field: Towards a derandomization? In *STOC*, pages 438–447, 1998.
- [MVV87] Ketan Mulmuley, Umesh V. Vazirani, and Vijay V. Vazirani. Matching is as easy as matrix inversion. *Combinatorica*, 7(1):105–113, 1987.
- [Sch80] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980.
- [Shp07] Amir Shpilka. Interpolation of depth-3 arithmetic circuits with two multiplication gates. In *STOC*, pages 284–293, 2007.
- [SS08] Nitin Saxena and C. Seshadhri. An almost optimal rank bound for depth-3 identities. *Electronic Colloquium on Computational Complexity (ECCC)*, 15(108), 2008.
- [SW01] Amir Shpilka and Avi Wigderson. Depth-3 arithmetic circuits over fields of characteristic zero. *Computational Complexity*, 10(1):1–27, 2001.
- [Zip90] R. Zippel. Interpolating polynomials from their values. *JSC*, 9(3):375–403, March 1990.

A The Sylvester–Gallai Theorem and its generalizations.

The following theorem was proved by Edelstein and Kelly [EK66].

Theorem A.1 *Let A, B and C be 3 nonempty finite subsets of points in \mathbb{R}^n such that $\text{affine-span}(A \cup B \cup C)$ has dimension at least 4 and $A \cap B \cap C = \emptyset$. Then there exists a line intersecting exactly 2 of the sets A, B, C .*

The proof is a clever application of the Sylvester–Gallai Theorem, and uses a special case of the hyperplane decomposition theorem proved in Section 8. By the correspondence of linear forms appearing in the circuit with colored points in \mathbb{R}^n , this almost immediately gives a rank bound for simple and minimal $\Sigma\Pi\Sigma(3)$ circuits that compute the zero polynomial. For the case of $k = 3$, the connection of the structure theorem to Sylvester–Gallai type theorems was even observed in [DS05]. However, this approach to proving the rank bound does not directly generalize and additional ideas are needed for $k > 3$.

Rank bound over other fields. Our proof the rank bound, Theorem 2.2, uses high-dimensional versions of the Sylvester–Gallai theorem. If the high dimensional Sylvester–Gallai theorem held over \mathbb{C} or over any other field, then using the techniques of our paper, they would translate to rank bounds for $\Sigma\Pi\Sigma(k)$ circuits over the corresponding field. For $k = 3$, it suffices just to prove a Sylvester–Gallai theorem for lines and fortunately this is known over \mathbb{C} [Kel86, EPS06]. The result by Edelstein and Kelly [EK66] essentially carries over for this case with a dimension bound of 5 instead of 4, and hence we get a rank bound for $\Sigma\Pi\Sigma(3)$ of 6 circuits over complex numbers. Furthermore proving a rank bound over complex numbers implies a rank bound over finite fields of characteristic significantly larger than the degree of the circuit.

In particular, the conjectures given in the conclusion would have the following implications.

Sylvester–Gallai over complex numbers, Conjecture (1): It will give a deterministic black-box identity testing algorithm for $\Sigma\Pi\Sigma(k)$ circuits (k fixed) over *all* fields of characteristic zero.

Sylvester–Gallai over large finite fields, Conjecture (2): It will give a deterministic blackbox identity testing algorithm for $\Sigma\Pi\Sigma(3, d, n)$ circuits over fields of characteristic $p > \text{poly}(d)$.

Furthermore, using a standard argument involving the Hilbert Nullstellensatz, it can be shown that Conjecture (1) implies high-dimensional Sylvester–Gallai over finite fields of characteristic $p > 2^{2^{O(d^2)}}$. Proving conjecture (2) will perhaps require a fundamentally new proof of the Sylvester–Gallai theorem which somehow manages to avoid the well-ordered property of the real field.