BLACKBOX IDENTITY TESTING FOR BOUNDED TOP-FANIN DEPTH-3 CIRCUITS: THE FIELD DOESN'T MATTER *

NITIN SAXENA† AND C. SESHADHRI‡

Abstract. Let C be a depth-3 circuit with n variables, degree d and top-fanin k (called $\Sigma\Pi\Sigma(k,d,n)$ circuits) over base field $\mathbb{F}$. It is a major open problem to design a deterministic polynomial time blackbox algorithm that tests if C is identically zero. Klivans & Spielman (STOC 2001) observed that the problem is open even when k is a constant. This case has been subjected to a serious study over the past few years, starting from the work of Dvir & Shpilka (STOC 2005).

We give the first polynomial time blackbox algorithm for this problem. Our algorithm runs in time poly$(n)d^k$, regardless of the base field. The *only* field for which polynomial time algorithms were previously known is $\mathbb{F} = \mathbb{Q}$ (Kayal & Saraf, FOCS 2009, and Saxena & Seshadhri, FOCS 2010). This is the first blackbox algorithm for depth-3 circuits that does not use the rank based approaches of Karnin & Shpilka (CCC 2008).

We prove an important tool for the study of depth-3 identities. We design a blackbox polynomial time transformation that reduces the number of variables in a $\Sigma\Pi\Sigma(k,d,n)$ circuit to k variables, but preserves the identity structure.

Key words. algebra homomorphism; blackbox; Chinese remaindering; depth-3 circuits; derandomization; identity testing; ideal theory

AMS subject classifications. 03D15, 68Q17, 68W30, 12Y05, 14Q20

1. Introduction. Polynomial identity testing (PIT) is a major open problem in theoretical computer science. The input is an arithmetic circuit that computes a polynomial $p(x_1, x_2, \ldots, x_n)$ over a base field $\mathbb{F}$. We wish to check if p is the zero polynomial, or in other words, is identically zero. We may be provided with an explicit circuit, or may only have blackbox access. In the latter case, we can only evaluate the polynomial p at various domain points. The main goal is to devise a *deterministic* blackbox polynomial time algorithm for PIT. One of the main reasons for interest in this problem is the connection between PIT algorithms and circuit lower bounds (Heintz & Schnorr [HS80], Kabanets & Impagliazzo [KI04] and Agrawal [Agr05, Agr06]). Refer to surveys for a detailed treatment of PIT [Sax09, AS09, SY10].

Since the problem of derandomizing PIT is very hard, restricted versions of it have been studied. One common and natural variant is that of *bounded depth circuits*. Results of Agrawal & Vinay [AV08] justify this restriction. They essentially show that an efficient blackbox identity test for depth-4 circuits leads to (almost) the complete resolution of PIT derandomization and also provides exponential lower bounds. Thus, understanding depth-3 identities seems to be a natural first step towards the goal of PIT and circuit lower bounds. Raz [Raz10] showed that strong lower bounds for depth-3 circuits imply super-polynomial lower bounds for general arithmetic formulas. Not surprisingly, the problem of derandomizing PIT is still wide open for the special case of depth-3 circuits.

A depth-3 circuit C over a field $\mathbb{F}$ is of the form $C(x_1, \ldots, x_n) = \sum_{i=1}^{k} T_i$, where T_i (*a multiplication term*) is a product of at most d linear polynomials with coefficients

*A preliminary version of this paper appeared as [SS11b].

†Hausdorff Center for Mathematics, Bonn, Germany `ns@hcm.uni-bonn.de`

‡Sandia National Laboratories, Livermore, USA `scomand@sandia.gov` (Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.)

in $\mathbb{F}$. The size of the circuit C can be expressed in three parameters: the number of variables n, the degree d, and the *top-fanin* k. Such a circuit is referred to as a $\Sigma\Pi\Sigma(k,d,n)$ circuit. This is the first result to give a blackbox polynomial time algorithm when the top-fanin k is constant, regardless of the field. This provides a complete resolution for the main open question of Klivan & Spielman [KS01].

The study of PIT algorithms for depth-3 circuits was initiated by Dvir & Shpilka [DS06], who gave a quasi-polynomial time non-blackbox algorithm. The first non-trivial *blackbox* algorithm was given by Karnin & Shpilka [KS08]. There have been many recent results in this area by Kayal & Saxena [KS07], Saxena & Seshadri [SS11a, SS10a], and Kayal & Saraf [KS09b]. Our main result is the first polynomial time blackbox tester for bounded top-fanin depth-3 circuits over *any* field.

THEOREM 1.1. *There is a deterministic blackbox* $\mathrm{poly}(nd^k)$ *time algorithm for PIT on* $\Sigma\Pi\Sigma(k,d,n)$ *circuits, regardless of the base field* $\mathbb{F}$.

Table 1.1 details the time complexities of previous algorithms. We also give the assumptions on the field required for the algorithm. Technically, the running times are polynomial in the stated times. We stress that the time complexities bound the number of bit operations of the hitting-set.

Table 1.1: Depth-3 blackbox PIT algorithms (actual running times are polynomial in stated times)

Paper	Time complexity	Field
[KS08]	$nd^{(2^{k^2}\log^{k-2} d)}$	Any
[SS11a]	$nd^{k^3\log d}$	Any
[KS09b]	nd^{k^k}	`Reals`
[SS10a]	$nd^{k^2\log d}$	Any
[SS10a]	nd^{k^2}	`Reals`
This paper	nd^k	Any

The only field for which such polynomial time algorithms were previously known was $\mathbb{R}$ (technically $\mathbb{Q}$ if we consider the overall bit complexity). This was a breakthrough result of Kayal & Saraf [KS09b], followed by improvements in [SS10a]. These used beautiful incidence geometry theorems for the reals, but analogues of these results are either unknown or false for other fields. Since the best running time of these algorithms is $\mathrm{poly}(nd^{k^2})$, we get an improved algorithm for this case as well.

As Table 1.1 shows, *even for the simple case of* $k = 3$ *and* $\mathbb{F} = \mathbb{F}_2$*, no deterministic polynomial-time blackbox PIT algorithm was known.* Kayal & Saxena [KS07] gave a non-blackbox algorithm (over all fields), which runs in $\mathrm{poly}(nd^k)$ time. Theorem 1.1 closes the gap between blackbox and non-blackbox algorithms.

Throughout the following discussion, we will think of k as a constant. Hence, when we refer to polynomial time, the dependence on k will be ignored.

1.1. Variable reduction. Dvir & Shpilka [DS06] introduced a powerful idea. They defined the notion of the *rank* of a $\Sigma\Pi\Sigma(k,d,n)$ circuit. We will not explain this precisely here but merely say that this is the number of "free variables" in a $\Sigma\Pi\Sigma$ circuit. They proved the remarkable fact that the rank of every identity[1] is

[1] A small caveat: there are some technical restrictions of simplicity and minimality.

small. This led to the reduction of PIT for general $\Sigma\Pi\Sigma(k, d, n)$ circuits to PIT on $\Sigma\Pi\Sigma$ circuits over few variables. They developed a non-blackbox quasi-polynomial time algorithm through their rank bounds. Karnin & Shpilka [KS08] used the low rank results to devise blackbox algorithms for $\Sigma\Pi\Sigma$ circuits. Their algorithms had a running time that depended exponentially in the rank. Hence, constant rank bounds would lead to polynomial time algorithms. Unfortunately, Kayal & Saxena [KS07] gave constructions (extended in [SS11a]) showing that for $\Sigma\Pi\Sigma$ identities over finite fields, the rank could be unbounded (as large as $k \log d$). This means that the best running time one could hope for over finite fields via this approach was $d^{k \log d}$. Tighter rank bounds from [SS11a, SS10a] gave algorithms that almost match this running time. For the special case when the field $\mathbb{F}$ is $\mathbb{R}$, Kayal & Saraf [KS09b] proved a constant rank bound, establishing the first polynomial time blackbox algorithm for this case. Refer to [SS10a] for a more detailed treatment of rank bounds.

Until this work, all blackbox algorithms relied solely on rank bounds. As the examples of [KS07] show, even for the case of $\mathbb{F}_2$, a new idea is required to get polynomial time algorithms. We provide the first blackbox algorithm that circumvents the problem of large rank identities. Interestingly, one of the main ideas has roots in the non-blackbox polynomial time algorithm of Kayal & Saxena [KS07]. This algorithm had a completely different idea and used generalizations of the *Chinese Remainder Theorem.* They effectively provided *low rank* certificates for non-identities. These algebraic ideas were further developed in a previous work of the authors [SS10a]. Karnin & Shpilka [KS08] used the extractors of Gabizon & Raz [GR08] to construct their blackbox algorithm. We combine these extractor ideas with the algebraic framework to develop a very useful algorithmic tool. Any $\Sigma\Pi\Sigma(k, d, n)$ circuit can be converted into a small family of $\Sigma\Pi\Sigma$ circuits over just k variables. The original circuit is an identity iff all circuits in the family are identities. This transformation has a running time $\mathrm{poly}(kdn)$ and only requires knowledge on the size bounds k, d, and n.

THEOREM 1.2. *Let $\mathbb{F}$ be an arbitrary field such that $|\mathbb{F}| > dnk^2$. There is a deterministic algorithm that takes as input a triple (k, d, n) of natural numbers and in time $\mathrm{poly}(kdn)$, outputs a set of linear maps (that are represented as matrices using $\mathrm{poly}(kdn)$ bits) $\Psi_i : \mathbb{F}[x_1, \ldots, x_n] \to \mathbb{F}[y_1, \ldots, y_k]$ ($1 \le i \le \mathrm{poly}(kdn)$). A $\Sigma\Pi\Sigma(k, d, n)$ circuit C is identically zero iff $\forall i, \Psi_i(C) = 0$.*

Remark: The field size restriction made in the theorem is really no loss of generality. In the blackbox model, it is standard to assume that we can query the given circuit on points in an algebraic extension field. If $\mathbb{F}$ is small, then we just need to move to a large enough extension field $\mathbb{F}'$. Such an extension can be found by a deterministic $\mathrm{poly}(\log |\mathbb{F}'|)$ time method of Adleman & Lenstra [AL86], or even by a slower brute-force method of finding a suitable irreducible polynomial over $\mathbb{F}$. We do all our computations in this extension field $\mathbb{F}'$. Henceforth, for convenience, we will assume that the field has size larger than dnk^2.

Regardless of n, d or $\mathbb{F}$, every $\Sigma\Pi\Sigma(3, d, n)$ non-identity can be converted to an "equivalent" non-identity involving just 3 variables. The time required to generate this transformation does not have the "usual" exponential dependence on k, that we see in all PIT algorithms (including ours). This theorem has a uniform treatment of all fields, and is hence stronger than rank bounds. We believe that this theorem will be useful in reaching the holy grail of a truly polynomial time PIT algorithm for $\Sigma\Pi\Sigma$ circuits.

The circuits $\Psi_i(C)$ only involve k variables. Using some standard PIT techniques, we can construct the following hitting-set. This proves Theorem 1.1.

THEOREM 1.3. *Given the triple of natural numbers (k, d, n), a hitting-set $\mathcal{H} \subseteq \mathbb{F}^n$ for $\Sigma\Pi\Sigma(k, d, n)$ circuits can be constructed in deterministic* $\mathrm{poly}(nd^k)$ *time. In other words, for every nonzero $\Sigma\Pi\Sigma(k, d, n)$ circuit C over $\mathbb{F}$, there is some vector $(\alpha_1, \dots, \alpha_n) \in \mathcal{H}$ such that $C(\alpha_1, \dots, \alpha_n) \neq 0$.*

1.2. History. The first randomized polynomial time PIT algorithm was given (independently) by Schwartz [Sch80], Zippel [Zip79] and DeMillo & Lipton [DL78]. Algorithms using less randomness were devised by Chen & Kao [CK97], Lewin & Vadhan [LV98], and Agrawal & Biswas [AB03]. For depth-2 circuits, there has been a long line of work studying blackbox PIT algorithms [BOT88, CDGK91, Wer94, KS96, SS96, GKS90, KS01]. Over algebras more general than fields, Saha *et al.* [SSS09] studied the complexity of depth-2 PIT. Raz & Shpilka studied non-blackbox algorithms for non-commutative formulas [RS05], initiating a line of work where the power of multiplication gates is restricted.

The history of depth-3 PIT has been explained quite a bit in the previous sections. Identity tests are known only for very special depth-4 circuits [AM10, Sax08, SV08, SV09, KMSV10, SV11, AvMV11, SSS12, BMS11]. Recently, Agrawal *et al.* [ASSS12] have shown how a higher-degree generalization of rank to *transcendence degree*, and the power of the *Jacobian*, unifies all known hitting-set generators. (It also proves related lower bounds.) This includes bounded top-fanin depth-3 circuits as well (and even relaxes bounded-ness). Nonetheless, we highlight two relative advantages of our work. Firstly, [ASSS12] requires a zero (or large) characteristic and hence could not prove our Theorem 1.1. Our techniques are, in this sense, truly field independent. Secondly, [ASSS12] reduces n variables to $O(k)$ many, hence, its variable reduction is not as strong as our Theorem 1.2.

At the end, we would just like to indicate how all the depth-3 PIT results are interrelated and how they collectively influenced progress in this problem. The rank notion of Dvir & Shpilka [DS06], the Chinese Remaindering of Kayal & Saxena [KS07], the rank preserving subspaces of Karnin & Shpilka [KS08], the series of improved rank bounds by the authors and Kayal & Saraf [SS11a, KS09b, SS10a]: Each paper built off previous results and provided enough food for thought for subsequent papers. This paper also builds on the edifice constructed so far.

1.3. Organization. In Section 2, we give some basic definitions and give an intuitive overview of our approach. Section 3 gives some of the tools that were developed in previous works. In Section 4, we give our main analysis and prove Theorems 1.2 and 1.3. To make the paper more self-contained, we have collected standard theorems and the proof-sketches in an appendix.

2. Preliminaries and intuition. We will denote the set $\{1, \dots, n\}$ by $[n]$. We fix the base field to be $\mathbb{F}$, so the circuits compute multivariate polynomials in the *polynomial ring* $\mathcal{R} := \mathbb{F}[x_1, \dots, x_n]$. We use $\mathbb{F}^*$ to denote $\mathbb{F} \setminus \{0\}$. A *linear form* is a linear polynomial in $\mathcal{R}$ with zero constant term. We will denote the set of all linear forms by $L(\mathcal{R}) := \{\sum_{i=1}^n a_i x_i \mid a_1, \dots, a_n \in \mathbb{F}\}$. The set $L(\mathcal{R})$ is a vector (or linear) space over $\mathbb{F}$, a fact that shall be repeatedly used. Much of what we do shall deal with *multi*-sets of linear forms (also products of linear forms) and equivalence classes inside them. A *list* of linear forms is a multi-set of forms with an arbitrary order associated with them.

DEFINITION 2.1. *We collect some important definitions from [SS11a]:*

[**Multiplication term and operators** $L(\cdot)$ **&** $M(\cdot)$] *Let* $c \in \mathbb{F}^*$ *and* S *be a list of nonzero linear forms. The* multiplication term of S*,* $M(S)$*, is* $\prod_{\ell \in S} \ell$*. The polynomial* $f := c \cdot \prod_{\ell \in S} \ell$ *is called a* multiplication term. *The* list of linear forms in f*,* $L(f)$*, is just the list* S *in the product. (Conventionally,* $L(c) = \emptyset$ *and* $M(\emptyset) = 1$*.)*

[$\Sigma\Pi\Sigma$ **circuits**] *A* $\Sigma\Pi\Sigma(k, d, n)$ *circuit* C *is a sum of* k *multiplication terms of degree* d*,* $C = \sum_{i=1}^{k} T_i$*. The list of* linear forms occurring in C *is* $L(C) := \bigcup_{i \in [k]} L(T_i)$*. Note that* $L(C)$ *is a list of size exactly* kd*.*

For any subset $S \subseteq [k]$*, the* sub-circuit C_S *is* $\sum_{s \in S} T_s$*. For an* $i \in \{0, \ldots, k-1\}$*, define* $[i]' := [k] \setminus [i]$*. (Conventionally,* $[0] := \emptyset$ *and* $C_\emptyset := 0$*.)*

[**Span** $\mathrm{sp}(\cdot)$ **and Rank** $\mathrm{rk}(\cdot)$] *For any* $S \subseteq L(\mathcal{R})$ *we let* $\mathrm{sp}(S) \subseteq L(\mathcal{R})$ *be the* linear span *of the linear forms in* S *over the field* $\mathbb{F}$*. (Conventionally,* $\mathrm{sp}(\emptyset) = \{0\}$*.) We use* $\mathrm{rk}(S)$ *to denote the rank of* S*, considered as vectors in* $\mathbb{F}^n$*.*

We always assume that the circuit is homogeneous, so each T_i is a product of d linear forms. This is a standard assumption, made rigorous in [DS06, Lemma 3.5]. Essentially, we introduce a dummy variable for the constant term in each linear form.

2.1. Intuition and main ideas. We give a high-level description of the main ideas used in this paper. Some notions are deliberately left vague, and others may even be formally incorrect. Nonetheless, this sketch is "morally" correct and, at some level, shows how the authors arrived at their conclusions.

How do we convert a high variate $\Sigma\Pi\Sigma(k, d, n)$ circuit C into a low variate one and still preserve the structure of C? We wish to do this by a linear transformation $\Psi : \mathbb{F}[x_1, \ldots, x_n] \to \mathbb{F}[y_1, \ldots, y_\ell]$, where ℓ is comparable to k. When the rank of the linear forms in C is itself comparable to k, this can be done quite directly. We will get a circuit $\Psi(C)$ that is essentially *isomorphic* to C. For identities of large rank, such a transformation seems impossible. Any linear transformation will necessarily destroy some of the structure of C. This is because forms that were independent in C are now dependent in $\Psi(C)$. But maybe we are trying too hard to preserve the dependencies in C? After all, we want a transformation Ψ that sends identities to identities. (And, of course, non-identities to non-identities. Surely, satisfying only the former condition is not too hard.) We are not particularly bothered about how well $\Psi(C)$ preserves the exact structure of C. Though this seems obvious, this is a very important point. All the previous approaches were trying to preserve the complete structure of C through a linear transformation. We show that relaxing this requirement allows for more efficient algorithms.

This is where the Chinese Remaindering techniques of [KS07] and the ideal framework of [SS10a] enters the picture. For any non-identity C, it was shown that there exists an ideal I generated by products of forms in $L(C)$ that "certifies" that C is nonzero. Essentially, the polynomial C is not in ideal I and hence, must be nonzero. In reality, it is much more complicated than that, but for the sake of explanation, it captures the main idea. The forms involved in generating I have rank at most k. This gives a low-dimensional certificate of the non-zeroness of C.

We argue that if Ψ can selectively preserve the forms generating I, then $\Psi(I)$ remains a certificate for $\Psi(C)$. In other words, $\Psi(C)$ will not be in $\Psi(I)$. All that is needed is to find such a Ψ that is independent of C (since we are interested in blackbox algorithms). Enter [KS08]. They develop the notion of *rank-preserving subspaces*. These can be viewed as linear transformations from a large-dimensional vector space $\mathcal{S}_1$ to a smaller dimensional one $\mathcal{S}_2$. These preserve the structure (in terms of linear independence) of specific low dimensional subspaces of $\mathcal{S}_1$. These were constructed in a blackbox manner using the extractors of [GR08]. By combining all the arguments properly, we can construct a transformation Ψ such that $\Psi(C)$ is a k-variate $\Sigma\Pi\Sigma$ circuit. (To make this idea work, we will actually need a *set* of transformations, and one Ψ will not suffice.)

The circuit $\Psi(C)$ is of the form $\Sigma\Pi\Sigma(k,d,k)$. This has only a constant number of variables, and the [Sch80, Zip79, DL78] lemma gives a simple blackbox algorithm for such circuits. This is combined with the transformation Ψ to construct the final hitting-set. Our hitting-set is structurally similar to the one constructed in [KS08], although various parameters differ and our analysis is completely different.

3. Necessary tools. In this section, we list out the basic tools that we need. In the first part, we explain the low-rank ideal certificates for the non-zeroness of C. This requires some technical definitions, before we can state the exact theorem. In the second part, we give the key lemma of the Vandermonde matrix transformation used in [KS08]. Once these tools are set in place, we will explain how the variable reduction works.

3.1. Low rank certificates. Theorems in [SS10a] (reinterpreting and generalizing the earlier work in [KS07]) essentially show that there is a low-rank subspace "certifying" the non-zeroness of C. We shall restate those results in a convenient form for this paper. We define the standard notion of an *ideal*.

DEFINITION 3.1. **[Ideal]** *An* ideal I *of* $\mathcal{R}$ *with* generators $f_i, i \in [m]$, *is the set* $\{\sum_{i\in[m]} q_i f_i | q_i\text{'s} \in \mathcal{R}\}$ *and is denoted by the notation* $\langle f_1,\ldots,f_m\rangle$. *For a set of polynomials* $\overline{f} = \{f_1,\ldots,f_m\}$, *let* $\langle\overline{f}\rangle$ *denote* $\langle f_1,\ldots,f_m\rangle$. *Two forms* ℓ *and* ℓ' *are* similar modulo I *if* $\ell \equiv \alpha\ell' (\mathrm{mod}\ I)$ *for some* $\alpha \in \mathbb{F}^*$.

Note that for an $f \in \mathcal{R}$, the following statements mean the same thing: $f \equiv 0(\mathrm{mod}\ I)$, $f \equiv 0(\mathrm{mod}\ f_1,\ldots,f_m)$ and $f \in I$. We now state the definition of a radical-span of the set of generators for an ideal I. This was introduced in [SS10a] and is crucial for formalizing various arguments about ideal generated by multiplication terms.

DEFINITION 3.2. **[Radical-span]** *Let* $S := \{f_1,\ldots,f_m\}$ *be multiplication terms generating an ideal* I. *The* radical-span of S *is the linear space* $\mathrm{radsp}(S) := sp(L(f_1)\cup \ldots \cup L(f_m))$.

When the set of generators S for ideal I are clear from the context we use the notation $\mathrm{radsp}(I)$. Similarly, $\mathrm{radsp}(I,f)$ is shorthand for $\mathrm{radsp}(S\cup\{f\})$. We can now state the theorem providing low-rank certificates for nonzero circuits. Appendix C gives more details about this theorem and how it follows from [SS10b, Theorem 25]. As we mentioned earlier, this can be seen as a formalization of an idea present in [KS07].

THEOREM 3.3 (Certificate for a non-identity). *Let* $C = \sum_{i\in[k]} T_i$ *be a nonzero* $\Sigma\Pi\Sigma(k,d,n)$ *circuit. Then there exists a set* $\overline{p} = \{p_1,p_2,\ldots,p_b\}$ *of multiplication terms, an* $i \in [k]$ *and an* $\alpha \in \mathbb{F}^*$ *such that:* $\mathrm{rk}(\mathrm{radsp}(\overline{p})) < k$ *and* $C \equiv \alpha\cdot T_i \not\equiv 0$ $(\mathrm{mod}\ \overline{p})$.

Remark: The $\overline{p}$ above, we call, a *certifying path*.

3.2. The Vandermonde linear transformation. Linear transformations based on the Vandermonde matrix are used widely, e.g. in [GR08] to construct linear seeded extractors for affine sources. This was adapted in [KS08] to design depth-3 blackbox PIT algorithms. We will need ideas from [GR08, Lemma 6.1] to construct linear transformations that reduce the dimension of a space but still preserve the structure of low rank subspaces.

We will design a linear transformation $\Psi_\beta : \mathbb{F}^n \to \mathbb{F}^k$ for any $\beta \in \mathbb{F}$ as follows. Let $V_{n,k,\beta}$ denote the $n \times k$ Vandermonde matrix. This is defined as $(V_{n,k,\beta})_{i,j} := \beta^{ij}$. We have

$$(b_1 \dots b_k) \;=\; (a_1 \dots a_n) \cdot V_{n,k,\beta} \qquad (b_r = \sum_{i=1}^{n} a_i \beta^{ir}) \tag{3.1}$$

So far, Ψ_β has been seen as a linear transformation. But we wish to eventually understand its action on ideals. For that reason, it is necessary to view Ψ_β as a *linear homomorphism* from $\mathcal{R} = \mathbb{F}[x_1, \dots, x_n]$ to $\mathcal{R}' := \mathbb{F}[y_1, \dots, y_k]$. This means that Ψ_β maps $\mathcal{R}$ to $\mathcal{R}'$ and preserves the ring operations of addition and multiplication. We can equivalently define Ψ_β as

$$\forall i \in [n], \quad \Psi_\beta : x_i \mapsto \sum_{j=1}^{k} \beta^{ij} y_j \tag{3.2}$$

We define $\Psi_\beta(\alpha) = \alpha$ for all $\alpha \in \mathbb{F}$. This (naturally) defines the action of Ψ_β, on *all* the elements of $\mathcal{R}$, that preserves the ring operations of $\mathcal{R}$. We now state a key property of Ψ [GR08, Lemma 6.1]. For completeness, a proof is provided in Appendix A. It is convenient to introduce a definition of *bad sets* of values.

DEFINITION 3.4. *Let $S \subseteq L(\mathcal{R})$ and ψ_β be defined as in Equation (3.2). The* bad set *$B(S)$ is the set of field values $\{\beta \in \mathbb{F} \mid \mathrm{rk}(\psi_\beta(S)) \neq \mathrm{rk}(S)\}$.*

LEMMA 3.5 (Ψ_β preserves k-rank). *Let $S \subseteq L(\mathcal{R})$ be a subset of linear forms with $\mathrm{rk}(S) \le k$. Then $|B(S)| \le nk^2$.*

The homomorphism Ψ_β depends solely on $\mathbb{F}, k, n$. If β has a bit representation of at most $O(kdn)$ bits, then Ψ_β is computable in $\mathrm{poly}(kdn)$ *time.*

4. The reduction algorithm. With all the basic definitions in place, we are now ready to convert any $\Sigma\Pi\Sigma(k,d,n)$ non-identity C into a $\Sigma\Pi\Sigma(k,d,k)$ non-identity. This will be done through a two-step process. We will use the properties of Ψ_β to prove a generalization of Lemma 3.5. Not only does Ψ_β preserve small subspaces, but it also maintains the structure of ideals having low radical-span. This leads to the second step. We deduce that Ψ_β must preserve all paths of C, since paths have a low radical-span. Since C has a certifying path $\overline{p}$, we show that $\Psi_\beta(\overline{p})$ must also be a certificate for $\Psi_\beta(C)$.

4.1. The moral nature of Ψ_β: It maintains ideals. We study the action of Ψ_β on ideals. Our main lemma is the following.

LEMMA 4.1 (Ψ_β preserve ideals). *Let $f_1, \dots, f_m, f$ be multiplication terms in $\mathcal{R}$. Define the ideal $I := \langle f_1, \dots, f_m \rangle$. Suppose the span* $\mathrm{radsp}(I, f)$*, over $\mathbb{F}$, is of rank at most k. Let S be any set of k linearly independent forms such that S contains a basis of* $\mathrm{radsp}(I, f)$.

For all $\beta \notin B(S)$: $f \in I$ iff $\Psi_\beta(f) \in \langle \Psi_\beta(f_1), \dots, \Psi_\beta(f_m) \rangle$.

To prove this, we need to show that the map Ψ_β is an isomorphism on small enough subrings of $\mathcal{R}$. This is a fairly direct consequence of Lemma 3.5. For $\ell_1, \dots, \ell_k \in L(\mathcal{R})$, $\mathbb{F}[\ell_1, \dots, \ell_k] \subset \mathcal{R}$ denotes the set of all polynomials $g(\ell_1, \dots, \ell_k)$, where $g \in \mathbb{F}[y_1, \dots, y_k]$. This is the *subalgebra* of $\mathcal{R}$ generated by $\{\ell_1, \dots, \ell_k\}$.

LEMMA 4.2. *Let $S = \{\ell_1, \dots, \ell_k\}$ be a set of k linearly independent forms (from $L(\mathcal{R})$). For any $\beta \notin B(S)$, Ψ_β induces an isomorphism between $\mathbb{F}[\ell_1, \dots, \ell_k]$ and $\mathcal{R}'$.*

Proof. Let $\mathcal{B}$ denote $\mathbb{F}[\ell_1, \dots, \ell_k]$ and set $S := \{\ell_1, \dots, \ell_k\} \subset L(\mathcal{R})$. It will be convenient to define Φ_β to be the map induced by Ψ_β on $\mathcal{B}$. We will show that $\Phi_\beta : \mathcal{B} \to \mathcal{R}'$ is an isomorphism. Since S is of rank k, $\mathcal{B}$ is isomorphic to $\mathcal{R}'$ (under an invertible linear transformation of variables). To show that the homomorphism Φ_β is an isomorphism, it suffices to prove that Φ_β is onto. We need to show that for any $g(y_1, \dots, y_k) \in \mathcal{R}'$, there exists $p \in \mathcal{B}$ such that $\Phi_\beta(p) = g$.

Since $\beta \notin B(S)$, $\mathrm{rk}(\Phi_\beta(S)) = \mathrm{rk}(S) = k$. Therefore, for each y_i, there exist constants $\alpha_j \in \mathbb{F}$ such that $y_i = \sum_j \alpha_j \Phi_\beta(\ell_j)$. By the linearity of Φ_β, $y_i = \Phi_\beta(\sum_j \alpha_j \ell_j)$. Hence, for each y_i, there is some linear form $t_i \in L(\mathcal{B})$, such that $\Phi_\beta(t_i) = y_i$. The polynomial $p := g(t_1, \dots, t_k)$ is certainly in $\mathcal{B}$. Since Φ_β is a homomorphism, $\Phi_\beta(p) = g$. □

Now we prove Lemma 4.1.

Proof. (of Lemma 4.1) If $f \in I$, then $f = \sum_{i \in [m]} g_i f_i$, where $g_i \in \mathcal{R}$. Since Ψ_β is a homomorphism, $\Psi_\beta(f) = \sum_{i \in [m]} \Psi_\beta(g_i)\Psi_\beta(f_i)$. So $\Psi_\beta(f) \in \langle \Psi_\beta(f_1), \dots, \Psi_\beta(f_m) \rangle$.

We now show the converse. Let the span $\mathrm{radsp}(I, f)$ over $\mathbb{F}$ be generated by linear forms $\ell_1, \dots, \ell_r \in S$. Note that $r \leq k$. Define subring (of $\mathcal{R}$) $\mathcal{B} := \mathbb{F}[\ell_1, \dots, \ell_k]$. Applying Lemma 4.2, we get that for all $\beta \notin B(S)$, Ψ_β induces an isomorphism $\Phi_\beta : \mathcal{B} \to \mathcal{R}'$. Fix the unique elements $t_1, \dots, t_k \in \mathcal{B}$ such that $\Phi_\beta(t_i) = y_i$ for all $i \in [k]$.

Suppose $\Psi_\beta(f) \in \langle \Psi_\beta(f_1), \dots, \Psi_\beta(f_m) \rangle$. Then there exist $g_1, \dots, g_m \in \mathcal{R}'$ such that,

$$\Psi_\beta(f) = \sum_{i=1}^{m} g_i \cdot \Psi_\beta(f_i) \tag{4.1}$$

Each g_i is a polynomial in $(y_1, \dots, y_k)$ over $\mathbb{F}$. So we can define the polynomial,

$$h := f - \sum_{i=1}^{m} g_i(t_1, \dots, t_k) \cdot f_i.$$

Note that f and f_i's are multiplication terms generated by forms in $\mathrm{sp}(\ell_1, \dots, \ell_k)$. Hence all of these are in $\mathcal{B}$. The polynomials $g_i(t_1, \dots, t_k)$ are also in $\mathcal{B}$, so $h \in \mathcal{B}$. Since Φ_β is a homomorphism,

$$\begin{aligned}\Phi_\beta(h) &= \Phi_\beta(f) - \sum_{i=1}^{m} g_i(\Phi_\beta(t_1), \dots, \Phi_\beta(t_k)) \cdot \Phi_\beta(f_i) \\ &= \Psi_\beta(f) - \sum_{i=1}^{m} g_i(y_1, \dots, y_k) \cdot \Psi_\beta(f_i) = 0\end{aligned}$$

But Φ_β is an isomorphism, so $h = 0$. This implies $f = \sum_{i=1}^{m} g_i(t_1, \dots, t_k) \cdot f_i$. Note that the evaluations of the g_i's are in $\mathcal{R}$. Thus, $f \in \langle f_1, \dots, f_m \rangle$. □

4.2. Variable reduction: Proof of Theorem 1.2. We come to the main part where we merge the path certificates with the properties of Ψ_β to prove the variable reduction. We will need a technical cancellation lemma from [SS10b], proven in Appendix B.

LEMMA 4.3. *Let $f_1, \ldots, f_m$ be multiplication terms generating an ideal I, let $\ell \in L(\mathcal{R})$ and $g \in \mathcal{R}$. If $\ell \notin \mathrm{radsp}(I)$ then: $\ell g \in I$ iff $g \in I$.*

We now state our main variable reduction lemma. This, combined with the polynomial time constructions of the Ψ_β's, completes the proof of Theorem 1.2.

LEMMA 4.4. *Let C be a $\Sigma\Pi\Sigma(k, d, n)$ circuit and $U \subseteq \mathbb{F}$ such that $|U| = dnk^2+1$. Then $C = 0$ iff $\forall \beta \in U, \Psi_\beta(C) = 0$.*

Proof. Since Ψ_β is a homomorphism, $C = 0$ implies $\forall \beta, \Psi_\beta(C) = 0$.

Suppose $C \neq 0$. Applying Theorem 3.3 on C yields a set $\overline{p} = \{p_1, p_2, \ldots, p_b\}$ of multiplication terms and an $i \in [k]$ such that

$$C \equiv \alpha \cdot T_i \not\equiv 0 \pmod{\overline{p}}, \text{ for some } \alpha \in \mathbb{F}^*. \tag{4.2}$$

Also, $\mathrm{rk}(\mathrm{radsp}(\overline{p})) < k$. Let $g := M(L(T_i) \cap \mathrm{radsp}(\overline{p}))$. This is just the product of all forms in T_i that are in $\mathrm{radsp}(\overline{p})$. So T_i/g is the product of forms *not* in $\mathrm{radsp}(\overline{p})$. By repeated applications of Lemma 4.3, since $T_i \notin \langle\overline{p}\rangle$, $g \notin \langle\overline{p}\rangle$. The rank of $\mathrm{radsp}(\overline{p}, g)$ is less than k. Indeed, for any linear form $\ell \in L(T_i)$, the rank of $\{\ell\} \cup \mathrm{radsp}(\overline{p})$ is at most k.

We collect a set B of bad β values. Let S'_ℓ be a set of forms that constitute a basis for $\{\ell\} \cup \mathrm{radsp}(\overline{p})$. We have $\mathrm{rk}(S'_\ell) \leq k$. The set S_ℓ is obtained by adding arbitrary independent linear forms to S'_ℓ (if necessary) to make it have rank k. Set $B = \bigcup_{\ell \in L(T_i)} B(S_\ell)$. By Lemma 3.5, $|B(S_\ell)| \leq nk^2$ and hence $|B| \leq dnk^2$. Since $|U| > dnk^2$, we can fix some good β in U.

Consider the multiplication terms $p_1, p_2, \ldots, p_b, g$. For any ℓ dividing g, S_ℓ contains a basis for $\mathrm{radsp}(\overline{p}, g)$. We already argued that $g \notin \langle\overline{p}\rangle$. Since $\beta \notin B(S_\ell)$, by Lemma 4.1, $\Psi_\beta(g) \notin \langle\Psi_\beta(\overline{p})\rangle$. We finish with the following easy claim.

CLAIM 4.5. *Choose $\beta \in U \setminus B$. If $\Psi_\beta(C) = 0$, then $\Psi_\beta(g) \in \langle\Psi_\beta(\overline{p})\rangle$.*

Proof of Claim 4.5. We first argue that $\Psi_\beta(T_i) \in \langle\Psi_\beta(\overline{p})\rangle$. By the claim assumption and Equation (4.2) we get:

$$0 = \Psi_\beta(C) \equiv \alpha \cdot \Psi_\beta(T_i) \pmod{\Psi_\beta(\overline{p})}.$$

Since $\alpha \neq 0$, we deduce that $\Psi_\beta(T_i) \in \langle\Psi_\beta(\overline{p})\rangle$. Consider any form $\ell \in L(T_i)$ such that $\ell \notin \mathrm{radsp}(\overline{p})$. We will show that $\Psi_\beta(T_i)/\Psi_\beta(\ell) \in \langle\Psi_\beta(\overline{p})\rangle$. By the choice of ℓ, $\mathrm{rk}(\{\ell\} \cup \mathrm{radsp}(\overline{p})) = \mathrm{rk}(\mathrm{radsp}(\overline{p})) + 1$. Since $\beta \notin B$, $\mathrm{rk}(\{\Psi_\beta(\ell)\} \cup \mathrm{radsp}(\Psi_\beta(\overline{p}))) = \mathrm{rk}(\mathrm{radsp}(\Psi_\beta(\overline{p}))) + 1$. This implies $\Psi_\beta(\ell) \notin \mathrm{radsp}(\Psi_\beta(\overline{p}))$. Since $\Psi_\beta(T_i) \in \langle\Psi_\beta(\overline{p})\rangle$, Lemma 4.3 tells us that $\Psi_\beta(T_i)/\Psi_\beta(\ell) \in \langle\Psi_\beta(\overline{p})\rangle$. We can iteratively repeat this process for all such forms ℓ. We will end up with $\Psi_\beta(g) \in \langle\Psi_\beta(\overline{p})\rangle$. □

Thus, for any good β, $\Psi_\beta(C) \neq 0$. □

4.3. The final hitting-set: Proof of Theorem 1.3. Let (k, d, n) be the triple of natural numbers given in the input. We design a simple hitting-set $\mathcal{H}$. We will generate a set of vectors $\overline{\delta} \in \mathbb{F}^n$ that make $\mathcal{H}$.

- Let $S \subseteq \mathbb{F}$ be an arbitrary set of size $dnk^2 + 1$.
- Let $T \subseteq \mathbb{F}$ be an arbitrary set of size $d + 1$.
- For each $\beta \in S$ and each vector $(\gamma_1, \ldots, \gamma_k) \in T^k$, define the vector $\overline{\delta}$ in $\mathbb{F}^n$ as follows:

$$\delta_i := \sum_{j \in [k]} \beta^{ij} \gamma_j.$$

We state the classic [Sch80, Zip79, DL78] lemma before finishing the proof of Theorem 1.3.

THEOREM 4.6. *[Sch80, Zip79, DL78] Let $f(y_1, \ldots, y_k)$ be a polynomial of degree d. Let T be a finite subset of $\mathbb{F}$. The probability that f is zero on a random point in T^k is at most $d/|T|$. Thus, for $|T| > d$, T^k is a hitting-set for all k-variate polynomials of degree d.*

THEOREM 4.7. *The set $\mathcal{H}$ is a hitting-set for $\Sigma\Pi\Sigma(k, d, n)$ circuits. It can be generated in* $\mathrm{poly}(nd^k)$ *time.*

Proof. The latter statement is quite clear, given the construction of $\mathcal{H}$. Consider a nonzero $\Sigma\Pi\Sigma(k, d, n)$ circuit C. We need to show the existence of some $\overline{\delta} \in \mathcal{H}$ such that $C(\overline{\delta}) \neq 0$. By Lemma 4.4, there exists a $\beta \in S$ such that $\Psi_\beta(C) \neq 0$. Since $\Psi_\beta(C)$ is a $\Sigma\Pi\Sigma(k, d, k)$ circuit, Theorem 4.6 tells us that there is some $\overline{\gamma} \in T^k$ such that $\Psi_\beta(C)(\overline{\gamma}) \neq 0$. Consider the $\overline{\delta}$ corresponding to this β and $\overline{\gamma}$. By construction of $\overline{\delta}$ and the definition of Ψ_β (Equation 3.2), $C(\overline{\delta}) = \Psi_\beta(C)(\overline{\gamma}) \neq 0$. □

5. Conclusion. We show that a $\Sigma\Pi\Sigma(k, d, n)$ identity is only as complicated as a $\Sigma\Pi\Sigma(k, d, k)$ identity. We prove this fact by observing that there is a "low rank" homomorphism that preserves the ideal structure in depth-3 circuits. Since this low rank homomorphism is easily computable, we get a $\mathrm{poly}(nd^k)$ time blackbox test. Can we identify properties of k-variate fanin k identities to develop faster PIT algorithms? Currently, no PIT algorithm is able to beat the exponential dependence on k.

This work also raises a question for depth-4 circuits: Are there analogous low rank homomorphisms for $\Sigma\Pi\Sigma\Pi(k)$ circuits? Such results would open the door for interesting PIT algorithms for higher depth circuits. Recently, in [BMS11, ASSS12], partial progress has been made to design such homomorphisms based on the notion of *algebraic independence*. But a general variable reduction for $\Sigma\Pi\Sigma\Pi(k)$ circuits (when k is constant) is still missing.

Can this approach be used beyond PIT? In particular, there are results known about learning $\Sigma\Pi\Sigma(k)$ circuits where PIT methods have turned out to be useful [KS09a]. The variable reduction techniques might have some utility for these problems.

Acknowledgements. We are grateful to Hausdorff Center for Mathematics, Bonn for its kind support, especially in hosting the second author when part of the work was done. The first author thanks Malte Beecken and Johannes Mittmann for several interesting discussions. We thank the anonymous referees for the detailed suggestions that improved the presentation.

REFERENCES

[AB03] M. Agrawal and S. Biswas. Primality and identity testing via Chinese remaindering. *J. ACM*, 50(4):429–443, 2003. (Conference version in FOCS 1999).

[Agr05] M. Agrawal. Proving lower bounds via pseudo-random generators. In *Proceedings of the 25th Annual Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 92–105, 2005.

[Agr06] M. Agrawal. Determinant versus permanent. In *Proceedings of the 25th International Congress of Mathematicians (ICM)*, volume 3, pages 985–997, 2006.

[AL86] L. M. Adleman and H. W. Lenstra. Finding irreducible polynomials over finite fields. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC)*, pages 350–355, 1986.

[AM10] V. Arvind and P. Mukhopadhyay. The ideal membership problem and polynomial identity testing. *Inf. Comput.*, 208(4):351–363, 2010. (Conference version in ISAAC 2007).

[AS09] M. Agrawal and R. Saptharishi. Classifying polynomials and identity testing. In *Current Trends in Science, Indian Academy of Sciences*, pages 149–162. 2009.

[ASSS12] M. Agrawal, C. Saha, R. Saptharishi, and N. Saxena. Jacobian hits circuits: Hitting-sets, lower bounds for depth-D occur-k formulas & depth-3 transcendence degree-k circuits. In *Proceedings of the Annual 44th ACM Symposium on Theory of Computing (STOC)*, 2012. http://eccc.hpi-web.de/report/2011/143/.

[AV08] M. Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *Proceedings of the 49th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 67–75, 2008.

[AvMV11] M. Anderson, D. van Melkebeek, and I. Volkovich. Derandomizing polynomial identity testing for multilinear constant-read formulae. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity (CCC)*, pages 273–282, 2011.

[BMS11] M. Beecken, J. Mittmann, and N. Saxena. Algebraic independence and blackbox identity testing. In *Proceedings of the 38th Annual International Colloquium on Automata, Languages and Programming (ICALP)*, pages 137–148, 2011.

[BOT88] M. Ben-Or and P. Tiwari. A deterministic algorithm for sparse multivariate polynominal interpolation. In *Proceedings of the 20th Annual ACM symposium on Theory of computing (STOC)*, pages 301–309, 1988.

[CDGK91] M. Clausen, A. W. M. Dress, J. Grabmeier, and M. Karpinski. On zero-testing and interpolation of k-sparse multivariate polynomials over finite fields. *Theoretical Computer Science*, 84(2):151–164, 1991.

[CK97] Z. Chen and M. Kao. Reducing randomness via irrational numbers. In *Proceedings of the 29th Annual ACM symposium on Theory of computing (STOC)*, pages 200–209, 1997.

[DL78] R. A. DeMillo and R. J. Lipton. A Probabilistic Remark on Algebraic Program Testing. *Inf. Process. Lett.*, 7(4):193–195, 1978.

[DS06] Z. Dvir and A. Shpilka. Locally decodable codes with 2 queries and polynomial identity testing for depth 3 circuits. *SIAM J. Comput.*, 36(5):1404–1434, 2006. (Conference version in STOC 2005).

[GKS90] D. Grigoriev, M. Karpinski, and M. F. Singer. Fast parallel algorithms for sparse multivariate polynomial interpolation over finite fields. *SIAM J. Comput.*, 19(6):1059–1063, 1990.

[GR08] A. Gabizon and R. Raz. Deterministic extractors for affine sources over large fields. *Combinatorica*, 28(4):415–440, 2008. (Conference version in FOCS 2005).

[HS80] J. Heintz and C. P. Schnorr. Testing polynomials which are easy to compute (extended abstract). In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing*, pages 262–272, 1980.

[KI04] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1):1–46, 2004. (Conference version in STOC 2003).

[KMSV10] Z. Karnin, P. Mukhopadhyay, A. Shpilka, and I. Volkovich. Deterministic identity testing of depth-4 multilinear circuits with bounded top fan-in. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pages 649–658, 2010.

[KS96] M. Karpinski and I. Shparlinski. On some approximation problems concerning sparse polynomials over finite fields. *Theoretical Computer Science*, 157(2):259–266, 1996.

[KS01] A. Klivans and D. Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the 33rd Symposium on Theory of Computing (STOC)*, pages 216–223, 2001.

[KS07] N. Kayal and N. Saxena. Polynomial identity testing for depth 3 circuits. *Computational Complexity*, 16(2):115–138, 2007. (Conference version in CCC 2006).

[KS08] Z. Karnin and A. Shpilka. Deterministic black box polynomial identity testing of depth-3 arithmetic circuits with bounded top fan-in. In *Proceedings of the 23rd Annual Conference on Computational Complexity (CCC)*, pages 280–291, 2008.

[KS09a] Z. Karnin and A. Shpilka. Reconstruction of generalized depth-3 arithmetic circuits with bounded top fan-in. In *Proceedings of the 24th Annual Conference on Computational Complexity (CCC)*, pages 274–285, 2009.

[KS09b] N. Kayal and S. Saraf. Blackbox polynomial identity testing for depth 3 circuits. In *Proceedings of the 50th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 198–207, 2009.

[LV98] D. Lewin and S. Vadhan. Checking polynomial identities over any field: Towards a derandomization? In *Proceedings of the 30th Annual Symposium on the Theory* *of Computing (STOC)*, pages 428–437, 1998.

[Raz10] R. Raz. Tensor-rank and lower bounds for arithmetic formulas. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 659–666, 2010.

[RS05] R. Raz and A. Shpilka. Deterministic polynomial identity testing in non-commutative models. *Computational Complexity*, 14(1):1–19, 2005. (Conference version in CCC 2004).

[Sax08] N. Saxena. Diagonal circuit identity testing and lower bounds. In *Proceedings of the 35th Annual International Colloquium on Automata, Languages and Programming (ICALP)*, pages 60–71, 2008.

[Sax09] N. Saxena. Progress on polynomial identity testing. *Bulletin of the European Association for Theoretical Computer Science (EATCS)- Computational Complexity Column*, (99):49–79, 2009.

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

[SS96] R. E. Schapire and L. M. Sellie. Learning sparse multivariate polynomials over a field with queries and counterexamples. *Journal of Computer and System Sciences*, 52(2):201–213, 1996.

[SS10a] N. Saxena and C. Seshadhri. From Sylvester-Gallai configurations to rank bounds: Improved black-box identity test for depth-3 circuits. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 21–29, 2010.

[SS10b] N. Saxena and C. Seshadhri. From Sylvester-Gallai configurations to rank bounds: Improved black-box identity test for depth-3 circuits. Technical Report TR10-013 (Revision #1), ECCC, http://eccc.hpi-web.de/report/2010/013/, 2010.

[SS11a] N. Saxena and C. Seshadhri. An almost optimal rank bound for depth-3 identities. *SIAM J. Comp.*, 40(1):200–224, 2011. (Conference version in CCC 2009).

[SS11b] N. Saxena and C. Seshadhri. Blackbox identity testing for bounded top fanin depth-3 circuits: the field doesn't matter. In *Proceedings of the 43rd annual ACM Symposium on Theory of Computing (STOC)*, pages 431–440, 2011.

[SSS09] C. Saha, R. Saptharishi, and N. Saxena. The power of depth 2 circuits over algebras. In *Proceedings of the 29th Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 371–382, 2009.

[SSS12] C. Saha, R. Saptharishi, and N. Saxena. A case of depth-3 identity testing, sparse factorization and duality. *Computational Complexity*, 2012. http://eccc.hpi-web.de/report/2011/021/.

[SV08] A. Shpilka and I. Volkovich. Read-once polynomial identity testing. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 507–516, 2008.

[SV09] A. Shpilka and I. Volkovich. Improved polynomial identity testing for read-once formulas. In *Proceedings of the 13th International Workshop on Randomization and Computation (RANDOM)*, pages 700–713, 2009.

[SV11] S. Saraf and I. Volkovich. Black-box identity testing of depth-4 multilinear circuits. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC)*, pages 421–430, 2011.

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

[Wer94] K. Werther. The complexity of sparse polynomial interpolation over finite fields. *Applicable Algebra in Engineering, Communication and Computing*, 5:91–103, 1994.

[Zip79] R. Zippel. Probabilistic algorithms for sparse polynomials. In *Proceedings of the International Symposium on Symbolic and Algebraic Manipulation (EUROSAM)*, pages 216–226, 1979.

Appendix A. A Vandermonde-inspired linear transformation.

Lemma 3.5. *Let $S \subseteq L(\mathcal{R})$ be a subset of linear forms with $\mathrm{rk}(S) \le k$. Then $|B(S)| \le nk^2$.*

The homomorphism Ψ_β depends solely on $\mathbb{F}, k, n$. If β has a bit representation of at most $O(kdn)$ bits, then Ψ_β is computable in $\mathrm{poly}(kdn)$ time.

Proof. We can assume wlog that S is of rank k, and has k linearly independent forms $\ell_1, \ldots, \ell_k \in L(\mathcal{R})$. Say Ψ_β maps $\ell_i = \sum_{j\in[n]} a_{i,j}x_j$ to $\sum_{j\in[k]} b_{i,j}y_j$, for all $i \in [k]$. Define matrices $B := ((b_{i,j}))_{i\in[k],j\in[k]}$ and $A := ((a_{i,j}))_{i\in[k],j\in[n]}$. By Equation (3.1) we deduce $B = A \cdot V_{n,k}$. We will show that B is invertible.

Since the rows of A are linearly independent over $\mathbb{F}$, we can apply *partial* Gaussian elimination (i.e. row operations) on A. This has the effect of *left*-multiplying A by an invertible matrix $E \in \mathbb{F}^{k\times k}$ ensuring: there are column indices $j_1 > \ldots > j_k \in [n]$ such that j_i is the *maximal* index with $(EA)_{i,j_i} \neq 0$, for all $i \in [k]$. Now we consider the matrix $B' := (EA) \cdot V_{n,k}$,

$$\det(B') = \sum_{\sigma\in S_k} \mathrm{sgn}(\sigma) \cdot P_\sigma(\beta), \quad \text{where } P_\sigma(\beta) := \prod_{i\in[k]} B'_{i,\sigma(i)}.$$

Note that we view $B'_{i,\sigma(i)}$ as a polynomial in $\mathbb{F}[\beta]$ which, by the assumption on EA, is of degree $j_i \cdot \sigma(i)$. Thus,

$$\deg(P_\sigma(\beta)) = \sum_{i\in[k]} j_i \cdot \sigma(i).$$

Since $j_1 > \ldots > j_k$, it can be easily shown that the expression above achieves its maxima (over $\sigma \in S_k$) only if $\sigma(1) > \ldots > \sigma(k)$. But this uniquely specifies σ, hence there is a unique $P_\sigma(\beta)$ of the largest degree ($\leq nk^2$). Thus, $\det(B')$ is a nonzero polynomial in $\mathbb{F}[\beta]$ of degree at most nk^2. This means that B' is invertible, hence $B = E^{-1}B'$ is invertible, for all but at most nk^2 values of β. ∎

Appendix B. A cancellation lemma.

An $f \in \mathcal{R}$ is called a *zerodivisor* of an ideal I (or mod I) if $f \notin I$ and there exists a $g \in \mathcal{R} \setminus I$ such that $fg \in I$. Let $f, g \in \mathcal{R}$. It is easy to see that if f is nonzero mod I and is a *non*-zerodivisor mod I then: $fg \in I$ iff $g \in I$. This can be seen as some sort of a "cancellation rule" for non-zerodivisors. We show such a cancellation rule in the case of ideals arising in $\Sigma\Pi\Sigma$ circuits.

Lemma 4.3. *Let $f_1, \ldots, f_m$ be multiplication terms generating an ideal I, let $\ell \in L(\mathcal{R})$ and $g \in \mathcal{R}$. If $\ell \notin \mathrm{radsp}(I)$ then: $\ell g \in I$ iff $g \in I$.*

Proof. Suppose $\ell \notin \mathrm{radsp}(I)$. If $I = \{0\}$, then the lemma is trivially true. Assume that $I \neq \{0\}$ and $\mathrm{rk}(\mathrm{radsp}(\ I)) =: r \in [n-1]$. Since $\ell \notin \mathrm{radsp}(I)$, there exists an invertible linear transformation $\tau : L(\mathcal{R}) \to L(\mathcal{R})$ that maps each form of $\mathrm{radsp}(I)$ to $\mathrm{sp}(x_1, \ldots, x_r)$ and maps ℓ to x_n. Now suppose that $\ell g \in I$. This means that there are $q_1, \ldots, q_m \in \mathcal{R}$ such that $\ell g = \sum_{i=1}^m q_i f_i$. Apply τ on this to get:

$$x_n g' = \sum_{i=1}^{m} q'_i \tau(f_i). \tag{B.1}$$

We know that $\tau(f_i)$'s are free of x_n. Express g', q'_i-s as polynomials wrt x_n, say

$$g' = \sum_{j\geq 0} a_j x_n^j, \ \text{ where } a_j \in \mathbb{F}[x_1, \ldots, x_{n-1}] \tag{B.2}$$

$$q'_i = \sum_{j\geq 0} b_{i,j} x_n^j, \ \text{ where } b_{i,j} \in \mathbb{F}[x_1, \ldots, x_{n-1}] \tag{B.3}$$

For $d \geq 1$, compare the coefficients of x_n^d on both sides of Equation (B.1). We get $a_{d-1} = \sum_{i=1}^m b_{i,d}\tau(f_i)$, thus a_{d-1} and $a_{d-1}x_n^{d-1}$ are in $\langle \tau(f_1), \ldots, \tau(f_m)\rangle$. Doing this for all $d \geq 1$, we get $g' \in \langle \tau(f_1), \ldots, \tau(f_m)\rangle$. Hence, $g = \tau^{-1}(g') \in \langle f_1, \ldots, f_m\rangle = I$. This finishes the proof. ▯

Appendix C. Details of Theorem 3.3.

Although it does not appear to be so, Theorem 3.3 is really a consequence of Chinese Remaindering for depth-3 circuits. We will provide an overview of the proof given in [SS10b, Theorem 25]. For more details, the reader should naturally look at the original paper. The technical heart of the proof is given in the following ideal decomposition statement.

THEOREM C.1 (Ideal Chinese remaindering). *Let $f_1, \ldots, f_m, f, g$ be multiplication terms. Define the ideal $I := \langle f_1, \ldots, f_m\rangle$. Assume that $L(g) \cap \mathrm{radsp}(I, f) = \emptyset$. Then, $\langle I, fg\rangle = \langle I, f\rangle \cap \langle I, g\rangle$.*

Proof sketch: Suppose that $L(f) \cap \mathrm{radsp}(I) = \emptyset$. If h is a polynomial in $\langle I, fg\rangle$ then clearly it is in each of the ideals $\langle I, f\rangle$ and $\langle I, g\rangle$. Suppose $h \in \langle I, f\rangle \cap \langle I, g\rangle$. Then $h = i_1 + bf = i_2 + cg$, for $i_1, i_2 \in I$ and $b, c \in \mathcal{R}$. Hence, $cg \in \langle I, f\rangle$. Since $L(g) \cap \mathrm{radsp}(I, f) = \emptyset$, we can repeatedly apply Lemma 4.3 and deduce that $c \in \langle I, f\rangle$. This implies that $cg \in g\langle I, f\rangle \subseteq \langle I, fg\rangle$, and we complete the proof. □

We can now split the ideal generated by multiplication terms into "constituent factors". Consider an ideal I and a multiplication term f. For any $\ell \in L(f)$, consider the set of forms $\{\ell' \in L(f) | \ell' \text{ is similar to } \ell \bmod \mathrm{radsp}(I)\}$. The product of these forms is a polynomial called a *node of f with respect to I*. Observe how f is the product of its nodes, and each node is a (maximal) product of forms that are similar modulo $\mathrm{radsp}(I)$. The set of nodes is denoted by $\mathrm{nod}_I(f)$. A repeated application of Theorem C.1 gives the following theorem, and its fairly obvious corollary.

THEOREM C.2. *Let I be an ideal and f a multiplication term. Let the set $\mathrm{nod}_I(f)$ be $\{g_1, \ldots, g_r\}$. Then*

$$\langle I, f\rangle = \bigcap_{i \in [r]} \langle I, g_i\rangle.$$

COROLLARY C.3. *Let $h \in \mathcal{R}$, f be a multiplication term, and let I be an ideal generated by some multiplication terms. Then, $h \notin \langle I, f\rangle$ iff $\exists g \in \mathrm{nod}_I(f)$ such that $h \notin \langle I, g\rangle$.*

We introduce yet another definition, that of *paths* of nodes in C. A path $\overline{p}$ with respect to an ideal I is a sequence of terms $\{p_1, p_2, \ldots, p_b\}$ with the following property. Each p_i divides T_i, and each p_i is a node of T_i with respect to the ideal $\langle I, p_1, p_2, \ldots, p_{i-1}\rangle$. So p_1 is a node of T_1 wrt I, p_2 is a node of T_2 wrt $\langle I, p_1\rangle$, etc.

Let us see an example of a path $(\langle 0\rangle, p_1, p_2, p_3)$ in Figure C.1. The oval bubbles represent the list of forms in a term, and the rectangles enclose forms in a node. The arrows show a path. Starting with I as the zero ideal, nodes $p_1 := x_1^2$, $p_2 := x_2(x_2 + 2x_1)$, and $p_3 := (x_4 + x_2)(x_4 + 4x_2 - x_1)(x_4 + x_2 + x_1)(x_4 + x_2 - 2x_1)$ form a path. Initially the path is just the zero ideal, so x_1^2 is a node. Note how p_2 is a power of x_2 modulo $\mathrm{radsp}(p_1)$, and p_3 is a power of x_4 modulo $\mathrm{radsp}\langle p_1, p_2\rangle$.

The non-identity certificate theorem states that for any non-identity C, there exists a path $\overline{p}$ such that modulo $\langle \overline{p}\rangle$, C reduces to a single nonzero multiplication term. The following directly implies Theorem 3.3. Simply set I to $\{0\}$ and note that $\mathrm{radsp}(\overline{p})$ has dimension at most $k - 1$.

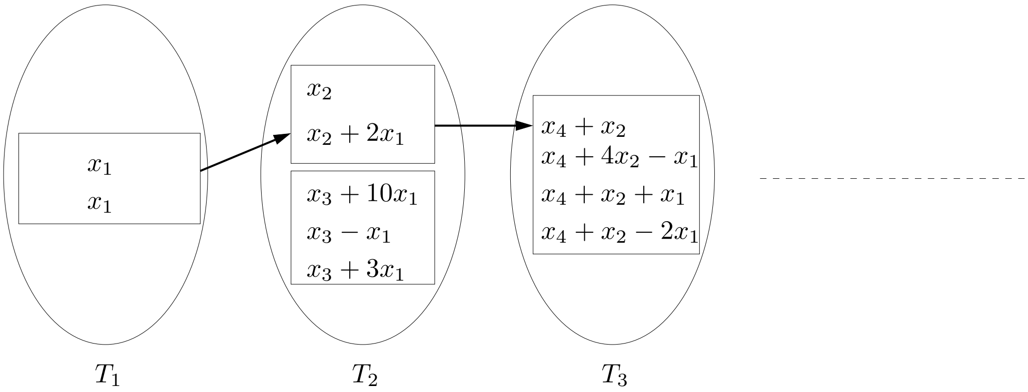

Fig. C.1: Nodes and paths in $C = T_1 + T_2 + T_3 + \dots$

Theorem C.4 (Certificate for a non-identity). *Let I be an ideal generated by some multiplication terms. Let $C = \sum_{i\in[k]} T_i$ be a $\Sigma\Pi\Sigma(k,d,n)$ circuit that is nonzero modulo I. Then $\exists i \in \{0,\dots,k-1\}$ such that $C_{[i]}$ mod I has a path $\overline{p}$ satisfying: $C_{[i]'} \equiv \alpha \cdot T_{i+1} \not\equiv 0 \pmod{\overline{p}}$ for some $\alpha \in \mathbb{F}^*$.*

Proof sketch: We start with the fact that $C = T_1 + \cdots + T_k$ is not in I. If $C_{[1]'} \in \langle I, T_1\rangle$, then C reduces to a single multiplication term modulo I (and we are done). So assume that $C_{[1]'} \notin \langle I, T_1\rangle$. By Corollary C.3, there is a node $g_1 \in \mathrm{nod}_I(T_1)$ such that $C_{[1]'} \notin \langle I, g_1\rangle$. This provides the first node in the final path. We now repeat this argument with the circuit $C_{[1]'}$ and the ideal $\langle I, g_1\rangle$, and get a second node g_2. This continues to yield the final path. $\square$