Electronic Colloquium on Computational Complexity, Report No. 132 (2022)

Fooling polynomials using invariant theory

Harm Derksen* Emanuele Viola†

September 18, 2022

Abstract

We revisit the problem of constructing explicit pseudorandom generators that fool with error ϵ degree-d polynomials in n variables over the field $\mathbb{F}_q$, in the case of large q. Previous constructions either have seed length $\geq 2^d \log q$, and thus are only non-trivial when $d < \log n$, or else rely on a seminal reduction by Bogdanov (STOC 2005). This reduction yields seed length not less than $d^4 \log n + \log q$ and requires fields of size $q \geq d^6/\epsilon^2$; and explicit generators meeting such bounds are known.

Departing from Bogdanov's reduction, we develop an algebraic analogue of the Bogdanov-Viola paradigm (FOCS 2007, SICOMP 2010) of summing generators for degree-one polynomials. Whereas previous analyses of the paradigm are restricted to degree $d < \log n$, we give a new analysis which handles large degrees. A main new idea is to show that the construction preserves indecomposability of polynomials. Apparently for the first time in the area, the proof uses invariant theory.

Our approach in particular yields several new pseudorandom generators. In particular, for large enough fields we obtain seed length $O(d \log n + \log q)$ which is optimal up to constant factors. We also construct generators for fields of size as small as $O(d^4)$. Further reducing the field size requires a significant change in techniques: Most or all generators for large-degree polynomials rely on Weil bounds; but such bounds are only applicable when $q > d^4$.

A *pseudorandom generator* for degree-d polynomials over the field $\mathbb{F}_q$ in n variables with error ϵ is an explicit map $P : S \to \mathbb{F}_q^n$ that "ϵ-fools" any such polynomial g, that is, the distributions $g(U)$ and $g(P(U))$ have *statistical distance* (or *error*) at most ϵ. Here U denotes the uniform distribution over the appropriate domain ($\mathbb{F}_q^n$ in the first occurrence and S in the second). The *seed length* of P is $\log_2 |S|$. The minimum possible seed length is $\Omega(d \log(n/d) + \log q + \log 1/\epsilon)$, at least when $d < n^{0.99}$ and q is prime [BV10, ABEK08]. Explicit constructions of generators (i.e., upper bounds on the seed length) have been intensely studied for at least 30 years. Two main lines of work exist. The first applies to any field [NN93, AGHP92, LVW93, Vio07, BV10, Lov09, Vio09]. The last paper gives seed length $O(\log n + 2^d \log q/\epsilon) \cdot d$ which is the best available for small fields such as $\mathbb{F}_2$. The corresponding generators are obtained within the Bogdanov-Viola paradigm [BV10]: to fool polynomials of degree d, sum $\ell \geq d$ independent copies of generators for degree-one polynomials. While the parameters given by

*Partially supported by NSF grant DMS 2147769.
†Supported by NSF grant CCF-2114116.

ISSN 1433-8092

the analysis in [Vio09] are non-trivial only for $d \leq \log n$, it is unknown whether the paradigm also works for larger degrees. If it did it would yield a breakthrough in complexity theory. For example, it would imply generators for small constant-depth circuits with parity gates, thanks to a well-known approximation due to Razborov [Raz87].

The second lines of works applies only to fields of large size $q \gg d$, but can handle much larger degrees. Here Bogdanov's seminal paper [Bog05] laid a paradigm that reduces constructing pseudorandom generators to constructing *hitting-set* generators for polynomials, an easier task. Bogdanov's paper was followed by a series of better and better constructions of hitting-set generators by Lu [Lu12], Cohen and Ta-Shma [CT13], and Guruswami and Xing [GX14]; see also [KS01] for earlier related work by Klivans and Spielman. Optimal hitting-set constructions are now known; in combination with Bogdanov's reduction they yield the following pseudorandom generators.

Theorem 1. *[[Bog05]+[GX14]+([Lu12] or [KS01])] There exist explicit pseudorandom generator that fool degree-d polynomials in n variables over $\mathbb{F}_q$ with seed length $O(d^4 \log n + \log q)$, provided $q \geq O(d^6/\epsilon^2)$.*

The notation $O(.)$ and $\Omega(.)$ denotes absolute constants. To connect with previous expressions for the seed length, note that adding a $\log 1/\epsilon$ term to the seed length in Theorem 1 does not change it since $q \geq 1/\epsilon$.

The parameters in Theorem 1 are essentially the best one can achieve using the reduction in [Bog05], as we now explain. That reduction proceeds by showing that restricting a polynomial g onto a "good" plane preserves its output distribution with high probability. Once a good plane is found, one can then just pick a uniform element from the plane, which only costs two field elements. To find a good plane, [Bog05] relies on results by Kaltofen [Kal95] showing that (the coefficients of) planes that are bad for g are zeroes of a low-degree polynomial K_g. One can then use a hitting set to find a good plane. A bottleneck in this reduction is that the degree of K_g is at least d^4. So one needs a hitting-set generator for polynomials of degree at least d^4, resulting in the d^4 factor in the final seed length. The same loss arises in earlier work dealing with polynomials over complex numbers, see [Kal95] for discussion. Over fields of large characteristic the degree can be improved from $O(d^4)$ to $O(d^2)$, which is known to be optimal, see [Lec07]. Thus, this approach does not yield seed length less than $d^2 \log n$. For related reasons, the reduction in [Bog05] requires the field size to be at least d^6.

Constructions of pseudorandom generators in the two lines of research above have followed different paradigms. By contrast, we shall prove that the [BV10] paradigm works also for large-degree polynomials, at least as long as the field is large enough. This in particular yields pseudorandom generators with improved parameters, stated next.

Theorem 2. *There are explicit pseudorandom generators that fool with error ϵ degree-d polynomials in n variables over $\mathbb{F}_q$ with seed length $O(d \cdot m \cdot \log(dk + dm) + \log q)$, provided that $q \geq O(dk)^4/\epsilon^2$, for any integers m and k such that $\binom{m+k-2}{m-1} \geq n$.*

In particular we can have either

(1) seed length $O(d \log(dn) + \log q)$ provided that $q \geq O(d^4 n^{0.001})/\epsilon^2$, or

(2) seed length $O(d \log n \cdot \log(d \log n) + \log q)$ provided that $q \geq O(d \log n)^4/\epsilon^2$.

Item (1) achieves optimal seed length up to constant factors, when $d < n^{0.99}$. In particular it improves on the $\Omega(d^4 \log n)$ seed lengths of previous constructions. The field size improves on the $\Omega(d^6/\epsilon^2)$ field size of previous constructions (Theorem 1) when say $d > n^{0.001}$. This item is obtained by suitably setting $m = O(1)$ and $k = n^{\Omega(1)}$.

Item (2) achieves optimal seed length up to the lower-order factor $\log(d \log n)$. The field size improves on previous constructions for $d \geq \omega(\log^2 n)$. This item is obtained by setting $m = O(\log n)$ and $k = O(\log n)$.

We also obtain pseudorandom generators with the same seed length as previous constructions, but that only require $q \geq O(d^4)$, see Theorem 21. This improves on the $\Omega(d^6)$ field size of previous constructions. Further reducing the field size will require a significant change in techniques: Most or all generators for large-degree polynomials rely on Weil bounds, cf. Fact 15 or [Sch04, Page 92]; but such bounds are only applicable when $q > d^4$.

Proof overview. A central concept in our proof, which was apparently not used before in the pseudorandomness literature, is that of *indecomposability*.

Definition 3. A polynomial g over a field $\mathbb{F}$ is *indecomposable* if it cannot be written as $c \circ h$ where c is a univariate polynomial of degree ≥ 2 and both c and h are over $\mathbb{F}$.

Let g be a polynomial we aim to fool. We begin by writing $g = c(h)$ where c is a univariate polynomial of maximal degree. We observe that the polynomial h is indecomposable, for else the degree of c is not maximal. A main technical contribution (discussed more below) is a universal (i.e., independent from g) construction of polynomials $f_1, f_2, \ldots, f_n$ that (i) are on few variables, (ii) have low degree, and (iii) *preserve indecomposability:* if $h(f_1, f_2, \ldots, f_n)$ is decomposable, then so is $h(x_1, x_2, \ldots, x_n)$. As observed above, the latter is not decomposable; hence the former is not decomposable either. We then prove (Lemma 12 in Section 2) that the output distribution of indecomposable polynomials is close to uniform. This proof combines several results in algebraic geometry, including Weil's bound and results about reducibility of shifts of indecomposable polynomials.

Putting the above together we conclude that the f_i fool g because

$$g(U) = c(h(U)) \approx c(U) \approx c(h(f_1, f_2, \ldots, f_n))(U) = g(f_1, f_2, \ldots, f_n)(U).$$

Hence we have reduced the problem of fooling g to that of fooling g composed with the f_i. The gain is that by (i) we have reduced the number of variables. The main cost is an increase in degree, but this increase is small by (ii). Overall we obtain the following result, which is a main technical contribution of this work.

Theorem 4. *For every positive integers n, d, k and field $\mathbb{F}_q$:*

There is an explicit family of degree-k polynomials $f_1, f_2, \ldots, f_n$ over $\mathbb{F}_q$ in $(d+1)m$ variables such that for any polynomial g over $\mathbb{F}_q$ of degree d in n variables the statistical distance between $g(U)$ and $g(f_1, f_2, \ldots, f_n)(U)$ is $O(d^2k^2/\sqrt{q})$, for any m and k as in Theorem 2.

If we plug uniform values for the variables of the f_i we obtain pseudorandom generators with seed length as in Theorem 2 except that the factor $\log(dk + dm)$ is replaced with $\log q$. This is sufficient to prove the theorem when q is polynomial in dn. If q is larger, for example $q \geq 2^d$, it is not sufficient, and we need to improve the dependence on q from multiplicative to

additive. To achieve this we combine Theorem 4 with another pseudorandom generator which we construct (Theorem 21). The latter generator combines Bogdanov's template [Bog05] discussed earlier with some of our proof ideas. Compared with [Bog05] and subsequent works, this generator has two main differences. First, we give a variant of Bogdanov's reduction of pseudorandom to hitting-set generators, again relying on preserving indecomposability. This allows us to improve the dependence on the field size. Note however that one can already obtain non-trivial generators over fields of size $O(d^4)$ from Theorem 4 (suitably set $k = O(1)$ and $m = n^{\Omega(1)}$). Second, we need to hit polynomials whose degree is larger than the number of variables, whereas in most previous work the degree is smaller. We note that such a hitting set can be obtained by combining [Lu12, GX14].

The construction of the f_i and its analysis using invariant theory. Let $M_1, M_2, \ldots$ be an enumeration of distinct monomials of degree k in m variables (in some cases we need some mild conditions on these monomials, discussed below). We take ℓ copies of the variables, and define $f_i := M_i^{[1]} + M_i^{[2]} + \ldots + M_i^{[\ell]}$ where $M_i^{[j]}$ is the monomial M_i where the variables are taken from copy j. Hence the construction is simple and very explicit.

The proof that the f_i preserve indecomposability uses *invariant theory*, apparently for the first time in this area, and proceeds as follows. Consider the polynomial $G := g(f_1, f_2, \ldots, f_n)$. First, note that G is *invariant* under permutation of the copies of variables (simply because the f_i are). Now assume that G can be decomposed as $G = c(H)$ for some univariate polynomial c. We show that H must be invariant as well. Next, we show that the f_i are a basis for the invariant polynomials; this allows us to write $H = h(f_1, f_2, \ldots, f_s)$ for some low-degree polynomial h, where note a priori s could be much larger than n. Hence we obtained

$$g(f_1, f_2, \ldots, f_n) = c(h(f_1, f_2, \ldots, f_s)).$$

Finally, we show that this implies $s = n$ and $g(x_1, x_2, \ldots, x_n) = c(h(x_1, x_2, \ldots, x_n))$ as desired.

Three results on preserving indecomposability. We give three formal versions of the analysis in the previous paragraph.

The first version (Theorem 6 in Section 1) has the easiest proof, requires fields of characteristic $> dk$, and takes $\ell > dk$ copies of variables. This version suffices to obtain generators with seed length $\tilde{O}(d \log^2 n) + O(\log q)$ over such fields, where $\tilde{O}(x)$ stands for $x \log^{O(1)} x$. Using the construction recursively, one can improve the seed length to $\tilde{O}(d \log n) + O(\log q)$, thus matching Item (2) in Theorem 2 up to lower-order factors for large characteristic, and in particular for prime fields. However these ideas do not suffice to obtain the optimal seed length in Item (1), for example. For this first version we can take any distinct monomials. This version also allows us to draw a close analogy with the Bogdanov-Viola paradigm [BV10]: We note that one can replace the M_i with any set of polynomials N_i of the same degree that fool degree-one polynomials. To verify this we can write the N_i as linear combinations of the M_i and use that the linear maps are full rank since the N_i fool degree-one polynomials.

The second version (Theorem 16 in Section 4) has a slightly more complicated proof, but requires only characteristic $> d$ and more importantly takes only $\ell = d + 1$ copies. This

essentially matches the number $\ell = d$ of copies in [BV10, Vio09]. For this we need a certain mild condition on the monomials. This version suffices to prove Theorem 2 for fields of characteristic $> d$, and in particular for prime fields.

The third version (Theorem 27 in Section 7) is the most complicated but works over any characteristic, and again takes only $\ell = d + 1$ copies. Here we need to avoid obvious counterexamples; for example over $\mathbb{F}_2$ we cannot take $M_1 = x^2$ because $g = x$ is trivially indecomposable but $g(f_1) = (x^{[1]})^2 + (x^{[2]})^2 + \cdots + (x^{[\ell]})^2 = (x^{[1]} + x^{[2]} + \cdots + x^{[\ell]})^2$ is decomposable. It turns out that it suffices to take any M_i that are indecomposable. This version can be used to prove Theorem 4 as stated, for fields of any characteristic. Besides this, the results in this section allow us to preserve indecomposability over any field, even small. The only restriction on the field size then comes from Weil's bound (cf. Fact 15).

Open problems. A natural goal is to reduce the field size in Item (1) in Theorem 2 to $O(d^4)$. This would yield a single generator that improves on all those in this paper. The current bounds on the field size arise from applying Weil's bound to polynomials of degree dk rather than d. However, these polynomials of degree dk have a special structure as they arise from the composition of an arbitrary polynomial of degree d with the M^{Σ}'s. It is conceivable that Weil's bound can be improved for such composed polynomials, perhaps to obtain bounds similar to those for degree-d polynomials. We raise this as an open problem.

1 Preserving indecomposability

In this section we give a first construction of polynomials that preserve indecomposability. We state the main theorem next after some notation. Then we proceed with the proof which involves several intermediate claims.

Let $\mathbb{F}_q$ be a field of characteristic p and let $R = \mathbb{F}_q[x_1, x_2, \ldots, x_m]$ be the polynomial ring in m variables. We define $R^{\otimes \ell} = \mathbb{F}_q[\{x_j^{[i]}\}]$ as the polynomial ring in the variables $x_j^{[i]}$ with $1 \le i \le \ell$ and $1 \le j \le m$. We can arrange the $\ell \cdot m$ variables in a matrix

$$X = \begin{pmatrix} x_1^{[1]} & x_2^{[1]} & \cdots & x_m^{[1]} \\ x_1^{[2]} & x_2^{[2]} & \cdots & x_m^{[2]} \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{[\ell]} & x_2^{[\ell]} & \cdots & x_m^{[\ell]} \end{pmatrix} \tag{1}$$

Definition 5. A *monomial* is a product of powers of variables (with leading coefficient 1). For a monomial $M = M(x_1, x_2, \ldots, x_m) \in R = \mathbb{F}_q[x_1, x_2, \ldots, x_m]$ we define $M^{[i]} = M(x_1^{[i]}, x_2^{[i]}, \ldots, x_m^{[i]})$ and $M^{\Sigma} = \sum_{i=1}^{\ell} M^{[i]}$.

Theorem 6. *Suppose that $M_1, M_2, \ldots, M_r \in R$ are distinct non-constant monomials of degree $\le k$, and let $g(x_1, x_2, \ldots, x_r)$ be a non-constant polynomial of degree d. Let $G := g(M_1^{\Sigma}, M_2^{\Sigma}, \ldots, M_r^{\Sigma})$ and assume that $p \ge dk + 1$ and $\ell \ge \max\{5, dk + 1\}$. If G is decomposable then g is decomposable.*

We remark that $\ell = d$ does not suffice for example for $d = 1$ and $k = 2$: take $M_1 = x_1^2$.

The rest of this section is devoted to proving the theorem. We view the *symmetric group* S_ℓ of permutations on ℓ elements as acting on $R^{\otimes \ell}$ by $\sigma(x_j^{[i]}) = x_j^{[\sigma(i)]}$ for all i, j. So the action of S_ℓ permutes the rows in X.

Definition 7. For a monomial M in $R^{\otimes \ell}$, the *diversity* of M is the smallest number d such that the variables in M come from d rows in X. For a nonzero polynomial $f \in R^{\otimes \ell}$, the diversity $\mathrm{div}(f)$ is the largest diversity over all monomials appearing in f.

For a subgroup G of S_ℓ and a polynomial $g \in R^{\otimes \ell}$ we say that g is G-invariant if $\sigma g = g$ for all $\sigma \in G$. Note that M^Σ is invariant under the action of S_ℓ and that $\mathrm{div}(M^\Sigma) = 1$ when M is not constant. In general we have the following proposition.

Proposition 8. *Suppose that $f \in R^{\otimes \ell}$ is an S_ℓ-invariant polynomial with $\mathrm{div}(f) = d$ and $p > d$. Then f can be written as a polynomial of degree d in the M^Σ's.*

Proof. The *orbit sum* of a monomial $M := M_1^{[1]} M_2^{[2]} \cdots M_\ell^{[\ell]}$ where the M_i are in R is the sum of all monomials in the S_ℓ orbit $\{M_1^{[\pi 1]} M_2^{[\pi 2]} \cdots M_\ell^{[\pi \ell]} : \pi \in S_\ell\}$ of M. We note that any S_ℓ-invariant polynomial f can be written as a linear combination of orbit sums of monomials. Using this fact we now prove the proposition by induction on $d = \mathrm{div}(f)$. If $d = 1$ then the orbit sums above are orbit sums of monomials that only involve one set of variables, so they are of the form M^Σ.

Now suppose $d > 1$. Without loss of generality, we may assume that f does not have monomials of diversity $< d$. Consider the orbit sum of a monomial $M_1^{[i_1]} M_2^{[i_2]} \cdots M_d^{[i_d]}$ where $M_1, M_2, \ldots, M_d \in \mathbb{F}_q[x_1, x_2, \ldots, x_m]$ are non-constant monomials. If the M_i are all distinct, then this orbit sum can be written as the sum

$$\sum_{i_1, i_2, \ldots, i_d} M_1^{[i_1]} M_2^{[i_2]} \cdots M_d^{[i_d]} \tag{2}$$

over all $(i_1, i_2, \ldots, i_d) \in \{1, 2, \ldots, \ell\}^d$ with $i_1, i_2, \ldots, i_d$ distinct.

If however some of the M_j's coincide, then in the sum (2) some of the monomials $M_1^{[i_1]} M_2^{[i_2]} \cdots M_d^{[i_d]}$ are summed more than once. (For example, if $d = \ell = 2$ and $M_1 = M_2 = x_1$ then the orbit sum of $M = M_1^{[1]} M_2^{[2]}$ has size 1, whereas in the sum above the same monomial would appear twice.) In the worst case, when all M_j's are the same, the same monomial is summed $d!$ times. This is not a problem, because the characteristic p of $\mathbb{F}_q$ is $> d$, so we can still write the orbit sum as the sum (2) multiplied by a non-zero field element.

So we can write f as a linear combination of sums (2) (for various choices of the M_i). Consider one such sum S. Note that the polynomial

$$S - M_1^\Sigma M_2^\Sigma \cdots M_d^\Sigma$$

has diversity $< d$. By the induction hypothesis, $S - M_1^\Sigma M_2^\Sigma \cdots M_d^\Sigma$ can be written as a polynomial of degree $< d$ in the M^Σ's. So f can be written as a polynomial of degree d in the M^Σ's. □

Let A_ℓ be the *alternating subgroup* of S_ℓ.

Lemma 9. *If a polynomial $f \in R^{\otimes \ell}$ is A_ℓ-invariant and $\deg(f) \leq \ell - 2$ then f is S_ℓ-invariant.*

Proof. First, assume that f is an A_ℓ orbit sum, i.e., there is a monomial N such that f is the sum of all elements in the set $\{\sigma \cdot N \mid \sigma \in A_\ell\}$. Because $\deg(N) \leq \ell - 2$, there exist two rows in (1), say i and j, such that N does not contain any variables from those rows. Then we have $(i\ j) \cdot N = N$, and since f was already A_ℓ-invariant we conclude that f is S_ℓ-invariant.

If f is arbitrary, we use the general fact that for any group G, if a polynomial is G-invariant, then f can be written as the sum of orbit sums polynomials. Hence we can apply the argument above to each orbit sum, and conclude the general case as well. □

Lemma 10. *If $f \in R^{\otimes \ell}$ is S_ℓ-invariant, $\deg(f) \leq \ell - 1$, $\ell \geq 5$ and $u \in R^{\otimes \ell}$ divides f, then u is S_ℓ-invariant.*

Proof. We can factor $f = f_1 f_2 \cdots f_s$ where f is irreducible. Factorization in the polynomial ring into irreducible factor is unique up to permuting factors and multiplying factors with nonzero constant scalars. From $f = \pi(f) = \pi(f_1)\pi(f_2)\cdots\pi(f_s)$ follows that for every i there exists a j and a nonzero constant $c \in \mathbb{F}_q^\times = \mathbb{F}_q - \{0\}$ such that $\pi(f_i) = cf_j$. In other words $\pi(L_i) = L_j$ where L_i is the span of f_i. Let $\mathcal{L} = \{L_1, L_2, \ldots, L_s\}$. Note that the set $\mathcal{L}$ may have less than s elements, because some factors may be the same up to a nonzero constant. Then S_ℓ acts on $\mathcal{L}$. Let $H_i \subseteq S_\ell$ be the stabilizer subgroup of L_i, that is, $\pi \in H_i$ if and only if $\pi(L_i) = L_i$. By the *orbit-stabilizer* theorem, the index $|S_\ell|/|H_i|$ of H_i in S_ℓ equals the size of the orbit of L_i. The latter is $\leq |\mathcal{L}| \leq s$. Moreover, $s \leq \deg(f) \leq \ell - 1$, where the second inequality is by assumption. Hence, the index of H_i is $< \ell$. It is known that the only proper subgroup of S_ℓ of index $< \ell$ is A_ℓ, see e.g. [Cla84, p. 84]. So it follows that $H_i = A_\ell$ or S_ℓ. This proves that $\pi(L_i) = L_i$ for all $\pi \in A_\ell$.

We now argue that in fact even $\pi(f_i) = f_i$ for all i and all $\pi \in A_\ell$. Fix i. From $\pi(L_i) = L_i$ for all $\pi \in A_\ell$ we know that for every $\pi \in A_\ell$ there exists a (unique) element $\chi_i(\pi) \in \mathbb{F}_q - \{0\}$ such that $\pi(f_i) = \chi_i(\pi) f_i$. Notice that $\chi_i : A_\ell \to \mathbb{F}_q^\times$ is a group homomorphism. Let K be its kernel. The kernel of any group homomorphism is a normal subgroup, so K is a normal subgroup of A_ℓ. On the other hand, A_ℓ is simple for $\ell \geq 5$, that is, it has no non-trivial normal subgroups. So either $K = A_\ell$ or $K = \{1\}$. We can exclude the latter possibility because it would imply that A_ℓ is commutative, which is not true. (We would have $\pi \cdot \pi' = \chi_i^{-1}\chi_i(\pi \cdot \pi') = \chi_i^{-1}(\chi_i(\pi) \cdot \chi_i(\pi')) = \chi_i^{-1}(\chi_i(\pi') \cdot \chi_i(\pi)) = \chi_i^{-1}\chi_i(\pi' \cdot \pi) = \pi' \cdot \pi$using that $\mathbb{F}_q^\times$ is commutative.) Hence $K = A_\ell$ and $\pi(f_i) = \chi_i(\pi) f_i = f_i$ for all $\pi \in A_\ell$.

Therefore, $f_1, f_2, \ldots, f_s$ are A_ℓ-invariant. If $s = 1$, then $f_t = f$ is S_ℓ-invariant. If $s > 1$, then $\deg(f_i) \leq \ell - 2$ for all i, and f_i is S_ℓ-invariant by lemma 9. Up to a constant, u is a product of the f_i's, so u is S_ℓ-invariant. □

Proposition 11. *Suppose that $M_1, M_2, \ldots, M_r \in \mathbb{F}_q[x_1, x_2, \ldots, x_m]$ are distinct non-constant monomials, $g(x_1, x_2, \ldots, x_r)$ is a polynomial of degree $d \leq \ell$ and $p > d$. If $g(M_1^\Sigma, M_2^\Sigma, \ldots, M_r^\Sigma) = 0$, then $g = 0$.*

Proof. Consider a monomial of maximal degree d in g, say $x_{i_1} x_{i_2} \cdots x_{i_d}$ with $i_1 \leq i_2 \leq \cdots \leq i_d$. Then the monomial $M_{i_1}^{[1]} M_{i_2}^{[2]} \cdots M_{i_d}^{[d]}$ appears in $M_{i_1}^\Sigma M_{i_2}^\Sigma \cdots M_{i_d}^\Sigma$. Here we use the assumption on the characteristic, needed for example if $i_1 = i_2 = \ldots = i_d$. Also, if $j_1 \leq j_2 \leq \cdots \leq j_d$ and $(i_1, i_2, \ldots, i_d) \neq (j_1, j_2, \ldots, j_d)$, then $M_{i_1}^{[1]} M_{i_2}^{[2]} \cdots M_{i_d}^{[d]}$ does not appear in $M_{j_1}^\Sigma M_{j_2}^\Sigma \cdots M_{j_d}^\Sigma$. Also, $M_{i_1}^{[1]} M_{i_2}^{[2]} \cdots M_{i_d}^{[d]}$ does not appear in $M_{j_1}^\Sigma M_{j_2}^\Sigma \cdots M_{j_{d'}}^\Sigma$ if $d' < d$ since the latter has diversity $\leq d'$ while the former has diversity d.

This shows that the monomial $M_{i_1}^{[1]} M_{i_2}^{[2]} \cdots M_{i_d}^{[d]}$ appears in $g(M_1^\Sigma, M_2^\Sigma, \ldots, M_r^\Sigma)$. In particular, $g(M_1^\Sigma, M_2^\Sigma, \ldots, M_r^\Sigma) \neq 0$. □

We can now prove the main theorem of this section.

Proof of Theorem 6. Suppose that G can be decomposed as $G = c(H)$ for some $H \in R^{\otimes \ell}$ and univariate polynomial $c \in \mathbb{F}_q[x]$ of degree $e \geq 1$. Note that G has degree $\leq dk$. Let $\alpha \in \overline{\mathbb{F}}_q$ be a root of $c(x)$. Then $x - \alpha$ divides $c(x)$, and so $H - \alpha$ divides $c(H)$. Then $H - \alpha$, and hence H, is S_ℓ-invariant by Lemma 10, using that $\ell \geq dk + 1$. Note that if $\alpha \in \overline{\mathbb{F}}_q$ does not lie in $\mathbb{F}_q$, then we have to apply Lemma 10 after replacing $\mathbb{F}_q$ with a finite field extension of $\mathbb{F}_q$ that contains α.

From the degree bounds on $G = c(H)$ and c it follows that H has degree $\leq dk/e$. In particular, $\mathrm{div}(H) \leq dk/e$. By Proposition 8 we can write H as a polynomial of degree $\leq dk/e$ in all M^Σ's, say $H = h(M_1^\Sigma, M_2^\Sigma, \ldots, M_s^\Sigma)$ for some s. Note that s may be larger than r.

If we set $u(x_1, x_2, \ldots, x_s) = g(x_1, x_2, \ldots, x_r) - c(h(x_1, x_2, \ldots, x_s))$, then we have

$$u(M_1^\Sigma, M_2^\Sigma, \ldots, M_s^\Sigma) = 0.$$

Proposition 11 implies that $u = 0$. So $g(x_1, x_2, \ldots, x_r) = c(h(x_1, x_2, \ldots, x_s))$. So $h(x_1, x_2, \ldots, x_s) = h(x_1, x_2, \ldots, x_r)$ only depends on $x_1, x_2, \ldots, x_r$ and the degree of h is $\leq d/e$. □

2 Indecomposability implies equidistribution

In this section we prove the following lemma.

Lemma 12. *Let h be a polynomial of degree d in n variables over $\mathbb{F}_q$. If h is indecomposable then $h(U)$ is $O(d^2/\sqrt{q})$-close to uniform over $\mathbb{F}_q$.*

For the proof we need several facts from the algebraic-geometry literature.

Fact 13. *[Naj05] Let h be a polynomial of degree d in n variables over an algebraically-closed field K. Suppose that h is indecomposable. Then the number of $\lambda \in K$ such that $h - \lambda$ is reducible in K is at most d.*

[Naj05] generalizes several previous works; we refer to [Naj05] for the history of this type of results. Our polynomials are over $\mathbb{F}_q$ which is not algebraically closed. However the following fact allows us to bypass this apparent obstacle. If K is a field the notation $\overline{K}$ denotes its *algebraic closure*.

Fact 14. *[BDN09, Theorem 4.2.] If a polynomial is indecomposable over $\mathbb{F}_q$ then it is also indecomposable over $\overline{\mathbb{F}_q}$.*

Finally, we use the following version of Weil's bound.

Fact 15. *[Bog05, Proposition 2.6] Let h be a non-constant polynomial of degree d in n variables over $\mathbb{F}_q$ that cannot be reduced in $\overline{\mathbb{F}_q}$. Then $|\mathbb{P}[h(U) = 0] - 1/q| \leq O(d^2 q^{-3/2})$, assuming $q > 5d^4$.*

Proof of Lemma 12. By Fact 14 h is also indecomposable over $\overline{\mathbb{F}_q}$. By Fact 13, $h - \lambda$ is not reducible in $\overline{\mathbb{F}_q}$ except for at most d values of $\lambda \in \mathbb{F}_q$. For each value λ for which it is not reducible, Fact 15 yields $|\mathbb{P}[h(U) = \lambda] - 1/q| \leq O(d^2 q^{-3/2})$. Note we can assume $q > 5d^4$ for else the conclusion of the lemma holds. For any other value of λ, by Schwartz-Zippel, $|\mathbb{P}[h(U) = \lambda] - 1/q| \leq d/q$. Combining these facts, the statistical distance between $h(U)$ and uniform is at most $O(d^2/\sqrt{q}) + d^2/q = O(d^2/\sqrt{q})$. □

3 Toy pseudorandom generators with what we have so far

In this section we derive "toy" pseudorandom generators with the results of the previous two sections, over fields of characteristic $> dk$. Define $f_i := M_i^\Sigma$ as in the introduction. The generator simply picks ℓm uniform values for the variables of the f_i and outputs $(f_1, f_2, \ldots, f_n)(U)$. The analysis goes as follows. Let g be a polynomial of degree d that we aim to fool. Let c be a univariate polynomial of maximal degree such that $g(x_1, x_2, \ldots, x_n) = c(h(x_1, x_2, \ldots, x_n))$. In particular we have $g(f_1, f_2, \ldots, f_n) = c(h(f_1, f_2, \ldots, f_n))$. Note that h has degree $\leq d$ and is indecomposable, for else the degree of c is not maximal. By Theorem 6, $h(f_1, f_2, \ldots, f_n)$ is indecomposable as well. By Lemma 12, $h(f_1, f_2, \ldots, f_n)(U)$ is $O(d^2k^2/\sqrt{q})$-close to uniform, and the same bound holds for $h(U)$.

Hence we obtained generators with seed length $O(\ell m \log q) = O(dkm \log q)$ and error $O(dk)^2/\sqrt{q}$. Here we just need $\binom{m+k}{m} \geq n$. For example, we can pick m and k to be $O(\log n)$. This gives seed length $O(d \log^2 n \log q)$. As mentioned earlier, one can improve the seed length to $O(d \log n \log^{O(1)} \log(dn))$ by applying the construction recursively.

4 Improving bounds for indecomposability

In this section we improve the bounds in Theorem 6 to get the preservation of indecomposability for $\ell \geq d + 1$ instead of $\ell \geq dk + 1$. The factor-k loss in the previous argument arises when bounding the diversity of H by the degree of H, where the latter is a priori as large as dk/e, see the proof of Theorem 6. In this section we consider a more constrained set Q of monomials, defined shortly. Using this, we can recoup a factor k when bounding the diversity of a polynomial in terms of its degree, see Lemma 17.

We fix a positive integer k and let $Q \subseteq R = \mathbb{F}_q[x_1, x_2, \ldots, x_m]$ be the subring spanned by all monomials of the form $x_1^{a_1} x_2^{a_2} \cdots x_m^{a_m} \in R$ where $a_1 + a_2 + \cdots + a_{m-1} = (k-1)a_m$. Note that the degree of a polynomial in Q is ka_m which is always divisible by k. Let $Q^{\otimes \ell} \subset R^{\otimes \ell}$ be the subring spanned by all monomials $M_1^{[1]} M_2^{[2]} \cdots M_\ell^{[\ell]}$ where $M_1, M_2, \ldots, M_\ell \in Q$.

We modify Theorem 6 by only considering monomials M^Σ where M is a monomial in the subring $Q \subset R$ rather than in R. By doing so, as we mentioned, we improve the parameters as follows.

Theorem 16. *Suppose that $M_1, M_2, \ldots, M_r \in Q$ are distinct non-constant monomials of degree k, and let $g(x_1, x_2, \ldots, x_r)$ be a non-constant polynomial of degree d. Let $G := g(M_1^\Sigma, M_2^\Sigma, \ldots, M_r^\Sigma)$ and assume that $p \geq d + 1$ and $\ell \geq \max\{5, d + 1\}$. If G is decomposable then g is decomposable.*

The rest of this section is devoted to the proof of this theorem. The proof follows the same outline of the proof of the corresponding Theorem 6 in Section 1, but some of the steps are more involved.

First, as mentioned above, we give a tighter connection between diversity and degree for polynomials in $Q^{\otimes \ell}$.

Lemma 17. *If $f \in Q^{\otimes \ell}$ has degree $\leq dk$, then $\mathrm{div}(f) \leq d$.*

Proof. The polynomial f is by definition a linear combination of monomials of the form $M_1^{[1]} M_2^{[2]} \cdots M_\ell^{[\ell]}$ where M_i is a monomial in Q of total degree $\leq dk$. If $M_i \neq 1$, then the degree of M_i is at least k. So $M_i \neq 1$ for at most d distinct indices i. This proves that $\mathrm{div}(f) \leq d$. □

We also modify Proposition 8 as follows.

Proposition 18. *Suppose that $f \in Q^{\otimes \ell}$ is an S_ℓ-invariant polynomial with $\mathrm{div}(f) = d$ and $p > d$. Then f can be written as a polynomial of degree d in the M^Σ's, where M ranges over monomials in Q.*

Proof. We follow the proof of Proposition 8 and note that all the monomials that appear can be chosen in Q and $Q^{\otimes \ell}$ instead of R and $R^{\otimes \ell}$ respectively. □

One difficulty that we face when generalizing the other statements in Section 1 such as Lemma 10 is that of arguing that the polynomials we encounter lie in $Q^{\otimes \ell}$ instead of $R^{\otimes \ell}$. For this purpose it is convenient to introduce the ring homomorphism

$$\Phi : R^{\otimes \ell} \to R^{\otimes \ell}[t^{[1]}, \ldots, t^{[\ell]}, (t^{[1]})^{-1}, \ldots, (t^{[\ell]})^{-1}]$$

with $\Phi(x_j^{[i]}) = t^{[i]} x_j^{[i]}$ for $j < m$, $\Phi(x_m^{[i]}) = (t^{[i]})^{1-k} x_m^{[i]}$ and $\Phi(\alpha) = \alpha$ for $\alpha \in \mathbb{F}_q$. Note that the image of Φ is a *Laurent polynomial*, that is, the exponents of the variables $t^{[i]}$ may be negative. If we work over the algebraically closed field $\overline{\mathbb{F}}_q$, then Φ corresponds to an action of the ℓ-dimensional torus group $T = (\overline{\mathbb{F}}_q^\times)^\ell$ on the ring $R^{\otimes \ell}$. This motivates the terminology that follows.

A polynomial $f \in R^{\otimes \ell}$ is *T-invariant* when $\Phi(f) = f$, i.e., the variables $t^{[i]}$ cancel out. Note that if $M \in Q$ then $M^{[i]}$ is T-invariant. More generally, $Q^{\otimes \ell}$ is the ring of all T-invariants. A polynomial $f \in R^{\otimes \ell}$ is called *T-semi-invariant* if $\Phi(f) = (\prod_{i=1}^{\ell} (t^{[i]})^{a^{[i]}}) f$ for some $a = (a^{[1]}, a^{[2]}, \ldots, a^{[\ell]}) \in \mathbb{Z}^\ell$, called *weight*. The monomials in $R^{\otimes \ell}$ are all T-semi-invariant.

As in Lemma 10, we need to argue about factors of invariant polynomials. We begin with the following lemma which will help us argue that these factors lie in $Q^{\otimes \ell}$ (as opposed to $R^{\otimes \ell}$).

Lemma 19. *Suppose that $u, f \in R^{\otimes \ell}$ and u divides f. If f is T-semi-invariant, then so is u.*

We shall only use this for T-invariant f, but the proof is the same.

Proof. Note that f is T-semi-invariant if and only if $\Phi(f)$, as a Laurent polynomial in $t^{[1]}, \ldots, t^{[\ell]}$, consists of a single monomial. If $f = uv$ then note $\Phi(f) = \Phi(u)\Phi(v)$. By assumption, $\Phi(f)$ consists of a single monomial as a Laurent polynomial. Then the same is also true for $\Phi(u)$ and $\Phi(v)$. Here we are using the general fact that if, say, $\Phi(u)$ has more than one term, then the product $\Phi(u)\Phi(v)$ has more than one term. To see this, consider the lexicographic order on monomials, and note that the product of the smallest monomial in $\Phi(u)$ with the smallest monomial in $\Phi(v)$ cannot be obtained by multiplying any other two monomials, and the same holds for the product of the largest monomials. Hence the product has at least two monomials. □

We modify Lemma 10 to the following statement:

Lemma 20. *If $f \in Q^{\otimes \ell}$ is S_ℓ-invariant, $\deg(f) \leq k\ell - 1$, $\ell \geq 5$ and $u \in R^{\otimes \ell}$ divides f, then u lies in $Q^{\otimes \ell}$ and is S_ℓ-invariant.*

Proof. We modify the proof of Lemma 10. We started with a factorization $f = f_1 f_2 \ldots f_s$ where $f_1, f_2, \ldots, f_s$ are irreducible. Since $f \in Q^{\otimes \ell}$ it is T-invariant, and therefore T-semi-invariant. We defined $L_i = \mathbb{F}_q f_i$ and considered the action of S_ℓ on $\mathcal{L} = \{L_1, L_2, \ldots, L_s\}$. As before H_i is the stabilizer of L_i. By Lemma 19, $f_1, f_2, \ldots, f_s$ are also semi-invariant. Let us fix some j, and let $(a^{[1]}, a^{[2]}, \ldots, a^{[\ell]})$ be the weight of the semi-invariant f_j.

We prove that $a^{[i]} \geq 0$ for all i. First recall that a monomial in Q is of the form $x_1^{a_1} x_2^{a_2} \cdots x_m^{a_m}$ with $a_1 + a_2 + \cdots + a_m = k a_m$. This means that the total degree of a polynomial $f \in Q$ in the variables $x_1, x_2, \ldots, x_m$ is exactly k times the degree of f as a polynomial in the variable x_m with coefficients in $\mathbb{F}_q[x_1, \ldots, x_{m-1}]$. Similarly, if $f \in Q^{\otimes \ell}$ then the total degree $\deg(f)$ of f is k times the degree $\deg_m(f)$ of f in the variables $x_m^{[1]}, x_m^{[2]}, \ldots, x_m^{[\ell]}$. Therefore for $f \in Q^{\otimes \ell}$ we have $\deg_m(f) = \deg(f)/k < \ell$.

Now suppose towards a contradiction that $a^{[i]} < 0$ for some i. Then $x_m^{[i]}$ must appear in f_j, so $\deg_m(f_j) \geq 1$. The orbit of L_j has $|S_\ell|/|H_j|$ elements, which correspond to as many irreducible factors of f that are distinct and have degree ≥ 1 with respect to $\deg_m$. This implies that $|S_\ell|/|H_j| \leq \deg_m(f) < \ell$. As in the proof of Lemma 10, this implies that $H_j = A_\ell$ or $H_j = S_\ell$.

For a permutation σ, $\sigma(f_j)$ is semi-invariant. Its weight is $\sigma(a) = (a^{[\sigma(1)]}, \ldots, a^{[\sigma(\ell)]})$. For $\sigma \in H_j$, $\sigma(f_j)$ and f_j are the same up to a constant, so $\sigma(a) = a$ for all $\sigma \in H_j$. If $H_j = S_\ell$ then $a^{[1]} = a^{[2]} = \cdots = a^{[\ell]} < 0$. The same holds if $H_j = A_\ell$ for $\ell \geq 3$ as we can again map i to any other value via $\sigma \in A_\ell$. This implies that all the monomials in f_j contain $\prod_{i=1}^m x_m^{[i]}$ and $\deg_m(f_j) \geq \ell$, contradicting the bound above.

We proved that $a^{[i]} \geq 0$ for all i, so $\Phi(f_j)$ lies in the polynomial ring $R^{\otimes \ell}[t^{[1]}, \ldots, t^{[\ell]}]$ for all j. From $\prod_{j=1}^s \Phi(f_j) = \Phi(f) = f \in R^{\otimes \ell}$ it follows that $\Phi(f_j) \in R^{\otimes \ell}$ for all j. This implies that $f_j \in Q^{\otimes \ell}$ for all j.

There remains to argue that the f_j are S_ℓ-invariant. As in the proof of Lemma 10, $\sigma(L_i) = L_i$ for all $\sigma \in A_\ell$ implies that $\sigma(f_i) = f_i$ for all $\sigma \in A_\ell$. Hence, the f_i are A_ℓ-invariant.

If $s = 1$, then $f = f_1$ is S_ℓ-invariant. Otherwise, $\deg_m(f_j) \leq \ell - 2$. Since $f_j \in Q^{\otimes \ell}$ we get $\deg(f_j) = k \deg_m(f_j) \leq (\ell - 2)k$. By Lemma 17, $\operatorname{div}(f_j) \leq \ell - 2$. Using Lemma 9 we conclude that f_j is S_ℓ-invariant. □

Proof of Theorem 16. Suppose that G can be decomposed as $G = c(H)$ for some $H \in R^{\otimes \ell}$ and univariate polynomial $c \in \mathbb{F}_q[x]$ of degree $e \geq 1$. We claim that in fact $H \in Q^{\otimes \ell}$. To verify this, note that from $G = c(H)$ it follows $\Phi(G) = c(\Phi(H))$. Since $\Phi(G)$ is T-invariant, i.e., constant in the variables $t^{[1]}, t^{[2]}, \ldots, t^{[\ell]}$, so is $\Phi(H)$ and $H \in Q^{\otimes \ell}$.

Note that G has degree $\leq dk$. Let $\alpha \in \overline{\mathbb{F}}_q$ be a root of $c(x)$. Then $x - \alpha$ divides $c(x)$, and so $H - \alpha$ divides $c(H)$. Because $\deg(H - \alpha) \leq dk < k\ell$, $H - \alpha$ lies in $Q^{\otimes \ell}$ and is S_ℓ-invariant by Lemma 20. (Possibly, we may have to replace $\mathbb{F}_q$ by a finite field extension.) It follows that $H \in Q^{\otimes \ell}$ and is S_ℓ-invariant.

From the degree bounds on $G = c(H)$ and c it follows that H has degree $\leq dk/e$. By Lemma 17, we get $\mathrm{div}(H) \leq d/e < \ell$. By Proposition 18 we can write H as a polynomial of degree $\leq d/e$ in all M^Σ's with $M \in Q$, say $H = h(M_1^\Sigma, M_2^\Sigma, \ldots, M_r^\Sigma, \ldots, M_s^\Sigma)$. If we set $u(x_1, x_2, \ldots, x_s) = g(x_1, x_2, \ldots, x_\ell) - c(h(x_1, x_2, \ldots, x_s))$, then we have

$$u(M_1^\Sigma, M_2^\Sigma, \ldots, M_s^\Sigma) = 0.$$

and the degree of u is $\leq d$. Proposition 11 implies that $u = 0$. So $g(x_1, x_2, \ldots, x_r) = c(h(x_1, x_2, \ldots, x_s))$. So $h(x_1, x_2, \ldots, x_s) = h(x_1, x_2, \ldots, x_r)$ only depends on $x_1, x_2, \ldots, x_r$ and the degree of h is $\leq d/e$. □

5 Bogdanov-style generators

In this section we prove the following theorem.

Theorem 21. *There are explicit pseudorandom generators that fool with error ϵ degree-d polynomials in n variables over $\mathbb{F}_q$, provided $q \geq O(d^4/\epsilon^2)$, with seed length either*

(1) $O(n \log(d + n) + \log q)$ or

(2) $O(d^4 \log n + \log q)$.

First we refine Bogdanov's reduction of pseudorandom generators to hitting-set generators. An explicit map $H : S \to \mathbb{F}_q^n$ is a *δ-hitting-set* generator for degree-d polynomials in n variables over $\mathbb{F}_q$ if for any such polynomial f, if $f \neq 0$ then $\mathbb{P}[f(H(U)) = 0] \leq \delta$. The seed length of H is $\log_2 |S|$.

We obtain the following refinement of Bogdanov's reduction:

Lemma 22. *Suppose there exists a δ-hitting-set generator with seed length s for polynomials of degree $3d^4$ in $2n$ variables over $\mathbb{F}_q$. Then there exists a pseudorandom generator for polynomials of degree d in n variables over $\mathbb{F}_q$ with seed length $2s + 2\log q$ and error $O(\delta + d^2/\sqrt{q})$.*

[Bog05, Theorem 3.1] proves the same but with error $O(\sqrt{\delta} d + d^2/\sqrt{q} + d^6/q)$. To prove Lemma 22 first we use the following result to relate indecomposability and *irreducibility*.

Fact 23. *[CN10, Lemma 7] Let $f \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a non-constant polynomial. Then f is indecomposable over $\overline{\mathbb{F}}$ iff $f - y$ is irreducible in $\overline{\mathbb{F}(y)}[x_1, x_2, \ldots, x_n]$.*

Here $\overline{\mathbb{F}(y)}$ is the algebraic closure of the function field $\mathbb{F}(y)$, where y is a variable.

We also need the following fact, mentioned already in [Bog05] when $\mathbb{E} = \overline{\mathbb{F}}$.

Fact 24. *Let $\mathbb{F} \subseteq \mathbb{E}$ be a field extension. Let H be a δ-hitting-set generator for degree-d polynomials over $\mathbb{F}$. Then H is also a δ-hitting-set generator for polynomials over $\mathbb{E}$.*

This fact follows because $\mathbb{E}$ is a vector space over $\mathbb{F}$.

Proof of Lemma 22. Let g be a polynomial that we aim to fool. As in Section 3, write $g = c(h)$ where c is a univariate polynomial of maximal degree, and h is indecomposable. It suffices to preserve the output distribution of h, which by Lemma 12 is close to uniform. We relate indecomposability to irreducibility via Fact 23, inspired by the proof of Theorem 8 in [CN10], then reason as in [Bog05], using Theorem 5 in [Kal95].

By Fact 14, h is indecomposable over $\overline{\mathbb{F}_q}$ as well. Hence we can apply Fact 23 to conclude that $h - y$ is irreducible in $\mathbb{E}[x_1, x_2, \ldots, x_n]$ where $\mathbb{E} := \overline{\mathbb{F}_q(y)}$. We now use Theorem 5 in [Kal95] over the field $\mathbb{E}$. For $v_{1..n} \in \mathbb{E}^n$ and $w_{2..n}, z_{2..n} \in \mathbb{E}^{n-1}$ define the following bivariate restriction of h:

$$h|_{v,w,z}[s,t] := h(s + v_1, w_2 s + z_2 t + v_2 \ldots, w_n s + z_n t + v_n).$$

Theorem 5 in [Kal95] shows that $h|_{v,w,z}$ is absolutely irreducible except when v, w are zeroes of a polynomial of degree $O(d^2)$ over $\mathbb{E}$, or z is the zero of a polynomial of degree $O(d^4)$ over $\mathbb{E}$ (where the latter polynomial may depend on the first).

For our generator, we pick (v, w) with a δ-hitting-set generator for polynomials of degree $O(d^2)$ and z with an independent δ-hitting-set generator with error ϵ for polynomials of degree $O(d^4)$. For the variables s and t we plug uniform values in $\mathbb{F}_q$.

By Fact 24, these hitting-set generators are also δ-hitting-set generators polynomials over $\mathbb{E}$. Hence, $h|_{v,w,z}$ is absolutely irreducible over $\mathbb{E}[s,t]$ with probability $\geq 1 - O(\delta)$. Then from Fact 23 we obtain that $h|_{v,w,z}$ is indecomposable with at least the same probability over the choice of v, w, z from the hitting-set generators. Whenever it is indecomposable, by Lemma 12 its output distribution is $O(d^2)/\sqrt{q}$-close to uniform. □

To prove Theorem 21 there remains to construct δ-hitting-set generators. Such for polynomials of degree d in n variables are known with optimal seed length $O(d \log n + \log 1/\delta)$, provided $q \geq O(d/\delta)$ [GX14]. In particular, for polynomials of degree d^4 in $O(n)$ variables we can set $\delta := \epsilon d^2/\sqrt{q}$ and have seed length $O(d^4 \log n + \log q)$, provided $q \geq O(d^4/(\epsilon d^2/\sqrt{q}))$. The last provision is equivalent to $q \geq O(d^4/\epsilon^2)$, which we can always assume for else the theorem is trivial. This gives Item (2) in Theorem 21.

Over fields of characteristic $\geq O(d^2)$ the d^4 factor can be improved to d^2 using Corollary 8 in [Lec07] – and that is the best possible, see Corollary 7 and the surrounding discussion in the same paper.

For Item (1) in Theorem 21 we need a different hitting-set generator, stated next.

Lemma 25. *[Implicit in [Lu12, GX14]] There is an explicit δ-hitting-set generator with seed length $O(n \log(n + d) + \log 1/\delta)$ for polynomials of degree d in n variables over $\mathbb{F}_q$, provided $q \geq O(d/\delta)$.*

This should be compared to the Schwartz-Zippel lemma, which yields a δ-hitting-set generator with seed length $n \log(d/\delta)$ provided $q \geq d/\delta$. As the above lemma is not stated in those works we quickly sketch how it follows from [Lu12, GX14]. Lu [Lu12] (Theorem 1)

gives a δ-hitting-set generator for polynomials with s terms with seed length $O(\log(sd/\delta))$ provided $q \geq d^{1.01}/\delta$. (Lu's proof focuses on constant δ, but as noted there and in [GX14] one can also obtain the stated parameters.) A degree-d polynomial in n variables has $s \leq \binom{n+d}{n}$ monomials. Hence we obtain seed length $O(n \log(n+d) + \log 1/\delta)$. Guruswami and Xing [GX14] use multiplication-friendly codes to bring down the bound on the field size to $q \geq O(d/\delta)$.

To prove Item (1) in Theorem 21, use the δ-hitting-set generator in Lemma 25 for polynomials of degree d^4 in $O(n)$ variables, setting $\delta := \epsilon d^2/\sqrt{q}$.

6 Proof of main results for fields of characteristic $> d$

In this section we prove our main results, Theorem 2 and Theorem 4, in the case of fields of characteristic $> d$ (for example, prime fields). The proofs over arbitrary characteristic are the same except that we use Theorem 27 of Section 7 instead of Theorem 16.

Proof of Theorem 4. Let Q and $M_1, M_2, \ldots$ be as in Theorem 16. The number of distinct monomials in Q is at least the number of positive integers $a_1, a_2, \ldots, a_{m-1}$ with sum equal to $k-1$ (corresponding to the setting $a_m = 1$ in Section 4). This number is $\binom{m-1+k-1}{m-1}$, which is $\geq n$ by assumption. Define $f_i := M_i^{\Sigma}$. The analysis is the same as in Section 3. □

Proof of Theorem 2. From Theorem 4 we reduce our task to that of fooling polynomials with degree $d' := dk$ in $n' := \ell m = (d+1)m$ variables, up to an error $O(d'^2/\sqrt{q})$. This error is $\leq \epsilon$ by our assumption that $q \geq O(dk)^4/\epsilon^2$.

Item (1) in Theorem 21 shows how to fool such polynomials with seed length $O(n' \log(d' + n') + \log q)$ and error β, provided $q \geq O(d'^4/\beta^2)$. This allows us to set $\beta := O(d'^2/\sqrt{q})$ and the provision is true. Again by our assumption that $q \geq O(dk)^4/\epsilon^2$, we have $\beta = O(\epsilon)$. Hence the combined error from the two steps is $O(\epsilon)$. The final seed length is $O(dm \log(dk+dm) + \log q)$, as desired. □

7 Preserving indecomposability over any characteristic

In this section we show how to preserve indecomposability over fields of arbitrary characteristic. The problem in small characteristic is that the M^{Σ}'s where M is a monomial no longer span the invariant ring $(R^{\otimes \ell})^{S_\ell}$. This is even a problem when $m = 1$. For example, if q is a power of 2, then $(x_1^2)^{\Sigma} = (x_1^{\Sigma})^2$ and the second elementary symmetric polynomial $\sum_{1 \leq i < j \leq \ell} x_1^{[i]} x_1^{[j]}$ does not lie in the ring generated by $(x_1^j)^{\Sigma}$, $j = 1, 2, 3, \ldots$. The solution is to avoid using M^{Σ}'s for which M is a power of another monomial of smaller degree.

The main results in this section are the following theorems, which can be used as a replacements to Theorem 6 and Theorem 16 :

Theorem 26. *Suppose that $M_1, M_2, \ldots, M_r \in R$ are distinct non-constant, indecomposable monomials of degree $\leq k$, and let $g(x_1, x_2, \ldots, x_r)$ be a non-constant polynomial of degree d.*

Let $G := g(M_1^\Sigma, M_2^\Sigma, \ldots, M_r^\Sigma)$ and assume that $\ell \geq \max\{5, dk+1\}$. If G is decomposable then g is decomposable.

Theorem 27. *Suppose that $M_1, M_2, \ldots, M_r \in Q$ are distinct non-constant monomials of degree k, and let $g(x_1, x_2, \ldots, x_r)$ be a non-constant polynomial of degree d. Let $G := g(M_1^\Sigma, M_2^\Sigma, \ldots, M_r^\Sigma)$ and assume that $\ell \geq \max\{5, d+1\}$. If G is decomposable then g is decomposable.*

The rest of this section is devoted to proving these theorems.

7.1 A first basis

In this subsection we obtain a first basis for $(R^{\otimes\ell})^{S_\ell}$. This is not the basis we will ultimately use. We start with some notation.

Let $\mathbb{N} = \{0, 1, 2, 3, \ldots\}$, m be a positive integer and $\mathcal{R} = \mathbb{N}^m \backslash \{0\}$. Suppose $(v_1, v_2, \ldots, v_m) \in \mathcal{R}$ and $k = \gcd(v_1, v_2, \ldots, v_m)$. If $k = 1$ then we call v *indivisible*. We can always write $v = kw$ where $w \in \mathcal{R}$ is indivisible. Let $\mathcal{R}_{\mathrm{i}} \subseteq \mathcal{R}$ be the set of indivisible vectors. We also define

$$\mathcal{Q} = \{(v_1, v_2, \ldots, v_m) \in \mathcal{R} \mid v_1 + v_2 + \cdots + v_{m-1} + (1-k)v_m = 0\}.$$

A basis of $R = \mathbb{F}_q[x_1, x_2, \ldots, x_m]$ is $\{1\} \cup \{x^v \mid v \in \mathcal{R}\}$ and a basis of the subring Q is $\{1\} \cup \{x^v \mid v \in \mathcal{Q}\}$.

For a set X, let $\mathbb{M}_j(X)$ be the set of all multisets that contain $\leq j$ elements of X, and let $\mathbb{M}(X) = \bigcup_{i=0}^\infty \mathbb{M}_j(X)$ be the set of all finite multisets of elements of X. If V and W are multisets, then we write $V \amalg W$ for their disjoint union. If k is a positive integer and $V = \{v_1, v_2, \ldots, v_n\}$ then we define $kV = \{kv_1, kv_2, \ldots, kv_n\}$ and

$$Vk = \coprod_{i=1}^k V = \{\underbrace{v_1, \ldots, v_1}_k, \underbrace{v_2, \ldots, v_2}_k, \ldots, \underbrace{v_n, \ldots, v_n}_k\}$$

is the disjoint union of k copies of V. We also use the convention $0V = V0 = \emptyset$.

Recall that the symmetric group S_ℓ acts on $R^{\otimes\ell}$ and we will study the invariant ring $(R^{\otimes\ell})^{S_\ell}$ and polynomials that span the invariant ring as an $\mathbb{F}_q$-vector space.

Definition 28. If $V = \{v_1, v_2, \ldots, v_r\} \in \mathbb{M}(\mathcal{R})$ then we define the invariant $\mathbf{x}^V \in (R^{\otimes\ell})^{S_\ell}$ as the sum of all monomials of the form

$$\prod_{j=1}^r (\mathbf{x}^{v_j})^{[i_j]},$$

where $i_1, i_2, \ldots, i_r \in \{1, 2, \ldots, \ell\}$ are distinct, each monomial is only summed once. In particular, $\mathbf{x}^\emptyset = 1$.

Note that in the case $r = 1$ this definition becomes $\mathbf{x}^{\{v\}} = \sum_{i \in \{1,2,\ldots,\ell\}} (\mathbf{x}^v)^{[i]}$. This is the same as M^Σ in the previous sections, if $M = \mathbf{x}^v$. Also note that if

$$V = \{\underbrace{v_1, v_1, \ldots, v_1}_{k_1}, \underbrace{v_2, v_2, \ldots, v_2}_{k_2}, \ldots, \underbrace{v_s, v_s, \ldots, v_s}_{k_s}\} = \coprod_{i=1}^s \{v_i\} k_i$$

where $v_1, v_2, \ldots, v_s$ are distinct, and $r = k_1 + k_2 + \cdots + k_s$ is the number of elements of the multiset V, then $\mathbf{x}^V$ is the sum of exactly

$$\binom{\ell}{\ell - r, k_1, k_2, \ldots, k_s} = \frac{\ell!}{(\ell - r)!k_1!k_2! \cdots k_s!}$$

distinct monomials, each having coefficient 1.

Example 29. If $\ell = 3$, $m = 2$, $V = \{(1,0),(1,0)\}$ and $W = \{(1,0),(0,1)\}$, then we have

$$\mathbf{x}^V = x_1^{[1]}x_1^{[2]} + x_1^{[1]}x_1^{[3]} + x_1^{[2]}x_1^{[3]}$$

and

$$\mathbf{x}^W = x_1^{[1]}x_2^{[2]} + x_1^{[1]}x_2^{[3]} + x_1^{[2]}x_2^{[1]} + x_1^{[2]}x_2^{[3]} + x_1^{[3]}x_2^{[1]} + x_1^{[3]}x_2^{[2]}.$$

Proposition 30. *A basis of the invariant ring $(R^{\otimes\ell})^{S_\ell}$ as an $\mathbb{F}_q$-vector space is given by all monomial sums $\mathbf{x}^V$, $V \in \mathbb{M}_\ell(\mathcal{R})$. A basis of the invariant ring $(Q^{\otimes\ell})^{S_\ell}$ as an $\mathbb{F}_q$-vector space is given by all monomial sums $\mathbf{x}^V$, $V \in \mathbb{M}_\ell(\mathcal{Q})$.*

Proof. Suppose $f \in (R^{\otimes\ell})^{S_\ell}$. As in the proof of Proposition 8, f is an $\mathbb{F}_q$-linear combination of orbit sums of monomials. But orbit sums of monomials are exactly the polynomials of the form $\mathbf{x}^V$ with $V \in \mathbb{M}_\ell(\mathcal{R})$. If $V \neq W$ then $\mathbf{x}^V$ and $\mathbf{x}^W$ have no monomial in common, so the $\mathbf{x}^V$ with $V \in \mathbb{M}_\ell(\mathcal{R})$ form a basis.

If $f \in (Q^{\otimes\ell})^{S_\ell}$ then the monomials that appear in f lie in $Q^{\otimes\ell}$. Orbits of such monomials are of the form x^V, $V \in \mathbb{M}_\ell(\mathcal{Q})$. □

A main problem with this basis is that it is not clear how to prove the corresponding of Proposition 11. In the next subsections we construct a different basis and prove such a result (Theorem 11).

7.2 A partial ordering on monomial sums

We will need a partial ordering on the basis elements $\mathbf{x}^V$, $V \in \mathbb{M}_\ell(\mathcal{R})$ in Proposition 30. We will do this by defining a partial ordering on $\mathbb{M}(\mathcal{R})$. For this we also need an ordering on partitions.

A *partition* of a nonnegative integer k is a tuple $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_r)$ of positive integers with $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_r \geq 1$ and $|\lambda| := \sum_{i=1}^r \lambda_i = k$. The set of all partitions of k is denoted by $\mathcal{P}_k$ and $\mathcal{P} = \bigcup_{k=0}^\infty \mathcal{P}_k$ is the set of all partitions. If λ is a partition, its *conjugate partition* λ' is obtained by transposing the Ferrers (or Young) diagram of λ. Since transposing twice gives us back the original diagram, we have $\lambda'' = \lambda$. Alternatively, if $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_r)$ is a partition, then we can define the conjugate partition λ' as $\lambda' = (\lambda'_1, \lambda'_2, \ldots, \lambda'_s)$, where $s = \lambda_1$, and λ'_j is the largest i for which $\lambda_i \geq j$. There is a *dominance ordering* $\trianglelefteq$ on $\mathcal{P}_k$. If $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_r)$ and $\mu = (\mu_1, \mu_2, \ldots, \mu_s)$ are partitions of k, then $\lambda \trianglelefteq \mu$ if and only if $r \geq s$ and $\sum_{i=1}^t \lambda_i \leq \sum_{i=1}^t \mu_i$ for $t = 1, 2, \ldots, s$.

Suppose that $V = \{v_1, v_2, \ldots, v_s\} \in \mathbb{M}(\mathcal{R})$. For a positive integer k we already defined kV and Vk. We generalize this to the case where k is a partition instead of a positive integer. For a partition $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_r)$, we define

$$\lambda V = \coprod_{i=1}^r \lambda_i V = \{\lambda_1 v_1, \lambda_1 v_2, \ldots, \lambda_1 v_s, \lambda_2 v_1, \lambda_2 v_2, \ldots, \lambda_2 v_s, \ldots, \lambda_r v_1, \lambda_r v_2, \ldots, \lambda_r v_s\}.$$

For the empty partition (), we define $()V = \emptyset$. Note that $\lambda(\mu v) = \mu(\lambda v)$ if μ is another partition. We also define right multiplication of V with a partition λ by $V\lambda := \lambda' V$. So for example, if k is a positive integer then we get

$$\{v\}(k) = (k)'\{v\} = (\underbrace{1, 1, \ldots, 1}_{k})\{v\} = \{\underbrace{v, v, \ldots, v}_{k}\} = \{v\}k.$$

In the following definition $\preceq$ is a priori a *preordering* on $\mathbb{M}(\mathcal{R})$, but we will prove that it is an ordering.

Definition 31. Let $\preceq$ be the preorder on $\mathbb{M}(\mathcal{R})$ generated by the following relations:

(I) For $v \in \mathcal{R}$ and $W \in \mathbb{M}(\mathcal{R})$ we have

$$W \preceq \{v\} \amalg W.$$

(II) For $v \in \mathcal{R}$, $W \in \mathbb{M}(\mathcal{R})$ and partitions λ, μ with $\lambda \trianglelefteq \mu$ we have

$$\lambda\{v\} \amalg W \preceq \mu\{v\} \amalg W.$$

(III) For linearly independent $v_1, v_2 \in \mathcal{R}$ and $W \in \mathbb{M}(\mathcal{R})$ we have

$$\{v_1 + v_2\} \amalg W \preceq \{v_1, v_2\} \amalg W.$$

Note that the $\preceq$ deals differently with dependent and independent vectors. For example, let $W = \emptyset$ and $v_1 = av$, $v_2 = bv$ for a vector v and integers a, b. Let $\lambda := (a, b)$ and $\mu := (a+b)$ be partitions. We have $\lambda \trianglelefteq \mu$. Also, $\lambda\{v\} = \{v_1, v_2\}$ and $\mu\{v\} = \{v_1 + v_2\}$. Hence in this case by (II) we have $\{v_1, v_2\} \preceq \{v_1 + v_2\}$. While if v_1, v_2 are linearly independent, then $\{v_1 + v_2\} \preceq \{v_1, v_2\}$ by (III).

Proposition 32. *The preorder $\preceq$ is a partial ordering on $\mathbb{M}(\mathcal{R})$, i.e., if $V \preceq W$ and $W \preceq V$ then $V = W$.*

Proof. Suppose that $V_1, V_2, \ldots, V_k \in \mathbb{M}(\mathcal{R})$ satisfy

$$V_1 \preceq V_2 \preceq \cdots \preceq V_k \preceq V_{k+1} := V_1$$

where for each i with $1 \leq i \leq k$, $V_i \preceq V_{i+1}$ follows from (I), (II) or (III) in Definition 31. We have to prove that $V_1 = V_2 = \cdots V_k$. We prove this by induction on $s = \min\{|V_1|, |V_2|, \ldots, |V_k|\}$.

If $s = 0$, then $V_i = \emptyset$ for some i. Now if $W \preceq \emptyset$ then $W = \emptyset$. So $V_j = \emptyset$ for all j and in particular, $V_1 = V_2 = \cdots = V_k$.

Suppose that $V_1, V_2, \ldots, V_k$ are nonempty and $s > 0$. For $W \in \mathbb{M}(\mathcal{R})$, let $\mathrm{cone}(W)$ be the convex cone in $\mathbb{R}^m$ spanned by the elements of W. It follows from the definition that if $W_1 \preceq W_2$, then $\mathrm{cone}(W_1) \subseteq \mathrm{cone}(W_2)$. So we get

$$\mathrm{cone}(V_1) \subseteq \mathrm{cone}(V_2) \subseteq \cdots \subseteq \mathrm{cone}(V_k) \subseteq \mathrm{cone}(V_1),$$

and it follows that

$$\mathrm{cone}(V_1) = \mathrm{cone}(V_2) = \cdots = \mathrm{cone}(V_k).$$

Let us call this cone $\mathcal{C}$. We choose an extremal ray γ of $\mathcal{C}$ and let $w \in \mathcal{R}$ be the indivisible vector that spans γ. We write $V_i = \lambda^{(i)}\{w\} \amalg Z_i$ where $\lambda^{(i)}$ is some partition and Z_i consists of elements in V_i that do not lie in the ray γ.

We claim that $\lambda^{(i)} \trianglelefteq \lambda^{(i+1)}$ for all i. If $V_i \preceq V_{i+1}$ because of (I), then V_{i+1} contains one additional vector, say v. If this vector v does not lie in the ray γ then $\lambda^{(i)} = \lambda^{(i+1)}$. If v does lie in γ, then $v = tw$ for some positive integer t, and $\lambda^{(i+1)}$ is obtained from $\lambda^{(i)}$ by inserting t. In particular, we have a strict inequality $\lambda^{(i)} \triangleleft \lambda^{(i+1)}$. If $V_i \preceq V_{i+1}$ because of (II), then it follows from the definition that $\lambda^{(i)} \preceq \lambda^{(i+1)}$. If $V_i \preceq V_{i+1}$ because of (III), then we have $V_i = \{v_1, v_2\} \amalg W \preceq \{v_1 + v_2\} \amalg W$ for some $W \in \mathbb{M}(\mathcal{R})$, where v_1, v_2 are linearly dependent. Now $v_1 + v_2$ cannot lie in any extremal ray, so it does not lie in γ, and at most one of the vectors v_1, v_2 will lie in γ. Similarly as in the case (I), we have $\lambda^{(i)} \trianglelefteq \lambda^{(i+1)}$ and the inequality is strict when v_1 or v_2 lies in γ. Since

$$\lambda^{(1)} \trianglelefteq \lambda^{(2)} \trianglelefteq \cdots \trianglelefteq \lambda^{(k)} \trianglelefteq \lambda^{(1)}$$

we get

$$\lambda^{(1)} = \lambda^{(2)} = \cdots = \lambda^{(k)}$$

From this we see that V_{i+1} is obtained from V_i by adding a vector not in γ in case (I), replacing some vectors not in γ by other vectors not in γ in case (II), or replacing a vector not in γ by two other vectors not in γ in case (3). It follows that

$$Z_1 \preceq Z_2 \preceq \cdots \preceq Z_k \preceq Z_1.$$

Since $|Z_i| < |V_i|$ for all i, we have $\min\{|Z_1|, |Z_2|, \ldots, |Z_k|\} < \min\{|V_1|, |V_2|, \ldots, |V_k|\}$. By induction, we have $Z_1 = Z_2 = \cdots = Z_k$ and therefore $V_1 = V_2 = \cdots = V_k$. □

7.3 Another basis

By Proposition 30, $(R^{\otimes \ell})^{S_\ell}$ is spanned by all $\mathbf{x}^V$ with $V \in \mathbb{M}_\ell(\mathcal{R})$. For V with more than ℓ elements we use the convention that $\mathbf{x}^V = 0$. This basis has a partial ordering, defined by $\mathbf{x}^V \preceq \mathbf{x}^W$ if and only if $V \preceq W$. We will study the multiplication in the invariant ring $(R^{\otimes \ell})^{S_\ell}$ in terms of this basis.

Suppose $U = \{u_1, u_2, \ldots, u_n\}$ is a multiset, then we define $\mathrm{stab}(U)$ as the cardinality of the stabilizer of $(u_1, u_2, \ldots, u_n)$ for the action of S_n on n-tuples. So if $v_1, v_2, \ldots, v_r$ are distinct, and U is the set that contains v_i with multiplicity k_i then $\mathrm{stab}(U) = k_1!k_2! \cdots k_r!$.

Definition 33. If $V, W \in \mathbb{M}(\mathcal{R})$ then a *partial matching* of V, W is a multiset

$$U = \{(v_1, w_1), (v_2, w_2), \ldots, (v_n, w_n)\}$$

such that V consists of all nonzero elements of the multiset $\{v_1, v_2, \ldots, v_n\}$ and W consists of all nonzero elements of the multiset $\{w_1, w_2, \ldots, w_n\}$ and $(0, 0)$ is not an element of U. We define $\Sigma U = \{v_1 + w_1, v_2 + w_2, \ldots, v_n + w_n\} \in \mathbb{M}(\mathcal{R})$.

Example 34. If $m = 1$ and $V = \{1, 1, 2\}$ and $W = \{3, 4\}$ then the possible partial matchings U are

$$\{(1,0),(1,0),(2,0),(0,3),(0,4)\}, \{(1,3),(1,0),(2,0),(0,4)\}, \{((1,4),(1,0),(2,0),(0,3)\},$$

$$\{(2,3),(1,0),(1,0),(0,4)\},\{(2,4),(1,0),(1,0),(0,3)\},$$
$$\{(1,3),(1,4),(2,0)\},\{(1,3),(2,4),(1,0)\},\{(1,4),(2,3),(1,0)\},$$

and ΣU is equal to

$$\{1,1,2,3,4\},\{4,1,2,4\},\{5,1,2,3\},\{5,1,1,4\},\{6,1,1,3\},\{4,5,2\},\{4,6,1\},\{5,5,1\}$$

respectively.

Proposition 35. *If $V, W \in \mathbb{M}(\mathcal{R})$ then we have*

$$\mathbf{x}^V \mathbf{x}^W = \sum_U \frac{\operatorname{stab}(\Sigma U)}{\operatorname{stab}(U)} \mathbf{x}^{\Sigma U} \tag{3}$$

where the sum is over all partial matchings U of V and W.

Proof. Suppose that $V = \{v_1, v_2, \ldots, v_r\}$, $W = \{w_1, w_2, \ldots, w_s\}$. It suffices to prove (3) over $\mathbb{Z}$ rather than over $\mathbb{F}_q$. From the definition, we have

$$\operatorname{stab}(V)\mathbf{x}^V = \sum_{i_1,i_2,\ldots,i_r} \prod_{\alpha=1}^{r} (\mathbf{x}^{v_\alpha})^{[i_\alpha]}$$

where the sum is over all $i_1, i_2, \ldots, i_r \in \{1, 2, \ldots, \ell\}$ that are distinct. We have

$$\operatorname{stab}(V)\operatorname{stab}(W)\mathbf{x}^V\mathbf{x}^W = \sum_{i_1,i_2,\ldots,i_r} \sum_{j_1,j_2,\ldots,j_s} \left(\prod_{\alpha=1}^{r} (\mathbf{x}^{v_\alpha})^{[i_\alpha]}\right)\left(\prod_{\beta=1}^{s} (\mathbf{x}^{w_\beta})^{[j_\beta]}\right).$$

Whenever $i_\alpha = j_\beta$, then the product of $(\mathbf{x}^{v_\alpha})^{[i_\alpha]}(\mathbf{x}^{w_\beta})^{[j_\beta]}$ gives a term $(\mathbf{x}^{v_\alpha+w_\beta})^{[i_\alpha]}$. A partial matching U comes from matching some of the indices in $\{i_1, i_2, \ldots, i_r\}$ with indices in $\{j_1, j_2, \ldots, j_s\}$. Because some of the v_i's may be the same, there may be several matchings of the indices that correspond to the matching U, namely there are exactly $\operatorname{stab}(V)\operatorname{stab}(W)/\operatorname{stab}(U)$ matchings of indices that correspond to the matching U and each of those result in a monomial sum $\operatorname{stab}(\Sigma U)\mathbf{x}^{\Sigma U}$. This yields

$$\operatorname{stab}(V)\operatorname{stab}(W)\mathbf{x}^V\mathbf{x}^W = \sum_U \frac{\operatorname{stab}(V)\operatorname{stab}(W)\operatorname{stab}(\Sigma U)}{\operatorname{stab}(U)} \mathbf{x}^{\Sigma U}$$

Dividing by $\operatorname{stab}(V)\operatorname{stab}(W)$ gives the desired result. □

Example 36. We go back to Example 34. We have

$$\mathbf{x}^{\{1,1,2\}}\mathbf{x}^{\{3,4\}} = \mathbf{x}^{\{1,1,2,3,4\}} + 2\mathbf{x}^{\{1,2,4,4\}} + \mathbf{x}^{\{1,2,3,5\}} + \mathbf{x}^{\{1,1,4,5\}} + \mathbf{x}^{\{1,1,3,6\}} + \mathbf{x}^{\{2,4,5\}} + \mathbf{x}^{\{1,4,6\}} + 2\mathbf{x}^{\{1,5,5\}}.$$

The coefficient of $\mathbf{x}^{\{1,5,5\}}$ is 2, because there is only 1 partial matching U with $\Sigma U = \{1, 5, 5\}$, namely $U = \{(1,4),(2,3),(1,0)\}$, and $\operatorname{stab}(\Sigma U)/\operatorname{stab}(U) = 2/1 = 2$.

Theorem 37. *The invariant ring $(R^{\otimes \ell})^{S_\ell}$ is generated by all invariants of the form $\mathbf{x}^{\{v\}j}$ where $v \in \mathcal{R}_{\mathrm{i}}$, and $1 \le j \le \ell$. Also, $(Q^{\otimes \ell})^{S_\ell}$ is generated by all $\mathbf{x}^{\{v\}j}$ with $v \in \mathcal{Q}_i$.*

Before we prove the theorem, we will need some other results.

Lemma 38. *Suppose $v \in \mathcal{R}_{\mathrm{i}}$ and $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_r)$ is a partition with $r \leq \ell$, and $\lambda' = (\lambda'_1, \lambda'_2, \ldots, \lambda'_s)$. Then*

$$\mathbf{x}^{\{v\}\lambda'_1}\mathbf{x}^{\{v\}\lambda'_2}\cdots\mathbf{x}^{\{v\}\lambda'_s}$$

is the sum of $\mathbf{x}^{\lambda\{v\}}$ and other polynomials of the form $\mathbf{x}^{\mu\{v\}}$ with $\mu \lhd \lambda$.

Proof. Suppose $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_r)$ is a partition, k is positive integer, and $v \in \mathcal{R}_{\mathrm{i}}$. If U is a partial matching of $\{v\}k = \{v, v, \ldots, v\}$ and $\lambda\{v\}$, then $\Sigma U = \mu\{v\}$ where μ is obtained from λ by increasing exactly k of the λ_i's. Here we use the convention that $\lambda_i = 0$ for $i > r$ and we can also increase a 0 to 1. There is a unique maximal μ we can get with the dominance ordering, namely if we increase the first k λ_i's then we get the maximum $\mu = (\lambda_1 + 1, \lambda_2 + 1, \ldots, \lambda_k + 1, \lambda_{k+1}, \lambda_{k+2}, \ldots)$. In terms of conjugate partitions, μ' is obtained from λ' by inserting k. If $\lambda_i \neq \lambda_j$, then we have $\mu_i \neq \mu_j$. This implies that $\mathrm{stab}(\mu\{v\}) = \mathrm{stab}(\Sigma\mu\{v\})$. So $\mathbf{x}^{\mu\{v\}}$ appears with coefficient 1 in the product $\mathbf{x}^{\{v\}k}\mathbf{x}^{\lambda\{v\}}$ and all the other $\mathbf{x}^U$ that appear satisfy $U \prec \mu\{v\}$.

From this it follows by induction on s that the largest term $\mathbf{x}^U$ that appears in

$$\mathbf{x}^{\{v\}\lambda'_1}\mathbf{x}^{\{v\}\lambda'_2}\cdots\mathbf{x}^{\{v\}\lambda'_s}$$

is $\mathbf{x}^{\{v\}(\lambda'_1,\ldots,\lambda'_s)} = \mathbf{x}^{\{v\}\lambda'} = \mathbf{x}^{\lambda\{v\}}$ and this term appears with coefficient 1. □

Lemma 39. *Suppose that $v_1, v_2, \ldots, v_r \in \mathcal{R}_{\mathrm{i}}$ are distinct (and therefore pairwise linearly independent), and $\lambda^{(1)}, \ldots, \lambda^{(r)}$ are partitions such that $\sum_{i=1}^r (\lambda^{(i)})'_1 \leq \ell$. If $V = \coprod_{i=1}^r \lambda^{(i)} v_i$ then*

$$\mathbf{x}^{\lambda^{(1)}\{v_1\}}\mathbf{x}^{\lambda^{(2)}\{v_2\}}\cdots\mathbf{x}^{\lambda^{(r)}\{v_r\}}$$

is the sum of $\mathbf{x}^V$ and other terms $\mathbf{x}^U$ with $U \prec V$.

Proof. Suppose that $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_r)$ is a partition, $W \in \mathbb{M}_{\ell-r}(\mathcal{R})$, and $v \in \mathcal{R}_{\mathrm{i}}$ is such that v and w are linearly independent for all $w \in W$. Let $Z = \lambda\{v\} \amalg W$. If U is any nontrivial contraction of $\lambda\{v\}$ and W then $U \prec Z$, because U is obtained from Z by replacing linearly independent vectors in U by their sum. If U is the trivial contraction, then $\Sigma U = Z$. Since $\lambda\{v\}$ and W are disjoint, we have $\mathrm{stab}(U) = \mathrm{stab}(\Sigma U)$. So $\mathbf{x}^Z$ appears in the product with coefficient 1.

Now the lemma follows by induction on r. □

Proposition 40. *Suppose that $V = \coprod_{i=1}^r \{v_i\}\lambda^{(i)}$ where $v_1, v_2, \ldots, v_r \in \mathcal{R}_{\mathrm{i}}$ are distinct, $\lambda^{(i)} = (\lambda_1^{(i)}, \lambda_2^{(i)}, \ldots, \lambda_{k_i}^{(i)})$ for $i = 1, 2, \ldots, r$, and $\sum_{i=1}^r \lambda_1^{(i)} \leq \ell$. Then*

$$g^V := \prod_{i=1}^{r}\prod_{j=1}^{k_i}\mathbf{x}^{\{v_i\}\lambda_j^{(i)}}$$

is the sum of $\mathbf{x}^V$ and other terms $\mathbf{x}^W$ with $W \prec V$.

Proof. By Lemma 38 (where $v = v_i$ and $\lambda = (\lambda^{(i)})'$) we see that

$$\prod_{j=1}^{k_i} \mathbf{x}^{\{v_i\}\lambda_j^{(i)}}$$

is the sum of $\mathbf{x}^{(\lambda^{(i)})'\{v_i\}} = \mathbf{x}^{\{v_i\}\lambda^{(i)}}$ and other terms $\mathbf{x}^{\mu\{v_i\}}$ with $\mu\{v_i\} \prec \{v_i\}\lambda^{(i)}$. Now the Proposition follows from Lemma 39. □

Theorem 41. *The polynomials g^V in Proposition 40 with $V \in \mathbb{M}_\ell(\mathcal{R})$ are a basis for $(R^{\otimes\ell})^{S_\ell}$ as an $\mathbb{F}_q$-vector space. The polynomials g^V with $V \in \mathbb{M}_\ell(\mathcal{Q})$ form a basis for $(Q^{\otimes\ell})^{S_\ell}$.*

Proof. Every $V \in \mathbb{M}_\ell(\mathcal{R})$ can uniquely be written in the form $V = \coprod_{i=1}^r \{v_i\}\lambda^{(i)}$ where $v_1, v_2, \ldots, v_r \in \mathcal{R}_\mathrm{i}$ are distinct, and $\sum_{i=1}^r \lambda_1^{(i)} \leq \ell$. Now g^V is the sum of $\mathbf{x}^V$ and other terms $\mathbf{x}^W$ with $W \prec V$. We claim that $\mathbf{x}^V$ can be uniquely expressed as an $\mathbb{F}_q$-linear combination of the g^W's. Suppose towards a contradiction that $\mathbf{x}^V$ is not in the span of the g^W's. Assume that V is minimal with respect to the ordering $\preceq$. Since $g^V - \mathbf{x}^V$ is a combination of $\mathbf{x}^W$ with $W \prec V$, we know by the minimality of V that $g^V - \mathbf{x}^V$ is an $\mathbb{F}_q$-linear combination of g^W's. Therefore, $\mathbf{x}^V$ is a linear combination of g^W's, which is a contradiction.

We claim that g^V's are free module generators. Suppose towards a contradiction that

$$\sum_{i=1}^r c_i g^{V_i} = 0.$$

with $r \geq 1$ and $c_1, c_2, \ldots, c_r$ nonzero. For some i, V_i is a maximal element in $\{V_1, V_2, \ldots, V_r\}$, i.e., there is no j with $V_i \prec V_j$. Now $\mathbf{x}^{V_i}$ appears in $\mathbf{g}^{V_i}$ but not in $\mathbf{g}^{V_j}$ for $j \neq i$. This implies $c_i = 0$, which is a contradiction.

The proof of the second statement in the Theorem is the same. We start with $V \in \mathbb{M}_\ell(\mathcal{R})$ (instead of $V \in \mathbb{M}_\ell(\mathcal{R})$) and note that the vectors $v_1, v_2, \ldots, v_r$ in the proof lie in $\mathcal{Q}_i$, and the W that appears lies in $\mathbb{M}_\ell(Q)$. □

Proof of Theorem 37. This follows from Theorem 41. □

To have finer control when composing polynomials it is useful to assign *weights* to the variables x_i, which are just positive integers c_i. Then the *weighted degree* of a monomial $x_1^{d_1} x_2^{d_2} \cdots x_r^{d_r}$ with respect to the c_i is defined as $\sum_{i=1}^r c_i d_i$.

Theorem 42. *Suppose that $v_1, v_2, \ldots, v_r \in \mathcal{R}_\mathrm{i}$ are distinct, $c_1, c_2, \ldots, c_r$ are positive integers, $g \in \mathbb{F}_q[x_1, x_2, \ldots, x_r]$ is a polynomial. Assign weight c_i to the variable x_i. If g has weighted degree $\leq \ell$ and $g(\mathbf{x}^{\{v_1\}c_1}, \mathbf{x}^{\{v_2\}c_2}, \ldots, \mathbf{x}^{\{v_r\}c_r}) = 0$. Then $g = 0$.*

Proof. If $x_1^{d_1} x_2^{d_2} \cdots x_r^{d_r}$ is a monomial of degree $\sum_{i=1}^r c_i d_i \leq \ell$, then $\prod_{i=1}^r (\mathbf{x}^{\{v_i\}c_i})^{d_i}$ is equal to g^V where $V = \coprod_{i=1}^r d_i\{v_i\}c_i \in \mathbb{M}(\mathcal{R})$. Note that the number of elements of V is equal to $\sum_{i=1}^r c_i \min\{d_i, 1\} \leq \sum_{i=1}^r c_i d_i \leq \ell$. So all the monomials in $\mathbf{x}^{\{v_1\}c_1}, \ldots, \mathbf{x}^{\{v_r\}c_r}$ that appear in $g(\mathbf{x}^{\{v_1\}c_1}, \mathbf{x}^{\{v_2\}c_2}, \ldots, \mathbf{x}^{\{v_r\}c_r})$ are linearly independent by Theorem 41. □

Using the above theorem we can finally prove Theorem 26.

Proof of Theorem 26. We follow the proof of Theorem 6 and modify it where necessary. We can write $M_i^\Sigma = \mathbf{x}^{\{v_i\}}$ for some vectors $v_1, v_2, \ldots, v_r \in \mathcal{R}_\mathrm{i}$. We get $G = c(H)$ for some $H \in (R^{\otimes \ell})^{S_\ell}$, where M_j has degree $\leq k$, G has degree $\leq kd < \ell$, c has degree e and H has degree $\leq kd/e$. By Theorem 37, we can write H as a polynomial in the generators $\mathbf{x}^{\{v\}j}$ with $1 \leq j \leq \ell$:

$$H = h(\mathbf{x}^{\{v_1\}}, \mathbf{x}^{\{v_2\}}, \ldots, \mathbf{x}^{\{v_r\}}, \mathbf{x}^{\{v_{r+1}\}c_{r+1}}, \ldots, \mathbf{x}^{\{v_s\}c_s})$$

for some polynomial $h(x_1, x_2, \ldots, x_s)$. If we give the variables $x_1, x_2, \ldots, x_r$ weight 1, and the variable x_j weight c_j for $j = r+1, r+2, \ldots, s$ then the weighted degree of h is $\leq kd/e < \ell/e$. We set $u(x_1, x_2, \ldots, x_s) = g(x_1, x_2, \ldots, x_r) - c(h(x_1, x_2, \ldots, x_s))$, then we have

$$u(\mathbf{x}^{\{v_1\}}, \mathbf{x}^{\{v_2\}}, \ldots, \mathbf{x}^{\{v_r\}}, \mathbf{x}^{\{v_{r+1}\}c_{r+1}}, \ldots, \mathbf{x}^{\{v_s\}c_s}) = 0.$$

Theorem 42 implies that $u = 0$. So $g(x_1, x_2, \ldots, x_r) = c(h(x_1, x_2, \ldots, x_s))$. So $h(x_1, x_2, \ldots, x_s) = h(x_1, x_2, \ldots, x_r)$ only depends on $x_1, x_2, \ldots, x_r$ and the degree of h is $\leq d/e$. □

Proof of Theorem 27. First note that the monomials are indecomposable since a monomial in Q of degree k has degree 1 in the variable x_m. We follow the proof of Theorem 16 and modify it where necessary. We can write $M_i^\Sigma = \mathbf{x}^{\{v_i\}}$ for some vectors $v_1, v_2, \ldots, v_r \in \mathcal{Q}_\mathrm{i}$. We get $G = c(H)$ for some $H \in (Q^{\otimes \ell})^{S_\ell}$, where M_j has degree k, G has degree $\leq kd < \ell$, c has degree e and H has degree $\leq kd/e$.

By Theorem 37, we can write H as a polynomial in the generators $\mathbf{x}^{\{v\}j}$ with $1 \leq j \leq \ell$ and $v \in \mathcal{Q}$:

$$H = h(\mathbf{x}^{\{v_1\}}, \mathbf{x}^{\{v_2\}}, \ldots, \mathbf{x}^{\{v_r\}}, \mathbf{x}^{\{v_{r+1}\}c_{r+1}}, \ldots, \mathbf{x}^{\{v_s\}c_s})$$

for some polynomial $h(x_1, x_2, \ldots, x_s)$ and $v_1, v_2, \ldots, v_s \in \mathcal{Q}$. We give the variables $x_1, x_2, \ldots, x_r$ weight 1, and the variable x_j weight c_j for $j = r+1, r+2, \ldots, s$. The monomial $\mathbf{x}^{\{v_i\}}$ has weight k for $i = 1, 2, \ldots, r$ and the monomial $\mathbf{x}^{\{v_i\}c_i}$ has weight kc_i for $i = r+1, r+2, \ldots, s$. It follows that the degree of h is $\leq d/e < \ell/e$. We set $u(x_1, x_2, \ldots, x_s) = g(x_1, x_2, \ldots, x_r) - c(h(x_1, x_2, \ldots, x_s))$, then we have

$$u(\mathbf{x}^{\{v_1\}}, \mathbf{x}^{\{v_2\}}, \ldots, \mathbf{x}^{\{v_r\}}, \mathbf{x}^{\{v_{r+1}\}c_{r+1}}, \ldots, \mathbf{x}^{\{v_s\}c_s}) = 0.$$

Theorem 42 implies that $u = 0$. So $g(x_1, x_2, \ldots, x_r) = c(h(x_1, x_2, \ldots, x_s))$. So $h(x_1, x_2, \ldots, x_s) = h(x_1, x_2, \ldots, x_r)$ only depends on $x_1, x_2, \ldots, x_r$ and the degree of h is $\leq d/e$. □

References

[ABEK08] Noga Alon, Ido Ben-Eliezer, and Michael Krivelevich. Small sample spaces cannot fool low degree polynomials. In *12th Workshop on Randomization and Computation (RANDOM)*, pages 266–275. Springer, 2008.

[AGHP92] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k-wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.

[BDN09] Arnaud Bodin, Pierre Dèbes, and Salah Najib. Indecomposable polynomials and their spectrum. *Acta Arith.*, 139(1):79–100, 2009.

[Bog05] Andrej Bogdanov. Pseudorandom generators for low degree polynomials. In *ACM Symp. on the Theory of Computing (STOC)*, pages 21–30, 2005.

[BV10] Andrej Bogdanov and Emanuele Viola. Pseudorandom bits for polynomials. *SIAM J. on Computing*, 39(6):2464–2486, 2010.

[Cla84] A. Clark. *Elements of Abstract Algebra*. Dover Books on Mathematics Series. Dover Publications, 1984.

[CN10] G. Chèze and S. Najib. Indecomposability of polynomials via Jacobian matrix. *J. Algebra*, 324(1):1–11, 2010.

[CT13] Gil Cohen and Amnon Ta-Shma. Pseudorandom generators for low degree polynomials from algebraic geometry codes. *Electron. Colloquium Comput. Complex.*, page 155, 2013.

[GX14] Venkatesan Guruswami and Chaoping Xing. Hitting sets for low-degree polynomials with optimal density. In *IEEE Conf. on Computational Complexity (CCC)*, pages 161–168. IEEE Computer Society, 2014.

[Kal95] Erich Kaltofen. Effective noether irreducibility forms and applications. *J. Comput. Syst. Sci.*, 50(2):274–295, 1995.

[KS01] Adam R. Klivans and Daniel A. Spielman. Randomness efficient identity testing of multivariate polynomials. In Jeffrey Scott Vitter, Paul G. Spirakis, and Mihalis Yannakakis, editors, *ACM Symp. on the Theory of Computing (STOC)*, pages 216–223. ACM, 2001.

[Lec07] Grégoire Lecerf. Improved dense multivariate polynomial factorization algorithms. *J. Symbolic Comput.*, 42(4):477–494, 2007.

[Lov09] Shachar Lovett. Unconditional pseudorandom generators for low degree polynomials. *Theory of Computing*, 5(1):69–82, 2009.

[Lu12] Chi-Jen Lu. Hitting set generators for sparse polynomials over any finite fields. In *IEEE Conf. on Computational Complexity (CCC)*, pages 280–286. IEEE Computer Society, 2012.

[LVW93] Michael Luby, Boban Veličković, and Avi Wigderson. Deterministic approximate counting of depth-2 circuits. In *2nd Israeli Symposium on Theoretical Computer Science (ISTCS)*, pages 18–24, 1993.

[Naj05] Salah Najib. Une généralisation de l'inégalité de Stein-Lorenzini. *Journal of Algebra*, 292(2):566–573, 2005.

[NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM J. on Computing*, 22(4):838–856, 1993.

[Raz87] Alexander Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Akademiya Nauk SSSR. Matematicheskie Zametki*, 41(4):598–607, 1987. English translation in Mathematical Notes of the Academy of Sci. of the USSR, 41(4):333-338, 1987.

[Sch04] Wolfgang Schmidt. *Equations Over Finite Fields: An Elementary Approach*. Kendrick Press, 2004.

[Vio07] Emanuele Viola. Pseudorandom bits for constant-depth circuits with few arbitrary symmetric gates. *SIAM J. on Computing*, 36(5):1387–1403, 2007.

[Vio09] Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d. *Computational Complexity*, 18(2):209–217, 2009.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il