THEORY OF COMPUTING, Volume 9 (21), 2013, pp. 665–683
www.theoryofcomputing.org

SPECIAL ISSUE: APPROX-RANDOM 2012

Extractors for Polynomial Sources over Fields of Constant Order and Small Characteristic*

Eli Ben-Sasson† Ariel Gabizon‡

Received November 13, 2012; Revised May 12, 2013; Published July 19, 2013

Abstract: A polynomial source of randomness over $\mathbb{F}_q^n$ is a random variable $X = f(Z)$ where f is a polynomial map and Z is a random variable distributed uniformly over $\mathbb{F}_q^r$ for some integer r. The three main parameters of interest associated with a polynomial source are the order q of the field, the (total) degree D of the map f, and the base-q logarithm of the size of the range of f over inputs in $\mathbb{F}_q^r$, denoted by k. For simplicity we call X a (q,D,k)-source.

Informally, an extractor for (q,D,k)-sources is a function $E : \mathbb{F}_q^n \to \{0,1\}^m$ such that the distribution of the random variable $E(X)$ is close to uniform over $\{0,1\}^m$ for any (q,D,k)-source X. Generally speaking, the problem of constructing extractors for such sources becomes harder as q and k decrease and as D increases. A rather large number of recent

*A conference version of this paper appeared in the Proceedings of RANDOM 2012 [1].

†Supported by funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 240258.

‡Supported by funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement number 240258.

ACM Classification: F.0, G.2.0, G.3

AMS Classification: 11T06, 11T23, 12F05

Key words and phrases: explicit construction, extractors, field extension, polynomials, polynomials - multivariate, Frobenius automorphisms, character sums

 DOI: 10.4086/toc.2013.v009a021

works in the area of derandomization have dealt with the problem of constructing extractors for $(q,1,k)$-sources, also known as "affine" sources. Constructing an extractor for non-affine sources, i. e., for $D>1$, is a much harder problem. Prior to the present work, only one construction was known, and that construction works only for fields of order much larger than n (Dvir et al., CCC 2009). In particular, even for $D=2$, no construction was known for any fixed finite field. In this work we construct extractors for (q,D,k)-sources for fields of constant order. Our proof builds on the work of DeVos and Gabizon (CCC 2010) on extractors for affine sources. Like the DeVos–Gabizon paper, our result makes crucial use of a theorem of Hou, Leung and Xiang (J. Number Theory 2002) which gives a lower bound on the dimension of products of subspaces.

1 Introduction

This paper is part of a long and active line of research devoted to the problem of "randomness extraction": Given a family of distributions all guaranteed to have a certain structure, devise a method that can convert a sample from any distribution in this family to a sequence of uniformly distributed bits—or at least a sequence *statistically close* to the uniform distribution. Usually, it is easy to prove that a random function is, with high probability, a good extractor for the given family, and the challenge is to give an explicit construction of such an extractor.

The first example of a randomness extraction problem was given by von Neumann [24], who gave an elegant solution[1] to the following problem: How can a biased coin with unknown bias be used to generate "fair" coin tosses? In this case the input distribution consists of independent identically distributed bits which makes the extraction task simpler. Since then many families of more complex distributions have been studied. Also, the concept of randomness extraction has proven to be useful for various applications. The reader is referred to the introduction of [10] for more details on the classes of distributions studied, references and motivation.

We now give a formal definition of extractors and related objects called dispersers.

Definition 1.1 (Extractors and dispersers)**.** Let Γ and Ω be some finite domains. Let $\mathcal{C}$ be a class of random variables taking values in Γ. We say a random variable P taking values in Ω is ε*-close to uniform* if for every $A \subseteq \Omega$,

$$\left|\Pr(P \in A) - \frac{|A|}{|\Omega|}\right| \le \varepsilon\,.$$

- Fix any $0 \le \varepsilon < 1$. A function $E:\Gamma \to \Omega$ is an ε*-extractor for* $\mathcal{C}$ if for every $X \in \mathcal{C}$, the random variable $E(X)$ is ε-close to uniform.

- A function $D:\Gamma \to \Omega$ is a *disperser for* $\mathcal{C}$ if for every $X \in \mathcal{C}$, the random variable $D(X)$ takes more than one value in Ω with nonzero probability.

[1] Von Neumann's algorithm as usually described does not precisely fit into the formal definition of a randomness extractor, as the algorithm is allowed not to produce an output.

1.1 Polynomial sources

In this paper we construct extractors for *polynomial sources*, which are distributions that are sampled by applying low-degree polynomials to uniform inputs as defined next. Throughout this paper, if Ω is a finite set, we let U_Ω denote the uniform distribution over Ω. By the *individual degree* of a multivariate polynomial f we mean the smallest d such that f has degree $\leq d$ in each variable.

Definition 1.2 (Polynomial sources)**.** Fix integers n,k,d with $k \leq n$ and a field $\mathbb{F}_q$. We define $\mathcal{M}[n,k,d]$ to be the set of mappings $f : \mathbb{F}_q^r \to \mathbb{F}_q^n$, where r is an integer counting the number of inputs to the source and

$$f(Z_1,\ldots,Z_r) = (f_1(Z_1,\ldots,Z_r),\ldots,f_n(Z_1,\ldots,Z_r))$$

such that

- for every $i \in [n]$, f_i is a polynomial in $\mathbb{F}_q[Z_1,\ldots,Z_r]$ of individual degree at most d.
- The range of f is of size at least q^k. Formally,

$$|\{f(z_1,\ldots,z_r) \mid (z_1,\ldots,z_r) \in \mathbb{F}_q^r\}| \geq q^k\,.$$

A (n,k,d)*-polynomial source* is a random variable of the form $f(U_{\mathbb{F}_q^r})$ for some and $f \in \mathcal{M}[n,k,d]$ with r inputs. (When the parameters n,k,d are clear from the context, we shall omit them, and simply use the term "polynomial source.")

Definition 1.3 (Polynomial-source extractors)**.** Let Ω be some finite set. A function $E : \mathbb{F}_q^n \to \Omega$ is a (k,d,D,ε)*-polynomial source extractor* if for every $f \in \mathcal{M}[n,k,d]$ of total degree at most D and r inputs, $E(f(U_{\mathbb{F}_q^r}))$ is ε-close to uniform (where $U_{\mathbb{F}_q^r}$ denotes the uniform distribution over $\mathbb{F}_q^r$).

Remark 1.4. A few words are in order regarding Definition 1.2.

- The number of inputs used by our source, denoted by r in Definition 1.2, does not affect the parameters of our extractors, and hence we omit this parameter from the definition of polynomial sources and extractors.
- In the context of extractors what might have seemed more natural is to require the random variable $f(U_{\mathbb{F}_q^r})$ to have *min-entropy*[2] at least $k \cdot \log q$. Our requirement on the size of the range of f is seemingly weaker, and suffices for our construction to work. (In particular, our result implies that, for some settings of field order and degree, when f has large range the random variable $f(U_{\mathbb{F}_q^r})$ is statistically close to a random variable that has at least a certain min-entropy.)
- Individual degree plays a larger role than total degree in our results. In fact, the first stage of our construction—constructing a non-constant polynomial over $\mathbb{F}_q$—requires a field of order depending *only* on individual degree. This is why it is more convenient to limit individual degree and not total degree in the definition of $\mathcal{M}[n,k,d]$.

[2]The min-entropy of a random variable X is the largest real number ℓ such that for every value x we have $\Pr(X = x) \leq 2^{-\ell}$. This is the standard measure of randomness in the context of extractors, originating from Chor and Goldreich [7].

Motivation To motivate our study of extractors for polynomial sources, we mention four distinct applications of such extractors for the simplest class of sources: affine ones, in which the degree of the source is 1 (see definition below). Demenkov and Kulikov [9] showed, using elementary methods, that any circuit over the full binary basis that computes an affine disperser for min-entropy rate $o(1)$ must contain at least $3n(1-o(1))$ gates, and this matches the previous best circuit lower bound of Blum from 1984 [4]. Another application of affine extractors was given by Viola [23] and independently by De and Watson [8] showing how to use them to construct extractors for bounded depth circuits. A third application was given by Ben-Sasson and Zewi [27] who showed how to construct two-source extractors and bipartite Ramsey graphs from affine extractors. Recent work of Guruswami [15] and of Dvir and Lovett [13] use "subspace evasive functions" which are closely related to affine extractors to get better algorithms for list-decoding of folded Reed-Solomon codes. These applications lead us to believe that extractors for general low-degree sources of the kind defined next will similarly be useful in other branches of computational complexity theory.

1.2 Previous work and our result

Polynomial-source extractors are a generalization of affine source extractors where the source is sampled by a degree-one map. There has been much work recently on affine-source extractors [2, 5, 26, 14, 10, 17] and related objects called affine-source dispersers [3, 22] where the output is required to be non-constant but not necessarily close to uniform.

Turning to extractors for non-affine, low-degree sources, the only previous work is by Dvir, Gabizon and Wigderson [12], and it requires large fields. In particular, to extract a single bit [12] needs a field of order at least n^c where $c > 1$ is a constant and n is number of inputs to the extractor, i. e., the number of outputs of the polynomial source.

(In a related albeit different vein, Dvir [11] constructed extractors for distributions that are uniform over low-degree algebraic varieties, which are sets of common zeros of a system of low-degree multivariate polynomials.)

In this work we construct polynomial-source extractors over much smaller fields than previously known, assuming the characteristic of the field is significantly smaller than the order of the field.

Theorem 1.5 (Main–Extractor)**.** *Fix a field $\mathbb{F}_q$ of characteristic p, integers $d, D, 4 \le k \le n$ where $n \ge 25$, and a positive integer $m < 1/2 \cdot \log_p q$. Let $\alpha = 3D \cdot (p \cdot d)^{3n/k}$. Assume that $q \ge 2 \cdot \alpha^2$. There is an explicit (k, d, D, ε)-polynomial source extractor $E : \mathbb{F}_q^n \to \mathbb{F}_p^m$ with error $\varepsilon = p^{m/2} \cdot \alpha \cdot q^{-1/2}$.*

In particular, when $D, n/k$, and p are constant, we get a polynomial-source extractor for fields of bounded order. We state such an instantiation.

Corollary 1.6 (Extractor for quadratic sources of min-entropy rate half over fields of characteristic 2)**.** *There is a universal constant C such that the following holds. For any $\varepsilon > 0$ and any $q > C/\varepsilon^2$ which is a power of 2, there is an explicit $(n/2, 2, 2, \varepsilon)$-polynomial source extractor $E : \mathbb{F}_q^n \to \{0, 1\}$.*

Non-Boolean dispersers for smaller fields Along the way to our proof we construct a weaker object called a *non-Boolean disperser*.

A non-Boolean disperser maps the source into a relatively small (but not $\{0,1\}$) domain and guarantees the output is non-constant. The advantage of this part of the construction is that it works for smaller fields than the extractor, and moreover, the field order for which it works depends only on the *individual* degrees of the source polynomials. In the theorem and corollary below we use an implicit isomorphism of $\mathbb{F}_q^n$ and $\mathbb{F}_{q^n}$. See an explanation of this in the beginning at the beginning of Section 3.

Theorem 1.7 (Main–Disperser)**.** *Fix a prime power $q = p^\ell$. Fix integers $k \le n$ and $d < s$ such that n is prime and s is a power of p. Fix a non-trivial $\mathbb{F}_q$-linear map $T : \mathbb{F}_q^n \to \mathbb{F}_q$. Let $u = \lceil (n-k)/(k-1) \rceil$. Define $P : \mathbb{F}_q^n \to \mathbb{F}_q$ by $P(x) \triangleq T(x^{1+s+s^2+\cdots+s^u})$. Assume that $q > d \cdot (s^{u+1}-1)/(s-1)$. Then, for any $f(\mathbf{Z}) = f(Z_1,\ldots,Z_r) \in \mathcal{M}[n,k,d]$, $P(f(\mathbf{Z}))$ is a non-constant function from $\mathbb{F}_q^r$ into $\mathbb{F}_q$.*

We instantiate this result for $\mathbb{F}_4$ which is the smallest field for which it works.

Corollary 1.8 (Disperser for min-entropy rate half over $\mathbb{F}_4$)**.** *Let n be prime. Define the function $P : \mathbb{F}_4^n \to \mathbb{F}_4$ as follows. Think of the input x as an element of $\mathbb{F}_{4^n}$ and compute x^3. Now output the first coordinate of the vector x^3. Then for any $f \in \mathcal{M}[n, \lceil n/2+1 \rceil, 1]$ (i. e., any multilinear $f \in \mathbb{F}_{4^n}[Z_1,\ldots,Z_r]$ that has range of size at least $4^{\lceil n/2+1 \rceil}$) the polynomial $P(f(Z_1,\ldots,Z_r))$ is a non-constant function from $\mathbb{F}_4^r$ into $\mathbb{F}_4$.*

Proof. We use Theorem 1.7, setting $k = \lceil n/2+1 \rceil$, $q = 4$, $d = 1$, $s = 2$ and $T : \mathbb{F}_4^n \to \mathbb{F}_4$ to be the map that projects to the first coordinate. This gives $u = 1$, and thus $P(x) = T(x^3)$ in this case. □

2 Overview of the proof

Our goal is to describe an explicit function $E : \mathbb{F}_q^n \to \{0,1\}^m$ such that for any (n,k,d)-polynomial source X we have that $E(X)$ is ε-close to the uniform distribution over $\{0,1\}^m$. We do this in two steps.

First we construct a function E_0, called a *non-Boolean disperser*, that is guaranteed to be non-constant on X, i. e., such that the random variable $Y = E_0(X)$ takes more than one value. This part is done in Section 4. Then we apply a second function E_1 to the output of E_0 and prove, using the fact that E_0 is a low-degree function in our case, that the distribution of $E_1(Y) = E_1(E_0(X))$ is ε-close to uniform. This "disperser–to–extractor" part is described in Sections 5 and 6. We now informally describe the two functions assuming for simplicity that the field $\mathbb{F}_q$ is of characteristic 2 and that n is prime. Before starting let us recall the notion of a Frobenius automorphism. If $\mathbb{K}$ is a finite field of characteristic 2 then the mapping

$$\sigma_i : \mathbb{K} \to \mathbb{K}, \qquad \sigma_i(z) = z^{2^i}$$

is a *Frobenius automorphism of* $\mathbb{K}$ *over* $\mathbb{F}_2$.

The three elementary properties of this mapping that we use below are

(i) its $\mathbb{F}_2$-*linearity:* $\sigma_i(a+b) = \sigma_i(a) + \sigma_i(b)$,

(ii) its *distinctness*: if $\mathbb{K}$ is an extension of $\mathbb{F}_2$ of degree at least t and $0 \le i < j \le t-1$ then σ_i and σ_j are different, and

(iii) its *dimension-preservation:* If $\mathbb{K} \supset \mathbb{F}_q \supset \mathbb{F}_2$ then $A \subset \mathbb{K}$ and $\sigma_i(A) \triangleq \{\sigma_i(a) \mid a \in A\}$ span spaces of equal dimension over $\mathbb{F}_q$ (see Claim 3.8).

A different view of low-degree sources The first part of our analysis uses a somewhat nonstandard view of low-degree sources that we need to highlight. The random variable X ranges over $\mathbb{F}_q^n$ and is the output of n degree-d polynomials over $\mathbb{F}_q$. Let

$$\mathbb{F}_q^{\leq d}[Z_1,\dots,Z_r]$$

denote the set of monomials over $\mathbb{F}_q$ of individual degree at most d where $d<q$. (We use Z variables to denote inputs of the polynomial source and X variables for its output.) Suppose the i-th coordinate of X is

$$X_i = P^{(i)}(Z_1,\dots,Z_r) = \sum_{M\in\mathbb{F}_q^{\leq d}[Z_1,\dots,Z_r]} a_M^{(i)}\cdot M(Z_1,\dots,Z_r)$$

where $a_M^{(i)}\in\mathbb{F}_q$ and $Z_1,\dots,Z_r$ are independent random variables distributed uniformly over $\mathbb{F}_q$. Applying an $\mathbb{F}_q$-linear bijection $\phi:\mathbb{F}_q^n\to\mathbb{F}_{q^n}$, let $a_M=\phi(a_M^{(1)},\dots,a_M^{(n)})$ denote the sequence of coefficients of the monomials M, viewed now as a single element in $\mathbb{F}_{q^n}$. Our nonstandard view is that our source is

$$X = P(Z_1,\dots,Z_r) = \sum_{M\in\mathbb{F}_q^{\leq d}[Z_1,\dots,Z_r]} a_M\cdot M(Z_1,\dots,Z_r) \tag{2.1}$$

where the coefficients a_M and the random variable X come from the "large" field $\mathbb{F}_{q^n}$ but the random variables $Z_1,\dots,Z_r$ still range over the "small" field $\mathbb{F}_q$. This large-field-small-field view will be important in what comes next. In particular, we shall use the following claim which reduces the problem of constructing a non-Boolean disperser to that of constructing a polynomial whose coefficients span $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$.

Claim 2.1 (Full-span polynomials are non-constant coordinate-wise)**.** *Suppose P has individual degree smaller than q. If the set of coefficients $A=\{a_M \mid \deg(M)>0\}$ appearing in (2.1) spans $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$ then $X_i=P^{(i)}(Z_1,\dots,Z_r)$ is a non-constant function on $\mathbb{F}_q^r$ for every $i\in\{1,\dots,n\}$.*

Proof. By way of contradiction. If $P^{(i)}$ is constant on $\mathbb{F}_q^r$ and has individual degrees smaller than q, then as a formal polynomial it is constant. This implies that all elements of A, as vectors in $\mathbb{F}_q^n$, are equal to zero in the i-th coordinate. Thus, A spans a strict subspace of $\mathbb{F}_{q^n}$ in contradiction to the assumption of the claim. $\square$

Non-Boolean disperser We start with the simplest nontrivial case to which our techniques apply and construct a non-Boolean disperser for homogeneous multilinear quadratic sources with min-entropy rate greater than half over the finite field with 4 elements (this is a special case of Corollary 1.8). Using $\binom{[r]}{2}$ to denote the set $\{(i,j)\mid 1\leq i<j\leq r\}$ and writing X as in (2.1) we get

$$X = \sum_{(i,j)\in\binom{r}{2}} a_{ij}Z_iZ_j, \qquad a_{ij}\in\mathbb{F}_{4^n} \tag{2.2}$$

where $Z_1,\dots,Z_r$ are uniformly and independently distributed over $\mathbb{F}_4$ and X takes more than $4^{n/2}$ distinct values. Let

$$A = \left\{ a_{ij} \;\middle|\; (i,j)\in\binom{[r]}{2} \right\} \tag{2.3}$$

denote the set of coefficients appearing in (2.2). In light of Claim 2.1 it suffices to construct E_0 such that $E_0(X)$, when written as a polynomial over $Z_1,\dots,Z_r$, has a set of coefficients that spans $\mathbb{F}_{4^n}$ over $\mathbb{F}_4$. (Then we "project" this polynomial onto, say, the first coordinate and get a non-constant function mapping into $\mathbb{F}_4$, i. e., a non-Boolean disperser.)

To do this we take the approach of DeVos and Gabizon [10] which uses the theorem of Hou, Leung and Xiang [16]. Assuming n is prime, this theorem implies that if $A,B \subset \mathbb{F}_{q^n}$ are sets spanning spaces of respective dimensions d_1,d_2 over $\mathbb{F}_q$, then the set of products

$$A\cdot B \triangleq \{a\cdot b \mid a\in A, b\in B\}$$

spans a subspace of $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$ of dimension at least $\min\{n, d_1+d_2-1\}$. Returning to our case and taking A as in (2.3), our first observation is that $\dim(\mathrm{span}(A)) > n/2$ because X is contained in $\mathrm{span}(A)$. So the theorem of [16] mentioned above implies that $\mathrm{span}(A\cdot A) = \mathbb{F}_{4^n}$. Consider what would happen if we could sample *twice* from X independently and take the product of the two samples in $\mathbb{F}_{4^n}$. Using $X', Z'_1,\dots,Z'_r$ to express the second sample we write this product as

$$X\cdot X' = \left(\sum_{(i,j)\in\binom{r}{2}} a_{ij}Z_iZ_j\right)\cdot\left(\sum_{(i',j')\in\binom{r}{2}} a_{i'j'}Z'_iZ'_j\right).$$

Opening the right-hand-side as a polynomial in $Z_1,\dots,Z_r,Z'_1,\dots,Z'_r$ we see that its set of coefficients is $A\cdot A$ which spans $\mathbb{F}_{4^n}$ over $\mathbb{F}_4$, as desired.[3]

Unfortunately we only have access to a *single* sample of X and have to make use of it. We use the fact that $\mathbb{F}_4$ is a degree 2 extension of a smaller field ($\mathbb{F}_2$) and hence has two distinct Frobenius automorphisms. And here comes our second observation: Taking the product of 2 distinct Frobenius automorphisms of a *single* sample of X has a similar effect to that of taking two independent samples of X! Indeed, take the product of $\sigma_0(X)$ and $\sigma_1(X)$ and, using the linearity of Frobenius mapping, expand as

$$\begin{aligned} X\cdot X^2 &= \left(\sum_{(i,j)\in\binom{r}{2}} a_{ij}Z_iZ_j\right)\cdot\left(\sum_{(i',j')\in\binom{r}{2}} a_{ij}^2 Z_i^2 Z_j^2\right) \\ &= \sum_{(i,j),(i',j')\in\binom{r}{2}} a_{ij}a_{i'j'}^2 Z_iZ_jZ_{i'}^2Z_{j'}^2 . \end{aligned}$$

The main point is that every element in the set of products of A and $A^2 \triangleq \left\{a^2 \mid a\in A\right\}$ appears as the coefficient of a monomial in the polynomial above and these monomials are distinct over $\mathbb{F}_4$. And the dimension-preservation of σ_1 implies that $\dim(\mathrm{span}(A^2)) = \dim(\mathrm{span}(A)) > n/2$. Consequently, the theorem of [16] implies that $A\cdot A^2$ spans $\mathbb{F}_{4^n}$ over $\mathbb{F}_4$, so by Claim 2.1 the function $E_0(X)$, which outputs the first coordinate of $X\cdot X^2$, is non-constant for X and this completes the sketch of our non-Boolean disperser for the special case of homogenous, quadratic, multilinear polynomials over $\mathbb{F}_4$.

[3] The same argument would work as well over the two-element field $\mathbb{F}_2$. The extension field is needed to deal with the case of a single source as explained next.

To extend this argument to general polynomial sources of individual degree $\leq d$ we carefully select a set of t distinct Frobenius automorphisms $\sigma_{i_0},\ldots,\sigma_{i_{t-1}}$ (assuming $\mathbb{F}_q$ is an extension-field of degree at least t) such that the mapping $f:(\mathbb{F}_q^{\leq d}[Z_1,\ldots,Z_r])^t \to \mathbb{F}_q[Z_1,\ldots,Z_r]$ given by

$$f(M_0,\ldots,M_{t-1}) = \prod_{j=0}^{t-1} \sigma_{i_j}(M_j) \mod (Z_1^q - Z_1,\ldots,Z_r^q - Z_r)$$

is injective. Then we argue, just as in the case above, that the function $g(X) \triangleq \prod_{j=0}^{t-1} \sigma_{i_j}(X)$ expands to a sum of distinct monomials with coefficients ranging over the product set $\hat{A} = \sigma_{i_0}(A)\cdots\sigma_{i_{t-1}}(A)$ where $\sigma(A) = \{\sigma(a) \mid a \in A\}$. The theorem of [16] is applied t times to conclude that $\hat{A}$ spans $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$. Now we apply Claim 2.1 and get that the first coordinate of $g(X)$ (viewing $g(X)$ as a tuple of n polynomials over $\mathbb{F}_q$) is a non-constant function. Details are provided in Section 4.

From dispersers to extractors This part is based on the work of Gabizon and Raz [14] and uses an important theorem of Weil [25]. This theorem implies the following. Suppose we evaluate a polynomial $g \in \mathbb{F}_q[Z_1,\ldots,Z_r]$ of small-enough degree $\deg(g) < \sqrt{q}$ on a uniformly random sample in $\mathbb{F}_q^r$ and then take the first bit of this evaluation (when viewing it as a vector over $\mathbb{F}_2$). Then, this bit will either be constant (in which case we then say that g is "degenerate," or close to the uniform distribution. Assuming our source is low-degree and the order q of the field is sufficiently large, we can argue that $\deg(E_0(X)) < \sqrt{q}$ because X is low-degree by assumption and E_0 is low-degree by construction. So to apply Weil's Theorem and get an extractor we only need to ensure that we have in hand a non-degenerate polynomial. Alas, we have relatively little control over the polynomial source so we need to transform it somehow into a non-degenerate one in a black-box manner. Here we apply another observation, proved by Swastik Kopparty, which says that $(E_0(X))^v$ is non-degenerate for odd[4] $v > 2$. This part is explained in Section 5. So we take $E_1(Y)$ to be the first[5] bit of Y^3 and using this observation and Weil's Theorem conclude that $E_1(E_0(X))$ is close to uniform. Analysis of the resulting extractor is given in the appendix.

3 Preliminaries

Notation: When we discuss identities between polynomials we only mean identities as *formal polynomials*. We will frequently alternate between viewing $\mathbf{x} \in \mathbb{F}_q^n$ as an element of either $\mathbb{F}_q^n$ or the field $\mathbb{F}_{q^n}$. When we do this we assume it is using an implicit bijective map $\phi : \mathbb{F}_q^n \to \mathbb{F}_{q^n}$ that is an isomorphism of vector spaces. That is, $\phi(t_1 \cdot a_1 + t_2 \cdot a_2) = t_1 \cdot \phi(a_1) + t_2 \cdot \phi(a_2)$ for any $t_1, t_2 \in \mathbb{F}_q$ and $a_1, a_2 \in \mathbb{F}_q^n$. Such ϕ is efficiently computable using standard representations of $\mathbb{F}_{q^n}$. (For details see for example the book of Lidl and Niederreiter [18].) For a set Ω we denote by U_Ω the uniform distribution over Ω.

3.1 Weil bounds for additive character sums

The seminal work of Weil [25] on the "Riemann hypothesis for curves over finite fields" implies very useful bounds on character sums. As we will see in this section, these bounds enable us to extract

[4]For characteristic $p > 2$ the criteria for v is a bit different: we need $p \nmid v$.

[5]In fact, we can output several bits. See Section 3.1 for details.

randomness from certain "low-degree distributions."

For background on characters of finite fields see [21] or Section 3.2 of [14]. The following version of the Weil bound was proved by Carlitz and Uchiyama [6].

Theorem 3.1 (Weil-Carlitz-Uchiyama bound). *Let $q = p^\ell$ for prime p and an integer ℓ. Let ψ be a non-trivial additive character of $\mathbb{F}_q$ (that is, not identically 1). Let $f(Z)$ be a polynomial in $\mathbb{F}_q[Z]$ of degree d. Suppose that f is not of the form $h^p + h + c$ for any $h \in \mathbb{F}_q[Z]$ and $c \in \mathbb{F}_q$. Then*

$$\left| \sum_{z \in \mathbb{F}_q} \psi(f(z)) \right| \leq (d-1) \cdot q^{1/2}.$$

We require the following generalization of Vazirani's XOR Lemma from Rao [20], appearing there as Lemma 4.2.

Lemma 3.2 (Rao's XOR lemma). *Let X be a distribution on a finite abelian group G s.t. $|\mathbf{E}(\psi(X))| \leq \varepsilon$ for any non-trivial character ψ of G. Then X is $\varepsilon \cdot \sqrt{|G|}$-close to uniform on G.*

The above lemma implies it suffices to bound additive character sums of a distribution over $\mathbb{F}_q$ in order to extract randomness. This is formalized in Lemma 3.4 below. To state the lemma we first define how to extract a few entries of an element in $\mathbb{F}_{p^\ell}$.

Definition 3.3 (Prefix projection). Let $q = p^\ell$ for prime p and an integer ℓ. Fix an isomorphism between $\mathbb{F}_q$ and $\mathbb{F}_p^\ell$ and view $x \in \mathbb{F}_q$ as $(x_1, \ldots, x_\ell) \in \mathbb{F}_p^\ell$. Fix an integer $m \leq \ell$. We define the *prefix projection function* $E_m : \mathbb{F}_q \to \mathbb{F}_p^m$ by $E_m(x) = E_m((x_1, \ldots, x_\ell)) \triangleq (x_1, \ldots, x_m)$.

Lemma 3.4 (XOR lemma for prefix projections). *Let $q = p^\ell$ for prime p and an integer ℓ. Let X be a distribution on $\mathbb{F}_q$ such that $|\mathbf{E}(\psi(X))| \leq \varepsilon$ for any non-trivial additive character ψ of $\mathbb{F}_q$. Then $E_m(X)$ is $p^{m/2} \cdot \varepsilon$-close to uniform.*

Proof. We claim that a function of the form $\psi'(a) \triangleq \psi(E_m(a))$ where ψ is a character of $\mathbb{F}_p^m$, is a character of $\mathbb{F}_q$: Let $\omega \in \mathbb{C}$ be a primitive p-th root of unity. The additive characters of $\mathbb{F}_q$ are exactly the functions $\psi : \mathbb{F}_q \to \mathbb{C}$ of the form $\psi(a) = \omega^{T(a)}$ where $T : \mathbb{F}_q \to \mathbb{F}_p$ is an $\mathbb{F}_p$-linear function and $T(a)$ is interpreted as an integer in $\{0, \ldots, p-1\}$. In particular, this includes such functions where T only looks at the first m coordinates of a (recall that we identify $\mathbb{F}_q$ with $\mathbb{F}_p^\ell$); and such functions in turn, are exactly those of the form $\psi(E_m(a))$ where ψ is a character of $\mathbb{F}_p^m$. Hence, from the assumption of the lemma $|\mathbf{E}(\psi(E_m(X)))| \leq \varepsilon$ for any non-trivial additive character of $\mathbb{F}_p^m$. From Lemma 3.2, we have that $E_m(X)$ is $p^{m/2} \cdot \varepsilon$-close to uniform.

□

Summing up the previous results we reach the statement that will be later used in analyzing our extractors.

Corollary 3.5 (Weil-Carlitz-Uchiyama for prefix projections). *Let $q = p^\ell$ for prime p and an integer ℓ. Let $f(Z)$ be a polynomial in $\mathbb{F}_q[Z]$ of degree d. Suppose that f is not of the form $h(Z)^p + h(Z) + c$ for any $h(Z) \in \mathbb{F}_q[Z]$ and $c \in \mathbb{F}_q$. Then $E_m(f(U_{\mathbb{F}_q}))$ is $p^{m/2} \cdot d/\sqrt{q}$-close to uniform.*

Proof. Follows immediately from Theorem 3.1 and Lemma 3.4. □

3.2 Dimension expansion of products

Recall that $\mathbb{F}_{q^n}$ is a vector space over $\mathbb{F}_q$ isomorphic to $\mathbb{F}_q^n$. For a set $A \subseteq \mathbb{F}_{q^n}$ we denote by $\dim(A)$ the dimension of the $\mathbb{F}_q$-span of A. For sets $A, B \subseteq \mathbb{F}_{q^n}$ let

$$A \cdot B \triangleq \{a \cdot b \mid a \in A, b \in B\}.$$

Hou, Leung and Xiang [16] show that such products expand in dimension. The following theorem is a corollary of Theorem 2.4 of [16].

Theorem 3.6 (Dimension expansion of products)**.** *Let $\mathbb{F}_q$ be any field, and let n be prime.*[6] *Let A and B be non-empty subsets of $\mathbb{F}_{q^n}$ such that $A, B \neq \{0\}$. Then*

$$\dim(A \cdot B) \geq \min\{n, \dim(A) + \dim(B) - 1\}.$$

In particular, if $A_1, \ldots, A_m$ are non-empty subsets of $\mathbb{F}_{q^n}$ such that for all $1 \leq i \leq m$, $\dim(A_i) \geq k$ for some $k \geq 1$. Then

$$\dim(A_1 \cdots A_m) \geq \min\{n, k \cdot m - (m-1)\}.$$

Remark 3.7. The definition of $A \cdot B$ is somewhat different from that in [16] where it is defined only for subspaces, and as the *span* of all possible products. The definition above will be more convenient for us. It is easy to see that Theorem 2.4 of [16] is equivalent to the theorem above with our definition. Still, we give a self-contained proof.[7]

Proof. First we note that it is enough to prove the theorem for linear subspaces A and B of dimension at least one: Given arbitrary sets A and B, let $A' \triangleq \operatorname{span}(A)$ and $B' \triangleq \operatorname{span}(B)$. If A and B both contain a non-zero element (as required in the theorem), then A' and B' are linear subspaces of dimension at least one. So we have that

$$\dim(A' \cdot B') \geq \min\{n, \dim(A') + \dim(B') - 1\} = \min\{n, \dim(A) + \dim(B) - 1\}.$$

Now, we observe that $\operatorname{span}(A' \cdot B') \subseteq \operatorname{span}(A \cdot B)$: An element of $A' \cdot B'$ has the form

$$\Big(\sum_i t_i \cdot a_i\Big) \cdot \Big(\sum_j s_j \cdot b_j\Big) = \sum_{i,j} t_i \cdot s_j \cdot a_i \cdot b_j\,,$$

where $a_i \in A, b_j \in B$ and $t_i, s_j \in \mathbb{F}_q$. This is obviously in $\operatorname{span}(A \cdot B)$. So $A' \cdot B' \subseteq \operatorname{span}(A \cdot B)$, and this implies $\operatorname{span}(A' \cdot B') \subseteq \operatorname{span}(A \cdot B)$. Therefore, the equation above implies

$$\dim(A \cdot B) \geq \min\{n, \dim(A) + \dim(B) - 1\}.$$

We now turn to proving the theorem for linear subspaces A and B of dimension at least one. We proceed by induction on $\dim(A)$. As a base, observe that the result holds trivially when $\dim(A) = 1$. For the inductive step, we may then assume that $\dim(A) > 1$. We may also assume that $B \neq \mathbb{F}_{q^n}$ as the theorem is immediate in this case.

[6] The theorem of [16] works also for non-prime n in which case the inequality involves the order of a certain subfield of $\mathbb{F}_{q^n}$.
[7] Also, see Section 3.2 of [10] for a self-contained proof using the definition of [16].

Note that we may freely replace A by $g\cdot A$ (or B by $g\cdot B$) for any $g\in\mathbb{F}_{q^n}$ as this has no effect on $\dim(A)$, $\dim(B)$, or $\dim(A\cdot B)$. By this operation, we may assume that $1\in A\cap B$. Since $\dim(A)>1$, we may choose $a\in A\setminus\mathbb{F}_q$. Let ℓ be the smallest nonnegative integer so that $a^\ell\notin B$. Note that such ℓ exists since $\mathbb{F}_{q^n}=\operatorname{span}(1,a,a^2,\dots,a^{n-1})$ for any $a\in\mathbb{F}_{q^n}\setminus\mathbb{F}_q$ as there are no non-trivial subfields $\mathbb{F}_q\subsetneq\mathbb{K}\subsetneq\mathbb{F}_{q^n}$ when n is prime, and $B\neq\mathbb{F}_{q^n}$. Furthermore, $\ell>0$ by the assumption that $1\in B$. Next, replace B by the set $a^{-(\ell-1)}\cdot B$. It now follows that $1\in B$ and $a\notin B$, so $A\cap B$ is a proper nonempty subset of A. In particular, $1\leq\dim(A\cap B)<\dim(A)$.

Consider the $\mathbb{F}_q$-linear subspaces $A\cap B$ and $A+B$ and observe that $(A\cap B)\cdot(A+B)\subseteq\operatorname{span}(A\cdot B)$. The next equation follows from this and the induction hypothesis applied to $A\cap B$ and $A+B$.

$$\begin{aligned}\dim(A\cdot B)&\geq\dim((A\cap B)\cdot(A+B))\\&\geq\min\{n,\dim(A\cap B)+\dim(A+B)-1\}\\&=\min\{n,\dim(A)+\dim(B)-1\}\,.\end{aligned}$$

This completes the proof. □

3.3 Frobenius automorphisms of $\mathbb{F}_q$

Let $q=p^\ell$ for prime p and let $i\geq 0$ be an integer. Raising to power p^i in $\mathbb{F}_q$ is known as a Frobenius automorphism of $\mathbb{F}_q$ over $\mathbb{F}_p$ and will play an important role. We record two useful and well-known properties of this automorphism that will be used in our proofs.

- **Linearity:** $\forall a,b\in\mathbb{F}_q$, $(a+b)^{p^i}=a^{p^i}+b^{p^i}$.
- **Bijection:** The map $x\to x^{p^i}$ over $\mathbb{F}_q$ is bijective. In particular, for $c\in\mathbb{F}_q$, c^{1/p^i} is always (uniquely) defined.

A useful fact following from these properties is that "taking the p-th power" of a set does not change its dimension.

Claim 3.8 (Dimension preservation). *Let $q=p^\ell$ from prime p and an integer ℓ. For an integer $i\geq 1$ and a set $A\subseteq\mathbb{F}_{q^n}$ let $A^{p^i}\triangleq\{a^{p^i}\mid a\in A\}$. Then* $\dim(A)=\dim(A^{p^i})$.

Proof. Let $\{a_1,\dots,a_k\}\subseteq A$ be a basis for the $\mathbb{F}_q$-span of A. Choose any $c_1,\dots,c_k\in\mathbb{F}_q$ that are not all zero. Then,

$$\sum_{j=1}^{k}c_j\cdot a_j^{p^i}=\left(\sum_{j=1}^{k}c_j^{1/p^i}\cdot a_j\right)^{p^i}\neq 0\,.$$

Thus $\{a_1^{p^i},\dots,a_k^{p^i}\}$ are independent over $\mathbb{F}_q$ and therefore $\dim(A^{p^i})\geq\dim(A)$. The reverse inequality is similar. □

4 The main construction

As before, we use r to denote the number of inputs of $f(Z_1,\ldots,Z_r) \in \mathcal{M}[n,k,d]$. We denote by $\mathcal{D}$ the product set $\{0,\ldots,d\}^r$. We use bold letters to denote vectors in $\mathbb{F}_q^r$. For example, $\mathbf{Z} = (Z_1,\ldots,Z_r)$. For an element $S = (s_1,\ldots,s_r) \in \mathcal{D}$ we use the notation

$$\mathbf{Z}^S \triangleq Z_1^{s_1} \cdots Z_r^{s_r}\,.$$

Fix $f = (f_1(\mathbf{Z}),\ldots,f_n(\mathbf{Z})) \in \mathcal{M}[n,k,d]$. For $1 \le j \le n$, we write

$$f_j(\mathbf{Z}) = \sum_{S\in\mathcal{D}} a_{j,S} \cdot \mathbf{Z}^S\,.$$

With the notation above, for $S \in \mathcal{D}$ let $a_S \triangleq (a_{1,S},\ldots,a_{n,S}) \in \mathbb{F}_q^n$. Using the isomorphism of the vectors spaces $\mathbb{F}_q^n$ and $\mathbb{F}_{q^n}$, we can view a_S as an element of $\mathbb{F}_{q^n}$ and write

$$f(\mathbf{Z}) = \sum_{S\in\mathcal{D}} a_S \cdot \mathbf{Z}^S. \tag{4.1}$$

That is, we view f as a multivariate polynomial with coefficients in $\mathbb{F}_{q^n}$. A crucial observation is that when f has large range the coefficients of f have large dimension.

Lemma 4.1 (Large range implies large span)**.** *Let $f \in \mathcal{M}[n,k,d]$. As in (4.1), write $f(\mathbf{Z}) = \sum_{S\in\mathcal{D}} a_S \cdot \mathbf{Z}^S$ where $a_S \in \mathbb{F}_{q^n}$. Then $\dim\{a_S\}_{S\in\mathcal{D}\setminus\{\mathbf{0}\}} \ge k$.*

Proof. The range of f over inputs in $\mathbb{F}_q^r$ is contained in an affine shift of the $\mathbb{F}_q$-linear span of $\{a_S\}_{S\in\mathcal{D}\setminus\{\mathbf{0}\}}$. Since this range is of size at least q^k, we must have $\dim\{a_S\}_{S\in\mathcal{D}\setminus\{\mathbf{0}\}} \ge k$. □

A simple but crucial observation from [10] is that a polynomial with coefficients in $\mathbb{F}_{q^n}$ whose non-constant coefficients span $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$ can be "projected" to a non-constant polynomial with coefficients in $\mathbb{F}_q$. We formalize this in the definition and lemma below.

Definition 4.2 (Full-span polynomial)**.** We say that a polynomial $G \in \mathbb{F}_{q^n}[\mathbf{Z}] = \mathbb{F}_{q^n}[Z_1,\ldots,Z_r]$ has *full span* if the coefficients of the non-constant monomials of G span $\mathbb{F}_{q^n}$ over $\mathbb{F}_q$.

Lemma 4.3 (Disperser for full-span polynomials)**.** *Suppose $G \in \mathbb{F}_{q^n}[\mathbf{Z}]$ has full span. Let $T : \mathbb{F}_{q^n} \to \mathbb{F}_q$ be a non-trivial $\mathbb{F}_q$-linear mapping. Then $T(G(\mathbf{Z}))$, as a function from $\mathbb{F}_q^r$ to $\mathbb{F}_q$, agrees with a non-constant polynomial in $\mathbb{F}_q[\mathbf{Z}]$ whose total and individual degrees are at most those of G.*

Proof. We write $G(\mathbf{Z}) = \sum_{S\in\mathcal{R}} a_S \cdot \mathbf{Z}^S$ for $a_S \in \mathbb{F}_{q^n}$, where $\mathcal{R} \subset \mathbb{N}^r$ denotes the set of tuples corresponding to the monomials of G. For every $\mathbf{x} = (x_1,\ldots,x_r) \in \mathbb{F}_q^r$, we have

$$T(G(\mathbf{x})) = T\left(\sum_{S\in\mathcal{R}} a_S \cdot \mathbf{x}^S\right) = \sum_{S\in\mathcal{R}} T(a_S)\cdot \mathbf{x}^S\,,$$

where the last inequality used the $\mathbb{F}_q$-linearity of T. Thus $T(G(\mathbf{Z}))$ agrees on all inputs in $\mathbb{F}_q^r$ with the polynomial $F(\mathbf{Z}) \triangleq \sum_{S\in\mathcal{R}} T(a_S)\cdot\mathbf{Z}^S$ which is in $\mathbb{F}_q[\mathbf{Z}]$. The full span of G means that $\dim\{a_S\}_{S\in\mathcal{R}\setminus\{\mathbf{0}\}} = n$.

Since T is a nontrivial linear map there is some $S \in \mathcal{R}$ such that $T(a_S) \neq 0$ and $S \neq \mathbf{0}$ and so F is a non-constant polynomial. As the monomials with non-zero coefficients in F are a subset of the monomials with non-zero coefficients in G, it is clear that the total and individual degrees of F are at most those of G. □

The previous lemma implies that to construct a disperser for polynomial sources it suffices to produce a function that increases the span of low-degree polynomials. We do this in the next theorem which is of paramount importance to this paper.

Theorem 4.4 (Product of distinct Frobenius automorphisms increases span)**.** *Fix a prime power $q = p^\ell$. Fix integers $k \leq n$ and $d < s$ such that n is prime and s is a power of p. (In particular, raising to power s^i is a Frobenius automorphism of $\mathbb{F}_q$ over $\mathbb{F}_p$.) Let $u = \lceil (n-k)/(k-1) \rceil$. Then for any $f(Z_1,\dots,Z_r) \in \mathcal{M}[n,k,d]$, the polynomial*

$$f^{1+s+s^2+\cdots+s^u}(Z_1,\dots,Z_r) = f(Z_1,\dots,Z_r) \cdot f^s(Z_1,\dots,Z_r) \cdots f^{s^u}(Z_1,\dots,Z_r)$$

has full span.

Proof. Fix $f \in \mathcal{M}[n,k,d]$. As in (4.1), write $f(\mathbf{Z}) = \sum_{S \in \mathcal{D}} a_S \cdot \mathbf{Z}^S$ with $a_S \in \mathbb{F}_{q^n}$.

$$f^{1+s+s^2+\cdots+s^u}(\mathbf{Z}) = \left(\sum_{S \in \mathcal{D}} a_S \cdot \mathbf{Z}^S \right)^{1+s+s^2+\cdots+s^u} = \prod_{i=0}^{u} \left(\sum_{S \in \mathcal{D}} a_S \cdot \mathbf{Z}^S \right)^{s^i}.$$

In what follows we use the notation $S_i = (S_{i,1},\dots,S_{i,r})$ and $S_i \cdot s^i = (S_{i,1} \cdot s^i, \dots, S_{i,r} \cdot s^i)$. Using the linearity of Frobenius automorphisms we continue the derivation and get

$$\begin{aligned} &= \prod_{i=0}^{u} \left(\sum_{S \in \mathcal{D}} a_S^{s^i} \cdot \mathbf{Z}^{S \cdot s^i} \right) = \sum_{S_0,\dots,S_u \in \mathcal{D}} \prod_{i=0}^{u} a_{S_i}^{s^i} \cdot \prod_{i=0}^{u} \mathbf{Z}^{S_i \cdot s^i} \\ &= \sum_{S_0,\dots,S_u \in \mathcal{D}} \prod_{i=0}^{u} a_{S_i}^{s^i} \cdot \prod_{i=0}^{u} \prod_{j=1}^{r} Z_j^{S_{i,j} \cdot s^i} = \sum_{S_0,\dots,S_u \in \mathcal{D}} A_{S_0,\dots,S_u} \cdot M_{S_0,\dots,S_u}(\mathbf{Z}), \end{aligned}$$

where

$$A_{S_0,\dots,S_u} = \prod_{i=0}^{u} a_{S_i}^{s^i} \quad \text{and} \quad M_{S_0,\dots,S_u}(\mathbf{Z}) = \prod_{i=0}^{u} \prod_{j=1}^{r} Z_j^{S_{i,j} \cdot s^i}.$$

The crucial observation is that if $(S_0,\dots,S_u)$ and $(S'_0,\dots,S'_u)$ are two distinct tuples of elements of $\mathcal{D}$ then the monomials $M_{S_0,\dots,S_u}(\mathbf{Z})$ and $M_{S'_0,\dots,S'_u}(\mathbf{Z})$ are distinct as well: Consider $j \in \{1,\dots,r\}$ such that $S_{i,j} \neq S'_{i,j}$ for some $0 \leq i \leq u$. Then Z_j is raised to power $\sum_{i=0}^{u} S_{i,j} \cdot s^i$ in $M_{S_0,\dots,S_u}(\mathbf{Z})$ and to power $\sum_{i=0}^{u} S'_{i,j} \cdot s^i$ in $M_{S'_0,\dots,S'_u}(\mathbf{Z})$. These powers are different as for all $0 \leq i \leq u$, $S_{i,j}, S'_{i,j} \leq d < s$, and there is only one way to write an integer in base s with "coefficients" smaller than s. Define

$$A \triangleq \{A_{S_0,\dots,S_u} \mid S_0,\dots,S_u \in \mathcal{D} \setminus \{\mathbf{0}\}\}.$$

For $0 \leq i \leq u$, define

$$B^{s^i} \triangleq \{a_S^{s^i} \mid S \in \mathcal{D} \setminus \{\mathbf{0}\}\}.$$

Note that $A = B^{s^0} \cdots B^{s^u}$. For all $0 \le i \le u$, by Lemma 4.1 and Claim 3.8 we have $\dim(B^{s^i}) \ge k$. Therefore, by Theorem 3.6 we get

$$\dim(A) \ge \min\{n, k \cdot (u+1) - u\} = n\,.$$

Our theorem follows by noticing that the coefficients of the non-constant monomials in $f^{1+s+s^2+\cdots+s^u}$ contain the set A, hence $f^{1+s+\cdots+s^u}$ has full span. □

Combining the lemma and theorem above we "project" into $\mathbb{F}_q$ and get a non-constant polynomial with coefficients in $\mathbb{F}_q$.

Theorem 4.5. *Fix a prime power $q = p^\ell$. Fix integers $k \le n$ and $d < s$ such that n is prime and s is a power of p. Fix a non-trivial $\mathbb{F}_q$-linear map $T : \mathbb{F}_{q^n} \to \mathbb{F}_q$. Let $u = \lceil (n-k)/(k-1) \rceil$. Define $P : \mathbb{F}_{q^n} \to \mathbb{F}_q$ by $P(x) \triangleq T(x^{1+s+s^2+\cdots+s^u})$. Fix any $f(Z_1,\ldots,Z_r) \in \mathcal{M}[n,k,d]$ of total degree D. Then $P(f(\mathbf{Z}))$, as a function on $\mathbb{F}_q^r$, agrees with a non-constant polynomial in $\mathbb{F}_q[\mathbf{Z}]$ of total degree at most $D \cdot (1+s+s^2+\cdots+s^u) < D \cdot s^{u+1}$ and individual degree at most $d \cdot (1+s+s^2+\cdots+s^u) = d \cdot (s^{u+1}-1)/(s-1)$.*

Proof. Follows immediately from Lemma 4.3 and Theorem 4.4. □

An immediate corollary is a construction of a "non-Boolean disperser" for polynomial sources.

Corollary 4.6. *Fix a prime power $q = p^\ell$. Fix integers $k \le n$ and $d < s$ such that n is prime and s is a power of p. Fix a non-trivial $\mathbb{F}_q$-linear map $T : \mathbb{F}_{q^n} \to \mathbb{F}_q$. Let $u = \lceil (n-k)/(k-1) \rceil$. Define $P : \mathbb{F}_{q^n} \to \mathbb{F}_q$ by $P(x) \triangleq T(x^{1+s+s^2+\cdots+s^u})$. Assume that $q > d \cdot (s^{u+1}-1)/(s-1)$. Then, for any $f(Z_1,\ldots,Z_r) \in \mathcal{M}[n,k,d]$ we have that $P(f(\mathbf{Z}))$ is a non-constant function from $\mathbb{F}_q^r$ into $\mathbb{F}_q$.*

Proof. Follows immediately from Theorem 4.5 by noticing that if $P(f)$ agrees with a non-constant polynomial whose individual degrees are smaller than q, then it is a non-constant function from $\mathbb{F}_q^r$ into $\mathbb{F}_q$. □

5 A useful criteria for the Weil bound

To get our main result we shall apply the Weil-Carlitz-Uchiyama bound for prefix projections (Corollary 3.5) to a certain polynomial $f \in \mathbb{F}_q[Z]$, and so we have to ensure that f is not of the "degenerate" form $h^p + h + c$ precluded by that bound. The common way to do this is to require $\gcd(\deg(f), p) = 1$ (cf., [14, 10]). However we have less control over the degree of the polynomial f we need to work with. For this reason, the following lemma will be very helpful to us. It gives us a simple way to "alter" f and get a polynomial that is not of the form $h^p + h + c$. The proof of the following lemma was shown to us by Swastik Kopparty.

Lemma 5.1 (Criteria for non-degenerateness). *Let $q = p^\ell$ for prime p and let $v \ge 2$ be an integer such that $p \nmid v$. Let $f \in \mathbb{F}_q[Z]$ be a non-constant polynomial. If f is of the form g^v for some $g \in \mathbb{F}_q[Z]$, it is not of the form $h^p + h + c$ for any $h \in \mathbb{F}_q[Z]$ and $c \in \mathbb{F}_q$.*

Proof. Suppose by way of contradiction there exists $f \in \mathbb{F}_q[Z]$ of degree $d \geq 1$ such that

$$f = g^v = h^p + h + c$$

for some $g, h \in \mathbb{F}_q[Z]$ and $c \in \mathbb{F}_q$. Fix such an f with minimal degree $d \geq 1$. It follows that $\deg(g) = d/v$ and $\deg(h) = d/p$. Taking a derivative in $\mathbb{F}_q[Z]$ of all 3 parts of the above equation we get

$$f'(Z) = v \cdot g^{v-1}(Z) \cdot g'(Z) = h'(Z),$$

where in the rightmost part we used the fact that the derivative of h^p is zero. Notice that $v \neq 0$ in $\mathbb{F}_q$ since $p \nmid v$. If $g' \not\equiv 0$ then this implies $\deg(h') \geq (v-1) \cdot \deg(g) = d \cdot (v-1)/v$. But $\deg(h') < d/p \leq d \cdot (v-1)/v$. (For the last inequality we use $p \geq 2$ and $v \geq 2$.) So g' and h' are the zero polynomial. It is not hard to see that this implies that all powers in g and h are multiples of p. So $g = g_1^p$ and $h = h_1^p$ for some $g_1, h_1 \in \mathbb{F}_q[Z]$. We now have

$$f = (g_1^p)^v = (h_1^p)^p + h_1^p + c.$$

This implies

$$g_1^v = h_1^p + h_1 + c^{1/p}.$$

(Recall that a p-th root always exists in $\mathbb{F}_q$.) Since g_1^v has positive degree smaller than $\deg(f) = d$, this contradicts the minimality of d and proves the lemma. $\square$

Reducing the multivariate case to the univariate case, we get the version of the Weil bound we need.

Lemma 5.2. *Let $q = p^\ell$ for a prime p and integer $\ell > 0$. Let $f(Z_1, \ldots, Z_r) \in \mathbb{F}_q[Z_1, \ldots, Z_r]$ be a non-constant polynomial of total degree $d < q$. Assume that $f = g^v$ for an integer $v \geq 2$ with $p \nmid v$ and some $g \in \mathbb{F}_q[Z_1, \ldots, Z_r]$. Let $m < \ell$ be a positive integer. Then $E_m(f(U_{\mathbb{F}_q^r}))$ is ε-close to uniform for $\varepsilon = p^{m/2} \cdot d \cdot q^{-1/2}$.*

Proof. We note first that there must be an $a = (a_1, \ldots, a_r) \in \mathbb{F}_q^r$ such that the univariate "line restriction" polynomial

$$f_a(Z) \triangleq f(a \cdot Z) = f(a_1 \cdot Z, \ldots, a_r \cdot Z)$$

has degree *exactly* d: The coefficient of Z^d in f_a is $f^d(a)$ where f^d is the d-homogeneous part of f, i.e., the sum of monomials of degree exactly d in f. By the Schwartz-Zippel lemma as $d < q$, there is an $a \in \mathbb{F}_q^r$ such that $f^d(a) \neq 0$ and therefore $f_a(Z)$ has degree d. Fix such an $a \in \mathbb{F}_q^r$. It follows that for all $b = (b_1, \ldots, b_r) \in \mathbb{F}_q^r$,

$$f_{a,b}(Z) \triangleq f(a \cdot Z + b) = f(a_1 \cdot Z + b_1, \ldots, a_r \cdot Z + b_r)$$

is non-constant, as the coefficient of Z^d in $f_{a,b}$ is also $f^d(a)$. Furthermore, for any $b \in \mathbb{F}_q^r$

$$f_{a,b} = f(a_1 \cdot Z + b_1, \ldots, a_r \cdot Z + b_r) = g^v(a_1 \cdot Z + b_1, \ldots, a_r \cdot Z + b_r),$$

and so $f_{a,b}$ is a v-th power of a polynomial in $\mathbb{F}_q[Z]$, and so by Lemma 5.1 is not of the form $h^p + h + c$ for any $h \in \mathbb{F}_q[Z]$ and $c \in \mathbb{F}_q$. As the distribution $f(U_{\mathbb{F}_q^r})$ is a convex combination of the distributions $f_{a,b}(U_{\mathbb{F}_q})$ for the different "shifts" $b \in \mathbb{F}_q^r$, the claim now follows from the Weil-Carlitz-Uchiyama bound for prefix projections (Corollary 3.5). $\square$

6 A polynomial-source extractor

We can now state and prove our main technical theorem, which immediately implies our main theorem on extractors for polynomial sources (Theorem 1.5).

Theorem 6.1 (Main–Extractors, parameterized version). *Fix a field $\mathbb{F}_q$ of characteristic p, integers $d, D, 2 \le k \le n$ where $n \ge 25$, and a positive integer $m < 1/2 \cdot \log_p q$. Let $\alpha = 3D \cdot (p \cdot d)^{\frac{1.2 \cdot n - k}{k-1} + 2}$. Assume that $q \ge 2 \cdot \alpha^2$. There is an explicit (k, d, D, ε)-polynomial source extractor $E : \mathbb{F}_q^n \to \mathbb{F}_p^m$ with error $\varepsilon = p^{m/2} \cdot \alpha \cdot q^{-1/2}$.*

Theorem 1.5 follows from the previous theorem by noticing that for $4 \le k \le n$,

$$\frac{1.2 \cdot n - k}{k-1} + 2 \le 3n/k \,.$$

Proof of Theorem 6.1. Choose a prime $n \le n' \le 1.2 \cdot n$ (which always exists for $n \ge 25$ according to Nagura's improvement of the Bertrand-Chebychev Theorem [19]). Given $f(Z_1, \dots, Z_r) \in \mathcal{M}[n, k, d]$ of total degree D we think of f as an element of $\mathcal{M}[n', k, d]$ by padding its output with zeros. Let s be the smallest power of p greater than d. Note that $s \le p \cdot d$. Let $P : \mathbb{F}_q^{n'} \to \mathbb{F}_q$ be the polynomial in Theorem 4.5 using s as above. If $p = 2$ let $v = 3$ and otherwise let $v = 2$. Let $E : \mathbb{F}_q^n \to \mathbb{F}_p^m$ be defined as $E(\mathbf{x}) \triangleq E_m(P^v(\mathbf{x}))$. From Theorem 4.5 we conclude that $P(f(\mathbf{Z}))$ is non-constant of degree at most $D \cdot s^{u+1}$ where

$$u = \lceil (n' - k)/(k-1) \rceil \le \frac{1.2 \cdot n - k}{k-1} + 1 \,.$$

Hence, from Lemma 5.2 we see that $E_m(P^v(f(U_{\mathbb{F}_q^r})))$ is ε-close to uniform for

$$\varepsilon = p^{m/2} \cdot v \cdot D \cdot s^{u+1} \cdot q^{-1/2} \le p^{m/2} \cdot 3D \cdot (p \cdot d)^{\frac{1.2 \cdot n - k}{k-1} + 2} \cdot q^{-1/2} = p^{m/2} \cdot \alpha \cdot q^{-1/2} \,. \qquad \square$$

Acknowledgements

We thank Swastik Kopparty for the proof of Lemma 5.1. We thank Swastik Kopparty and Shubhangi Saraf for helpful discussions. We thank Zeev Dvir for reading a previous version of this paper. The first author thanks Emanuele Viola for raising the question addressed in this paper. We thank the anonymous referees for a careful reading and helpful comments.

References

[1] ELI BEN-SASSON AND ARIEL GABIZON: Extractors for polynomials sources over constant-size fields of small characteristic. In *Proc. 16th Internat. Workshop on Randomization and Computation (RANDOM'12)*, pp. 399–410. Springer, 2012. [doi:10.1007/978-3-642-32512-0_34] 665

[2] ELI BEN-SASSON, SHLOMO HOORY, EYAL ROZENMAN, SALIL VADHAN, AND AVI WIGDERSON: Extractors for affine sources. Unpublished Manuscript, 2001. 668

[3] ELI BEN-SASSON AND SWASTIK KOPPARTY: Affine dispersers from subspace polynomials. *SIAM J. Comput.*, 41(4):880–914, 2012. Preliminary version in STOC'09. [doi:10.1137/110826254] 668

[4] NORBERT BLUM: A Boolean function requiring $3n$ network size. *Theoret. Comput. Sci.*, 28(3):337–345, 1983. [doi:10.1016/0304-3975(83)90029-4] 668

[5] JEAN BOURGAIN: On the construction of affine extractors. *Geometric and Functional Analysis*, 17(1):33–57, 2007. [doi:10.1007/s00039-007-0593-z] 668

[6] LEONARD CARLITZ AND SABURÔ UCHIYAMA: Bounds for exponential sums. *Duke Math. J.*, 24(1):37–41, 1957. [doi:10.1215/S0012-7094-57-02406-7] 673

[7] BENNY CHOR AND ODED GOLDREICH: Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988. Preliminary version in FOCS'85. [doi:10.1137/0217015] 667

[8] ANINDYA DE AND THOMAS WATSON: Extractors and lower bounds for locally samplable sources. *ACM Trans. Computation Theory*, 4(1):3, 2012. Preliminary version in RANDOM'11. [doi:10.1145/2141938.2141941] 668

[9] EVGENY DEMENKOV AND ALEXANDER S. KULIKOV: An elementary proof of a $3n - o(n)$ lower bound on the circuit complexity of affine dispersers. In *36th Internat. Symp. on Mathematical Foundations of Computer Science (MFCS'11)*, pp. 256–265, 2011. [doi:10.1007/978-3-642-22993-0_25] 668

[10] MATT DEVOS AND ARIEL GABIZON: Simple affine extractors using dimension expansion. In *Proc. 25th IEEE Conf. on Computational Complexity (CCC'10)*, pp. 50–57. IEEE Comp. Soc. Press, 2010. [doi:10.1109/CCC.2010.14] 666, 668, 671, 674, 676, 678

[11] ZEEV DVIR: Extractors for varieties. *Comput. Complexity*, 21(4):515–572, 2012. Preliminary version in CCC'09. [doi:10.1007/s00037-011-0023-3] 668

[12] ZEEV DVIR, ARIEL GABIZON, AND AVI WIGDERSON: Extractors and rank extractors for polynomial sources. *Comput. Complexity*, 18(1):1–58, 2009. Preliminary version in FOCS'07. [doi:10.1007/s00037-009-0258-4] 668

[13] ZEEV DVIR AND SHACHAR LOVETT: Subspace evasive sets. In *Proc. 44th STOC*, pp. 351–358. ACM Press, 2012. See also at ECCC. [doi:10.1145/2213977.2214010] 668

[14] ARIEL GABIZON AND RAN RAZ: Deterministic extractors for affine sources over large fields. *Combinatorica*, 28(4):415–440, 2008. Preliminary version in FOCS'05. [doi:10.1007/s00493-008-2259-3] 668, 672, 673, 678

[15] VENKATESAN GURUSWAMI: Linear-algebraic list decoding of folded Reed-Solomon codes. In *Proc. 26th IEEE Conf. on Computational Complexity (CCC'11)*, pp. 77–85. IEEE Comp. Soc. Press, 2011. [doi:10.1109/CCC.2011.22] 668

[16] XIANG-DONG HOU, KA HIN LEUNG, AND QING XIANG: A generalization of an addition theorem of Kneser. *J. Number Theory*, 97(1):1–9, 2002. [doi:10.1006/jnth.2002.2793] 671, 672, 674

[17] XIN LI: A new approach to affine extractors and dispersers. In *Proc. 26th IEEE Conf. on Computational Complexity (CCC'11)*, pp. 137–147. IEEE Comp. Soc. Press, 2011. [doi:10.1109/CCC.2011.27] 668

[18] RUDOLF LIDL AND HARALD NIEDERREITER: *Introduction to Finite Fields and Their Applications*. Cambridge Univ. Press, Cambridge, 1994. 672

[19] JITSURO NAGURA: On the interval containing at least one prime number. *Proc. Japan Acad.*, 28(4):177–181, 1952. [doi:10.3792/pja/1195570997] 680

[20] ANUP RAO: An exposition of Bourgain's 2-source extractor. *Electron. Colloq. on Comput. Complexity (ECCC)*, 14(034), 2007. ECCC. 673

[21] WOLFGANG M. SCHMIDT: *Equations over Finite Fields: An Elementary Approach*. Volume 536 of *Lecture Notes in Mathematics*. Springer, 1976. [doi:10.1007/BFb0080437] 673

[22] RONEN SHALTIEL: Dispersers for affine sources with sub-polynomial entropy. In *Proc. 52nd FOCS*, pp. 247–256. IEEE Comp. Soc. Press, 2011. Full version available at author's home page. [doi:10.1109/FOCS.2011.37] 668

[23] EMANUELE VIOLA: Extractors for circuit sources. In *Proc. 52nd FOCS*, pp. 220–229. IEEE Comp. Soc. Press, 2011. See also at ECCC. [doi:10.1109/FOCS.2011.20] 668

[24] JOHN VON NEUMANN: Various techniques used in connection with random digits. *Applied Math Series*, 12:36–38, 1951. 666

[25] ANDRÉ WEIL: On some exponential sums. *Proc. Nat. Acad. Sci. USA*, 34(5):204–207, 1948. PNAS. 672

[26] AMIR YEHUDAYOFF: Affine extractors over prime fields. *Combinatorica*, 31(2):245–256, 2011. [doi:10.1007/s00493-011-2604-9] 668

[27] NOGA ZEWI AND ELI BEN-SASSON: From affine to two-source extractors via approximate duality. In *Proc. 43rd STOC*, pp. 177–186. ACM Press, 2011. [doi:10.1145/1993636.1993661] 668

AUTHORS

Eli Ben-Sasson
associate professor
Technion, Haifa, Israel
eli@cs.technion.ac.il
http://eli.net.technion.ac.il/

Ariel Gabizon
visiting researcher
Technion, Haifa, Israel
ariel.gabizon@gmail.com
https://sites.google.com/site/arielgabizon1/

ABOUT THE AUTHORS

ELI BEN-SASSON graduated from the Hebrew University in 2001. His advisor was Avi Wigderson. He believes that the internet has killed the ritual of "telling a joke" (as opposed to forwarding it). He is sometimes described as "relaxed" though feels stressed, and enjoys the company of his wife and four kids.

ARIEL GABIZON graduated from the Weizmann Institue in 2008. His advisors were Ran Raz and Ronen Shaltiel. He is interested in using nice algebraic techniques for computer science problems, and in figuring out how powerful the randomized complexity classes are. He is a big supporter of practicing Vipassana meditation, and humanity gradually becoming vegan. He loves anything to do with creativity and free expression, like theater improv, singing, writing songs and dancing.