

Unbalanced Expanders and Randomness Extractors from Parvaresh-Vardy Codes*

VENKATESAN GURUSWAMI[†]
Dept. of Computer Science & Engineering
University of Washington
Seattle, WA 98195
venkat@cs.washington.edu

CHRISTOPHER UMANS[‡]
Computer Science Department
California Institute of Technology
Pasadena, CA 91125
umans@cs.caltech.edu

SALIL VADHAN[§]
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138
salil@eecs.harvard.edu

Abstract

We give an improved explicit construction of highly unbalanced bipartite expander graphs with expansion arbitrarily close to the degree (which is polylogarithmic in the number of vertices). Both the degree and the number of right-hand vertices are polynomially close to optimal, whereas the previous constructions of Ta-Shma, Umans, and Zuckerman (STOC '01) required at least one of these to be quasipolynomial in the optimal. Our expanders have a short and self-contained description and analysis, based on the ideas underlying the recent list-decodable error-correcting codes of Parvaresh and Vardy (FOCS '05).

Our expanders can be interpreted as near-optimal “randomness condensers,” that reduce the task of extracting randomness from sources of arbitrary min-entropy rate to extracting randomness from sources of min-entropy rate arbitrarily close to 1, which is a much easier task. Using this connection, we obtain a new construction of randomness extractors that is optimal up to constant factors, while being much simpler than the previous construction of Lu et al. (STOC '03) and improving upon it when the error parameter is small (e.g. $1/\text{poly}(n)$).

Keywords: expander graphs, randomness extractors,

*A preliminary version of this paper appeared on *ECCC* under the title “Extractors and Condensers from Univariate Polynomials” [10].

[†]Supported by NSF CCF-0343672, a Sloan Research Fellowship, and a David and Lucile Packard Foundation Fellowship.

[‡]Supported by NSF CCF-0346991, BSF 2004329, a Sloan Research Fellowship, and an Okawa Foundation research grant.

[§]Supported by NSF CCF-01133096, ONR N00014-04-1-0478, and US-Israel BSF 2002246.

error-correcting codes, list decoding, condensers.

1 Introduction

One of the exciting developments in the theory of pseudorandomness has been the discovery of intimate connections between a number of fundamental and widely studied objects — expander graphs, randomness extractors, list-decodable error-correcting codes, pseudorandom generators, and randomness-efficient samplers. Indeed, substantial advances have been made in our understanding of each of these objects by translating intuitions and techniques from the study of one to the study of another. In this work, we continue this in tradition. Specifically, we use ideas from recent breakthrough constructions of list-decodable codes, due to Parvaresh and Vardy [22], to give improved and simplified constructions of both unbalanced bipartite expander graphs and randomness extractors.

1.1 Unbalanced expander graphs

Expanders are graphs that are sparse yet very highly connected. They have a wide variety of applications in theoretical computer science, and there is a rich body of work on constructions and properties of expanders. (See the survey [11]). The classic measure of the connectivity of an expander is *vertex expansion*, which asks that every set S of vertices that is not too large has significantly more than $|S|$ neighbors. This property is formalized for bipartite graphs through the following definitions.

Definition 1.1. A bipartite (multi)graph with N left-vertices, M right-vertices, and left-degree D is specified by a function $\Gamma : [N] \times [D] \rightarrow [M]$, where $\Gamma(x, y)$ denotes the y 'th neighbor of x . For a set $S \subseteq [N]$, we write $\Gamma(S)$ to denote its set of neighbors $\{\Gamma(x, y) : x \in S, y \in [D]\}$.

Definition 1.2. A bipartite graph $\Gamma : [N] \times [D] \rightarrow [M]$ is a (K, A) expander if for every set $S \subseteq [N]$ of size K , we have $|\Gamma(S)| \geq A \cdot K$. It is a $(\leq K_{max}, A)$ expander if it is a (K, A) expander for all $K \leq K_{max}$.

The typical goals in constructing expanders are to maximize the expansion factor A and minimize the degree D . In this work, we are also interested minimizing the size M of the right-hand side, so $M \ll N$ and the graph is highly unbalanced. Intuitively, this makes expansion harder to achieve because there is less room in which to expand. Using the probabilistic method, it can be shown that very good expanders exist — with expansion $A = (1 - \varepsilon) \cdot D$, degree $D = O(\log(N/M)/\varepsilon)$, and $M = O(K_{max}D/\varepsilon) = O(K_{max}A/\varepsilon)$ right vertices. Thus, if $M \leq N^c$ for some constant $c < 1$, then the degree is logarithmic in N , and logarithmic degree is in fact necessary if $M = O(K_{max}A)$.¹ However, applications of expanders require *explicit constructions* — ones where the neighbor function Γ is computable in polynomial time (in its input length, $\log N + \log D$) — and the best known explicit constructions still do not match the ones given by the probabilistic method.

Most classic constructions of expanders, such as [19, 5, 18, 20], focus on the balanced (or non-bipartite) case (i.e. $M = N$), and thus are able to achieve constant degree $D = O(1)$. The expansion properties of these constructions are typically proven by bounding the second-largest eigenvalue of the adjacency matrix of the graph. While such ‘spectral’ expansion implies various combinatorial forms of expansion (e.g., vertex expansion) and many other useful properties, it seems insufficient for deducing vertex expansion beyond $D/2$ [14] or for obtaining highly imbalanced expanders with polylogarithmic degree [38]. This is unfortunate, because some applications of expanders require these properties. A beautiful example of such an application was given by Buhrman et. al. [1]. They showed that a $(\leq K_{max}, A)$ expander with N left-vertices, M right-vertices, and expansion $A = (1 - \varepsilon)D$ yields a method for storing any set $S \subseteq [N]$ of size at most $K_{max}/2$ in an M -bit data structure so that membership in S can be probabilistically tested by reading only *one bit* of the data structure. An optimal expander would give $M = O(K_{max} \log N)$, only a constant factor more than what is needed to represent an arbitrary set of size $K_{max}/2$

¹More generally, the degree must be at least $\Omega(\log(N/K_{max})/\log(M/(K_{max}A)))$, as follows from the lower bounds on the degree of dispersers [23].

(even without supporting efficient membership queries).²

Explicit constructions of expanders with expansion $A = (1 - \varepsilon)D$ were obtained by Ta-Shma, Umans, and Zuckerman [33] for the highly imbalanced (and nonconstant-degree) case and Capalbo et al. [2] for the balanced (and constant-degree) case. The constructions of Ta-Shma et al. [33] can make either one of the degree or right-hand side polynomially larger than the nonconstructive bounds mentioned above, at the price of making the other quasipolynomially larger. That is, one of their constructions gives $D = \text{poly}(\log N)$ and $M = \text{quasipoly}(K_{max}D) \stackrel{\text{def}}{=} \exp(\text{poly}(\log(K_{max}D)))$, whereas the other gives $D = \text{quasipoly}(\log N)$ and $M = \text{poly}(K_{max}D)$. The quasipolynomial bounds were improved recently in [32], but remained superpolynomial.

We are able to simultaneously achieve $D = \text{poly}(\log N)$ and $M = \text{poly}(KD)$, in fact with a good tradeoff between the degrees of these two polynomials.

Theorem 1.1. For all constants $\alpha > 0$, every $N \in \mathbb{N}$, $K_{max} \leq N$, and $\varepsilon > 0$, there is an explicit $(\leq K_{max}, (1 - \varepsilon)D)$ expander $\Gamma : [N] \times [D] \rightarrow [M]$ with degree $D = O((\log N)(\log K_{max})/\varepsilon)^{1+1/\alpha}$ and $M \leq D^2 \cdot K_{max}^{1+\alpha}$. Moreover, D is a power of 2.

The construction of our expanders is based on the recent list-decodable codes of Parvaresh and Vardy [22], and can be described quite simply. The proof of the expansion property is inspired by the list-decoding algorithm for the PV codes, and is short and self-contained. An overview of this ‘list-decoding approach’ to proving expansion is provided in Section 2.1.

1.2 Randomness extractors

One of the main motivations and applications of our expander construction is the construction of *randomness extractors*. These are functions that convert weak random sources, which may have biases and correlations, into almost-perfect random sources. For general models of weak random sources, this is impossible, so the extractor is also provided with a short ‘seed’ of truly random bits to help with the extraction [21]. This seed can be so short (e.g. of logarithmic length), that one can often eliminate the need for any truly random bits by enumerating all choices for the seed. For example, this allows extractors to be used for efficiently simulating randomized algorithms using only a weak random source [39, 21]. Extractors have also found a

²We note that to implement the data structure of [1], it is not sufficient that the expander be explicit in terms of its neighbor function Γ being efficiently computable, but it is also necessary that the expander have efficient ‘decoding algorithms’. Such expanders were constructed by Ta-Shma [31]. Our expanders also have efficient decoding algorithms, but they only provide improvements over [1, 31] for this application when the set size is relatively small, e.g. $(\log N)^{\omega(1)} \leq K_{max} \leq \exp((\log \log N)^3)$.

wide variety of other applications in theoretical computer science beyond their original motivating application, and thus a long body of work has been devoted to providing efficient constructions of extractors. (See the survey of Shaltiel [26].)

To formalize the notion of an extractor, we need a few definitions. Following [3, 39], the randomness in a source is measured by *min-entropy*: a random variable \mathbf{X} has min-entropy at least k iff $\Pr[\mathbf{X} = x] \leq 2^{-k}$ for all x . A random variable \mathbf{Z} is ε -close to a distribution D if for all events A , $\Pr[\mathbf{Z} \in A]$ differs from the probability of A under the distribution D by at most ε . Then an extractor is defined as follows:

Definition 1.3 ([21]). *A function $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a (k, ε) extractor if for every \mathbf{X} with min-entropy at least k , $E(\mathbf{X}, \mathbf{Y})$ is ε -close to uniform, when \mathbf{Y} is uniformly distributed on $\{0, 1\}^d$. An extractor is explicit if it is computable in polynomial time.*

The competing goals when constructing extractors are to obtain a short seed length and to obtain a long output length. Nonconstructively, it is possible to simultaneously have a seed length $d = \log n + 2\log(1/\varepsilon) + O(1)$ and an output length of $m = k + d - 2\log(1/\varepsilon) - O(1)$. It remains open to match these parameters with an explicit construction.

Building on a long line of work, Lu et al. [17] achieved seed length and output length that are within constant factors of optimal, provided that the error parameter ε is not too small. More precisely, they achieve seed length $d = O(\log n)$ and output length $m = (1 - \alpha)k$ for $\varepsilon \geq n^{-1/\log^{(c)} n}$, where α and c are any two positive constants. For general ε , they pay with either a larger seed length of $d = O((\log^* n)^2 \log n + \log(1/\varepsilon))$, or a smaller output length of $m = k/\log^{(c)} n$ for any constant c .

In this work, we also achieve extractors that are optimal up to constant factors, but are able to handle an error parameter ε that is even exponentially small:

Theorem 1.2 (extractor). *For every constant $\alpha > 0$, and all positive integers n, k and all $\varepsilon > \exp(-n/2^{O(\log^* n)})$, there is an explicit construction of a (k, ε) extractor $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\log n + \log(1/\varepsilon))$ and $m \geq (1 - \alpha)k$.*

Our extractor is also substantially simpler than that of [17], which is a complex recursive construction involving many tools. The key component in our construction is the interpretation of our expander graph as a *randomness condenser*:

Definition 1.4. *A function $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is an $k \rightarrow_\varepsilon k'$ condenser if for every \mathbf{X} with min-entropy at least k , $C(\mathbf{X}, \mathbf{Y})$ is ε -close to a distribution with min-entropy k' , when \mathbf{Y} is uniformly distributed on $\{0, 1\}^d$. A*

condenser is explicit if it is computable in polynomial time. A condenser is called lossless if $k' = k + d$.

Observe that a $k \rightarrow_\varepsilon k'$ condenser with output length $m = k'$ is an extractor, because the unique distribution on $\{0, 1\}^m$ with min-entropy m is the uniform distribution. Condensers are a natural stepping-stone to constructing extractors, as they can be used to increase the *entropy rate* (the ratio of the min-entropy in a random variable to the length of the strings over which it is distributed), and it is often easier to construct extractors when the entropy rate is high. Condensers have also been used extensively in less obvious ways to build extractors, often as part of complex recursive constructions (e.g., [12, 25, 17]). Non-constructively, there exist *lossless* condensers with seed length $d = \log n + \log(1/\varepsilon) + O(1)$, and output length $m = k + d + \log(1/\varepsilon) + O(1)$.

As shown by [33], lossless condensers are equivalent to bipartite expanders with expansion close to the degree. Applying this connection to Theorem 1.1, we obtain the following condenser:

Theorem 1.3. *For all constants $\alpha > 0$, and every $n \in \mathbb{N}$, $k \leq n$, and $\varepsilon > 0$, there is an explicit $(k \rightarrow_\varepsilon k + d)$ (lossless) condenser $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = (1 + 1/\alpha) \cdot (\log n + \log k + \log(1/\varepsilon)) + O(1)$ and $m \leq 2d + (1 + \alpha)k$.*

Consider the case that α is a constant close to 0. Then the condenser has seed length $O(\log(n/\varepsilon))$ and output min-entropy rate roughly $1/(1 + \alpha)$. Thus, the task of constructing extractors for arbitrary min-entropy is reduced to that of constructing extractors for min-entropy rate close to 1, which is a much easier task. Indeed, when ε is constant, we can use a well-known and simple extractor based on expander walks. When ε is sub-constant, we use Zuckerman's extractor [40] to obtain the proper dependence on ε . Thus, we obtain Theorem 1.2.

1.3 Organization

We begin in Section 2 with a high level overview of our construction and proof method. In Section 3 we describe and analyze our expander construction. This expander construction implies a lossless condenser construction, which is discussed in Section 4. By applying extractors for high min-entropy to the output of this condenser, we obtain our new extractors, also in Section 4. In Section 5, we analyze a lossy version of our main construction, which allows us to minimize the seed lengths of the resulting condensers and extractors. In Section 6 we analyze a variant of the construction that utilizes Reed-Solomon codes and is a univariate analogue of [27], and whose analysis is based on [8]. Finally we conclude in Section 7 with some open problems.

1.4 Notation

Throughout this paper, we use boldface capital letters for random variables (e.g., “ \mathbf{X} ”), capital letters for indeterminates, and lower case letters for elements of a set. Also throughout the paper, \mathbf{U}_t is the random variable uniformly distributed on $\{0, 1\}^t$. The *support* of a random variable \mathbf{X} is $\text{supp}(\mathbf{X}) \stackrel{\text{def}}{=} \{x : \Pr[\mathbf{X} = x] > 0\}$. The *statistical distance* between random variables (or distributions) \mathbf{X} and \mathbf{Y} is $\max_T |\Pr[\mathbf{X} \in T] - \Pr[\mathbf{Y} \in T]|$. We say \mathbf{X} and \mathbf{Y} are ε -close if their statistical distance is at most ε . All logs are base 2.

2 Overview of our approach

In this section we give a high level overview of our construction and the proof technique.

2.1 Expansion via list-decoding

Before explaining our approach, we briefly review the basics of list-decodable codes. A *code* is mapping $C : [N] \rightarrow [M]^D$, encoding messages of bit-length $n = \log_2 N$ to D symbols over the alphabet $[M]$. (Contrary to the usual convention in coding theory, we use different alphabets for the message and the encoding.) The *rate* of such a code is $\rho = n/(D \log_2 M)$. We say that C is (ε, K) *list-decodable* if for every $r \in [M]^D$, the set $\text{LIST}(r, \varepsilon) \stackrel{\text{def}}{=} \{x : \Pr_y[C(x)_y = r_y] \geq \varepsilon\}$ is of size at most K . We think of r as a *received word* obtained by corrupting all but an ε fraction of symbols in some codeword. The list-decodability property says that there are not too many messages x that could have led to the received word r . The goal in constructing list-decodable codes is to optimize the tradeoff between the agreement ε and the rate ρ , which are typically constants independent of the message length n . Both the alphabet size M and the list-size K should be relatively small (e.g. constant or $\text{poly}(n)$). Computationally, we would like efficient algorithms both for computing $C(x)$ given x and for enumerating the messages in $\text{LIST}(r, \varepsilon)$ given a received word r .

The classic Reed-Solomon codes were shown to achieve these properties with polynomial-time list-decoding in the seminal work of Sudan [29]. The tradeoff between ε and ρ was improved by Guruswami and Sudan [9], and no better result was known for a number of years. Recently, Parvaresh and Vardy [22] gave an ingenious variant of Reed-Solomon codes for which the agreement-rate tradeoff is even better, leading finally to the optimal tradeoff achieved by Guruswami and Rudra [8] (namely, $\rho = \varepsilon - o(1)$).

Our expanders are based on the Parvaresh-Vardy codes. Specifically, for a left-vertex $x \in [N]$ and $y \in [D]$, we

define the y 'th neighbor of x to be $\Gamma(x, y) = (y, C(x)_y)$, where $C : [N] \rightarrow [M]^D$ is a Parvaresh-Vardy code with a somewhat unusual setting of parameters. (Note that here we take the right-hand vertex set to be $[D] \times [M]$.) To prove that this graph is an expander, we adopt a ‘list-decoding’ view of expanders. Specifically, for a right-set $T \subseteq [D] \times [M]$, we define

$$\text{LIST}(T) \stackrel{\text{def}}{=} \{x \in [N] : \Gamma(x) \subseteq T\}.$$

Then the property of Γ being a (K, A) expander can be reformulated as follows:

for all right-sets T of size less than AK , we have

$$|\text{LIST}(T)| < K.$$

Let us compare this to the standard list-decodability property for error-correcting codes. Notice that for a received word $r \in [M]^D$,

$$\begin{aligned} \text{LIST}(r, \varepsilon) &= \{x : \Pr_y[C(x)_y = r_y] \geq \varepsilon\} \\ &= \{x : \Pr_y[\Gamma(x, y) \in T_r] \geq \varepsilon\}, \end{aligned}$$

where $T_r = \{(y, r_y) : y \in [D]\}$. Thus, the two list-decoding problems are related, but have the following key differences:

- In the coding setting, we only need to consider sets T of the form T_r . In particular, these sets are all very small — containing only D of the possible DM right vertices.
- In the expander setting, we only need to bound the number of left-vertices whose neighborhood is entirely contained in T , whereas in the coding setting we need to consider left-vertices for which even an ε fraction of neighbors are in T_r .
- In the coding setting, it is desirable for the alphabet size M to be small (constant or $\text{poly}(n)$), whereas our expanders are most interesting and useful when M is in the range between, say, $n^{\omega(1)}$ and $2^{n/2}$.
- In the coding setting, the exact size of $\text{LIST}(r, \varepsilon)$ is not important, and generally any $\text{poly}(n/\varepsilon)$ bound is considered sufficient. Here, however, the relation between the list size and the size of T is crucial. A factor of 2 increase in the list size (for T of the same size) would change our expansion factor A from $(1 - \varepsilon)D$ to $(1 - \varepsilon)D/2$.

For these reasons, we cannot use the analysis of Parvaresh and Vardy [22] as a black box. Indeed, in light of the last item, it is somewhat of a surprise that we can optimize the bound on list size to yield such a tight relationship between $|T|$ and $|\text{LIST}(T)|$ and thereby provide near-optimal expansion.

This list-decoding view of expanders is related to the list-decoding view of randomness extractors that was implicit in Trevisan’s breakthrough extractor construction [36] and was crystallized by Ta-Shma and Zuckerman [34]. There one considers *all* sets $T \subseteq [D] \times [M]$ (not just ones of bounded size) and bounds the size of $\text{LIST}(T, \mu(T) + \varepsilon) = \{x : \Pr_y[\Gamma(x, y) \in T] \geq \mu(T) + \varepsilon\}$, where $\mu(T) \stackrel{\text{def}}{=} |T|/(DM)$ is the density of T . Indeed, our work began by observing a strong similarity between a natural ‘univariate’ analog of the Shaltiel–Umans extractor [27] and the Guruswami–Rudra codes [8], and by hoping that the list-decoding algorithm for the Guruswami–Rudra codes could be used to prove that the univariate analog of the Shaltiel–Umans construction is indeed a good extractor (as conjectured in [15]). However, we were only able to bound $|\text{LIST}(T, \varepsilon)|$ for “small” sets T , which led to constructions of *lossy* condensers, as in the preliminary version of our paper [10]. In this version, we instead bound the size of $\text{LIST}(T) = \text{LIST}(T, 1)$, and this bound is strong enough to yield expanders with expansion $(1 - \varepsilon) \cdot D$ and thus directly implies lossless condensers, as discussed above. (We still consider lossy condensers in Section 5 of this paper for the purpose of getting improved bounds on some other parameters.)

It is also interesting to compare our construction and analysis to recent constructions of extractors based on algebraic error-correcting codes, namely those of Ta-Shma, Zuckerman, and Safra [35] and Shaltiel and Umans [27]. Both of those constructions use multivariate polynomials (Reed–Muller codes) as a starting point, and rely on the fact that these codes are *locally decodable*, in the sense that any bit of the message can be recovered by reading only a small portion of the received word. While the advantage of local decodability is clear in the computational setting (i.e., constructions of pseudorandom generators [30, 27, 37]), where it enables efficient reductions, it is less clear why it is needed in the information-theoretic setting of extractors, where the ‘decoding’ only occurs in the analysis. Indeed, Trevisan’s extractor [36] corresponds to the pseudorandom generator construction of [30], but with the locally list-decodable code replaced by a standard list-decodable code. However, the extractor analyses of [35] and [27] seem to rely essentially on multivariate polynomials and local (list-)decodability. Our construction works with univariate polynomials and the analysis does not require any local decoding – indeed, univariate polynomial codes are not locally decodable.

2.2 Parvaresh-Vardy codes and the proof technique

Our constructions are based on Parvaresh-Vardy codes [22], which in turn are based on Reed-Solomon codes. A

Reed-Solomon codeword is a univariate degree $n - 1$ polynomial $f \in \mathbb{F}_q[Y]$, evaluated at all points in the field. A Parvaresh-Vardy codeword is a bundle of several related degree $n - 1$ polynomials $f_0, f_1, f_2, \dots, f_{m-1}$, each evaluated at all points in the field. The evaluations of the various f_i at a given field element are packaged into a symbol from the larger alphabet \mathbb{F}_{q^m} . The purpose of this extra redundancy is to enable a better list-decoding algorithm than is possible for Reed-Solomon codes.

The main idea in [22] is to view degree $n - 1$ polynomials as elements of the extension field $\mathbb{F} = \mathbb{F}_q[Y]/E(Y)$, where E is some irreducible polynomial of degree n . The f_i (now viewed as elements of \mathbb{F}) are chosen so that $f_i = f_0^{h^i}$ for $i \geq 1$, and a positive integer parameter h . As explained in Section 2.1, our expander is constructed directly from Parvaresh-Vardy codes as follows:

$$\Gamma(f_0, y) = [y, f_0(y), f_1(y), \dots, f_{m-1}(y)].$$

In the analysis, our task is to show that for any set T of size L , the set $\text{LIST}(T) = \{f_0 : \Gamma(f_0) \subseteq T\}$ is small. To do this we follow the list-decoding analysis of [22], which in turn has the same general structure as the list-decoding algorithms for Reed–Solomon codes [29, 9]. We first produce a non-zero polynomial $Q : \mathbb{F}_q^{1+m} \rightarrow \mathbb{F}_q$ that vanishes on T . Now, for every $f_0 \in \text{LIST}(T)$, we have

$$Q(y, f_0(y), \dots, f_{m-1}(y)) = 0 \quad \forall y \in \mathbb{F}_q,$$

and by ensuring that Q has small degree (which is possible because T is not too large), we will be able to argue that the univariate polynomial $Q(Y, f_0(Y), \dots, f_{m-1}(Y))$ is the zero polynomial. Recalling the definition of the f_i , and viewing the f_i as elements of the extension field $\mathbb{F} = \mathbb{F}_q[Y]/E(Y)$, we observe that f_0 is a *root* of the univariate polynomial

$$Q^*(Z) \stackrel{\text{def}}{=} Q(Y, Z, Z^h, Z^{h^2}, \dots, Z^{h^{m-1}}) \bmod E(Y).$$

This is because when simplifying the formal polynomial $Q^*(f_0(Y)) \bmod E(Y)$, we can first take each $f_0(Y)^{h^i}$ term modulo $E(Y)$, resulting in $f_i(Y)$, and we have just argued that $Q(Y, f_0(Y), \dots, f_{m-1}(Y))$ is the zero polynomial, so it is still the zero polynomial modulo $E(Y)$. This argument holds for every $f_0 \in \text{LIST}(T)$, and so we can upper-bound $|\text{LIST}(T)|$ by the degree of Q^* .

3 Expander Graphs

We first formally develop the list-decoding view of expanders described in Section 2.1.

Definition 3.1. For a bipartite graph $\Gamma : [N] \times [D] \rightarrow [M]$ and a set $T \subseteq [M]$, define

$$\text{LIST}(T) = \{x \in [N] : \Gamma(x) \subseteq T\}.$$

The proof of the next lemma follows from the definitions:

Lemma 3.1. *A graph Γ is a (K, A) expander iff for every set T of size at most $AK - 1$, $\text{LIST}(T)$ is of size at most $K - 1$.*

3.1 The construction

Fix the field \mathbb{F}_q and let $E(Y)$ be an irreducible polynomial of degree n over \mathbb{F}_q . We identify elements of \mathbb{F}_q^n with univariate polynomials over \mathbb{F}_q with degree at most $n - 1$. Fix an integer parameter h .

Our expander is the bipartite graph $\Gamma : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ defined as:

$$\Gamma(f, y) \stackrel{\text{def}}{=} [y, f(y), (f^h \bmod E)(y), (f^{h^2} \bmod E)(y), \dots, (f^{h^{m-1}} \bmod E)(y)]. \quad (1)$$

For ease of notation, we will refer to $(f^{h^i} \bmod E)$ as “ f_i .”

Theorem 3.2. *The graph $\Gamma : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ defined in (1) is a $(\leq K_{\max}, A)$ expander for $K_{\max} = h^m$ and $A = q - (n - 1)(h - 1)m$.*

Proof. Let K be any integer less than or equal to $K_{\max} = h^m$, and let $A = q - (n - 1)(h - 1)m$. By Lemma 3.1, it suffices to show that for every set $T \subseteq \mathbb{F}_q^{m+1}$ of size at most $AK - 1$, we have $|\text{LIST}(T)| \leq K - 1$. Fix such a set T .

Our first step is to find a nonzero “low-degree” polynomial $Q(Y, Y_1, \dots, Y_m)$ that vanishes on T . Specifically, Q will only have nonzero coefficients on monomials of the form $Y^i M_j(Y_1, \dots, Y_m)$ for $0 \leq i \leq A - 1$ and $0 \leq j \leq K - 1 \leq h^m - 1$, where $M_j(Y_1, \dots, Y_m) = Y_1^{j_0} \dots Y_m^{j_{m-1}}$ and $j = j_0 + j_1 h + \dots + j_{m-1} h^{m-1}$ is the base- h representation of j . (For simplicity, one may think of $K = h^m$, in which case we are simply requiring that Q has degree at most $h - 1$ in each variable Y_i .) For each $z \in T$, requiring that $Q(z) = 0$ imposes a homogeneous linear constraint on the AK coefficients of Q . Since the number of constraints is smaller than the number of unknowns, this linear system has a nonzero solution. Moreover, we may assume that among all such solutions, Q is the one of smallest degree in the variable Y . This implies that if we write Q in the form $Q(Y, Y_1, \dots, Y_m) = \sum_{j=0}^{K-1} p_j(Y) \cdot M_j(Y_1, \dots, Y_m)$ for univariate polynomials $p_0(Y), \dots, p_{K-1}(Y)$, then at least one of the p_j 's is not divisible by $E(Y)$. Otherwise $Q(Y, Y_1, \dots, Y_m)/E(Y)$ would have smaller degree in Y and would still vanish on T (since E is irreducible and thus has no roots in \mathbb{F}_q).

Consider any polynomial $f(Y) \in \text{LIST}(T)$. By the definition of $\text{LIST}(T)$ and our choice of Q , it holds that

$$Q(y, f_0(y), f_1(y), \dots, f_{m-1}(y)) = 0 \quad \forall y \in \mathbb{F}_q.$$

That is, the univariate polynomial $R_f(Y) \stackrel{\text{def}}{=} Q(f_0(Y), \dots, f_{m-1}(Y))$ has q zeroes. Since the degree of $R_f(Y)$ is at most $(A - 1) + (n - 1)(h - 1)m < q$, it must be identically zero. So

$$Q(Y, f_0(Y), \dots, f_{m-1}(Y)) = 0$$

as a formal polynomial. Now recall that $f_i(Y) \equiv f(Y)^{h^i} \pmod{E(Y)}$. Thus,

$$\begin{aligned} Q(Y, f(Y), f(Y)^h, \dots, f(Y)^{h^{m-1}}) \\ \equiv Q(Y, f_0(Y), \dots, f_{m-1}(Y)) \equiv 0 \pmod{E(Y)}. \end{aligned}$$

So if we interpret $f(Y)$ as an element of the extension field $\mathbb{F} = \mathbb{F}_q[Y]/E(Y)$, then $f(Y)$ is a root of the univariate polynomial Q^* over \mathbb{F} defined by

$$\begin{aligned} Q^*(Z) &\stackrel{\text{def}}{=} Q(Y, Z, Z^h, Z^{h^2}, \dots, Z^{h^{m-1}}) \bmod E(Y) \\ &= \sum_{j=0}^{K-1} (p_j(Y) \bmod E(Y)) \cdot M_j(Z, Z^h, \dots, Z^{h^{m-1}}) \\ &= \sum_{j=0}^{K-1} (p_j(Y) \bmod E(Y)) \cdot Z^j. \end{aligned}$$

Since this holds for every $f(Y) \in \text{LIST}(T)$, we deduce that Q^* has at least $|\text{LIST}(T)|$ roots in \mathbb{F} . On the other hand, Q^* is a non-zero polynomial, because at least one of the $p_j(Y)$'s is not divisible by $E(Y)$. Thus, $|\text{LIST}(T)|$ is bounded by the degree of Q^* , which is at most $K - 1$. \square

3.2 Setting parameters

The following theorem differs from Theorem 1.1 only by allowing α to be non-constant (and then making the dependence of D on α explicit).

Theorem 3.3 (Thm. 1.1, generalized). *For every $N \in \mathbb{N}$, $K_{\max} \leq N$, and $\varepsilon, \alpha > 0$, there is an explicit $(\leq K_{\max}, (1 - \varepsilon)D)$ expander $\Gamma : [N] \times [D] \rightarrow [M]$ with degree $D = 2^{2+\alpha}((\log N)(\log K_{\max})/\varepsilon)^{1+1/\alpha}$ and $M \leq D^2 \cdot K_{\max}^{1+\alpha}$. Moreover, D is a power of 2.*

Proof. Let $n = \log N$ and $k = \log K_{\max}$. Let $h = \lceil (nk/\varepsilon)^{1/\alpha} \rceil$ and let q be the power of 2 in the interval $(h^{1+\alpha}, 2h^{1+\alpha}]$.

Set $m = \lceil (\log K_{\max})/(\log h) \rceil$, so that $h^{m-1} \leq K_{\max} \leq h^m$. Then, by Theorem 3.2, the graph $\Gamma : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ defined in (1) is a $(\leq h^m, A)$ expander for $A = q - (n - 1)(h - 1)m$. Since $K_{\max} \leq h^m$, it is also a $(\leq K_{\max}, A)$ expander.

Note that the number of left-vertices in Γ is $q^n \geq N$, and the number of right-vertices is

$$M = q^{m+1} \leq q^2 \cdot h^{(1+\alpha) \cdot (m-1)} \leq q^2 \cdot K_{\max}^{1+\alpha}.$$

The degree is

$$D \stackrel{\text{def}}{=} q \leq 2h^{1+\alpha} \leq 2(2(nk/\varepsilon)^{1/\alpha})^{1+\alpha} = 2^{2+\alpha}(nk/\varepsilon)^{1+1/\alpha} \leq 2^{2+\alpha} \cdot ((\log N)(\log K_{\max})/\varepsilon)^{1+1/\alpha}.$$

To see that the expansion factor $A = q - (n-1)(h-1)m \geq q - nhk$ is at least $(1-\varepsilon)D = (1-\varepsilon)q$, note that

$$nhk \leq \varepsilon \cdot h^{1+\alpha} \leq \varepsilon q,$$

where the first inequality holds because $h^\alpha \geq nk/\varepsilon$.

Finally, the construction is explicit because a representation of \mathbb{F}_q for q a power of 2 (i.e. an irreducible polynomial of degree $\log q$ over \mathbb{F}_2) as well as an irreducible polynomial $E(Y)$ of degree n over \mathbb{F}_q can be found in time $\text{poly}(n, \log q) = \text{poly}(\log N, \log D)$ [28]. \square

Remark 1. *In this proof we work in a field \mathbb{F}_q of characteristic 2, which has the advantage of yielding a polynomial-time construction even when we need to take q to be super-polynomially large (which occurs when $\varepsilon(n) = n^{-\omega(1)}$). When $\varepsilon \geq 1/\text{poly}(n)$, then we could use any prime power q instead, with some minor adjustments to the construction and the parameters claimed in the theorem.*

4 Lossless condensers and extractors

We now interpret the expanders constructed in the previous section as lossless condensers (see Definition 1.4). This connection, due to Ta-Shma, Umans, and Zuckerman [33], is based on viewing a function $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ as the neighbor function of a bipartite graph with 2^n left-vertices, 2^d right-vertices, and left-degree 2^d . It turns out that this graph has expansion close to the degree if and only if C is a lossless condenser.

Lemma 4.1 ([33]). *$C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ is a $k \rightarrow_\varepsilon k + d$ condenser iff the corresponding bipartite graph is a $(2^k, (1-\varepsilon) \cdot 2^d)$ expander.*

Thus the following is an immediate consequence of Theorem 3.3.

Theorem 4.2 (Theorem 1.3, generalized). *For every $n \in \mathbb{N}$, $k_{\max} \leq n$, and $\alpha, \varepsilon > 0$, there is an explicit function $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = (1+1/\alpha) \cdot (\log n + \log k_{\max} + \log(1/\varepsilon) + \alpha + 2)$ and $m \leq 2d + (1+\alpha)k_{\max}$ such that for all $k \leq k_{\max}$, C is a $k \rightarrow_\varepsilon k + d$ (lossless) condenser.*

Once we have condensed almost all of the entropy into a source with entropy rate close to 1 (as in Theorem 4.2), extracting (most of) that entropy is not that difficult. All we need to do is to compose the condenser with an extractor that works for entropy rates close to 1. The following standard fact makes the composition formal:

Proposition 4.3. *Suppose $C : \{0, 1\}^n \times \{0, 1\}^{d_1} \rightarrow \{0, 1\}^{n'}$ is a $k \rightarrow_{\varepsilon_1} k'$ condenser, and $E : \{0, 1\}^{n'} \times \{0, 1\}^{d_2} \rightarrow \{0, 1\}^m$ is a (k', ε_2) -extractor; then $E \circ C : \{0, 1\}^n \times \{0, 1\}^{t_1+t_2} \rightarrow \{0, 1\}^m$ defined by $(E \circ C)(x, y_1, y_2) \stackrel{\text{def}}{=} E(C(x, y_1), y_2)$ is a $(k, \varepsilon_1 + \varepsilon_2)$ -extractor.*

For the best dependence on the error parameter ε , the extractor we will use is due to Zuckerman:

Theorem 4.4 ([40]). *For all constants $\alpha, \delta > 0$: for all positive integers n, k and all $\varepsilon > \exp(-n/2^{O(\log^* n)})$, there is an explicit construction of a $(k = \delta n, \varepsilon)$ extractor $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\log n + \log(1/\varepsilon))$ and $m \geq (1-\alpha)k$.*

We now prove our main extractor theorem, restated here:

Theorem 4.5 (Thm. 1.2, restated). *For every constant $\alpha > 0$, and all positive integers n, k and all $\varepsilon > \exp(-n/2^{O(\log^* n)})$, there is an explicit construction of a (k, ε) extractor $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\log n + \log(1/\varepsilon))$ and $m \geq (1-\alpha)k$.*

Proof. Consider the condenser of Theorem 4.2, with its parameter ε set to one half the present ε , and its parameter α set to 1. Its seed length is $d_1 = O(\log(n/\varepsilon))$, and its output length is $n' \leq 2d_1 + 2k$, while its output min-entropy is $k' \geq k + d_1$. Applying Proposition 4.3 to this condenser and the extractor of Theorem 4.4 (with its parameter ε set to half the present ε , and $\delta = 1/2$) gives the claimed extractor. \square

In the fairly common case that ε is a constant, we can use the much simpler ‘‘expander-walk’’ extractor which extracts almost all of the entropy for entropy rates close to 1 (in place of the extractor of Theorem 4.4). Note that our condenser from Theorem 4.2 achieves a constant entropy rate arbitrarily close to 1, and so can be combined with any extractor for such high min-entropy rates. A standard construction achieving this is based on expander walks [6, 40, 41]. Specifically, such an extractor can be obtained by combining the equivalence between extractors and ‘averaging samplers’ [40], and the fact that expander walks are an averaging sampler, as established by the Chernoff bound for expander walks [6].³

Theorem 4.6. *For all constants $\alpha, \varepsilon > 0$, there is a constant $\delta < 1$ for which the following holds: for all positive integers n , there is an explicit construction of a $(k = \delta n, \varepsilon)$ extractor $E : \{0, 1\}^n \times \{0, 1\}^t \rightarrow \{0, 1\}^m$ with $t \leq \log(\alpha n)$ and $m \geq (1-\alpha)n$.*

For completeness, we present the short proof:

³The papers [13, 4] prove hitting properties of expander walks, and observe that these imply objects related to (but weaker than) extractors, known as dispersers.

Proof. Let $m = \lceil (1 - \alpha)n \rceil$, and for some absolute constants $c > 1$ and $\lambda < 1$, let G be an explicit 2^c -regular expander on 2^m vertices (identified with $\{0, 1\}^m$) and second eigenvalue $\lambda = \lambda(G) < 1$. Let L be the largest power of 2 at most $(n - m)/c$ (so $L > (n - m)/(2c)$), and let $t = \log L \leq \log(\alpha n)$. The extractor E is constructed as follows. Its first argument x is used to describe a walk v_1, v_2, \dots, v_L of length L in G by picking v_1 based on the first m bits of x , and each further step of the walk from the next c bits of x — so in all, L must satisfy $n = m + (L - 1)c$. The seed y is used to pick one of the vertices of the walk at random. The output $E(x, y)$ of the extractor is the m -bit label of the chosen vertex.

Let \mathbf{X} be a random variable with min-entropy $k = \delta n$. We wish to prove that for any $S \subseteq \{0, 1\}^m$, the probability that $E(\mathbf{X}, \mathbf{U}_t)$ is a vertex in S is in the range $\mu \pm \varepsilon$ where $\mu = |S|/2^m$. Fix any such subset S . Call an $x \in \{0, 1\}^n$ “bad” if

$$\left| \Pr_y[E(x, y) \in S] - \mu \right| > \varepsilon/2.$$

The known Chernoff bounds for random walks on expanders [6] imply that the number of bad x 's is at most

$$2^n \cdot e^{-\Omega(\varepsilon^2(1-\lambda)L)} = 2^n \cdot e^{-\Omega(\varepsilon^2(1-\lambda)\alpha n/c)} = 2^n \cdot 2^{-\Omega(\varepsilon^2\alpha n)}$$

(since c, λ are absolute constants). Therefore the probability that \mathbf{X} is bad is at most $2^{-\delta n} \cdot 2^n \cdot 2^{-\Omega(\varepsilon^2\alpha n)}$, which is exponentially small for large enough $\delta < 1$. Therefore

$$|\Pr[E(\mathbf{X}, \mathbf{U}_t) \in S] - \mu| \leq \varepsilon/2 + 2^{-\Omega(n)} \leq \varepsilon,$$

implying that E is a (k, ε) -extractor. \square

Combining Theorem 4.2 with Theorem 4.6 via Proposition 4.3, as in the proof of Theorem 4.5, we obtain the following extractor, which has the advantage that its proof is short and self-contained (except for the Chernoff bound for expander walks [6]):

Theorem 4.7. *For every constant $\alpha > 0$, for all positive integers n, k , and all constant $\varepsilon > 0$, there is an explicit construction of a (k, ε) extractor $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with $d = O(\log n + \log(1/\varepsilon))$ and $m \geq (1 - \alpha)k$.*

5 Lossy condensers

In this section we show how minor modifications to the proof allow us to optimize the seed length or the output entropy. We also show that a small modification to the construction yields condensers from Reed-Solomon codes. The price for both of these modifications is that the resulting objects are no longer *lossless* condensers, but instead just ordinary (lossy) condensers.

5.1 The list-decoding viewpoint

First, we record some standard facts about min-entropy:

Proposition 5.1. *For $K \in \mathbb{N}$, a distribution D has min-entropy at least $\log K$ iff D is a convex combination of flat distributions on sets of size exactly K .*

Proposition 5.2. *For any $k > 0$, the distance from a distribution D to a closest distribution with min-entropy k is exactly $\sum_{a:D(a) \geq 2^{-k}} (D(a) - 2^{-k})$.*

Proposition 5.3. *A distribution D with min-entropy $\log(K - c)$ is c/K -close to some distribution with min-entropy $\log K$.*

Proof. By Proposition 5.2, the distance from D to the closest distribution with min-entropy $\log K$ is

$$\sum_{a:D(a) \geq 1/K} (D(a) - 1/K) \leq 1 - (K - c) \cdot 1/K = c/K. \quad \square$$

The next lemma gives a useful sufficient condition for a distribution to be close to having large min-entropy:

Lemma 5.4. *Let \mathbf{Z} be a random variable. If for all sets T of size K , $\Pr[\mathbf{Z} \in T] \leq \varepsilon$ then \mathbf{Z} is ε -close to having min-entropy at least $\log(K/\varepsilon)$.*

Proof. Let T be a set of the K heaviest elements x (under the distribution of \mathbf{Z}). Let $2^{-\ell}$ be the average probability mass of the elements in T . Then $\varepsilon \geq \Pr[\mathbf{Z} \in T] = 2^{-\ell}K$, so $\ell \geq \log(K/\varepsilon)$. But every element outside T has weight at most $2^{-\ell}$, and with all but probability ε , \mathbf{Z} hits elements outside T . \square

Now we can develop a ‘list-decoding’ view of lossy condensers, analogous to the one we have used for expanders (Lemma 3.1) and the one known for extractors [34]. The following definition should be compared to Definition 3.1:

Definition 5.1. *For a function $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ and a set $T \subseteq \{0, 1\}^m$, define*

$$\text{LIST}(T, \varepsilon) \stackrel{\text{def}}{=} \left\{ x : \Pr_y[C(x, y) \in T] > \varepsilon \right\}.$$

Similar to the situation with expanders, if we can bound the size of $\text{LIST}(T, \varepsilon)$ for all sets T that are not too large, then we have a condenser:

Lemma 5.5. *Fix a function $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$. If for every set $T \subseteq \{0, 1\}^m$ of size at most L , we have $|\text{LIST}(T, \varepsilon)| \leq H$, then C is a*

$$\log(H/\varepsilon) \rightarrow_{2\varepsilon} \log(L/\varepsilon) - 1$$

condenser.

Proof. We have a random variable \mathbf{X} with min-entropy $\log(H/\varepsilon)$. For a fixed T of size at most L , the probability that \mathbf{X} is in $\text{LIST}(T, \varepsilon)$ is at most ε ; if that does not happen, then the probability $C(\mathbf{X}, \mathbf{U}_t)$ lands in T is at most ε . Altogether the probability $C(\mathbf{X}, \mathbf{U}_t)$ falls in T is at most 2ε . Now apply Lemma 5.4. \square

5.2 An analysis for minimizing the seed length

The condenser of Theorem 4.2 is lossless and achieves an entropy rate of $1/(1+\alpha)$ for any desired $\alpha > 0$, but its seed length is $(1+1/\alpha)(\log(n/\varepsilon) + \log k + O(1))$. By picking α to be large, say $\alpha = 1/\gamma$ for a small constant $\gamma > 0$, we can reduce the seed length to $(1+\gamma)(\log(n/\varepsilon) + \log k + O(1))$, at the expense of a worse output entropy rate of $\Omega(\gamma)$.

We now show how one can improve the seed length further, to $(1+\gamma)(\log(n/\varepsilon) + O(1))$ — that is, save the $\log k$ term. The new condenser, while not lossless, still retains a fraction $(1 - O(1/\log(n/\varepsilon)))$ of the input entropy, and the entropy rate is $\Omega(\gamma)$.

The improved analysis that permits us to optimize the seed length is in the following theorem, which exploits the “multiple-roots” idea in [9] (compare to Theorem 3.2):

Theorem 5.6. *Define $\Gamma : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ as in (1) and define $\text{LIST}(T, \varepsilon)$ with respect to Γ as in Definition 5.1. Fix positive integer parameters $s \geq 1$, $H \leq h^m$. Then for every set $T \subseteq \mathbb{F}_q^{m+1}$ of size at most*

$$L = \left\lfloor \frac{AH - 1}{\binom{m+s}{s-1}} \right\rfloor,$$

we have $|\text{LIST}(T, \varepsilon)| \leq H - 1$, where $A = \varepsilon sq - (n - 1)(h - 1)m$.

Some intuition about the parameters above may be in order. In Theorem 3.2, the lower bound on q (implicit in the demand that $A > 0$) needed in order to ensure expansion by a $(1 - \varepsilon)q$ factor was $q \geq nmh/\varepsilon$. In the above theorem, the lower bound requirement is weaker by a factor s , and this turns into an improvement in the seed length (which is $\log q$). When viewed as a condenser, the price we pay is that the input entropy is larger by about $\log \binom{m+s}{s-1}$ (which is $\Theta(m)$ when we pick $s = m$) than the output entropy, and thus the condenser incurs an entropy loss.

Proof. Let $T \subseteq \mathbb{F}_q^{m+1}$ be an arbitrary set of size at most $L = \lfloor (AH - 1) / \binom{m+s}{s-1} \rfloor$. The proof follows along the lines of the proof of Theorem 3.2, with the main change being that we make sure that the interpolated polynomial $Q(Y, Y_1, Y_2, \dots, Y_m)$ has a root of multiplicity at least s at

each element $\alpha = (\alpha_0, \alpha_1, \alpha_2, \dots, \alpha_m) \in S$. By a “root of multiplicity at least s ,” we mean that the polynomial

$$Q_\alpha(Y, Y_1, \dots, Y_m) \stackrel{\text{def}}{=} Q(\alpha_0 + Y, \alpha_1 + Y_1, \dots, \alpha_m + Y_m)$$

has no monomials of degree $s - 1$ or smaller with nonzero coefficients, which amounts to $\binom{m+s}{s-1}$ homogeneous linear constraints on the coefficients of Q . The polynomial Q will only have nonzero coefficients on AH monomials of the form $Y^i M_j(Y_1, \dots, Y_m)$ for $0 \leq i \leq A - 1$ and $0 \leq j \leq H - 1$, where $M_j(Y_1, \dots, Y_m) = Y_1^{j_0} \dots Y_m^{j_{m-1}}$ and $j = j_0 + j_1 h + \dots + j_{m-1} h^{m-1}$ is the base- h representation of j . Since $AH > |T| \binom{m+s}{s-1}$, we have more unknowns than the number of homogeneous linear constraints and such a nonzero polynomial Q exists. In the following, we fix Q to be any such nonzero polynomial, and if several such polynomials exist, we choose the one with smallest Y -degree.

Suppose $f(Y) \in \text{LIST}(T, \varepsilon)$. Let $y \in \mathbb{F}_q$ be such that $\Gamma(f, y) \in T$. Then, by the choice of Q ,

$$Q(y, f_0(y), f_1(y), \dots, f_{m-1}(y)) = Q(\Gamma(f, y)) = 0.$$

In fact, since $\Gamma(f, y)$ is a root of multiplicity s , we can show that the the polynomial

$$R_f(Y) \stackrel{\text{def}}{=} Q(Y, f_0(Y), f_1(Y), \dots, f_{m-1}(Y))$$

has a root of multiplicity s at y . To see this, note that

$$\begin{aligned} R_f(y + Y) &= Q(y + Y, f_0(y + Y), \dots, f_{m-1}(y + Y)) \\ &= Q(y + Y, f_0(y) + Y g_0(Y), \dots, f_{m-1}(y) + Y g_{m-1}(Y)) \\ &= Q_{\Gamma(f, y)}(Y, Y \cdot g_0(Y), Y \cdot g_1(Y), \dots, Y \cdot g_{m-1}(Y)) \end{aligned}$$

for some polynomials g_0, \dots, g_{m-1} . Since every monomial in $Q_{\Gamma(f, y)}$ has degree at least s , when we substitute $Y \cdot g_i(Y)$ for the variables we get a univariate polynomial divisible by Y^s . Thus $Y^s | R_f(y + Y)$, i.e. R_f has a root of multiplicity s at y . Equivalently, $(Y - y)^s | R_f(Y)$. We conclude that if $f(Y) \in \text{LIST}(T, \varepsilon)$, i.e., if

$$\Pr_y [Q(y, f_0(y), f_1(y), \dots, f_{m-1}(y)) = 0] > \varepsilon,$$

then $R_f(Y)$ has more than εsq roots counting multiplicities. On the other hand the degree of $R_f(Y)$ is at most $(A - 1) + (n - 1)(h - 1)m$. Therefore, since εsq exceeds this degree, we must have $R_f(Y) = 0$.

From this point on, the proof proceeds identically to that of Theorem 3.2 (with H playing the role of K), leading to the desired conclusion $|\text{LIST}(T, \varepsilon)| \leq H - 1$. \square

5.3 Setting parameters

Picking parameters suitably we obtain the following condenser:

Theorem 5.7. For every $n \in \mathbb{N}$, $\ell \leq n$ such that 2^ℓ is an integer, and $\alpha, \varepsilon > 0$, there is an explicit function $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^{n'}$ that is a

$$(k = \ell t + \log(1/\varepsilon)) \rightarrow_{3\varepsilon} k + d - (2\ell + \log(1/\varepsilon) + O(1))$$

condenser with $d \leq (1 + 1/\alpha)t$ and $n' \leq (1 + 1/\alpha)k + d$, where $t = \lceil \alpha \log(4n/\varepsilon) \rceil$, provided $\ell(t - 2) \geq \log(2/\varepsilon)$.

Proof. Set $h = 2^t$ and note that $h^{1/\alpha} \geq 4n/\varepsilon$. Let q be the power of 2 in $(h^{1+1/\alpha}/2, h^{1+1/\alpha}]$. Set $m = s = \ell$. Note that

$$A \stackrel{\text{def}}{=} \varepsilon s q - (n - 1)(h - 1)m \geq \varepsilon s q - n h m \geq \varepsilon s q / 2,$$

because $q \geq h^{1+1/\alpha}/2 \geq 2hn/\varepsilon$, and $s = m$.

Consider the function $\Gamma : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ defined in (1). By Theorem 5.6, for every $T \subseteq \mathbb{F}_q^{m+1}$ of size at most

$$L = \left\lfloor \frac{Ah^m - 1}{\binom{2m}{m+1}} \right\rfloor,$$

we have $|\text{LIST}(T, \varepsilon)| \leq h^m - 1$. Applying Lemma 5.5, we find that Γ is a

$$\log((h^m - 1)/\varepsilon) \rightarrow_{2\varepsilon} \log(L/\varepsilon) - 1$$

condenser. Now

$$L \geq \frac{Ah^m - 1}{2^{2m}} - 1 \geq \frac{Ah^m}{2^{2m}} - 2.$$

By Proposition 5.3, the output distribution of the condenser Γ is within statistical distance

$$\frac{2^{2m+1}}{Ah^m} \leq \frac{2^{2\ell+1}}{2^{\ell t}} \leq \varepsilon \quad (2)$$

of a distribution with min-entropy at least $\log(Ah^m/2^{2m+1}) + \log(1/\varepsilon)$, where we used the hypothesis $\ell(t - 2) \geq \log(2/\varepsilon)$ to conclude the last inequality in (2). Together with the lower bound $A \geq \varepsilon \ell q / 2$, we can conclude that Γ is a

$$\ell t + \log(1/\varepsilon) \rightarrow_{3\varepsilon} \log q + \log \ell + \ell t - 2\ell - 2$$

condenser. This is the claimed condenser; the upper bounds on d and n' follow from the fact that $q \leq 2^{(1+1/\alpha)t}$.

Finally, the construction is explicit because a representation of \mathbb{F}_q for q a power of 2 as well as an irreducible polynomial $E(Y)$ of degree n over \mathbb{F}_q can be found in time $\text{poly}(n, \log q)$ [28]. \square

In the previous theorem, α may be subconstant, and in the following corollary we show that it can be set to produce a seed length of $d = \log n + O(1)$ (for constant ε), which would be optimal up to an additive constant if our

condenser produced an output that is almost perfectly condensed (i.e., if the output length exceeded the output min-entropy by only an additive $O(1)$ bits). We can achieve such a short seed at the expense of an output entropy rate of $\Omega(1/\log(n/\varepsilon))$, which is subconstant, but still quite good.

Corollary 5.8. For every constant integer $c \geq 2$, and for every $n \in \mathbb{N}$, $k \leq n$, and $\varepsilon \geq 2^{-k+3}$, there is an explicit construction of a

$$k \rightarrow_{3\varepsilon} (1 - 2/c)k + d - \log(1/\varepsilon) - O(1)$$

condenser $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^{n'}$ with $d = \log(n/\varepsilon) + O(1)$ and $n' = \left(1 + \frac{\log(4n/\varepsilon)}{c}\right)k + d$.

Proof. Set $\alpha = c/\log(4n/\varepsilon)$ in Theorem 5.7. \square

5.4 Extractors with short seed length

We now combine the condenser of Theorem 5.7 with Zuckerman's recent extractor. (This extractor in turn starts by applying a condenser due to Raz [24] that has constant seed length and can increase the entropy rate from δ to $1 - \delta$ for any constant $\delta > 0$, while retaining a constant fraction of the min-entropy.)

Theorem 5.9 ([41]). For all constants $\alpha, \delta, \varepsilon > 0$: for all positive integers n , there is an explicit construction of a $(k = \delta n, \varepsilon)$ extractor $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = \log n + O(1)$ and output length $m \geq (1 - \alpha)k$.

Combining Theorem 5.7 with Theorem 5.9 via Proposition 4.3, as in the proof of Theorem 4.5, we obtain the following extractor, which has a near-optimal seed length:

Theorem 5.10. For all constants $\alpha, \gamma, \varepsilon > 0$: for all positive integers n, k , there is an explicit construction of a (k, ε) extractor $E : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^m$ with seed length $d = (1 + \gamma) \log n + \log k + O(1)$ and output length $m \geq (1 - \alpha)k$, provided $k \geq cd/\alpha$ for a universal constant c .

6 Reed-Solomon version

We use one of the main ideas from [8] to argue that a small modification to our construction gives a good condenser from Reed-Solomon codes, answering a question raised in [15].

Let q be an arbitrary prime power, and let $\zeta \in \mathbb{F}_q$ be a generator of the multiplicative group \mathbb{F}_q^* . Then the polynomial $E(Y) = Y^{q-1} - \zeta$ is irreducible over \mathbb{F}_q [16, Chap. 3, Sec. 5]. The following identity holds for all $f(Y) \in \mathbb{F}_q[Y]$:

$$f(Y)^q \equiv f(Y^q) \equiv f(Y^{q-1}Y) \equiv f(\zeta Y) \pmod{E(Y)}.$$

seed length d	output length	output entropy	Thm.
$(1 + 1/\gamma)(\log(n/\varepsilon) + \log k) + O(1)$	$(1 + \gamma)k + 2d$	$k + d$	4.2
$(1 + \gamma)(\log(n/\varepsilon) + \log k) + O(1)$	$(1 + 1/\gamma)k + 2d$	$k + d$	4.2
$(1 + \gamma)\log(n/\varepsilon) + O(1)$	$(1 + 1/\gamma)k + d$	$(1 - O(1/\log(n/\varepsilon)))k + d$	5.7
$\log(n/\varepsilon) + O(1)$	$(1 + \gamma\log(4n/\varepsilon))k + d$	$(1 - 2\gamma)k + d - O(\log(1/\varepsilon))$	5.7

Figure 1. Condensers in this paper for k min-entropy. Above, $\gamma > 0$ is an arbitrarily small constant. Note that the first two constructions condense all entropy thresholds less than k simultaneously.

In this case, if we modify our basic function Γ (see (1)) slightly so that we raise f to successive powers of q rather than h , we obtain the function $C : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ defined by:

$$\begin{aligned} C(f, y) &\stackrel{\text{def}}{=} [y, f(y), (f^q \bmod E)(y), (f^{q^2} \bmod E)(y), \\ &\quad \dots, (f^{q^{m-1}} \bmod E)(y)] \\ &= [y, f(y), f(\zeta y), \dots, f(\zeta^{m-1}y)]. \end{aligned} \quad (3)$$

In other words, our function interprets its first argument as describing a univariate polynomial over \mathbb{F}_q of degree at most $n-1$ (i.e., a Reed-Solomon codeword), it uses the seed to select a random location in the codeword, and it outputs m successive symbols of the codeword, together with the seed. This is precisely the analog of the Shaltiel-Umans q -ary extractor construction [27], for univariate polynomials rather than multivariate polynomials.

With a minor modification to the proof of Theorem 3.2, we show that this is good condenser:

Theorem 6.1. *Define C as in (3) and $\text{LIST}(T, \varepsilon)$ with respect to C as in Definition 5.1. Then for every $T \subseteq \mathbb{F}_q^{m+1}$ of size at most $L = Ah^m - 1$, we have*

$$|\text{LIST}(T, \varepsilon)| \leq (h-1) \cdot \frac{q^m - 1}{q-1},$$

where $A = \varepsilon q - (n-1)(h-1)m$.

Proof. Let $T \subseteq \mathbb{F}_q^{m+1}$ with $|T| \leq Ah^m - 1$. The proof follows along the lines of Theorem 3.2. We interpolate a nonzero polynomial $Q(Y, Y_1, Y_2, \dots, Y_m)$ that vanishes on T , and has degree at most $A-1$ in Y and at most $(h-1)$ in each Y_j . The number of coefficients of such a Q equals Ah^m which exceeds $|T|$, and therefore such a nonzero polynomial Q indeed exists. We can also ensure that $E(Y)$ does not divide Q . For every $f(Y) \in \text{LIST}(T, \varepsilon)$, the polynomial $R_f(Y) \stackrel{\text{def}}{=} Q(Y, f(Y), f(\zeta Y), \dots, f(\zeta^{m-1}Y))$ has more than εq roots, and degree at most $(A-1) + (n-1)(h-1)m$, and therefore must be the zero polynomial. We define Q^* slightly differently:

$$Q^*(Z) \stackrel{\text{def}}{=} Q(Y, Z, Z^q, Z^{q^2}, \dots, Z^{q^{m-1}}) \bmod E(Y).$$

As before, Q^* is a nonzero polynomial over the extension field $\mathbb{F} = \mathbb{F}_q[Y]/(E(Y))$. Further, every $f(Y) \in \text{LIST}(T, \varepsilon)$, viewed as an element of the extension field \mathbb{F} , is a root of Q^* . It follows that $|\text{LIST}(T, \varepsilon)| \leq \deg(Q^*)$. The degree of Q^* is at most

$$(h-1)(1 + q + q^2 + \dots + q^{m-1}) = (h-1) \cdot \frac{q^m - 1}{q-1},$$

and this proves the claimed bound. \square

By picking parameters suitably in the above construction, we obtain the following condenser. Unlike our basic condenser (Theorem 4.2), this condenser is no longer lossless. Instead, the ratio of the input and output min-entropies is roughly $d/t \approx (1 + 1/\alpha)$, which means that we retain only a $\alpha/(1 + \alpha)$ fraction of the min-entropy (compare with Theorem 5.7).

Theorem 6.2 (Reed-Solomon condenser). *For every $n \in \mathbb{N}$, $\ell \leq n$ such that 2^ℓ is an integer, and $\alpha, \varepsilon > 0$, there is an explicit function $C : \{0, 1\}^n \times \{0, 1\}^d \rightarrow \{0, 1\}^{n'}$ defined in (3) that is a*

$$(1 + 1/\alpha)\ell t + \log(1/\varepsilon) \rightarrow_{3\varepsilon} \ell t + d - 2$$

condenser with $d \leq (1 + 1/\alpha)t$ and $n' \leq (1 + 1/\alpha)\ell t + d$, where $t = \lceil \alpha \log(4n\ell/\varepsilon) \rceil$, provided $\ell t \geq \log(1/\varepsilon)$.

Proof. Set $h = 2^t$ and note that $h^{1/\alpha} \geq 4n/\varepsilon$. Let q be the power of 2 in $(h^{1+1/\alpha}/2, h^{1+1/\alpha}]$. Set $m = \ell$. Note that

$$A \stackrel{\text{def}}{=} \varepsilon q - (n-1)(h-1)m \geq \varepsilon q - nhm \geq \varepsilon q/2,$$

because $q \geq h^{1+1/\alpha}/2 \geq 2nh\ell/\varepsilon$, and $m = \ell$.

Consider the function $C : \mathbb{F}_q^n \times \mathbb{F}_q \rightarrow \mathbb{F}_q^{m+1}$ defined in (3). By Theorem 6.1, for every $T \subseteq \mathbb{F}_q^{m+1}$ of size at most $L = Ah^m - 1$ we have $|\text{LIST}(T, \varepsilon)| \leq q^m - 1$. Applying Lemma 5.5, we find that C is a

$$\log\left(\frac{q^m - 1}{\varepsilon}\right) \rightarrow_{2\varepsilon} \log\left(\frac{Ah^m - 1}{2\varepsilon}\right)$$

condenser. By Proposition 5.3, the output distribution of the condenser C is within statistical distance $\frac{1}{Ah^m} \leq 2^{-\ell t} \leq \varepsilon$

of a distribution with min-entropy at least

$$\log\left(\frac{Ah^m}{2\varepsilon}\right) \geq \log q + \ell t - 2 = \ell t + d - 2.$$

We can thus conclude that C is a

$$(1 + 1/\alpha)\ell t + \log(1/\varepsilon) \rightarrow_{3\varepsilon} \ell t + d - 2$$

condenser. This is the claimed condenser; the upper bounds on d and n' follow from the fact that $q = 2^d \leq 2^{(1+1/\alpha)t}$.

Finally, the construction is explicit because a representation of \mathbb{F}_q for q a power of 2 as well as a generator of \mathbb{F}_q^* can be found in time $\text{poly}(\log q)$ [28]. \square

6.1 Limitation of the Reed-Solomon condenser

For the Reed-Solomon-based construction, a relatively simple argument shows that the entropy rate must in general be a constant less than 1. The example below comes from [7, 34]:

Lemma 6.3. *Define C as in (3). For every positive integer $p < n$ such that $p|(q-1)$, there is a source \mathbf{X} with minentropy at least $\lfloor n/p \rfloor \cdot \log q$ for which the support of $C(\mathbf{X}, \mathbf{U}_{\log q})$ is entirely contained within a set of size w^m , where $w = (q-1)/p + 1$.*

Proof. Take the source to be p -th powers of all degree $\lfloor (n-1)/p \rfloor$ polynomials. Every output symbol of C is an evaluation of such a polynomial, and therefore must be a p -th power, or 0. There are thus only $w = (q-1)/p + 1$ possible output symbols, so the output is contained within a set of size w^m . \square

For such a source \mathbf{X} , the output minentropy of C is at most $m \log w$ and the output length is $m \log q$. Thus the entropy rate is at most

$$\frac{\log w}{\log q} \approx 1 - \frac{\log p}{\log q}.$$

So for example, for a source obtained when $p \approx \sqrt{n}$, the Reed-Solomon condenser C has a constant entropy rate less than 1 unless the seed length $\log q$ is $\omega(\log n)$.

This implies that the entropy rate obtained in Theorem 6.2 is not an artifact of the analysis. That is, it is not possible to improve the entropy rate (e.g., to 1) simply by giving a different, improved analysis for the generic Reed-Solomon construction.

7 Conclusions

The “list-decoding” view of expanders and condensers used in this paper seems to be quite powerful, leading to

constructions that are more direct, achieve improved parameters. It is thus natural to ask how far this approach can be pushed. Constructing unbalanced expanders with expansion close to the degree where the degree and/or size of the right-hand side are within *constant factors* of optimal is a natural next goal. This is closely related to question of constructing truly optimal extractors, ones that are optimal up to *additive* constants in the seed length and/or output length. Towards this end, we wonder if there is some variant of our construction with a better entropy rate – the next natural threshold is to have entropy *deficiency* only $k^{o(1)}$. Another interesting question is whether some variant of these constructions can give a block-wise source directly. Depending on the actual parameters, either of these two improvements have the potential to lead to extractors with optimal output length (i.e. ones extract all the min-entropy). Alternatively, if we can find an extractor with optimal output length for high min-entropy (say $.99n$), then, by composing it with our condenser, we would get one for arbitrary min-entropy.

We also wonder whether these new techniques can help in other settings. For example, can we use them to argue about *computational* analogues of the objects in this paper – pseudorandom generators and pseudoentropy generators? Or, can variants of our constructions yield so-called “2-source” objects, in which both the source and the seed are only weakly random?

Acknowledgements. This paper began with a conversation at the BIRS workshop “Recent Advances in Computation Complexity.” We would like to thank the organizers for inviting us, and BIRS for hosting the workshop. We also thank Oded Goldreich, Prahladh Harsha, Farzad Parvaresh, Jaikumar Radhakrishnan, Omer Reingold, Ronen Shaltiel and Dieter van Melkebeek for helpful comments.

References

- [1] H. Buhrman, P. B. Miltersen, J. Radhakrishnan, and S. Venkatesh. Are bitvectors optimal? *SIAM Journal on Computing*, 31(6):1723–1744 (electronic), 2002.
- [2] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson. Randomness conductors and constant-degree expansion beyond the degree/2 barrier. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 659–668, 2002.
- [3] B. Chor and O. Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM Journal on Computing*, 17(2):230–261, Apr. 1988.
- [4] A. Cohen and A. Wigderson. Dispersers, deterministic amplification, and weak random sources (extended abstract). In *Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science*, pages 14–19, 1989.

- [5] O. Gabber and Z. Galil. Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, June 1981.
- [6] D. Gillman. A Chernoff bound for random walks on expander graphs. *SIAM J. Comput.*, 27(4):1203–1220 (electronic), 1998.
- [7] V. Guruswami, J. Hastad, M. Sudan, and D. Zuckerman. Combinatorial bounds for list decoding. *IEEE Transactions on Information Theory*, 48(5):1021–1035, 2002.
- [8] V. Guruswami and A. Rudra. Explicit capacity-achieving list-decodable codes. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 1–10, 2006.
- [9] V. Guruswami and M. Sudan. Improved decoding of Reed-Solomon and Algebraic-Geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, 1999.
- [10] V. Guruswami, C. Umans, and S. Vadhan. Extractors and condensers from univariate polynomials. Technical Report TR06-134, Electronic Colloquium on Computational Complexity, October 2006.
- [11] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43(4):439–561 (electronic), 2006.
- [12] R. Impagliazzo, R. Shaltiel, and A. Wigderson. Extractors and pseudo-random generators with optimal seed length. In *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing*, pages 1–10, 2000.
- [13] R. Impagliazzo and D. Zuckerman. How to recycle random bits. In *Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science*, pages 248–253, 1989.
- [14] N. Kahale. Eigenvalues and expansion of regular graphs. *Journal of the ACM*, 42(5):1091–1106, Sept. 1995.
- [15] S. Kalyanaraman and C. Umans. On obtaining pseudorandomness from error-correcting codes. In S. Arun-Kumar and N. Garg, editors, *FSTTCS*, volume 4337 of *Lecture Notes in Computer Science*, pages 105–116. Springer, 2006.
- [16] R. Lidl and H. Niederreiter. *Introduction to Finite Fields and their applications*. Cambridge University Press, 1986.
- [17] C.-J. Lu, O. Reingold, S. Vadhan, and A. Wigderson. Extractors: Optimal up to constant factors. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pages 602–611, 2003.
- [18] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.
- [19] G. A. Margulis. Explicit constructions of expanders. *Problemy Peredachi Informacii*, 9(4):71–80, 1973.
- [20] G. A. Margulis. Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators. *Problemy Peredachi Informatsii*, 24(1):51–60, 1988.
- [21] N. Nisan and D. Zuckerman. Randomness is linear in space. *Journal of Computer and System Sciences*, 52(1):43–52, 1996.
- [22] F. Parvaresh and A. Vardy. Correcting errors beyond the Guruswami-Sudan radius in polynomial time. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 285–294, 2005.
- [23] J. Radhakrishnan and A. Ta-Shma. Bounds for dispersers, extractors, and depth-two superconcentrators. *SIAM Journal on Discrete Mathematics*, 13(1):2–24 (electronic), 2000.
- [24] R. Raz. Extractors with weak random seeds. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 11–20, 2005.
- [25] O. Reingold, R. Shaltiel, and A. Wigderson. Extracting randomness via repeated condensing. *SIAM J. Comput.*, 35(5):1185–1209, 2006.
- [26] R. Shaltiel. Recent developments in explicit constructions of extractors. *Bulletin of the European Association for Theoretical Computer Science*, 77:67–, June 2002. Columns: Computational Complexity.
- [27] R. Shaltiel and C. Umans. Simple extractors for all min-entropies and a new pseudorandom generator. *Journal of the ACM*, 52(2):172–216, 2005. Conference version appeared in FOCS 2001.
- [28] V. Shoup. New algorithms for finding irreducible polynomials over finite fields. *Mathematics of Computation*, 54(189):435–447, 1990.
- [29] M. Sudan. Decoding of Reed Solomon codes beyond the error-correction bound. *J. Complexity*, 13(1):180–193, 1997.
- [30] M. Sudan, L. Trevisan, and S. Vadhan. Pseudorandom generators without the xor lemma. *J. Comput. Syst. Sci.*, 62(2):236–266, 2001.
- [31] A. Ta-Shma. Storing information with extractors. *Inform. Process. Lett.*, 83(5):267–274, 2002.
- [32] A. Ta-Shma and C. Umans. Better lossless condensers through derandomized curve samplers. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, 2006. To appear.
- [33] A. Ta-Shma, C. Umans, and D. Zuckerman. Loss-less condensers, unbalanced expanders, and extractors. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 143–152, 2001.
- [34] A. Ta-Shma and D. Zuckerman. Extractor codes. *IEEE Transactions on Information Theory*, 50(12):3015–3025, 2004.
- [35] A. Ta-Shma, D. Zuckerman, and S. Safra. Extractors from Reed-Muller codes. *J. Comput. Syst. Sci.*, 72(5):786–812, 2006.
- [36] L. Trevisan. Extractors and pseudorandom generators. *Journal of the ACM*, 48(4):860–879, 2001.
- [37] C. Umans. Pseudo-random generators for all hardnesses. *J. Comput. Syst. Sci.*, 67(2):419–440, 2003.
- [38] A. Wigderson and D. Zuckerman. Expanders that beat the eigenvalue bound: Explicit construction and applications. *Combinatorica*, 19(1):125–138, 1999.
- [39] D. Zuckerman. Simulating BPP using a general weak random source. *Algorithmica*, 16(4-5):367–391, 1996.
- [40] D. Zuckerman. Randomness-optimal oblivious sampling. *Random Struct. Algorithms*, 11(4):345–367, 1997.
- [41] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 681–690, 2006.