Dimension Expanders via Rank Condensers*

Michael A. Forbes[†1] **and Venkatesan Guruswami**[‡2]

1 **School of Mathematics, Institute for Advanced Study**
Princeton, NJ, USA
miforbes@csail.mit.edu

2 **Computer Science Department, Carnegie Mellon University**
Pittsburgh, PA, USA
guruswami@cmu.edu

Abstract

An emerging theory of "linear algebraic pseudorandomness" aims to understand the linear algebraic analogs of fundamental Boolean pseudorandom objects where the rank of subspaces plays the role of the size of subsets. In this work, we study and highlight the interrelationships between several such algebraic objects such as subspace designs, dimension expanders, *seeded rank condensers*, *two-source rank condensers*, and rank-metric codes. In particular, with the recent construction of near-optimal subspace designs by Guruswami and Kopparty [12] as a starting point, we construct good (seeded) rank condensers (both *lossless* and *lossy* versions), which are a small collection of linear maps $\mathbb{F}^n \to \mathbb{F}^t$ for $t \ll n$ such that for every subset of $\mathbb{F}^n$ of small rank, its rank is preserved (up to a constant factor in the lossy case) by at least one of the maps.

We then compose a tensoring operation with our lossy rank condenser to construct constant-degree dimension expanders over polynomially large fields. That is, we give $O(1)$ explicit linear maps $A_i : \mathbb{F}^n \to \mathbb{F}^n$ such that for any subspace $V \subseteq \mathbb{F}^n$ of dimension at most $n/2$, $\dim\left(\sum_i A_i(V)\right) \geqslant (1 + \Omega(1)) \dim(V)$. Previous constructions of such constant-degree dimension expanders were based on Kazhdan's property T (for the case when $\mathbb{F}$ has characteristic zero) or monotone expanders (for every field $\mathbb{F}$); in either case the construction was *harder* than that of usual vertex expanders. Our construction, on the other hand, is *simpler*.

For two-source rank condensers, we observe that the lossless variant (where the output rank is the product of the ranks of the two sources) is equivalent to the notion of a linear rank-metric code. For the lossy case, using our seeded rank condensers, we give a reduction of the general problem to the case when the sources have high ($n^{\Omega(1)}$) rank. When the sources have $O(1)$ rank, combining this with an "inner condenser" found by brute-force leads to a two-source rank condenser with output length nearly matching the probabilistic constructions.

1998 ACM Subject Classification F.2.1 Numerical Algorithms and Problems

Keywords and phrases dimension expanders, rank condensers, rank-metric codes, subspace designs, Wronskians

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2015.800

* A full version of this work can be found at http://arxiv.org/abs/1411.7455.

† This work was performed when the author was a graduate student at MIT CSAIL (which was supported by Scott Aaronson's Waterman award, NSF CCF-1249349), and when the author was a Google Research Fellow at the Simons Institute for the Theory of Computing.

‡ Some of this work was done when the author was a visiting researcher at Microsoft Research New England, Cambridge, MA. Research supported in part by NSF grant CCF-0963975.

18th Int'l Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'15) /
19th Int'l Workshop on Randomization and Computation (RANDOM'15).
Editors: Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim; pp. 800–814

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The broad area of pseudorandomness deals with efficiently generating objects that exhibit the desirable properties of "random-like" objects despite being constructed either explicitly or with limited randomness. Pseudorandomness is a central and influential theme in many areas such as complexity theory, derandomization, coding theory, cryptography, high-dimensional geometry, graph theory, and additive combinatorics. The topic has witnessed much progress over the years and continues to be intensively studied. We now have non-trivial constructions of various pseudorandom objects such as expander graphs, randomness extractors and condensers, Ramsey graphs, list-decodable codes, compressed sensing matrices, Euclidean sections, and pseudorandom generators for various concrete models. Despite the seemingly different definitions and contexts of these objects, insights in pseudorandomness have uncovered intimate connections between them, and this has led to a rich theory of "Boolean pseudorandomness" drawing a common pool of broadly useful techniques (see for instance the recent survey by Vadhan [26].)

Recently, there is an emerging theory of "linear-algebraic pseudorandomness" aimed at understanding the linear-algebraic analogs of fundamental Boolean pseudorandom objects where the dimension of subspaces plays the role analogous to min-entropy. Examples of such algebraic objects include dimension expanders, subspace-evasive sets, subspace designs, rank-preserving condensers, etc. In addition to their intrinsic interest, these notions also have surprising applications; for instance, subspace-evasive sets to the construction of Ramsey graphs [22] and list-decodable codes [13, 15], subspace designs to list decoding both in the Hamming metric and the rank metric [16, 14], and rank-preserving condensers to affine extractors [10] [1] and polynomial identity testing [18, 9].

In this work, we study several interesting pseudorandom objects in the linear-algebraic world, such as subspace evasive sets, subspace designs, dimension expanders, *seeded rank condensers*, and *two-source rank condensers*. The last two notions are also introduced in this work, though closely related concepts were studied earlier in the literature. We briefly and informally define these notions now, with more precise statements appearing in later sections. A subspace evasive set is a (large) subset of $\mathbb{F}^n$ that has small intersection with every low-dimensional subspace of $\mathbb{F}^n$. Subspace designs are a (large) collection of subspaces such that every low-dimensional subspace intersects few of them. Dimension expanders are a (small) collection of linear maps $A_i : \mathbb{F}^n \to \mathbb{F}^n$ such that for every subspace $V \subseteq \mathbb{F}^n$ of bounded dimension, the dimension of $\sum_i A_i(V)$ is at least $\alpha \cdot \dim(V)$ for a constant $\alpha > 1$ (so that the dimension grows, or *expands*). Rank condensers are a (small) collection of linear maps $\mathbb{F}^n \to \mathbb{F}^t$ (for $t \ll n$) such that for every subspace of dimension r, its image under at least one of the maps has large dimension. That is, the ambient dimension n is *condensed* to t while roughly the preserving the rank (to r in the *lossless* case (so that no rank is lost), and to $\Omega(r)$ in the *lossy* case). A two-source rank condenser is a map $E : \mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F}^t$ such that for every pair $A, B \subseteq \mathbb{F}^n$ with rank r each, $f(A \times B)$ has rank $\Omega(r^2)$ (or even r^2 in the lossless case) – the tensor product construction is lossless but requires $t = n^2$, so the challenge here is to "derandomize" the tensor product and achieve $t \ll n^2$ (and even $t \ll n$

[1] Despite the usage of rank condensers in the Gabizon-Raz [10] construction of affine extractors, affine extractors seem to not quite fit the restricted notion of a "linear-algebraic pseudorandom object" in the sense of this paper. That is, the objects we consider focus on functions and their interactions with the rank of certain sets of vectors. In contrast, affine extractors (maps which convert uniform distributions over large-enough subspaces of the input into uniform distributions over full-dimensional subspaces) require further statistical properties. The weaker notion of an affine disperser (a map which is almost surjective on its range when applied to a large-enough subspace of the input) similarly requires one-sided statistical guarantees.

for the lossy case for $r \ll \sqrt{n}$).

We remark that there are two perspectives on the above objects. The first is that of *subspaces*, so that we only consider subspaces and their dimension. The second is that of *sets of vectors*, where we consider arbitrary sets of vectors measured by their rank (the dimension of their span). When the underlying functions are linear (or multilinear) these viewpoints are equivalent. For example, one can equally discuss dimension expanders as expanding the dimension of subspaces or as increasing the rank of matrices through matrix multiplication. In this work, we take both views, using "dimension" to refer to subspaces and "rank" to refer to the dimension of the span of a set of vectors.

Conceptually, our work highlights close interconnections between the above pseudorandomness notions. In particular, we show that subspace designs (which were introduced in the context of list decoding variants of algebraic-geometric codes of Guruswami and Xing [16]) are the *same* concept as lossless rank condensers while emphasizing a different regime of parameters. This connection also highlights that a strong variant of subspace designs yields lossy rank condensers. The near-optimal explicit construction of (strong) subspace designs of Guruswami-Kopparty [12] then yields lossless and lossy rank condensers with parameters close to the existential constructions. Our main technical application is an explicit construction of constant-degree dimension expanders over polynomially large fields, that expands all subspaces of $\mathbb{F}^n$ of dimension $n/2$ (say) by a factor $\alpha > 1$. We achieve this construction by first increasing the rank in a trivial way by increasing the dimension of the ambient space, and then using a lossy rank condenser to reduce the ambient space back to $\mathbb{F}^n$ while preserving the rank up to a constant factor. While previous constructions of dimension expanders were at least as complicated as constructions of standard expander graphs (or more so), our construction and analysis is rather elementary. Unfortunately, unlike previous work, our techniques are currently best suited to large fields due to connections with Reed-Solomon codes. However, we do obtain dimension expanders over small fields by paying various logarithmic penalties.

Turning to two-source rank condensers, our original motivation to propose them was a possible route to iteratively construct subspace-evasive sets that might offer some way around the exponential dependence on intersection size that seems inherent to constructions based on algebraic varieties. While there appears to be serious obstacles to such an approach, the notion seems a fundamental one to study regardless. In this work, we focus on two-source rank condensers $f : \mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F}^t$ where the map f is bilinear as this seems like a natural class of constructions to study. We observe that the lossless variant is *equivalent* to the notion of a linear rank-metric code. Known optimal constructions of rank-metric codes such as the Gabidulin codes thereby yield lossless two-source condensers with optimal output length (equal to $\Theta(nr)$ for rank-r subsets of $\mathbb{F}^n$). For lossy two-source rank condensers, we can enumerate over the seeds of our seeded lossy condenser, applying it to both sources separately and condensing the sources to $r^{\Theta(1)}$ dimensions (from the original n). For small r (e.g., constant), we can "concatenate" this construction with a near-optimal lossy two-source condenser found by brute-force to obtain output length $\Theta(n/r)$, matching the non-constructive bound. In general, our method reduces the problem to the case of relatively high "rate" (when $r \approx n^{1/3}$), which is typically easier to tackle.

Organization: In the next three sections, we state (informal versions of) our results, all of ideas behind them, and brief discussions of prior work for seeded rank condensers (section 2), dimension expanders (section 3), and two-source rank condensers (section 4). An expanded treatment with formal statements and proofs can be found in the full version of this work (arXiv:1411.7455).

2 Subspace Designs and Rank Condensers

We begin by discussing the notion of a *subspace design*, as recently defined by Guruswami and Xing [16], and contrast this with the notion of a *seeded (single source) rank condenser* to which we add the qualifier of *lossless*, as defined by Forbes, Saptharishi and Shpilka [8]. We will describe how these objects are essentially the same notion, where the rank condenser can be considered the "primal" object and the subspace design the "dual" object. We then introduce *lossy rank condensers*, a new notion that is key to our construction of dimension expanders see section 3) and describe how the construction of subspace designs of Guruswami and Kopparty [12] implies nearly optimal lossy rank condensers.

2.1 Subspace Designs

We begin with the definition of a subspace design, which is a collection of subspaces $\{H_i\}_i$ such that small-dimensional subspaces V intersect few of the H_i.

▶ **Definition 1** (Guruswami-Xing [16] and Guruswami-Kopparty [12]). Let $\mathbb{F}$ be a field. A collection $\mathcal{H} = \{H_i\}_i$ of subspaces $H_i \subseteq \mathbb{F}^n$ is a **weak** (r, L)**-subspace design** if for every subspace $V \subseteq \mathbb{F}^n$ with $\dim V = r$,

$$|\{i \mid \dim(H_i \cap V) > 0\}| \leqslant L \, .$$

The collection $\mathcal{H}$ is a **strong** (r, L)**-subspace design** if for every subspace $V \subseteq \mathbb{F}^n$ with $\dim V = r$,

$$\sum_i \dim(H_i \cap V) \leqslant L \, .$$

The collection $\mathcal{H}$ is **explicit** if given an index $i \in [|\mathcal{H}|]$ a basis for the i-th subspace in $\mathcal{H}$ can be constructed in $\mathsf{poly}(n, \log |\mathcal{H}|)$ operations in $\mathbb{F}$.

We note here that the above subspaces H_i are not constrained to be of equal dimension. Allowing the dimension of the H_i to vary could conceivably allow for improved constructions, but no construction so far uses this freedom. As such, we will primarily concern ourselves with the case when the dimensions are equal.

Guruswami-Xing [16] defined subspace designs as a way to prune list-decodable codes to ensure a small list-size while maintaining high rate. As such, one wishes for the size $|\mathcal{H}|$ of the design to be large while maintaining L of moderate size. In particular, they showed that large designs exist non-constructively.

▶ **Proposition** (Guruswami-Xing [16]). *Let $\mathbb{F}_q$ be a finite field. Let $\epsilon > 0$, $n \geqslant 8/\epsilon$ and $s \leqslant \epsilon n/2$. Then there is a strong $(s, 8s/\epsilon)$-subspace design $\mathcal{H}$ of $(1-\epsilon)n$-dimensional subspaces in $\mathbb{F}_q^n$ with $|\mathcal{H}| = q^{\epsilon n/8}$.*

Note that the *co*-dimension of the subspaces in $\mathcal{H}$ is ϵn, which is twice that of the maximum dimension $s \approx \epsilon n/2$. We now further remark on the variations of this definition. The following relation between the weak and strong versions is immediate.

▶ **Lemma 2** (Guruswami-Kopparty [12]). *Let $\mathbb{F}$ be a field, and let $\mathcal{H}$ be a collection of subspaces in $\mathbb{F}^n$. Then if $\mathcal{H}$ is a strong (r, L)-subspace design, then $\mathcal{H}$ is a weak (r, L)-subspace design. If $\mathcal{H}$ is a weak (r, L)-subspace design, then $\mathcal{H}$ is a strong (r, rL)-subspace design.*

We also observe that as every dimension $\leqslant r$ subspace can be padded to a dimension r subspace, we immediately can see that subspace designs apply to smaller subspaces as well.

▶ **Lemma 3.** *Let $\mathbb{F}$ be a field, and let $\mathcal{H}$ be a weak/strong (r, L)-subspace design in $\mathbb{F}^n$. Then $\mathcal{H}$ is a (s, L)-subspace design over $\mathbb{F}^n$ for every $1 \leqslant s \leqslant r$.*

While the above seems to allow one to focus on dimension r as opposed to dimension $\leqslant r$, this is not strictly true as one can achieve a better list size L for dimension $s \ll r$. Similarly, the above lemma relating strong and weak designs seems to suggest that qualitatively (up to polynomial factors) these notions are the same. However, as described in the full version, obtaining the appropriate (strong) list size simultaneously for all $s \leqslant r$ will be crucial for our application to constant-degree dimension expanders.

2.2 Seeded Lossless Rank Condensers

Strong subspace designs ask that for any small subspace V there is some $H_i \in \mathcal{H}$ so that $H_i \cap V$ is *small* (that is, by averaging, $\dim H_i \cap V \leqslant L/|\mathcal{H}|$). Equivalently, the amount of dimension in V that is outside H_i is *large* so that in some sense the dimension of V is preserved. This perspective is more naturally phrased in the language of *(seeded) rank condensers*, as defined by Forbes, Saptharishi and Shpilka [8]. The definition we use here is tuned to the equivalence with subspace designs, and we recover their definition as the lossless version of what we term here a *lossy seeded rank condenser* (see Theorem 6). We will discuss prior work and motivation for rank condensers that is less immediately relevant in the full version.

We begin with the definition of rank condensers, which are a collection of linear maps $\varphi : \mathbb{F}^n \to \mathbb{F}^t$ (given explicitly as matrices $E \in \mathbb{F}^{t \times n}$) such that for any small-dimensional subspace V, most of the maps have $\dim \varphi(V) = \dim V$.

▶ **Definition 4.** Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$. A collection of matrices $\mathcal{E} \subseteq \mathbb{F}^{t \times n}$ is a **weak (seeded) (r, L)-lossless rank condenser** if for all matrices $M \in \mathbb{F}^{n \times r}$ with $\operatorname{rank} M = r$,

$$|\{E \mid E \in \mathcal{E}, \operatorname{rank} EM < \operatorname{rank} M\}| \leqslant L \, .$$

The collection $\mathcal{E}$ is a **strong (seeded) (r, L)-lossless rank condenser** if for all matrices $M \in \mathbb{F}^{n \times r}$ with $\operatorname{rank} M = r$,

$$\sum_{E \in \mathcal{E}} (\operatorname{rank} M - \operatorname{rank} EM) \leqslant L \, .$$

The collection $\mathcal{E}$ is **explicit** if given an index $i \in [|\mathcal{E}|]$ the i-th matrix of $\mathcal{E}$ can be constructed in $\mathsf{poly}(t, n, \log |\mathcal{E}|)$ operations in $\mathbb{F}$.

As we have many types of condensers in this paper (weak, strong, lossless, lossy, two-source, etc.) we will often just refer to them as "condensers" (perhaps with some relevant parameters such as "(r, ϵ)") when the relevant adjectives are clear from context.

As it can only increase the quality of the condenser, one naturally considers the case when $\operatorname{rank} E = t$ for all $E \in \mathcal{E}$. However, we do not impose this restriction just as we do not impose the condition that subspaces in subspace designs all have the same dimension. In fact, by the equivalence of subspace designs and lossless rank condensers one can see that these two restrictions are equivalent.

We briefly remark that as all of the pseudorandom objects we consider in this work are linear (or in the case of two-source condensers, bilinear) we will often freely pass between subspaces $V \subseteq \mathbb{F}^n$ of dimension r and matrices $M \in \mathbb{F}^{n \times r}$ of rank r, using that we can choose a basis for V so that $\text{col-span}\, M = V$. As such, we will often treat a matrix $M \in \mathbb{F}^{n \times r}$ as a list of r vectors in $\mathbb{F}^n$.

We now note that subspace designs are equivalent to lossless rank condensers.

▶ **Proposition 5.** *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$. Let $\mathcal{H} = \{H_i\}_{i\in[N]}$ be a collection of subspaces $H_i \subseteq \mathbb{F}^n$ and let $\mathcal{E} = \{E_i\}_{i\in[N]} \subseteq \mathbb{F}^{t\times n}$ be a collection of matrices, where we have that* row-span $E_i = (H_i)^\perp$ *for $i \in [N]$. Then $\mathcal{H}$ is a weak/strong (r, L)-subspace design iff $\mathcal{E}$ is a weak/strong (r, L)-lossless rank condenser.*

While the above proposition is quite simple, it offers a unifying perspective of these different objects which was key to obtaining further results.

2.3 Seeded Lossy Rank Condensers

While the above seeded lossless rank condensers already have applications to list-decodable codes, rank condensers were defined in Forbes, Saptharishi and Shpilka [8] for quite different reasons. We now give a definition closer to their motivation.

▶ **Definition 6.** Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$ and $\epsilon \geqslant 0$. A collection of matrices $\mathcal{E} \subseteq \mathbb{F}^{t\times n}$ is a **(seeded)** (r, ϵ)**-lossy rank condenser** if for all matrices $M \in \mathbb{F}^{n\times r}$ with $\operatorname{rank} M = r$,

$$\operatorname{rank} EM \geqslant (1-\epsilon)\operatorname{rank} M \ ,$$

for some $E \in \mathcal{E}$. The collection $\mathcal{E}$ is a **(seeded)** $(\leqslant r, \epsilon)$**-lossy rank condenser** if it a (s, ϵ)-lossy condenser for all $1 \leqslant s \leqslant r$.

The collection $\mathcal{E}$ is **explicit** if given an index $i \in [|\mathcal{E}|]$ the i-th matrix of $\mathcal{E}$ can be constructed in $\mathsf{poly}(t, n, \log|\mathcal{E}|)$ operations in $\mathbb{F}$.

This notion is a natural linear-algebraic analogue of condensers for *min-entropy* from the realm of Boolean pseudorandomness. One contrast is that we do not require that *most* $E \in \mathcal{E}$ have the desired condensing property as this does not seem important for our applications, although we note that our constructions can meet this stronger requirement with the appropriate modifications[2].

In is worthwhile to contrast this object with subspace designs or lossless rank condensers. The goal of subspace designs was (due to connections with list-decodable codes) to construct a *large* design while less focus was on the exact list-size bound. Here, we have the somewhat different goal of obtaining a *small* collection of matrices, which is akin to obtaining a very small list size in a subspace design. The focus on the collection being small is from the use of such condensers in derandomization, as we will need to enumerate over each matrix in the collection.

In particular, the notion of a $(r, 0)$-lossy rank condenser is of interest because it is *lossless*, which is important for many applications. In particular, this notion was previously defined as a "*rank condenser (hitting set)*" in the work of Forbes, Saptharishi and Shpilka [8], but the construction and usage of these objects predates them[3]. In particular, Gabizon and Raz [10] constructed a $(r, 0)$-condenser with size nr^2, and they used this to construct affine extractors over large fields. Karnin and Shpilka [18] named the construction of Gabizon and Raz [10] to be "rank preserving subspaces" and used this construction to make a *polynomial*

[2] More precisely, this stronger definition requiring most E to condense rank is closer to the definition of a min-entropy condenser. Only requiring *some* E to condense rank is more akin to the notion of a *somewhere condenser* as defined by Barak-Kindler-Shaltiel-Sudakov-Wigderson [3].

[3] We note that the works we highlight are not necessarily the first or last in their respective lines of research, and rather we only highlight those that (to the best of our knowledge) had results concerning lossless rank condensers.

identity testing[4] algorithm of Dvir and Shpilka [7] work in the *black box* model. Forbes and Shpilka [9] later gave an improved construction of a rank condenser with only nr size, and showed how they can be used to make another polynomial identity testing algorithm of Raz and Shpilka [24] work in the black-box model. Forbes, Saptharishi and Shpilka [8], building on the work of Agrawal, Saha, and Saxena [1], analyzed "multivariate" lossless rank condensers as they arose naturally in a polynomial identity testing algorithm.

Beyond applications to polynomial identity testing, Lokshtanov, Misra, Panolan and Saurabh [19] used these condensers to derandomize a fixed-parameter-tractable algorithm of Marx [21] for ℓ-matroid intersection. Cheung, Kwok and Lau [6] rediscovered the rank condenser of Gabizon and Raz [10] and (among other things) used this to give faster randomized algorithms for exact linear algebra. Forbes, Saptharishi and Shpilka [8] showed a generic recipe to construct such rank condensers from *any* error-correcting code (over large fields). Given these applications and connections present in (r, ϵ)-lossy rank condensers for $\epsilon = 0$, we expect the $\epsilon > 0$ version will similarly have many applications.

We now quote the parameters given by the probabilistic method.

▶ **Proposition 7.** *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$, $\epsilon \geqslant 0$ and $t > (1-\epsilon)r$. Then there is a collection $\mathcal{E}$ of k matrices $\mathcal{E} \subseteq \mathbb{F}_q^{t \times n}$ that is a (r, ϵ)-lossy rank condenser whenever*

$$k \geqslant \frac{rn + o_q(1)}{(t-(1-\epsilon)r)(\lfloor \epsilon r \rfloor + 1) - o_q(1)} \,. \tag{2.1}$$

For $\epsilon > 0$, there is a collection $\mathcal{E}$ of size k that is a $(\leqslant r, \epsilon)$-lossy rank condenser whenever

$$k \geqslant \frac{n + o_q(1)}{\epsilon(t-(1-\epsilon)r) - o_q(1)} \,.$$

Thus we can make the output size t of the condenser to be almost equal to the guaranteed dimension bound of $(1-\epsilon)r$. Further, we see that there is essentially no penalty in (existentially) insisting for a $(\leqslant r, \epsilon)$-condenser over a (r, ϵ)-condenser. However, we show in the full version that the notion of $(\leqslant r, \epsilon)$-condenser is provably stronger.

2.4 Our Results

We now turn to our constructions of condensers. We begin with the following construction, which is the rank condenser of Forbes and Shpilka [9] and was named the *folded Wronskian* by Guruswami-Kopparty [12].

▶ **Construction 8** (Folded Wronskian). *Let $\mathbb{F}$ be a field. Let $\omega \in \mathbb{F}$ be an element of multiplicative order $\geqslant n$. Define the matrix $\mathrm{W}_{t,\omega}(x) \in \mathbb{F}[x]^{[\![t]\!] \times [\![n]\!]}$ by $(\mathrm{W}_{t,\omega}(x))_{i,j} := (\omega^i x)^j$.*

Identifying $\mathbb{F}^{[\![n]\!]}$ with the vector space of degree $< n$ polynomials $\mathbb{F}[x]^{<n}$, the matrix $\mathrm{W}_{t,\omega}(x)$ defines the linear map $\mathrm{W}_{t,\omega}(x) : \mathbb{F}[x]^{<n} \to \mathbb{F}[x]^t$ given by

$$f(x) \mapsto (f(x), f(\omega x), \ldots, f(\omega^{t-1}x)) \,.$$

That is, we define $[\![n]\!] := \{0, \ldots, n-1\}$ so that in the above the indices i and j are indexed from zero. When the value of ω is clear from context we will just write "W_t". Note

[4] The *polynomial identity testing problem* is when given a algebraic circuit C (perhaps from a restricted class of circuits) to *deterministically* decide whether the circuit C computes the identically zero polynomial. The *black box* version is where we only allow access to C by evaluating the polynomial it computes. See Shpilka and Yehudayoff [25] for more on this problem.

that the fact that ω has large multiplicative order means that we require a large field, in particular that $|\mathbb{F}| > n$.

The key result that forms the starting point for our constructions is the following analysis of the folded Wronskian by Guruswami and Kopparty [12]. While their analysis was originally in the context of subspace designs, we state their result here in the language of lossless rank condensers as it is more natural in our context.

▶ **Theorem 9** (Guruswami-Kopparty [12]). *Assume the setup of Theorem 8 where we take $t \geqslant r \geqslant 1$. Let $S \subseteq \{(\omega^\ell)^j \mid j \geqslant 0\}$ where $\ell \geqslant t - r + 1$. Then $\{\mathrm{W}_t(\alpha) \mid \alpha \in S\} \subseteq \mathbb{F}^{t \times n}$ is a strong $(r, \frac{r(n-r)}{t-r+1})$-lossless rank condenser.*

We note here that the above parameters are slightly stronger than what Guruswami and Kopparty [12] obtain, as they only obtain a list bound of $\frac{r(n-1)}{t-r+1}$. This improved bound follows by using some of the analysis from Forbes, Saptharishi and Shpilka [8] as explained in the full version. Note that this construction essentially matches the non-constructive bound (2.1) when $\epsilon = 0$.

The above analysis indicates that for a matrix $M \in \mathbb{F}^{n \times r}$ of rank r that the total rank loss over all maps in $\mathcal{E}$ is at most $\frac{r(n-r)}{t-r+1}$. Thus, by an averaging argument, at most $1/k \cdot \frac{r(n-r)}{t-r+1}$ such maps can have a rank loss of $\geqslant k$. This observation thus shows that the above construction is not just a *lossless* rank condenser but also a *lossy* condenser (with different parameters).

▶ **Corollary 10.** *Let $\mathbb{F}$ be a field. Let $n, t \geqslant r \geqslant 1$ and $\epsilon > 0$, where $\omega \in \mathbb{F}$ is an element of multiplicative order $\geqslant \mathsf{poly}(n)$. Define $\mathcal{E} := \{\mathrm{W}_{t,\omega}((\omega^t)^j) \mid 0 \leqslant j < \frac{n}{\epsilon(t-r+1)}\}$, that is, the folded Wronskian evaluated at $\frac{n}{\epsilon(t-r+1)}$ distinct powers of ω^t. Then $\mathcal{E}$ is an explicit $(\leqslant r, \epsilon)$-lossy rank condenser.*

To motivate our below application to dimension expanders, suppose that $r = n/3$, $t = n/2$ and $\epsilon > 0$. This says then that we construct a rank condenser that maps $\mathbb{F}^n$ to $\mathbb{F}^{n/2}$ that maps rank $n/3$ subspaces to rank $(1-\epsilon)n/3$ subspaces. Further, this condenser is a collection of at most

$$\frac{n}{\epsilon(n/2 - n/3)} = 6/\epsilon$$

maps such that one map from the collection always preserves the desired rank. To obtain these parameters, it is key to the analysis that we have a *strong* lossless condenser and that it obtains the (near-optimal) bound given by Guruswami and Kopparty [12]. Note that these condensing parameters are similar to the min-entropy condensers of Raz [23] and Barak-Kindler-Shaltiel-Sudakov-Wigderson [3], which use a constant number of random bits to condense a source with a constant-rate of min-entropy.

3 Dimension Expanders

We now turn to our main object of interest, *dimension expanders*. Dimension expanders were defined by Barak, Impagliazzo, Shpilka and Wigderson [2] in an attempt to translate challenges in the explicit construction of objects in Boolean pseudorandomness into the regime of linear algebra. Indeed, in combinatorics there is a well-established analogy between subsets of $[n]$ and subspaces of vector spaces over finite fields. In the context of pseudorandomness, we can then translate questions that manipulate the *size* of subsets $S \subseteq \{0,1\}^n$ (or more generally, the min-entropy of distributions over $\{0,1\}^n$) into questions about manipulating the *dimension* of subspaces $V \subseteq \mathbb{F}^n$. While these regimes seem different, it is conceivable that such

linear algebraic constructions could yield new constructions in Boolean pseudorandomness (such as how the inner-product function is a two-source extractor). Indeed, as in the work of Guruswami and Wang [13], this idea has borne fruit (if in a perhaps unexpected way) by showing how linear-algebraic pseudorandom objects can improve list-decodable codes. We now define dimension expanders.

▶ **Definition 11.** Let $\mathbb{F}$ be a field, $n \geqslant 1$, $\epsilon > 0$ and $\alpha \in \mathbb{R}$ with $\alpha \geqslant 1$. A collection of matrices $\mathcal{A} = \{A_1, \ldots, A_d\} \subseteq \mathbb{F}^{n \times n}$ is a (ϵ, α)**-dimension expander of degree** d if for all subspaces $V \subseteq \mathbb{F}^n$ of dimension $\leqslant \epsilon n$ that

$$\dim \sum_{i=1}^{d} A_i(V) = \dim \operatorname{span}\{A_i(V)\}_{i=1}^{d} \geqslant \alpha \cdot \dim V \ .$$

The collection $\mathcal{A}$ is **explicit** if given an index $i \in [|\mathcal{A}|]$ the i-th matrix in $\mathcal{A}$ can be constructed in $\mathsf{poly}(n, \log |\mathcal{A}|)$ operations in $\mathbb{F}$.

We remark that in the above definition one can generally assume that all of the maps A_i are of full-rank, as that can only increase $\dim \sum_{i=1}^{d} A_i(V)$. Similarly, one can assume that A_1 equals the identity matrix I_n as we can use the transform $A_i \mapsto A_1^{-1} A_i$ as again this does not affect the size of the outputted dimension. While these assumptions are thus without loss of generality, we will not impose them.

In general we will be most interested in $(\Omega(1), 1 + \Omega(1))$-dimension expanders of constant degree, which we shall thus call "dimension expanders" without any quantification. This parameter regime is of interest because it matches that of the probabilistic method, which we quote the results of below.

▶ **Proposition 12.** *Let $\mathbb{F}_q$ be a finite field, $n \geqslant 1$, $\epsilon > 0$ and $\alpha \in \mathbb{R}$ with $\alpha \geqslant 1$. Then there exist a collection matrices $\mathcal{A} = \{A_1, \ldots, A_d\} \subseteq \mathbb{F}^{n \times n}$ which is a (ϵ, α)-dimension expander of degree d whenever*

$$d \geqslant \alpha + \frac{1}{1 - \alpha\epsilon} + o_q(1) \ .$$

Put into more concrete terms, we see that one can existentially obtain $(1/2d, d - O(1))$-dimension expansion with degree d. That we have an expansion of $(1 - \epsilon)d$ in a degree d expander is akin to *lossless (vertex) expanders* which have a similar degree/expansion relation, and these expanders have applications beyond those of normal expanders (see Capalbo, Reingold, Vadhan and Wigderson [5] and references therein). While previous work focused on obtaining constant-degree dimension expanders, our work raises the questions of obtaining *lossless* dimension expanders so that we match the above bound. Our work, as discussed below, lends itself to being particularly quantitative with regards to the size and parameters of the construction. However, we do not obtain lossless dimension expanders, and to the best of our knowledge, neither do the other previous constructions of dimension expanders discussed below.

While we discuss prior work in depth in the full version, we briefly summarize the state of art in dimension expanders in the following theorems.

▶ **Theorem** (Lubotzky and Zelmanov [20] and Harrow [17])**.** *Let $\mathbb{F}$ be a field of characteristic zero and $n \geqslant 1$. There exists an explicit $O(1)$-sized collection $\mathcal{A} \subseteq \mathbb{F}^{n \times n}$ such that $\mathcal{A}$ is a $(1/2, 1 + \Omega(1))$-dimension expander over $\mathbb{F}^n$.*

This construction requires characteristic zero as it uses a notion of distance that lacks a good definition in finite characteristic.

▶ **Theorem** (Bourgain and Yehudayoff [4]). *Let $n \geqslant 1$. There exists an explicit $O(1)$-sized collection $\mathcal{A} \subseteq \{0,1\}^{n \times n}$ such that $\mathcal{A}$ is a $(1/2, 1 + \Omega(1))$-dimension expander over $\mathbb{F}^n$, over every field $\mathbb{F}$.*

Note that the above construction is only a function of n, and not of the field, so that this *single* construction is a dimension expander over *all* fields.

As explained in the full version of this work, both of the above constructions in some way attempt to extend existing ideas about expander graphs into the world of dimension expanders. The first replicates the representation theory approach to constructing expanding Cayley graphs, and the second shows how bipartite expanders (with the strong requirement of *monotonicity*) extend to also be dimension expanders.

Our Work: In our work we take a different approach to constructing dimension expanders that treats such expanders as part of an emerging theme of *linear-algebraic* pseudorandomness as seen by recent linear-algebraic approaches to list-decoding [11, 15, 13, 16, 14] and linear-algebraic derandomization of subclasses of polynomial identity testing [18, 9]. The first consequence of this perspective is that we work in fields that are at least polynomially large as this is the setting of Reed-Solomon codes. To obtain dimension expanders over smaller fields, a natural solution within this theory is to use "code concatenation" ideas from coding theory. Unfortunately the idea of code concatenation is somewhat subtle in our setting and so only supplies a concatenation (based on converting Reed-Solomon codes to BCH codes) that incurs a logarithmic loss in the parameters. The second consequence is that we build our dimension expanders out of the existing linear-algebraic pseudorandom objects that have emerged from prior work. That is, just how in Boolean pseudorandomness the notions of expanders, extractors and list-decodable codes are all related (see for example Vadhan [26]), we leverage such connections to construct our expanders from the above mentioned rank condensers.

We now explain our construction, which while ultimately was motivated by the connections between two-source rank condensers and dimension expanders, can be explained in a self-contained manner. The first observation is that one can easily obtain "$(1, d)$-expanders" of degree $d \in \mathbb{N}$ if one is willing to allow the ambient space to grow. That is, consider the tensor product $\mathbb{F}^n \otimes \mathbb{F}^d = \mathbb{F}^{nd}$. By properties of the tensor product, for $V \subseteq \mathbb{F}^n$ of rank $r \leqslant n$ we know that $V \otimes \mathbb{F}^d$ is of rank rd in $\mathbb{F}^{nd}$. Further, $V \otimes \mathbb{F}^d$ can be seen as the image of d maps $T_i : \mathbb{F}^n \to \mathbb{F}^{nd}$ where the i-th map places the space $\mathbb{F}^n$ into the "i-th block" of $(\mathbb{F}^n)^d = \mathbb{F}^{nd}$. In analogy to bipartite expander graphs, this is akin to giving each left vertex its own "private neighborhood" of right vertices into which it expands.

While trivial, the above step now allows us to convert a question of *expansion* to a question of *condensing*. That is, tensoring achieves expansion only because the output of the maps are larger than the input, while the non-trivial aspect of dimension expanders is to expand while keeping the output size the *same*. However, tensoring *has* expanded dimension and thus we can now focus on reducing the output size. Specifically, suppose that we consider $V \subseteq \mathbb{F}^n$ of rank $r = n/2d$. Then its image under the above tensoring is $W := \sum_i T_i(V)$ of dimension $n/2$. This subspace W lies in an nd-dimensional space and we wish return it to an n-dimensional space while not losing too much in the dimension. However, this last problem is exactly the question of *lossy rank condensing*. For constant d, the above discussion shows that we can condense such constant-rate dimension in a lossy way using a *constant* number of maps. In this example, we can condense W to $\mathbb{F}^n$ using $\frac{dn}{\epsilon(n-n/2)} = \frac{2d}{\epsilon}$ maps, at least one of which produces a $(1-\epsilon)n/2$ dimensional space. Thus, this expands $V \subseteq \mathbb{F}^n$ of rank $\frac{n}{2d}$ to be of dimension $(1-\epsilon)n/2$ within $\mathbb{F}^n$, all while using $d \cdot \frac{2d}{\epsilon} = \frac{2d^2}{\epsilon}$ maps (we multiply the

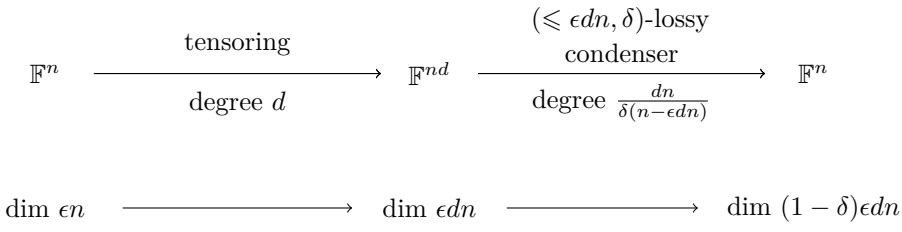

Figure 1 Constructing dimension expanders from tensoring and lossy rank condensers.

number of maps due to the composition). We summarize this composition in Figure 1.

We note that the above discussion has only discussed constant-rate rank, that is, subspaces of $\mathbb{F}^n$ with rank $\Omega(n)$. Dimension expanders however are required to expand *all* small subspaces. Our construction also handles this case as the lossy rank condensers we use will preserve a $(1-\delta)$ fraction of the input rank, as long as that rank is small enough. In the above sketch there is also the technicality that we must tensor with $\mathbb{F}^d$ with d being *integral*, which restricts $d \geqslant 2$ as $d = 1$ does not yield expansion. With this construction alone one would only obtain expansion in $\mathbb{F}^n$ for rank $< n/d \leqslant n/2$, but we manage to sidestep this restriction by a simple truncation argument. Putting the above pieces together we obtain the following theorem.

▶ **Theorem 13** (Main Theorem). *Let $n, d \geqslant 1$ and let $0 < \epsilon \leqslant \eta < 1$ be constants. Let $\mathbb{F}$ be a field with $|\mathbb{F}| \geqslant \mathsf{poly}(n)$. There is an explicit $(\epsilon, \eta/\epsilon)$-dimension expander in $\mathbb{F}^n$ of degree $\Theta\left(\frac{1}{\epsilon^2(1-\eta)^2}\right)$. If $\epsilon < 1/d$ then for any $\delta > 0$ there is an explicit $(\epsilon, (1-\delta)d)$-dimension expander in $\mathbb{F}^n$ with degree $\frac{d^2}{\delta(1-\epsilon d)}$.*

These expanders yield an expansion of α with degree $\approx \alpha^2$, and thus are not lossless. In particular, existential methods show that there are $(\epsilon, \eta/\epsilon)$-dimension expanders with degree $\approx 1/\epsilon + \frac{1}{1-\eta}$. In remains an interesting challenge to obtain such lossless dimension expanders. In particular, we note that we get "all of the dimension" from the tensoring step using only *one* map from the condenser. This occurs despite the fact that *most* maps in the condenser preserve all of the dimension (assuming we double the seed length). It seems natural to hope that an integrated analysis of the tensoring and condensing stages would show that the construction has a better expansion than what we obtain.

Over small fields our results are comparatively weaker as we simulate a larger field within the small field (as how one transforms Reed-Solomon codes to BCH codes), so that we pay various logarithmic penalties.

▶ **Corollary 14.** *Let $\mathbb{F}_q$ be finite and $n, d \geqslant 1$. Then there are explicit $\left(\Theta\left(\frac{1}{d \log_q dn}\right), \Theta(d)\right)$-dimension expanders in $\mathbb{F}_q^n$ of degree $\Theta(d^2 \log_q dn)$.*

4 Two-Source Rank Condensers

In the context of Boolean pseudorandomness, it is well known (see for example Vadhan [26]) that strong min-entropy seeded extractors (extractors that output the entropy of the source *plus* the entropy of the seed) are equivalent to a form of vertex expansion. Such extractors are a special case of (seedless) two-source min-entropy extractors where one of the sources is very small and of full entropy. Thus, as a generalization of the dimension expanders we have already defined, we can thus define the notion of a *(seedless) two-source rank condenser*.

While it is often most natural to consider the two sources to be of equal dimension, to highlight the connection to dimension expanders we consider sources with unbalanced dimension.

▶ **Definition 15.** Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$ and $m \geqslant s \geqslant 1$. A function $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ is a **(seedless)** (r, s, ϵ)**-two-source rank condenser** if for all sets $A \subseteq \mathbb{F}^n$ and $B \subseteq \mathbb{F}^m$ with $\operatorname{rank} A = r$ and $\operatorname{rank} B = s$,

$$\operatorname{rank} f(A \times B) = \operatorname{rank}\{f(\overline{v}, \overline{w})\}_{\overline{v} \in A, \overline{w} \in B} \geqslant (1 - \epsilon) \operatorname{rank} A \cdot \operatorname{rank} B \ .$$

The function f is a $(\leqslant r, s, \epsilon)$-condenser if it is a (r', s, ϵ)-condenser for all $1 \leqslant r' \leqslant r$, and $(\leqslant r, \leqslant s, \epsilon)$-condensers are defined similarly. If $\epsilon = 0$ we say the rank condenser is **lossless** and it is otherwise **lossy**. The function f is **bilinear** if $f(\overline{v}, \overline{w}) = (\overline{v}^{\mathrm{tr}} E_i \overline{w})_{i=1}^t$ for $E_i \in \mathbb{F}^{n \times m}$. The function f is **explicit** if it can be evaluated in $\mathsf{poly}(n, m, t)$ steps.

While this definition is naturally motivated as a generalization of dimension expanders, we originally were motivated to study these objects due to potential applications for constructing *subspace evasive sets*, as we describe in the full version.

Note that in general we allow the function f to be arbitrary, but in this work we will restrict ourselves to bilinear functions f as they are the most natural. In this case, as discussed after Theorem 4, we see that the function f acts on *subspaces* so that we ask that for subspaces $V \subseteq \mathbb{F}^n$ and $W \subseteq \mathbb{F}^m$ that $\dim f(V, W) \geqslant (1 - \epsilon) \dim V \cdot \dim W$. In this way, f can be thought of as a *derandomized tensor product*.

We now quote the parameters as given by the probabilistic method.

▶ **Proposition 16.** *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$ and $m \geqslant s \geqslant 1$ and $\epsilon \geqslant 0$. Then there exists a function $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ which is a bilinear (r, s, ϵ)-two-source rank condenser, assuming that*

$$t \geqslant \frac{n}{\epsilon s} + \frac{m}{\epsilon r} + (1 - \epsilon) rs + o_q(1) \ .$$

for $\epsilon > 0$. Further, there exists an f which is a $(\leqslant r, s, \epsilon)$-condenser assuming that

$$t \geqslant \frac{n}{\epsilon s} + \frac{m}{\epsilon} + (1 - \epsilon) rs + o_q(1) \ .$$

If $\epsilon = 0$, then there exists an f which is a $(r, s, 0)$-condenser assuming that

$$t \geqslant rn + sm + rs + o_q(1) \ .$$

In particular, in the balanced case of $n = m$ and $r = s$ this shows that any $t \geqslant \frac{2n}{\epsilon r} + (1 - \epsilon) r^2 + o_q(1)$ suffices. Note that unlike the single-source setting, there is a large penalty for condensing all small enough sources. Thus, the above gives $(r, r, 1/2)$-condensers with output $\approx \frac{n}{r} + r^2$ but to obtain a $(\leqslant r, r, 1/2)$-condenser the resulting output size is $\approx n + r^2$ (and the full version shows that a linear dependence on n is needed in this case).

Note that in our definitions of seeded rank condensers there was no analogue of *strong* min-entropy extractors, which are extractors that also recover the entropy of the seed in addition to the entropy of the source. That is, in our setting, there is no "rank of the seed" to recover as the seed is simply an index into the collection $\mathcal{E}$. The notion of a two-source rank condenser in some sense allows the second source to be a "seed" in that we can associate elements of $\mathcal{E}$ with elements in a basis for $\mathbb{F}^m$. However, we do not pursue this analogy further as two-source rank condensers meeting the probabilistic method do not seem to yield good lossy rank condensers in all regimes as two-source condensers can require an output size which is linear in the input size.

However, the connection between two-source extractors and expanders does hold tightly for the notion of rank, as we show. Note that for this connection it suffices to have condensers that work when one of the two sources has full rank.

▶ **Proposition 17.** *Let $\mathbb{F}_q$ be a finite field. For large n and all other parameters constant, constructions of bilinear $(\leqslant \delta n, m, \epsilon)$-two-source rank condensers $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ that meet the probabilistic method bound yield constructions of (δ, α)-dimension expanders in $\mathbb{F}^n$ meeting the probabilistic method bound.*

We also give constructions of two-source condensers using seeded rank condensers. That is, for two sources we use a seeded rank condenser to condense each source and use a union bound to show that the seed-length only doubles. We then enumerate over seeds and for each seed we then tensor the two condensed sources together. While this approach seems wasteful, we show that it yields *optimal* lossless two-source rank condensers by appropriate pruning. In particular, we observe that this is the same construction as given by Forbes and Shpilka [9] for an object known as a *rank-metric code*. We push this observation further to see that bilinear lossless two-source rank condensers are *equivalent* to rank-metric codes. Using this connection, we obtain optimal such condensers over *any* field using known constructions of rank-metric codes.

▶ **Theorem 18.** *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$ and $m \geqslant s \geqslant 1$. Then there is an explicit bilinear $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ which is a $(r, s, 0)$-two-source rank condenser with $t \leqslant O(\min\{r, s\} \cdot (n + m))$.*

We then turn to constructions of *lossy* two-source condensers, where our results are considerably weaker. However, we are able to give near-optimal results for *constant* r by using a brute force "inner condenser" and using our condense-then-tensor results as an "outer condenser".

▶ **Proposition 19.** *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$, where $|\mathbb{F}| \geqslant \mathsf{poly}(n)$ and $r \leqslant O(1)$. Then there is an explicit bilinear $(r, r, 1 - (1 - \epsilon)^3)$-two source rank condenser $f : \mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F}^t$ with $t \leqslant O(n/\epsilon^2 r)$.*

5 Open Questions

This work leaves several directions for future work.

1. Can one obtain (r, ϵ)-lossy seeded rank *extractors*, where the output is $\approx (1 - \epsilon)r$? Our methods require the output to be $\geqslant r$.
2. Can one develop of theory of "code concatenation" to improve our results for small fields?
3. Can one obtain lossy two-source rank condensers with output size $o(nr)$ for $r = \omega(1)$?
4. Can one obtain *lossless* dimension expanders, where the degree/expansion relationship matches the probabilistic method?
5. What is the complexity of computing dimension expansion? That is, given matrices $A_1, \ldots, A_d \in \mathbb{F}^{n \times n}$, compute the largest α so that $\mathcal{A} := \{A_i\}_{i=1}^d$ is a $(1/2, \alpha)$-dimension expander.

Acknowledgments. We would like to thank Swastik Kopparty, Prasad Raghavendra, Amir Shpilka, Amir Yehudayoff, Avi Wigderson, and anonymous reviewers for helpful comments.

References

1 Manindra Agrawal, Chandan Saha, and Nitin Saxena. Quasi-polynomial hitting-set for set-depth-Δ formulas. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC 2013)*, pages 321–330, 2013. Full version at arXiv:1209.2333.

2 Boaz Barak, Russell Impagliazzo, Amir Shpilka, and Avi Wigderson. Personal Communication to Dvir-Shpilka [**?**], 2004.

3 Boaz Barak, Guy Kindler, Ronen Shaltiel, Benny Sudakov, and Avi Wigderson. Simulating independence: New constructions of condensers, Ramsey graphs, dispersers, and extractors. *J. ACM*, 57(4), 2010. Preliminary version in the *37th Annual ACM Symposium on Theory of Computing (STOC 2005)*.

4 Jean Bourgain and Amir Yehudayoff. Expansion in $\mathrm{SL}_2(\mathbb{R})$ and monotone expanders. *Geometric and Functional Analysis*, 23(1):1–41, 2013. Preliminary version in the *44th Annual ACM Symposium on Theory of Computing (STOC 2012)*. This work is the full version of [**?**].

5 Michael R. Capalbo, Omer Reingold, Salil P. Vadhan, and Avi Wigderson. Randomness conductors and constant-degree lossless expanders. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC 2002)*, pages 659–668, 2002.

6 Ho Yee Cheung, Tsz Chiu Kwok, and Lap Chi Lau. Fast matrix rank algorithms and applications. *J. ACM*, 60(5):31, 2013. Preliminary version in the *44th Annual ACM Symposium on Theory of Computing (STOC 2012)*.

7 Zeev Dvir and Amir Shpilka. Locally decodable codes with two queries and polynomial identity testing for depth 3 circuits. *SIAM J. Comput.*, 36(5):1404–1434, 2007. Preliminary version in the *37th Annual ACM Symposium on Theory of Computing (STOC 2005)*.

8 Michael A. Forbes, Ramprasad Saptharishi, and Amir Shpilka. Hitting sets for multilinear read-once algebraic branching programs, in any order. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC 2014)*, pages 867–875, 2014. Full version at arXiv:1309.5668.

9 Michael A. Forbes and Amir Shpilka. On identity testing of tensors, low-rank recovery and compressed sensing. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC 2012)*, pages 163–172, 2012. Full version at arXiv:1111.0663.

10 Ariel Gabizon and Ran Raz. Deterministic extractors for affine sources over large fields. *Combinatorica*, 28(4):415–440, 2008. Preliminary version in the *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005)*.

11 Venkatesan Guruswami. Linear-algebraic list decoding of folded reed-solomon codes. In *Proceedings of the 26th Annual IEEE Conference on Computational Complexity (CCC 2011)*, pages 77–85, 2011. The full version of this paper is merged into Guruswami-Wang [13].

12 Venkatesan Guruswami and Swastik Kopparty. Explicit subspace designs. *Combinatorica*, pages 1–25, 2014. Preliminary version in the *54th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2013)*.

13 Venkatesan Guruswami and Carol Wang. Linear-algebraic list decoding for variants of reed-solomon codes. *IEEE Transactions on Information Theory*, 59(6):3257–3268, 2013. Preliminary versions appeared in Proceedings of the 26th Annual IEEE Conference on Computational Complexity (CCC 2011) and Proceedings of the 15th International Workshop on Randomization and Computation (RANDOM 2011).

14 Venkatesan Guruswami and Carol Wang. Evading subspaces over large fields and explicit list-decodable rank-metric codes. In *Proceedings of the 18th International Workshop on Randomization and Computation (RANDOM 2014)*, pages 748–761, 2014. Full version at arXiv:1311.7084.

15 Venkatesan Guruswami and Chaoping Xing. Folded codes from function field towers and improved optimal rate list decoding. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC 2012)*, pages 339–350, 2012. Full version at arXiv:1204.4209.

16 Venkatesan Guruswami and Chaoping Xing. List decoding Reed-Solomon, algebraic-geometric, and Gabidulin subcodes up to the Singleton bound. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC 2013)*, pages 843–852, 2013. Full version in the Electronic Colloquium on Computational Complexity (ECCC), Technical Report TR12-146.

17 Aram W. Harrow. Quantum expanders from any classical cayley graph expander. *Quantum Information & Computation*, 8(8–9):715–721, 2008.

18 Zohar S. Karnin and Amir Shpilka. Black box polynomial identity testing of generalized depth-3 arithmetic circuits with bounded top fan-in. *Combinatorica*, 31(3):333–364, 2011. Preliminary version in the *23rd Annual IEEE Conference on Computational Complexity (CCC 2008)*.

19 Daniel Lokshtanov, Pranabendu Misra, Fahad Panolan, and Saket Saurabh. Deterministic truncation of linear matroids. *arXiv*, 1404.4506, 2014.

20 Alexander Lubotzky and Efim Zelmanov. Dimension expanders. *J. Algebra*, 319(2):730–738, 2008.

21 Dániel Marx. A parameterized view on matroid optimization problems. *Theor. Comput. Sci.*, 410(44):4471–4479, 2009. Preliminary version in the *33rd International Colloquium on Automata, Languages and Programming (ICALP 2006)*.

22 Pavel Pudlák and Vojtěch Rödl. Pseudorandom sets and explicit constructions of Ramsey graphs. In *Complexity of computations and proofs*, volume 13 of *Quad. Mat.*, pages 327–346. Dept. Math., Seconda Univ. Napoli, Caserta, 2004.

23 Ran Raz. Extractors with weak random seeds. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, pages 11–20, 2005. Full version in the Electronic Colloquium on Computational Complexity (ECCC), Technical Report TR04-099.

24 Ran Raz and Amir Shpilka. Deterministic polynomial identity testing in non-commutative models. *Comput. Complex.*, 14(1):1–19, April 2005. Preliminary version in the *19th Annual IEEE Conference on Computational Complexity (CCC 2004)*.

25 Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5(3-4):2070–388, 2010.

26 Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012.