Dimension Expanders via Rank Condensers

Michael A. Forbes [*]   Venkatesan Guruswami [†]

November 28, 2014

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 [GK13] 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.

Contents

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 [Vad12].)

Recently, there is an emerging theory of "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 [PR04] and list-decodable codes [GW13, GX12], subspace designs to list decoding both in the Hamming metric and the rank metric [GX13, GW14], and rank-preserving condensers to affine extractors [GR08a] and polynomial identity testing [KS11, FS12].

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 $\alpha > 1$. 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 (equal to r in the *lossless* case, and $\Omega(r)$ 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 $t \ll n$ for the lossy case for $r \ll \sqrt{n}$).

Conceptually, our work highlights close interconnections between these notions. In particular, we show that subspace designs (which were introduced in the context of list decoding variants of algebraic-geometric codes in [GX13]) are the *same* concept as lossless rank condensers but that they emphasize 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 in [GK13] 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). We then give an expanded treatment with formal statements and proofs, as well as detailed comparison with related work in Sections 6, 7, 8 and 9. The connection between rank-metric codes and two-source lossless rank condensers is described in Appendix B. We also include detailed results on the parameters achieved by random constructions in Appendix C.

2 Subspace Designs and Rank Condensers

We begin by discussing the notion of a *subspace design*, as recently defined by Guruswami and Xing [GX13], 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 [FSS14]. 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 [GK13] implies nearly optimal lossy rank condensers.

2.1 Subspace Designs

We begin with the definition of a subspace design.

Definition 2.1 (Guruswami-Xing [GX13] and Guruswami-Kopparty [GK13]). *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 [GX13] 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 [GX13]). *Let $\mathbb{F}_q$ be a finite field. Let $\varepsilon > 0$, $n \geqslant {}^8\!/\varepsilon$ and $s \leqslant \varepsilon n/2$. Then there is a strong $(s, {}^{8s}\!/\varepsilon)$-subspace design $\mathcal{H}$ of $(1-\varepsilon)n$-dimensional subspaces in $\mathbb{F}_q^n$ with $|\mathcal{H}| = q^{\varepsilon n/8}$.* □

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

Lemma 2.2 (Guruswami-Kopparty [GK13]). *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 2.3. *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 we will later (Section 7), 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

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*. 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 [FSS14]. 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 Definition 2.5). We will discuss prior work and motivation for rank condensers that is less immediately relevant in Section 6. We begin with the definition.

Definition 2.4. *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 $\mathrm{rank}\, M = r$,*

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

*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 $\mathrm{rank}\, M = r$,*

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

*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 $\mathrm{rank}\, E = t$ 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 (Proposition 6.1) 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 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 (Proposition 6.1). *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$. Let $\mathcal{H} = \{H_i\}_{i \in [M]}$ be a collection of subspaces $H_i \subseteq \mathbb{F}^n$ and let $\mathcal{E} = \{E_i\}_{i \in [M]} \subseteq \mathbb{F}^{t \times n}$ be a collection of matrices, where we have that* row-span $E_i = (H_i)^{\perp}$ *for $i \in [M]$. Then $\mathcal{H}$ is a weak/strong (r,L)-subspace design iff $\mathcal{E}$ is a weak/strong (r,L)-lossless rank condenser.* $\qquad\square$

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 [FSS14] for quite different reasons. We now give a definition closer to their motivation.

Definition 2.5. *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$ and $\varepsilon \geqslant 0$. 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 $\mathrm{rank}\, M = r$,*

$$\mathrm{rank}\, EM \geqslant (1 - \varepsilon)\, \mathrm{rank}\, M \ ,$$

*for some $E \in \mathcal{E}$. The collection $\mathcal{E}$ is a (**seeded**) $(\leqslant r,\varepsilon)$-**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 $\mathrm{poly}(t,n,\log|\mathcal{E}|)$ operations in $\mathbb{F}$.* $\qquad\diamond$

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 one can also obtain this stronger requirement with our methods.

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 [FSS14], but the construction and usage of

these objects predates them[1]. In particular, Gabizon and Raz [GR08a] constructed a $(r,0)$-condenser with size nr^2, and they used this to construct affine extractors over large fields. Karnin and Shpilka [KS11] named the construction of Gabizon and Raz [GR08a] to be "rank preserving subspaces" and used this construction to make a *polynomial identity testing*[2] algorithm of Dvir and Shpilka [DS07] work in the *black box* model. Forbes and Shpilka [FS12] 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 [RS05] work in the black-box model. Forbes, Saptharishi and Shpilka [FSS14], building on the work of Agrawal, Saha, and Saxena [ASS13], 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 [LMPS14] used these condensers to derandomize a fixed-parameter-tractable algorithm of Marx [Mar09] for ℓ-matroid intersection. Cheung, Kwok and Lau [CKL13] rediscovered the rank condenser of Gabizon and Raz [GR08a] and (among other things) used this to give faster randomized algorithms for exact linear algebra. Forbes, Saptharishi and Shpilka [FSS14] 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 $\varepsilon = 0$, we expect the $\varepsilon > 0$ version will similarly have many applications.

We now quote the parameters given by the probabilistic method.

Proposition (Informal version of Proposition C.11 and Proposition C.12)**.** *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$, $\varepsilon \geqslant 0$ and $t > (1-\varepsilon)r$. 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*

$$k \geqslant \frac{rn + o_q(1)}{(t - (1-\varepsilon)r)(\lfloor \varepsilon r \rfloor + 1) - o_q(1)} \ . \tag{2.6}$$

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

$$k \geqslant \frac{n + o_q(1)}{\varepsilon(t - (1-\varepsilon)r) - o_q(1)} \ . \qquad \square$$

Thus we can make the output size t of the condenser to be almost equal to the guaranteed dimension bound of $(1-\varepsilon)r$. Further, we see that there is essentially no penalty in (existentially) insisting for a $(\leqslant r, \varepsilon)$-condenser over a (r, ε)-condenser. However, we show in Lemma 6.6 that the notion of $(\leqslant r, \varepsilon)$-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 [FS12] and was named the *folded Wronskian* by Guruswami-Kopparty [GK13].

Construction 2.7 (Folded Wronskian)**.** *Let $\mathbb{F}$ be a field. Let $\omega \in \mathbb{F}$ be an element of multiplicative order $\geqslant n$. Define $W_{t,\omega}(x) \in \mathbb{F}[x]^{[\![t]\!] \times [\![n]\!]}$ by $(W_{t,\omega}(x))_{i,j} := (\omega^i x)^j$.*

That is, identifying $\mathbb{F}^{[\![n]\!]}$ with the degree $< n$ polynomials $\mathbb{F}[x]^{<n}$, we see that $W_{t,\omega}(x) : \mathbb{F}[x]^{<n} \to \mathbb{F}[x]^t$ defined by

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

---

[1] 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.

[2] 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 [SY10] for more on this problem.

That is, we define $[\![n]\!] := \{0, \ldots, n-1\}$ so that in the above i and j are indexed from zero. When the value of ω is clear from context we will just write "W_t". Note 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 [GK13]. 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 2.8 (Guruswami-Kopparty [GK13]). *Assume the setup of Construction 2.7 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 $\{W_t(\alpha) \mid \alpha \in S\} \subseteq \mathbb{F}^{t \times n}$ is a strong $(r, \frac{r(n-r)}{t-r+1})$-lossless rank condenser.* $\qquad\square$

We note here that the above parameters are slightly stronger than what Guruswami and Kopparty [GK13] 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 [FSS14] as explained in Section 6. Note that this construction essentially matches the non-constructive bound (2.6) when $\varepsilon = 0$.

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 (Corollary 6.9). *Let $\mathbb{F}$ be a field. Let $n, t \geqslant r \geqslant 1$ and $\varepsilon > 0$, where $\omega \in \mathbb{F}$ is an element of multiplicative order $\geqslant \mathrm{poly}(n)$. Define $\mathcal{E} := \{W_{t,\omega}((\omega^t)^j) \mid 0 \leqslant j < \frac{n}{\varepsilon(t-r+1)}\}$, that is, the folded Wronskian evaluated at $\frac{n}{\varepsilon(t-r+1)}$ distinct powers of ω^t. Then $\mathcal{E}$ is an explicit $(\leqslant r, \varepsilon)$-lossy rank condenser.* $\qquad\square$

To motivate our below application to dimension expanders, suppose that $r = n/3$ and $t = n/2$ and some $\varepsilon > 0$. 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 $(1-\varepsilon)n/3$ subspaces. Further, this condenser is a collection of at most

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

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 [GK13]. Note that these condensing parameters are very much similar to the min-entropy condenser of Raz [Raz05], who uses a constant number of random bits to condense a source with constant-rate 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 [BISW04] 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 [GW13], 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 3.1. *Let $\mathbb{F}$ be a field, $n \geqslant 1$, $\varepsilon > 0$ 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 $\leqslant \varepsilon n$ that*

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

*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 $\mathrm{poly}(n, \log|\mathcal{A}|)$ 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 (Informal version of Proposition C.10)**.** *Let $\mathbb{F}_q$ be a finite field, $n \geqslant 1$, $\varepsilon > 0$ 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*

$$d \geqslant \alpha + \frac{1}{1 - \alpha\varepsilon} + o_q(1) \ . \qquad \square$$

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 $(1 - \varepsilon)d$ 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 [CRVW02] 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 Section 7, we briefly summarize the state of art in dimension expanders in the following theorems.

Theorem (Lubotzky and Zelmanov [LZ08] and Harrow [Har08])**.** *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 $\mathbb{F}^n$.* $\square$

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

Theorem (Bourgain and Yehudayoff [BY13])**.** *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 $\mathbb{F}$.* $\square$

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 Section 7, 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

$$\mathbb{F}^n \xrightarrow[\text{degree } d]{\text{tensoring}} \mathbb{F}^{nd} \xrightarrow[\text{degree } \frac{dn}{\delta(n-\varepsilon dn)}]{\substack{(\leqslant \varepsilon dn, \delta)\text{-lossy} \\ \text{condenser}}} \mathbb{F}^n$$

$$\text{dim } \varepsilon n \xrightarrow{\hspace{3cm}} \text{dim } \varepsilon dn \xrightarrow{\hspace{3cm}} \text{dim } (1-\delta)\varepsilon dn$$
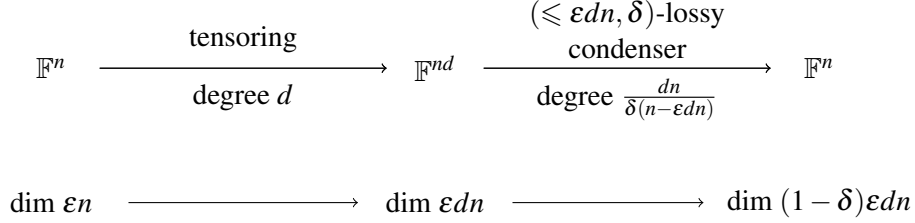
Figure 1: Constructing dimension expanders from tensoring and lossy rank condensers.

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 [Gur11, GX12, GW13, GX13, GW14] and linear-algebraic derandomization of subclasses of polynomial identity testing [KS11, FS12]. 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 [Vad12]), 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 (Proposition 9.1), 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 loosing too much in the dimension. However, this last problem is exactly the question of *lossy rank condensing*. As shown above after the informal version of Corollary 6.9, 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 $\frac{dn}{\varepsilon(n-n/2)} = \frac{2d}{\varepsilon}$ maps, at least one of which produces a $(1-\varepsilon)n/2$ dimensional space. Thus, this expands $V \subseteq \mathbb{F}^n$ of rank $\frac{n}{2d}$ to be of dimension $(1-\varepsilon)n/2$ within $\mathbb{F}^n$, all while using $d \cdot \frac{2d}{\varepsilon} = \frac{2d^2}{\varepsilon}$ maps (we multiply the 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 (Main Theorem, Informal version of [Corollary 7.3](#) and [Corollary 7.4](#)). *Let $n, d \geqslant 1$ and let $0 < \varepsilon \leqslant \eta < 1$ be constants. Let $\mathbb{F}$ be a field with $|\mathbb{F}| \geqslant \mathrm{poly}(n)$. There is an explicit $(\varepsilon, \eta/\varepsilon)$-dimension expander in $\mathbb{F}^n$ of degree $\Theta\left(\frac{1}{\varepsilon^2(1-\eta)^2}\right)$. If $\varepsilon < 1/d$ then there is an explicit $(\varepsilon, (1-\delta)d)$-dimension expander in $\mathbb{F}^n$ with degree $\frac{d^2}{\delta(1-\varepsilon d)}$.* □

These expanders yield an expansion of α with degree $\approx \alpha^2$, and thus are not lossless. In particular, the existential bound of [Proposition C.10](#) shows that there are $(\varepsilon, \eta/\varepsilon)$-dimension expanders with degree $\approx 1/\varepsilon + \frac{1}{1-\eta}$. 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 (Informal version of [Corollary 8.8](#)). *Let $\mathbb{F}_q$ be a finite field. Let $n, d \geqslant 1$. 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 [[Vad12](#)]) 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 ([Proposition 9.1](#)) we consider sources with unbalanced dimension.

Definition 4.1. *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 $\mathrm{rank}\,A = r$ and $\mathrm{rank}\,B = s$,*

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

*The function f is a $(\leqslant r, s, \varepsilon)$-condenser if it is a (r', s, ε)-condenser for all $1 \leqslant r' \leqslant r$, and $(\leqslant r, \leqslant s, \varepsilon)$-condensers are defined similarly. If $\varepsilon = 0$ 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 $\mathrm{poly}(n, m, t)$ 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 [Appendix A](#).

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 [Definition 2.4](#), 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 $\dim f(V,W) \geqslant (1-\varepsilon) \dim V \cdot \dim W$. 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 (Informal version of Proposition C.13 and Proposition C.14)**.** *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$ and $m \geqslant s \geqslant 1$ and $\varepsilon \geqslant 0$. 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*

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

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

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

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

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

In particular, in the balanced case of $n = m$ and $r = s$ this shows that any $t \geqslant \frac{2n}{\varepsilon r} + (1-\varepsilon)r^2 + o_q(1)$ 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 Proposition 9.4 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 (Proposition C.14) 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 (Proposition 9.4).

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 (Informal version of Proposition 9.1)**.** *Let $\mathbb{F}_q$ be a finite field. For large n and all other parameters constant, constructions of bilinear $(\leqslant \delta n, m, \varepsilon)$-two-source rank condensers $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ that meet the probabilistic method bound of Proposition C.14 yield constructions of (δ, α)-dimension expanders in $\mathbb{F}^n$ meeting the probabilistic method bound of Proposition C.10.* $\qquad \Box$

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 [FS12] for an object known as a *rank-metric code*. We push this observation further to see (in Appendix B) 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 (Informal version of Proposition 9.8 and Corollary B.6). *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$ and $m \geqslant s \geqslant 1$. Then there is an explicit $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ which is a $(r, s, 0)$ with $t \leqslant O(\max\{r, s\}(n + m))$.* $\square$

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 (Informal version of Proposition 9.9). *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$, where $\mathbb{F}$ is polynomially large and $r \leqslant O(1)$. Then there is an explicit bilinear $(r, r, 1 - (1 - \varepsilon)^3)$-two source rank condenser $f :$ $\mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F}^t$ with $t \leqslant O(n/\varepsilon^2 r)$.* $\square$

5 Notation

We briefly summarize some notation we use. We will apply functions not just to inputs, but to sets of inputs. Thus, if $f : S \to T$ and $U \subseteq S$, then $f(U) := \{f(x)\}_{x \in U}$. We denote $[n] := \{1, \ldots, n\}$ and $[\![n]\!] = \{0, \ldots, n-1\}$. We sometimes use $[\![n]\!]$ to index matrices from 0, so that $M \in \mathbb{F}^{[\![n]\!] \times m}$ has its rows indexed from zero but its columns indexed from 1. The i-th row of a matrix M will be denoted by $M_{i, \bullet}$ and the j-th column denoted $M_{\bullet, j}$. The transpose of a matrix M will be denoted by M^{tr}.

6 Constructions of Subspace Designs and Rank Condensers

In this section we relate subspace designs to lossless rank condensers, and use this relation to construct lossless and lossy seeded rank condensers. We begin with this relation, as given by taking dual spaces. We then construct lossless rank condensers using the folded Wronskian from Construction 2.7 and the analysis of Guruswami-Kopparty [GK13] quoted in Theorem 2.8 (which we slightly improve using the analysis of Forbes, Saptharishi and Shpilka [FSS14]). We then turn to lossy condensers, first noting that condensing "dimension $\leqslant r$" and "dimension r" is more subtle in this regime. We then show how to convert any strong lossless condenser to a lossy condenser and in particular do so for our explicit lossless condenser. As a result we obtain a lossy condenser nearly matching the probabilistic method.

6.1 Lossless Rank Condensers

We first show that subspace designs and lossless seeded rank condensers are the *same* object.

Proposition 6.1. *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$. Let $\mathcal{H} = \{H_i\}_{i \in [M]}$ be a collection of subspaces $H_i \subseteq \mathbb{F}^n$ and let $\mathcal{E} = \{E_i\}_{i \in [M]} \subseteq \mathbb{F}^{t \times n}$ be a collection of matrices, where we have that row-span $E_i = (H_i)^\perp$ for $i \in [M]$. Then $\mathcal{H}$ is a weak/strong (r, L)-subspace design iff $\mathcal{E}$ is a weak/strong (r, L)-lossless rank condenser.*

Proof: There is a natural surjection from matrices $M \in \mathbb{F}^{n \times r}$ with rank r to r-dimensional subspaces V given by $M \mapsto \mathrm{col\text{-}span}\, M$. We now show that rank condensers act on matrices in the same way that subspace designs act on subspaces. Each matrix E_i naturally defines a map $\varphi_i : \mathbb{F}^n \to \mathbb{F}^t$ where by definition $\ker \varphi_i = (\mathrm{row\text{-}span}\, E_i)^\perp = H_i$. Basic linear algebra shows that $\mathrm{rank}\, M = \dim V = \dim \varphi_i(V) + \dim(\ker \varphi_i \cap V)$. Using that $\varphi_i(V) = \mathrm{rank}\, E_i M$, we have that $\mathrm{rank}\, M - \mathrm{rank}\, E_i M = \dim(H_i \cap V)$.

Thus, summing over i we see that for M with $V = \mathrm{col\text{-}span}\, M$,

$$\sum_i (\mathrm{rank}\, M - \mathrm{rank}\, E_i M) = \sum_i \dim(H_i \cap V) \, .$$

13

Thus, the left summation is bounded by L for all M iff the right summation is bounded by L for all V, so that $\mathcal{E}$ is (r,L)-strong iff $\mathcal{H}$ is (r,L)-strong.

Observing that $\operatorname{rank} E_i M < \operatorname{rank} M$ iff $\dim(H_i \cap V) > 0$, the same conclusion holds for (r,L)-weak designs and condensers. $\qquad\square$

Given the above equivalence, we now phrase the construction of Guruswami and Kopparty [GK13] of subspace designs in the language of lossless condensers. To do so, we first need the following result of Forbes, Saptharishi and Shpilka [FSS14] which gives an analysis of the folded Wronskian.

Proposition 6.2 (Implicit in Forbes, Saptharishi and Shpilka [FSS14], given explicitly by Forbes [For14, Theorem 5.4.3]). *Assume the setup of Construction 2.7. Let $M \in \mathbb{F}^{n\times r}$ be of rank r. Then the $r \times r$ $\mathrm{W}_{r,\omega}(x)M \in \mathbb{F}[x]^{r\times r}$ has determinant $\det \mathrm{W}_{r,\omega}(x)M \in \mathbb{F}[x]$ which is a non-zero polynomial such that after dividing out powers of x has degree $\leqslant r(n-r)$.*

There are roughly three components to the above result: the non-zeroness of the polynomial $\det \mathrm{W}_{r,\omega}(x)M$, the degree bound on $\det \mathrm{W}_{r,\omega}(x)M$, and the multiplicative order lower bound of the distinguished element ω needed for these facts.

The non-zeroness of $\det \mathrm{W}_{r,\omega}(x)M$ has been explicitly established by various works. Forbes and Shpilka [FS12] gave a proof based on the Cauchy-Binet formula, Guruswami and Wang [GW13] via analyzing a certain linear system, Guruswami and Kopparty [GK13] via an analysis akin to that of the folded Reed-Solomon codes of Guruswami and Rudra [GR08b], and Lokshtanov, Misra, Panolan and Saurabh [LMPS14] via a reduction to the $n = 2$ case. We remark that these proofs have some similarity to analyses of the (classical) Wronskian, as discussed in Guruswami-Kopparty [GK13] and Forbes-Saptharishi-Shpilka [FSS14], so that similar proofs can be extracted from that literature.

The above proofs of non-zeroness yield different degree upper bounds on $\det \mathrm{W}_{r,\omega}(x)M$ and the multiplicative order of ω. The obvious degree bound for $\det \mathrm{W}_{r,\omega}(x)M$ is $r(n-1)$ as this is an $r \times r$ determinant with entries who are polynomials of degree $< n$. Forbes and Shpilka [FS12] refined this bound to $nr - \binom{r+1}{2}$, and this is tight. However, Forbes, Saptharishi and Shpilka [FSS14] observed [3] (and this is explicitly stated in Forbes [For14]) that one can clear powers of x to derive the degree bound of $r(n-r)$, and again this is tight. Regarding the multiplicative order of ω, most of the above proofs show that order $\geqslant n$ is sufficient, and this is also tight.

Given the above analysis, we now derive the weak lossless condenser of Forbes and Shpilka [FS12] (with an updated degree bound).

Corollary 6.3 (Forbes and Shpilka [FS12]). *Assume the setup of Construction 2.7 where we take $t = r$. Let $S \subseteq \mathbb{F} \setminus \{0\}$. Then $\{\mathrm{W}_r(\alpha) \mid \alpha \in S\} \subseteq \mathbb{F}^{r\times n}$ is a weak $(r, r(n-r))$-lossless rank condenser.*

Proof: As $\operatorname{rank} \mathrm{W}_r(\alpha)M = \operatorname{rank} M = r$ iff $\det \mathrm{W}_r(\alpha)M \neq 0$, we see that by Proposition 6.2 there are at most $r(n-r)$ such $\alpha \in \mathbb{F} \setminus \{0\}$ where $\operatorname{rank} \mathrm{W}_r(\alpha)M < \operatorname{rank} M$, giving the desired bound. $\qquad\square$

Note that in the above condenser the output size t is equal to the rank bound r so that this might appropriately be called a rank *extractor*[4]. In some applications this equality is important, and indeed for our

---

[3]Forbes, Saptharishi and Shpilka [FSS14] noted more generally that a similar construction follows from any error-correcting code, and that the code in the folded Wronskian is the dual Reed-Solomon code. However, the dual Reed-Solomon code seems special in that it gives rise to "folding" of the resulting construction, which is crucial for Theorem 2.8.

[4]The paper of Dvir, Gabizon and Wigderson [DGW09] also defined a notion of a rank extractor, but that was for *algebraic* rank (or transcendence degree) of polynomials which generalizes linear-algebraic rank to polynomials of higher degree.

construction of dimension expanders we would ideally have lossy rank extractors (which we do not achieve, as discussed after Corollary 6.9).

While the above degree and order analysis is tight, Guruswami-Kopparty [GK13] extended it in two ways. The first is to study the *multiplicities* and the second is to study the case when $t > r$. They observed that the multiplicity of vanishing of $\det W_r(\alpha)M$ upper bounds the rank deficiency $(\operatorname{rank} M - \operatorname{rank} W_r(\alpha)M)$ so that instead of a weak condenser we can obtain a *strong* condenser with the same list bound. For $t > r$, they observed that taking the values of α as powers of ω introduces a certain *redundancy* between different α-values which is akin to the folding of folded Reed-Solomon codes of Guruswami-Rudra [GR08b] and this allows one to obtain a smaller list bound by pruning this redundancy. Updating their argument with the above degree bound (as powers of ω are non-zero, so we can safely ignore powers of x in the determinant with respect to the determinant being non-zero) we get the following analysis of the folded Wronskian as a lossless rank condenser.

Theorem 2.8 (Guruswami-Kopparty [GK13]). *Assume the setup of Construction 2.7 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 $\{W_t(\alpha) \mid \alpha \in S\} \subseteq \mathbb{F}^{t \times n}$ is a strong $(r, \frac{r(n-r)}{t-r+1})$-lossless rank condenser.* $\square$

We now conclude with an explicit instantiation of the above when taking all points in a finite field.

Corollary 6.4. *Let $\mathbb{F}_q$ be a finite field and let $n, t \geqslant r \geqslant 1$ such that $q > n$. Given a generator ω of $\mathbb{F}_q$, there is an explicit $\mathcal{E} \subseteq \mathbb{F}_q^{t \times n}$ with $|\mathcal{E}| = \lfloor \frac{q-1}{t} \rfloor$, such that $\mathcal{E}$ is a strong $\left(r, \frac{r(n-r)}{t-r+1}\right)$-lossless rank condenser.*

Proof: As ω generates $\mathbb{F}_q$, we can apply Theorem 2.8 as ω has the desired multiplicative order of $q - 1 \geqslant n$. Now take $S = \left\{1, \omega^t, (\omega^t)^2, \ldots, (\omega^t)^{\lfloor \frac{q-1}{t} \rfloor}\right\}$ and $\mathcal{E} := \{W_t(\alpha) \mid \alpha \in S\}$. For any $1 \leqslant r \leqslant t$, we have that $t \geqslant t - r + 1$ so that $\mathcal{E}$ has the desired condensing properties by Theorem 2.8. That $|\mathcal{E}| = \lfloor \frac{q-1}{t} \rfloor$ follows from construction by the multiplicative order of ω. That $\mathcal{E}$ is explicit is also clear as we can index $\mathcal{E}$ by $\left[\lfloor \frac{q-1}{t} \rfloor\right]$ and then given an $i \in \left[\lfloor \frac{q-1}{t} \rfloor\right]$ produce $W_t(\omega^i)$ in $\operatorname{poly}(n, t, \log q)$ operations in $\mathbb{F}_q$ via repeated squaring. $\square$

6.2   Lossy Rank Condensers

We now turn from lossless condensers to *lossy* condensers, as defined in Section 2. The motivation for studying objects that can lose a small amount of rank is to obtain a comparatively smaller seed length. Before turning to constructions, we study the issue of condensing "rank r" as compared to the stronger notion of condensing "rank $\leqslant r$", where the latter notion is provably stronger. We then relate lossless condensers to lossy condensers, observing that an averaging argument converts strong lossless condensers to lossy condensers with a *smaller* list bound. This leads us to constructing lossy condensers from our lossless constructions, and this will achieve condensing of "rank $\leqslant r$" as needed for our application to dimension expanders (Section 7) as they must expand *all* small subspaces.

We begin by recalling that a (r, ε)-lossy condenser should condenser rank r to rank $(1 - \varepsilon)r$, and that a $(\leqslant r, \varepsilon)$-condenser should condense rank s to rank $(1 - \varepsilon)s$ for *all* $s \leqslant r$. Insisting on the latter notion is somewhat out of line with the usual definition of min-entropy condensers, that ask for min-entropy $\geqslant k$ being condensed to some $\geqslant k'$ min-entropy, without any (stated) guarantee on sources with input min-entropy $\ll k$. As such, $(\leqslant r, \varepsilon)$-condenser notion is more akin to the *conductors* of Capalbo, Reingold, Vadhan and Wigderson [CRVW02] (which are maps that have min-entropy output guarantees for any input min-entropy) than condensers. However, when $\varepsilon = 0$, we see that condensing for rank r implies condensing $\leqslant r$ with the same list bound, similar to Lemma 2.3.

Lemma 6.5. *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$. Let $\mathcal{E} \subseteq \mathbb{F}^{t \times n}$ be a weak/strong (r, L)-lossless rank condenser. Then $\mathcal{E}$ is a $(\leqslant r, L)$-lossless rank condenser.* $\square$

However, we now note that this connection breaks for $\varepsilon > 0$.

Lemma 6.6. *Let $\mathbb{F}$ be a field, let $n \geqslant 1$. Let $\mathcal{E} \subseteq \mathbb{F}^{t \times 3n}$ be a $(2n, 1/2)$-lossy rank condenser. Let $\pi : \mathbb{F}^{4n} \to \mathbb{F}^{3n}$ be the projection map onto the first $3n$ coordinates and let $P \in \mathbb{F}^{3n \times 4n}$ be the associated projection matrix. Then $\mathcal{E}' := \{EP \mid E \in \mathcal{E}\} \subseteq \mathbb{F}^{t \times 4n}$ is a $(3n, 2/3)$-lossy condenser but not a $(s, 1 - \delta)$-condenser for any $1 \leqslant s \leqslant n$ and $\delta < 1$.*

Proof: For a vector space $V \subseteq \mathbb{F}^{4n}$, the dimension of the projection $\pi(V)$ has $\dim \pi(V) = \dim V - \dim(V \cap \ker \pi) \geqslant \dim V - n$. Thus, if $\dim V \geqslant 3n$ then $\dim \pi(V) \geqslant 2n$, so that there is some $E \in \mathcal{E}$ so that $E(\pi(V))$ has dimension $\geqslant n$. Thus $\mathcal{E}'$, the composition of $\mathcal{E}$ and π, will map spaces of dimension $3n$ to spaces of dimension $\geqslant n$ so that $\mathcal{E}'$ is a $(3n, 2/3)$-condenser.

Now consider $V \subseteq \mathbb{F}^{4n}$ of dimension $s \leqslant n$ which is a subspace of the kernel of the projection map π (which has dimension n). Then clearly $\pi(V) = 0$, so that the rank of the condenser $\mathcal{E}'$ on $\pi(V)$ is always zero, so $\mathcal{E}'$ preserves none of the rank of V, so that $\mathcal{E}'$ is not a $(s, 1 - \delta)$-condenser for any $1 \leqslant s \leqslant n$ and $\delta < 1$. $\square$

This example embeds itself into many examples of "manipulating rank r" versus "manipulating rank $\leqslant r$" by pseudorandom objects. Thus, it shows that to obtain the latter guarantee one needs to explicitly consider dimension $\leqslant r$, and indeed our techniques will work in this regime.

We now give constructions of good lossy rank condensers, using *strong* lossless rank condensers. As strong lossless condensers bound the sum of *all* rank losses $(\operatorname{rank} M - \operatorname{rank} EM)$ it follows from an averaging argument that an at most $1/k$ fraction of the maps can have a rank deficiency $\geqslant k$. Thus we can take a seed length of the lossy condenser that is $1/k$ of the list bound of the lossless condenser. For this reduction, a weak lossless condensers would not suffice as the resulting seed length would not be smaller than the original list bound.

Proposition 6.7. *Let $\mathbb{F}$ be a field and let $n \geqslant r \geqslant 1$ and $\varepsilon \geqslant 0$. Let $\mathcal{E} \subseteq \mathbb{F}^{t \times n}$ be a strong (r, L)-lossless rank condenser. Then for any $\mathcal{E}' \subseteq \mathcal{E}$ with $|\mathcal{E}'| > \frac{L}{\lfloor \varepsilon r \rfloor + 1}$, $\mathcal{E}'$ is a (r, ε)-lossy rank condenser.*

Proof: We argue the contrapositive. As $\mathcal{E}'$ is not a lossy condenser there is a matrix $M \in \mathbb{F}^{n \times r}$ with rank r so that for all $E \in \mathcal{E}'$,
$$\operatorname{rank} EM < (1 - \varepsilon) \operatorname{rank} M ,$$
and thus
$$\operatorname{rank} M - \operatorname{rank} EM > \varepsilon \operatorname{rank} M = \varepsilon r ,$$
so that
$$\operatorname{rank} M - \operatorname{rank} EM \geqslant \lfloor \varepsilon r \rfloor + 1 .$$
Thus, using that $\mathcal{E}'$ is a strong (r, L)-design (as it is a subset of $\mathcal{E}$ is such a design, and this is preserved under taking subsets),
$$\begin{aligned} L &\geqslant \sum_{E \in \mathcal{E}'} (\operatorname{rank} M - \operatorname{rank} EM) \\ &\geqslant \sum_{E \in \mathcal{E}'} (\lfloor \varepsilon r \rfloor + 1) \\ &= |\mathcal{E}'| \cdot (\lfloor \varepsilon r \rfloor + 1) , \end{aligned}$$
and thus $|\mathcal{E}'| \leqslant \frac{L}{\lfloor \varepsilon r \rfloor + 1}$ as desired. $\square$

We now use this lemma, along with the results on lossless condensers, to obtain our desired lossy condenser.

Proposition 6.8. *Assume the setup of Construction 2.7 where we take $n,t \geqslant r \geqslant 1$ and $\varepsilon > 0$. Let $S \subseteq \{(\omega^\ell)^j \mid j \geqslant 0\}$ where $\ell \geqslant t$ and $|S| \geqslant \min\left\{\frac{n}{\varepsilon(t-r+1)}, n^2\right\}$. Then $\mathcal{E} := \{W_t(\alpha) \mid \alpha \in S\} \subseteq \mathbb{F}^{t \times n}$ is a $(\leqslant r, \varepsilon)$-lossy rank condenser.*

Proof: If $|S| \geqslant n^2 \geqslant r(n-r)+1$ then we have a *lossless* $(\leqslant r, 0)$-condenser by Corollary 6.3 (and Lemma 6.5).

Thus consider $|S| \geqslant \frac{n}{\varepsilon(t-r+1)}$. By Theorem 2.8 and Proposition 6.7 it follows that $\mathcal{E}$ is a (s, ε)-lossy condenser as long as $|S| > \frac{s(n-s)}{(\lfloor \varepsilon s \rfloor + 1)(t-s+1)}$ (as $\ell \geqslant t \geqslant t-s+1$ as $s \geqslant 1$). In particular, it suffices if

$$|S| \geqslant \frac{sn}{(\lfloor \varepsilon s \rfloor + 1)(t-s+1)}$$

and thus as $\lfloor \varepsilon s \rfloor + 1 \geqslant \varepsilon s$ and $s \leqslant r$, it suffices if

$$|S| \geqslant \frac{sn}{\varepsilon s \cdot (t-r+1)} = \frac{n}{\varepsilon \cdot (t-r+1)} \,,$$

where this last bound is independent of s, so that we get the desired result. $\qquad\square$

Note that in the above it is crucial that not only does the condenser of Theorem 2.8 condense all small ranks $s \leqslant r$, but also that the list bound is smaller as s decreases. Now we take this construction with explicit values for the α's.

Corollary 6.9. *Let $\mathbb{F}$ be a field. Let $n,t \geqslant r \geqslant 1$ and $\varepsilon > 0$. Define $N := \left\lceil \frac{n}{\varepsilon(t-r+1)} \right\rceil$ and let $M \geqslant N$. Suppose $|\mathbb{F}|$ is of size $> tn^2$. Then there is an explicit $(\leqslant r, \varepsilon)$-lossy rank condenser $\mathcal{E} \subseteq \mathbb{F}^{t \times n}$ of size $|\mathcal{E}| = M$.*

Proof: As $|\mathbb{F}| > tn^2$ it follows that we can find an element $\omega \in \mathbb{F}$ with multiplicative order $\geqslant tn^2$ in $\mathrm{poly}(n,t)$ steps (see for example Forbes [For14, Lemma A.0.5]). Thus, the set $S := \{(\omega^t)^j \mid 0 \leqslant j < \min\{N, n^2\}\}$ has size $\min\{N, n^2\}$ and these explicit elements are all distinct as ω has order $\geqslant t\min\{N, n^2\}$.

We now appeal to Proposition 6.8, seeing that our set S is sufficiently large and our setup matches that of Construction 2.7 as ω has order $\geqslant n$. Further, we see that the resulting matrices $\mathcal{E}$ are explicit as Construction 2.7 is explicit. Padding the result with zero-matrices yields an explicit set $\mathcal{E}$ with size M. $\qquad\square$

Thus, we see that the above essentially almost matches the list bound of $\frac{n}{\varepsilon(t-(1-\varepsilon)r)}$ of the probabilistic method for lossy condensers condensing dimension $\leqslant r$, as given by Proposition C.12. However, there are two gaps. The first is that this construction requires the use of polynomially large fields, while the existential bound also holds for constant-sized fields. Second is that this construction requires that the output t have "$t \geqslant r$" even though we only require the rank to be $\geqslant (1-\varepsilon)r$. The probabilistic method of Proposition C.12 allows us to take $t \approx (1-\varepsilon)r$. While this does not seem dramatic, it will cause a slight complication in our construction of dimension expanders (as discussed after Corollary 7.3).

7 Constructions of Dimension Expanders

In this section we construct constant-degree dimension expanders by composing a tensoring operation with our construction of a lossy rank condenser, as constructed in in Section 6. We first discuss previous constructions of dimension expanders, then turn to our own construction.

7.1 Previous Constructions

There have been two main types of constructions. The first is to use Cayley graphs from groups with Kazhdan's *property T*, and the second is to use *monotone expanders*.

Property T: This approach to constructing dimension expanders is rooted in their similarities to expanding Cayley graphs, which we now discuss. Specifically, dimension expanders can be seen as a certain type of (vertex) expander in the usual sense, as observed by Dvir and Shpilka [DS11]. That is, given a $(\Omega(1), 1 + \Omega(1))$-dimension expander $\mathcal{A} = \{A_1, \ldots, A_d\} \subseteq \mathbb{F}_q^n$ (where we assume without loss of generality that all A_i are invertible, as discussed after Definition 3.1), consider the graph G with vertex set $\mathbb{F}_q^n$ such that $\bar{v} \in \mathbb{F}_q^n$ is connected to $\{A_i \bar{v}\}_i$. Note that this is a Schreier graph, as we have a subgroup of $\mathrm{GL}_n(\mathbb{F}_q)$ (the subgroup generated by the A_i) acting on the set $\mathbb{F}_q^n$. While in general this graph G is directed, one could assume (as is common in Cayley and Schreier graphs) that $\mathcal{A}$ is symmetric so that $A \in \mathcal{A}$ iff $A^{-1} \in \mathcal{A}$, in which case G is undirected. That the graph G expands (as a vertex expander) would mean that whenever $S \subseteq \mathbb{F}_q^n$ has $|S| \leqslant (1 - \Omega(1))|\mathbb{F}_q^n|$, that the neighborhood of S has size at least $1 + \Omega(1)$ times the size of S. That is, that $|\cup_i A_i(S)| \geqslant (1 + \Omega(1))|S|$. That $\mathcal{A}$ is a *dimension* expander asks for an expansion property of G that is both weaker and stronger than that of vertex expansion in some respects. It is weaker in that we only care about when the set S is a *subspace* of $\mathbb{F}_q^n$. However, it is also stronger as vertex expansion only yields that $|\cup_i A_i(S)| \geqslant (1 + \Omega(1))|S|$ which only implies $\dim \mathrm{span} \cup_i A_i(S) \geqslant \dim S + \Omega(1)$, while dimension expansion yields that $\dim \mathrm{span} \cup_i A_i(S) \geqslant (1 + \Omega(1)) \dim S$.

As seen by the above connection, dimension expanders can be seen as a type of Schreier graph. Thus, to understand the construction of dimension expanders it is first helpful to recall the construction of expanding Cayley graphs in particular (as these are Schreier graphs). In particular, if you have a group G with generators Γ, then the corresponding Cayley graph is an expander if the *Kazhdan constant* is strictly bounded away from zero (in which case it is said that G has *property T* with respect to Γ). The Kazhdan constant being bounded away from zero roughly means that each irreducible unitary representation of G must "move" each non-zero vector a non-trivial amount via some generator in Γ.

Given that the above notions are inherently linear algebraic, Wigderson [Wig04] made a conjecture (see Dvir and Wigderson [DW10, Conjecture 7.1]) that any expanding Cayley graph would yield a dimension expander. Specifically, that any irreducible representation $\rho : G \to \mathbb{F}^{n \times n}$ of that group G with generators $\Gamma \subseteq G$ would have that $\rho(\Gamma)$ is a dimension-expander. This collection will be of constant-size if the original expander was constant degree, and would intuitively expand dimension by analogy to how a positive Kazhdan constant "moves" any vector so these matrices must "move" a subspace to obtain a non-trivial amount of additional dimension.

Lubotzky and Zelmanov [LZ08] proved Wigderson's conjecture in characteristic zero by exploiting the connection of expansion in Cayley graphs to the underlying group having property T. Thus, they established explicit constant-degree $(\Omega(1), 1 + \Omega(1))$-dimension expanders in any field of characteristic zero. Unfortunately, as property T relies on the representations being over the characteristic zero (so that distance is a well-defined notion), the conjecture remains open in finite characteristic. Harrow [Har08] independently obtained this result in the context of *quantum expanders*, which imply dimension expanders in characteristic zero. We summarize this in the following theorem.

Theorem (Lubotzky and Zelmanov [LZ08] and Harrow [Har08])**.** *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 $\mathbb{F}^n$.* $\qquad\square$

Monotone Expanders: This second approach to constructing dimension expanders exploits another similarity to expander graphs, but now to bipartite (vertex) expanders. Specifically, as observed in Bourgain

and Yehudayoff [BY13], bipartite vertex expanders can be seen as special cases of dimension expanders, where bipartite vertex expanders only expand subspaces spanned by basis vectors and dimension expanders expand all subspaces. Indeed, suppose the graph G is on the vertex set with bipartition $[n] \sqcup [n]$ and we partition the edges E of G into partial matchings $E = E_1 \sqcup \cdots \sqcup E_d$ so that d is an upper bound on the degree of G. Then we can view the sets E_i as defining (partial) maps $E_i : [n] \to [n]$ which we can then view as matrices $A_i \in \{0,1\}^{n \times n}$ by using these maps to act on the standard basis vectors and then extending linearly. That G is a good vertex expander means that for any $S \subseteq [n]$ with $|S| \leqslant (1 - \Omega(1))n$ that the neighborhood of S in G is slightly larger, that is, $|\cup_i E_i(S)| \geqslant (1 + \Omega(1))|S|$. However, from this we see that the vector space $V := \mathrm{span}\{\overline{e}_i\}_{i \in S}$ (where $\overline{e}_i$ is the i-th standard basis vector) thus expands under the collection $\{A_i\}$, as $\dim \sum_i A_i(V) = |\cup_i E_i(S)| \geqslant (1 + \Omega(1))|S| = (1 + \Omega(1)) \dim V$.

Given the above connection, one can then ask in which contexts do the above matrices A_i also expand *any* subspace, as opposed to just those spanned by basis vectors. Dvir and Shpilka [DS11] implicitly observed that dimension expansion occurs when these (partial) maps $E_i : [n] \to [n]$ are *monotone*. That is, if each edge set E_i defines a partial map $E_i : [n] \to [n]$ so that if $E_i(j)$ and $E_i(k)$ are defined for $j, k \in [n]$, then $j < k \implies E_i(j) < E_i(k)$. When this monotonicity occurs, the resulting matrices $\{A_i\}_i$ are a dimension expander, and this statement was made explicit by Dvir and Wigderson [DW10].

Thus, to construct explicit constant-degree dimension expanders it then suffices to construct explicit constant-degree monotone expanders (where the partition of the edges into monotone maps must also be explicit). Unfortunately, monotone expanders seem a more delicate object than unrestricted expanders. Indeed, the standard probabilistic method arguments cannot demonstrate even the *existence* of constant-degree monotone expanders (see [DW10, BY13]). However, using this connection along with Cayley expanders over $\mathbb{Z}_n$, Dvir and Shpilka [DS11] were able to construct monotone expanders (and thus dimension expanders) with logarithmic degree, as well as constant-degree expanders with inverse-logarithmic expansion. We formally state their result here to contrast with our results over $\mathbb{F}_2$ (Corollary 8.8), as our results are much weaker (only achieving $(\Omega(1/\lg n), 1 + \Omega(1))$-expanders of logarithmic degree).

Theorem (Dvir and Shpilka [DS11]). *Let $n \geqslant 1$. There exists an explicit $O(\lg n)$-sized collection $\mathcal{A} \subseteq \{0,1\}^{n \times n}$ such that $\mathcal{A}$ is a $(\Omega(1), 1 + \Omega(1))$-dimension expander over $\mathbb{F}^n$, for every field $\mathbb{F}$.*

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 $(\Omega(1), 1 + \Omega(1/\lg n))$-dimension expander over $\mathbb{F}^n$, for every field $\mathbb{F}$. $\square$

Dvir and Wigderson [DW10] gave an iterative construction of monotone expanders in spirit of the zig-zag product of Reingold, Vadhan and Wigderson [RVW02]. Using this approach they were able to give monotone expanders (and thus dimension expanders) of degree $\lg^{(c)}(n)$ (the c-th iterated logarithm) for any constant c. Assuming a base construction of a constant-degree monotone expander (which is *not* known to exist via the probabilistic method), they were able to produce constant-degree dimension expanders.

In a more sophisticated work, Bourgain and Yehudayoff [BY13] used expansion in the group $\mathrm{SL}_2(\mathbb{R})$ to obtain explicit constant degree monotone expanders, and thus dimension expanders.

Theorem (Bourgain and Yehudayoff [BY13]). *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 $\mathbb{F}^n$, for every field $\mathbb{F}$.* $\square$

A remarkable feature of the above works is that they achieve (for each n) a *single* collection of matrices that is a dimension expander for *every* field, in particular by only using $\{0,1\}$-values. In contrast, our constructions will very much depend on the underlying field by using elements of large multiplicative order.

7.2 Our Construction

We now proceed to give the details of our construction, following the outline given in Section 3. That is, we apply a tensoring operation to yield expansion but increasing the ambient dimension, and then use a lossy rank condenser to preserve this expansion while reducing the ambient dimension to its original size. As mentioned in Section 3, this approach is somewhat limited to expanding rank $\leqslant n/2$ subspaces as we can only tensor with integral-dimensional spaces. To circumvent this, we observe that we can simply "forget" some of the rank we obtained by tensoring and this allows the construction to expand any rank. We start with a generic analysis, where we parameterize the "forgetfullness" by γ.

Proposition 7.1. *Let $\mathbb{F}$ be a field. Let $n \geqslant r \geqslant 1$, $d \geqslant 1$, $\delta \in (0,1]$, and $\gamma \in [0,1]$. Let $\mathcal{E} \subseteq \mathbb{F}^{n \times nd}$ be a $(\leqslant \lceil (1-\gamma)rd \rceil, \delta)$-lossy rank condenser. For $i \in [d]$, define $T_i \in \mathbb{F}^{nd \times n}$ to be the matrix of the map $\bar{v} \mapsto \bar{v} \otimes \bar{e}_i$, where $\bar{e}_i \in \mathbb{F}^d$ is the i-th standard basis vector. Define $\mathcal{A} := \{ET_i \mid E \in \mathcal{E}, i \in [d]\}$. Then $\mathcal{A} \subseteq \mathbb{F}^{n \times n}$ is a $(r/n, (1-\gamma)(1-\delta)d)$-dimension expander of degree $d \cdot |\mathcal{E}|$.*

Proof: That $\mathcal{A} \subseteq \mathbb{F}^{n \times n}$ is clear from construction, as $E \in \mathbb{F}^{n \times nd}$ and $T_i \in \mathbb{F}^{nd \times n}$. That $|\mathcal{A}| = d \cdot |\mathcal{E}|$ is also clear. We now argue the expansion property.

Consider some $V \subseteq \mathbb{F}^n$ with $\dim V = s \leqslant r$. Then we have that $V \otimes \mathbb{F}^d \subseteq \mathbb{F}^{nd}$ has rank sd by the properties of the tensor product, and using that $V \otimes \mathbb{F}^d = \sum_i T_i(V)$, it follows that $\dim \sum_i T_i(V) = sd$. In particular, we have that there is some subspace $W \subseteq \sum_i T_i(V)$ with $\dim W = \lceil (1-\gamma)sd \rceil \leqslant \lceil (1-\gamma)rd \rceil$. By the hypothesis on $\mathcal{E}$, it follows that for some $E \in \mathcal{E}$ that $\dim E(W) \geqslant (1-\delta)\lceil (1-\gamma)sd \rceil \geqslant (1-\delta)(1-\gamma)sd = (1-\delta)(1-\gamma)d \cdot \dim V$. Thus, as $\dim \sum_{A \in \mathcal{A}} A(V) \geqslant \dim E(W) \geqslant (1-\delta)(1-\gamma)d \cdot \dim V$ we see that there is the desired expansion. $\quad \square$

We now instantiate this recipe with our explicit construction of a lossy rank condenser (Corollary 6.9) to deduce the following.

Theorem 7.2. *Let $\mathbb{F}$ be a field. Let $n, d \geqslant 1$. Let $\varepsilon \in (0,1)$, $\delta \in (0,1]$ and $\gamma \in [0,1]$, subject to*

$$(1-\gamma)\varepsilon d < 1 .$$

Then there is an explicit $(\varepsilon, (1-\gamma)(1-\delta)d)$-dimension expander of degree

$$d \cdot \left\lceil \frac{d}{\delta(1-(1-\gamma)\varepsilon d)} \right\rceil ,$$

whenever $|\mathbb{F}| > d^2 n^3$.

Proof: For $r = \lfloor \varepsilon n \rfloor$, Corollary 6.9 yields an explicit $(\leqslant \lceil (1-\gamma)rd \rceil, \delta)$-lossy rank condenser of size $M := \left\lceil \frac{d}{\delta(1-(1-\gamma)\varepsilon d)} \right\rceil$ as

$$\left\lceil \frac{nd}{\delta(n - \lceil (1-\gamma)rd \rceil + 1)} \right\rceil \leqslant \left\lceil \frac{nd}{\delta(n - (1-\gamma)rd)} \right\rceil \leqslant \left\lceil \frac{d}{\delta(1-(1-\gamma)\varepsilon d)} \right\rceil .$$

Note that Corollary 6.9 requires that $n \geqslant \lceil (1-\gamma)rd \rceil$, which is true iff $(1-\gamma)rd \leqslant n$, which follows from $(1-\gamma)\varepsilon d \leqslant 1$, and we have by hypothesis that $(1-\gamma)\varepsilon d < 1$. We now use this condenser in Proposition 7.1, noting that this construction is also explicit and multiplies the degree by d. $\quad \square$

We now note one particularly natural set of parameters for the above construction, namely when $\gamma = 0$.

Corollary 7.3. *Let $\mathbb{F}$ be a field. Let $n, d \geqslant 1$. Let $\varepsilon \in (0,1)$, $\delta \in (0,1]$, subject to $\varepsilon < 1/d$. Then there is an explicit $(\varepsilon, (1-\delta)d)$-dimension expander of degree $d \cdot \left\lceil \frac{d}{\delta(1-\varepsilon d)} \right\rceil$ whenever $|\mathbb{F}| > d^2 n^3$.* $\qquad\square$

Now consider this construction when $\varepsilon \leqslant 1/2d$ and $\delta = 1/2$, so that we have that the degree of the expander is $\approx 4d^2$, when the expansion α has $\alpha = d/2$. However, the existential construction of Proposition C.10 yields that the degree could be $\approx \alpha + 2$ here. Thus, our construction suffers from a quadratic loss as compared to the existential bounds. In particular, we do not achieve lossless dimension expanders. Intuitively, this quadratic loss is because we compose a tensor step with a condensing step, so that the degree of each step multiplies.

While the above corollary with $\gamma = 0$ is sufficient for expanding small rank subspaces, it intrinsically cannot yield, for example, $(2/3, 1 + \Omega(1))$-expanders, because the parameter d must be an *integer*. Thus, after our tensoring step, from rank r we get rank rd and we must take $d > 1$ to achieve *any* expansion. But as there is no way to tensor with $\mathbb{F}^d$ for d fractional, we must have that $d \geqslant 2$. However, our construction of a lossy rank condenser (Proposition 6.8) *requires* that the output size be at least the rank bound of $rd \geqslant 2r$. Thus, since the output size must be n so that we achieve $n \times n$ matrices, it follows that this method does not work for $2r > n$. Note that if we had lossy condensers meeting the existential bound (Proposition C.12) then we would only need that $n \geqslant (1-\delta)rd$ and thus taking δ sufficiently close to 1 would remedy the $r \leqslant n/2$ limitation seen here.

However, as we currently cannot match the above existential bound, the above construction introduces the γ parameter. When $\gamma > 0$, we simply "forget" that the tensoring yields dimension rd and simply work with the smaller rank bound of $(1-\gamma)rd$, thus allowing us to take any r where $(1-\gamma)rd \leqslant n$. In particular, by letting γ close to 1 we can take *any* rank r. We now implement the above by choosing γ carefully so that we can obtain constant degree expanders in $\mathbb{F}^n$ that expand rank εn to rank ηn for any constants $\varepsilon \leqslant \eta < 1$.

Corollary 7.4. *Let $\mathbb{F}$ be a field. Let $n \geqslant 1$. Let $0 < \varepsilon \leqslant \eta < 1$. Then there is an explicit $(\varepsilon, \eta/\varepsilon)$-dimension expander of degree*

$$\left\lceil \frac{1+\eta}{2\varepsilon} \right\rceil \cdot \left\lceil \frac{2(1+\eta)\left\lceil \frac{1+\eta}{2\varepsilon} \right\rceil}{(1-\eta)^2} \right\rceil \leqslant \lceil 1/\varepsilon \rceil \cdot \left\lceil \frac{4\lceil 1/\varepsilon \rceil}{(1-\eta)^2} \right\rceil ,$$

whenever $|\mathbb{F}| > d^2 n^3$.

Proof: This follows from Theorem 7.2 by choosing parameters carefully. In particular, we will to choose d, δ, γ so that $(1-\delta)(1-\gamma)\varepsilon d = \eta$ but that $(1-\gamma)\varepsilon d$ is bounded away from one.

In particular, choose $d = \left\lceil \frac{1+\eta}{2\varepsilon} \right\rceil$. Note that as $\varepsilon < 1$ we have that $d \geqslant 2$. Now choose γ so that $(1-\gamma)\varepsilon d = \frac{1+\eta}{2}$ so that $(1-\gamma)\left\lceil \frac{1+\eta}{2\varepsilon} \right\rceil = \frac{1+\eta}{2\varepsilon}$ from which it follows that $\gamma \in [0,1)$. Now choose δ so that $(1-\delta)(1-\gamma)\varepsilon d = \eta$, in particular that $(1-\delta) = \frac{2\eta}{1+\eta}$, so that $\delta = \frac{1-\eta}{1+\eta} \in (0,1)$. Thus, plugging these values into Corollary 7.3 yields the desired parameters. $\qquad\square$

The above constructions all require large fields. In Section 8 we show how to simulate these results in small fields by paying certain logarithmic penalties.

8   Constructions over Small Fields

The main constructions of this paper rely on the folded Wronskian (Construction 2.7) which requires a polynomially large field. In this section, we discuss to what extent we can extend our techniques to smaller

fields. Guruswami and Kopparty [GK13] gave a way to convert subspace designs over large fields to subspace designs over small fields with comparable parameters. However, their method was not able to preserve the *strong-ness* of the subspace design and as we saw this strong-ness is essential for our construction of constant degree dimension expanders (Section 7). We give in this section an alternate method for simulating large fields that preserves strong-ness but is slightly worse in other parameters than the method of Guruswami and Kopparty [GK13]. Our method is based on the conversion of Reed-Solomon codes to BCH codes and can be seen as a basic form of "code concatenation" from coding theory (although more sophisticated versions of that idea do not seem to work in our setting, see Remark 8.9). As a consequence, we obtain over *any* field constructions of strong lossless rank condensers with "inverse logarithmic output rate" and logarithmic-degree $(\Omega(1/\lg n), 1 + \Theta(1))$-dimension expanders.

We begin by showing how to transform a matrix over an extension field $\mathbb{K}$ of $\mathbb{F}$ to a matrix over $\mathbb{F}$, while preserving rank. This operation will increase the rows of this matrix (which is undesirable, but tolerable) and is akin to converting Reed-Solomon codes to BCH codes.

Construction 8.1. *Let $\mathbb{F}$ be a subfield of $\mathbb{K}$, where $\dim_{\mathbb{F}} \mathbb{K} = k$ so that $\varphi : \mathbb{K} \cong \mathbb{F}^k$ is an $\mathbb{F}$-vector space isomorphism. Define $\varphi^n : \mathbb{K}^n \cong \mathbb{F}^{kn}$ by applying φ coordinate-wise, so that $(\overline{\alpha}_1, \dots, \overline{\alpha}_n) \mapsto (\varphi(\overline{\alpha}_1), \dots, \varphi(\overline{\alpha}_n))$. For a matrix $M \in \mathbb{K}^{n \times m}$, define $\varphi^n(M) \in \mathbb{F}^{kn \times m}$ as the result of applying φ^n to each column of M.* □

The key point about this map φ^n is that it is $\mathbb{F}$-linear so that it composes nicely with matrix multiplication.

Lemma 8.2. *Assume the setup of Construction 8.1. Let $E \in \mathbb{K}^{t \times n}$ and $M \in \mathbb{F}^{n \times r}$. Then $\varphi^t(EM) = \varphi^t(E)M$.*

Proof: Let E have columns $\overline{v}_1, \dots, \overline{v}_n \in \mathbb{K}^t$. Then the i-th column is $(EM)_{\bullet,i} = \sum_{j=1}^n M_{j,i} \overline{v}_j$. Thus by $\mathbb{F}$-linearity of φ^t, $\varphi^t((EM)_{\bullet,i}) = \sum_{j=1}^n M_{j,i} \varphi^t(\overline{v}_j) = (\varphi^t(E)M)_{\bullet,i}$. Thus, each column of $\varphi^t(EM)$ and $\varphi^t(E)M$ agree, so they are equal. □

We now show that this map φ^n preserves rank of matrices as we switch fields.

Lemma 8.3. *Assume the setup of Construction 8.1. Let $M \in \mathbb{K}^{n \times m}$ be a matrix. Then $\varphi^n(M)$ has $\operatorname{rank}_{\mathbb{F}} \varphi^n(M) \geqslant \operatorname{rank}_{\mathbb{K}} M$.*

Proof: Consider the kernel of the matrix M over $\mathbb{K}$, $\ker_{\mathbb{K}} M = \{\overline{v} \mid \overline{v} \in \mathbb{K}^m, M\overline{v} = \overline{0}\}$, so that this represents the $\mathbb{K}$-dependencies among the columns of M. Likewise, consider the $\mathbb{F}$-dependencies of $\varphi^n(M)$, $\ker_{\mathbb{F}} \varphi^n(M) = \{\overline{v} \mid \overline{v} \in \mathbb{F}^m, \varphi^n(M)\overline{v} = \overline{0}\}$. By Lemma 8.2 we have $\varphi^n(M\overline{v}) = \varphi^n(M)\overline{v}$ for $\overline{v} \in \mathbb{F}^m$. As φ^n is an isomorphism, we conclude that for $\overline{v} \in \mathbb{F}^m$, $M\overline{v} = \overline{0}$ iff $\varphi^n(M)\overline{v} = \overline{0}$. Thus, it follows that $\ker_{\mathbb{F}} \varphi^n(M) \subseteq \ker_{\mathbb{K}} M$. In particular, we see that the $\mathbb{F}$-dependencies between the columns of M are a subset of the $\mathbb{K}$-dependencies of the columns of M, from which the claim follows. □

Thus, we arrive at the following.

Corollary 8.4. *Assume the setup of Construction 8.1. Let $E \in \mathbb{K}^{t \times n}$ and $M \in \mathbb{F}^{n \times r}$. Then $\operatorname{rank}_{\mathbb{F}} \varphi^t(E)M \geqslant \operatorname{rank}_{\mathbb{K}} EM$.* □

Note that in the above the results are in general neither tight nor improvable. To see that the rank can increase, consider $\alpha_1, \dots, \alpha_k \in \mathbb{K}$ which are an $\mathbb{F}$-basis for $\mathbb{K}$. Then the matrix $M = [\alpha_1, \dots, \alpha_k] \in \mathbb{K}^{1 \times k}$ has $\mathbb{K}$-rank equal to 1, while $\varphi^1(M)$ has $\mathbb{F}$-rank equal to k. To see that the rank may not increase, consider the $n \times n$ identity matrix $I_n \in \mathbb{K}^n$. The resulting matrix $\varphi^n(I_n) \in \mathbb{F}^{kn \times n}$ still has rank n as its rank cannot exceed the number of columns.

In general, one can "get more" $\mathbb{F}$-rank from a matrix over $\mathbb{K}$ by expanding the number of rows *and columns* of this matrix. That is, for a $\mathbb{K}$-vector space V we use that $\dim_{\mathbb{F}} V = k \dim_{\mathbb{K}} V$ instead of just

that $\dim_{\mathbb{F}} V \geqslant \dim_{\mathbb{K}} V$. Indeed, Guruswami-Kopparty [GK13] used this to simulate large fields in their construction of subspace designs. The downside of this approach is that it does not act as the composition of linear maps (so we do not get Lemma 8.2), and thus does not preserve strong-ness of the design.

We now apply the above results to convert strong lossless rank condensers over large fields to condensers over small fields.

Proposition 8.5. *Assume the setup of Construction 8.1 with $n \geqslant r \geqslant 1$. Let $\mathcal{E} \subseteq \mathbb{K}^{t \times n}$ be a weak/strong (r, L)-lossless rank condenser. Then $\varphi^n(\mathcal{E}) := \{\varphi^n(E) \mid E \in \mathcal{E}\} \subseteq \mathbb{F}^{kt \times n}$ is a weak/strong (r, L)-lossless rank condenser.*

Proof: Consider a matrix $M \in \mathbb{F}^{n \times r}$ with $\operatorname{rank}_{\mathbb{F}} M = r$. Then by Corollary 8.4, we have that for any $E \in \mathcal{E}$ that $\operatorname{rank}_{\mathbb{F}} \varphi^t(E)M \geqslant \operatorname{rank}_{\mathbb{K}} EM$. Noting that $\operatorname{rank}_{\mathbb{F}} M = \operatorname{rank}_{\mathbb{K}} M$ as M is an $\mathbb{F}$-matrix, we have the inclusion of sets $\{E \mid \operatorname{rank}_{\mathbb{K}} EM < \operatorname{rank}_{\mathbb{K}} M\} \supseteq \{E \mid \operatorname{rank}_{\mathbb{F}} \varphi^t(E)M < \operatorname{rank}_{\mathbb{F}} M\}$, so since the former set has size $\leqslant L$ then so does the latter. As this holds for any M, it follows that $\varphi^t(\mathcal{E})$ has the desired weak condensing.

Similarly, we see that $\operatorname{rank}_{\mathbb{F}} M - \operatorname{rank}_{\mathbb{F}} \varphi^t(E)M \leqslant \operatorname{rank}_{\mathbb{K}} M - \operatorname{rank}_{\mathbb{K}} EM$ and since summing over $E \in \mathcal{E}$ of the right-hand side yields $\leqslant L$ (assuming now $\mathcal{E}$ is strong), then so does summing over the left-hand side. Thus, $\varphi^t(\mathcal{E})$ has the desired strong condensing as well. $\qquad\square$

We now apply the above simulation of a large field to translate Corollary 6.4 to small fields.

Corollary 8.6. *Let $\mathbb{F}_q$ be a finite field and let $n, t \geqslant r \geqslant 1$ such that $q^k > n$. Given an explicit presentation of $\mathbb{F}_{q^k} \cong \mathbb{F}_q^k$ and a generator ω of $\mathbb{F}_{q^k}$, there is an explicit $\mathcal{E} \subseteq \mathbb{F}_q^{kt \times n}$ with $|\mathcal{E}| = \lfloor \frac{q^k - 1}{t} \rfloor$, such that for all $1 \leqslant r \leqslant t$, $\mathcal{E}$ is a strong $(r, \frac{r(n-r)}{t-r+1})$-lossless rank condenser (and thus strong $(r, \frac{r(n-r)}{t-r+1})$-subspace design).*

Proof: Corollary 6.4 implies that we have such a $\mathcal{E}' \in \mathbb{F}_{q^k}^{t \times n}$ with $|\mathcal{E}'| = \lfloor \frac{q^k - 1}{t} \rfloor$ that has the desired condensing properties (and thus design properties by our equivalence (Proposition 6.1)). We now apply our large field simulation (Proposition 8.5) to obtain $\mathcal{E}$. Clearly $\mathcal{E}$ has the desired size and condensing properties. That $\mathcal{E}$ can be indexed in $\operatorname{poly}(n, t, k \log q)$ is also clear given the explicit presentation $\mathbb{F}_{q^k} \cong \mathbb{F}_q^k$. $\qquad\square$

In comparison, Guruswami and Kopparty [GK13] obtain a similar construction of weak designs where the list bound has *decreased* by a factor of k. However, for our applications the strong-ness of the above is more important.

Similarly, we can obtain explicit lossy rank condensers where the output size is logarithmically larger than it was previously.

Corollary 8.7. *Let $\mathbb{F}_q$ be a finite field. Let $n, t \geqslant r \geqslant 1$ and $\varepsilon > 0$. Define $N := \left\lceil \frac{n}{\varepsilon(t-r+1)} \right\rceil$ and let $M \geqslant N$. Then there is an explicit $(\leqslant r, \varepsilon)$-lossy rank condenser $\mathcal{E}' \subseteq \mathbb{F}^{kt \times n}$ of size $|\mathcal{E}'| = M$, where $k := \lceil \log_q(tn^2 + 1) \rceil = \Theta(\log_q tn)$.*

Proof: First, observe that $tn^2 + 1 \leqslant q^k \leqslant q \cdot (tn^2 + 1) \leqslant (tn^2 + 1)^2$. Now, recall that we can construct an explicit presentation of $\mathbb{F}_{q^k}$ extending $\mathbb{F}_q$ in time $\operatorname{poly}(q^k) \leqslant \operatorname{poly}(n, t)$ (see for example Forbes [For14, Lemma 3.2.5]). Now by Corollary 6.9 we have a $(\leqslant r, \varepsilon)$-lossy rank condenser $\mathcal{E} \subseteq \mathbb{F}_{q^k}^{t \times n}$ of size $|\mathcal{E}| = M$. We then convert this to be over $\mathbb{F}_q$ following the logic of Proposition 8.5. $\qquad\square$

Likewise we obtain logarithmic-degree dimension expanders, but we can only expand rank of inverse logarithmic rate.

Corollary 8.8. *Let $\mathbb{F}_q$ be a finite field. Let $n, d \geqslant 1$ and define $k := \lceil \log_q(d^2 n^3 + 1) \rceil$. Let $\varepsilon \in (0,1)$ and $\delta \in (0,1]$, subject to $1 - \varepsilon dk < 1$. Then there is an explicit $(\varepsilon, (1-\delta)d)$-dimension expander of degree $d \cdot \left\lceil \frac{dk}{\delta(1-\varepsilon dk)} \right\rceil$.*

Proof: As in Theorem 7.2, we will instantiate Proposition 7.1 with an explicit lossy condenser. However, here we take the parameter γ of Theorem 7.2 to have $\gamma = 0$, as $\gamma > 0$ was only needed to expand in the high-rate regime (which this proof will not achieve).

Thus, for $r = \lfloor \varepsilon n \rfloor$ we use Corollary 8.7 to obtain an explicit $(\leqslant rd, \delta)$-condenser in $\mathbb{F}_q^{n \times nd}$ with resulting seed length of $\left\lceil \frac{dk}{\delta(1-\varepsilon dk)} \right\rceil$ as

$$
\left\lceil \frac{nd}{\delta(\lfloor n/k \rfloor - rd + 1)} \right\rceil \leqslant \left\lceil \frac{nd}{\delta(n/k - rd)} \right\rceil = \left\lceil \frac{ndk}{\delta(n - rdk)} \right\rceil \leqslant \left\lceil \frac{dk}{\delta(1 - \varepsilon dk)} \right\rceil ,
$$

noting that we are restricted in needing that $\lfloor n/k \rfloor \cdot k \leqslant n$ (and we pad matrices from $k \lfloor n/k \rfloor$ rows to n rows). Plugging this into Proposition 7.1 thus yields the desired dimension expander. $\qquad\square$

Put more informally, this yields for every integer d, a $\left(\Theta\left(\frac{1}{d \log_q dn} \right), \Theta(d) \right)$-dimension expanders in $\mathbb{F}_q^n$ of degree $\Theta(d^2 \log_q dn)$.

We now briefly remark on the difficulty of using more sophisticated code concatenation ideas from coding theory to obtain better results.

Remark 8.9. The above constructions mimic the conversion of Reed-Solomon codes of block-length n (over the large alphabet of $\mathbb{F}_{q^k}$, with $k = \lceil \log_q n \rceil$) to BCH codes of block-length n (over the small alphabet of $\mathbb{F}_q$). This conversion preserves the distance of the code, but multiplies the number of parity checks of the code by factor of $\lceil \log_q n \rceil$ for codes of block-length n. As such, the relative distance of the resulting BCH code cannot be better than $1/\lceil \log_q n \rceil$ without the code becoming trivial.

Coding theory has the method of *code concatenation* for reducing the alphabet size without incurring the above logarithmic loss. In particular, one can reduce Reed-Solomon codes to a code over $\mathbb{F}_2$ while only incurring a constant factor loss in the distance and rate. Unfortunately, it is unclear how to implement code concatenation in the context of this paper.

Specifically, in code concatenation for the Hamming metric, one first uses the isomorphism that $\mathbb{F}_q^{nk} = \mathbb{F}_{q^k}^n$, one then uses the outer-code $E_1 : \mathbb{F}_{q^k}^n \to \mathbb{F}_{q^k}^m$, to which one then applies an inner-code $E_2 : \mathbb{F}_{q^k} = \mathbb{F}_q^k \to \mathbb{F}_q^{k'}$ component-wise. In the Hamming metric, the isomorphism in the first step of the composition will preserve (or increase) the relative Hamming weight, so that a vector $\bar{v} \in \mathbb{F}_q^{nk}$ with weight $\frac{d}{nk}$ will have weight at least $\frac{d/k}{n} = \frac{d}{nk}$ when considered as a vector in $\mathbb{F}_{q^k}^n$. Thus, we have not lost in the distance by appealing to this isomorphism, which intuitively follows from the fact that the Hamming weight of a vector is defined coordinate-wise.

However, this isomorphism does not seem to play well with rank (as the rank of a matrix is not defined column-wise), as discussed in the comments after Corollary 8.4. That is, the rank can drop by a factor of k when treating a $\mathbb{F}_q$-vector space $V \subseteq \mathbb{F}_q^{nk}$ as a $\mathbb{F}_{q^k}$-vector space in $\mathbb{F}_{q^k}^n$. One can recover this factor of k as done in the large-field simulation of Guruswami-Kopparty [GK13] in the context of subspace designs but this construction seems to lose the strong-ness of the construction which is essential for our work. $\qquad\diamond$

9 Constructions of Two-Source Rank Condensers

In this section, we relate two-source rank condensers to the pseudorandom objects we have already discussed. First, we show that two-source rank condensers with good parameters (even for the special case when one of the sources has full rank) yield constructions of dimension expanders with good parameters. We then show that *seeded* (single-source) rank condensers can be used to construct two-source rank condensers with good parameters. In particular, for $(r, r, 0)$-condensers we obtain an output length of $\Theta(nr)$, which is essentially optimal. However, we note that this is essentially the construction of Forbes and Shpilka [FS12] for constructing *rank-metric codes*. We show in Appendix B that (bilinear) lossless two-source rank condensers are *equivalent* to (linear) rank-metric codes, and as a consequence derive *optimal* such condensers from rank-metric code constructions. For $(r, r, 1 - (1 - \varepsilon)^3)$-lossy two-source condensers, we show how to obtain output size $\Theta(n/\varepsilon^2 r)$ for *constant r* by using our ideas as an "outer condenser" and finding via brute-force an "inner condenser".

We begin by showing how two-source rank condensers imply dimension expanders.

Proposition 9.1. *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$, $m \geqslant 1$, and $\varepsilon > 0$. Let $E_1, \ldots, E_t \in \mathbb{F}^{n \times m}$ be such that $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ with $f(\bar{v}, \bar{w}) := (\bar{v}^{\mathrm{tr}} E_i \bar{w})_{i \in [t]}$ is a bilinear $(\leqslant r, m, \varepsilon)$-two-source rank condenser.*

Then, for $i \in [m]$, define $A_i \in \mathbb{F}^{t \times n}$ to the be matrix of the linear transformation $\bar{v} \mapsto f(\bar{v}, \bar{e}_i)$, where $\bar{e}_i \in \mathbb{F}^m$ is the i-th standard basis vector. Then $\mathcal{A} := \{A_i\}_{i=1}^{m}$ has the following property: for all $V \subseteq \mathbb{F}^n$ of dimension $\leqslant r$,

$$\dim \sum_i A_i(V) \geqslant (1 - \varepsilon) m \dim V .$$

In particular, consider $\delta := r/n$, $\alpha := (1 - \varepsilon)m$ with $\alpha \delta < 1$, where m, δ, α and ε are constant. For $n \geqslant \Omega_{m,\alpha,\delta,\varepsilon}(1)$ whenever

$$t \leqslant \frac{n}{\varepsilon m} + \frac{m}{\varepsilon} + (1 - \varepsilon) rm + o_q(1) + O(1) ,$$

we have that the collection $\mathcal{A}$ is a degree-m (δ, α)-dimension expander in $\mathbb{F}^n$ whenever

$$m > \alpha + \frac{1}{1 - \alpha \delta} .$$

Proof: First note that for each $i \in [m]$, the map $\bar{v} \mapsto (\bar{v}^{\mathrm{tr}} E_j \bar{e}_i)_{j \in [t]}$ is indeed a linear map in $\bar{v}$ from $\mathbb{F}^n \to \mathbb{F}^t$, so that A_i is well defined. Now consider a subspace $V \subseteq \mathbb{F}^n$ of dimension $\leqslant r$. By bilinearity of f, we have that $\dim f(V, \mathbb{F}^m) = \dim f(V, I_m)$, where I_m is the $m \times m$ identity matrix so that col-span $I_m = \mathbb{F}^m$. The rank condenser property guarantees that $\dim f(V, I_m) \geqslant (1 - \varepsilon) m \dim V$. In particular, as $\{f(V, \bar{e}_i)\}_{i \in [m]}$ spans $\sum_i A_i(V)$, it follows that $\dim \sum_i A_i(V) \geqslant (1 - \varepsilon) m \dim V$ as desired.

To see the second part of the claim, note that for $\mathcal{A}$ to be a dimension expander we only need $t \leqslant n$ as we can then pad the matrices in $\mathcal{A}$ to be $n \times n$ matrices. Thus it suffices that

$$\frac{n}{\varepsilon m} + \frac{m}{\varepsilon} + (1 - \varepsilon) rm + o_q(1) + O(1) \leqslant n ,$$

for which it suffices that

$$\frac{1}{\varepsilon m} + o_n(1) + (1 - \varepsilon) \delta m + o_n(1) \leqslant 1 ,$$

which is equivalent to

$$\frac{1}{\varepsilon m} \leqslant 1 - \alpha \delta - o_n(1) .$$

Using that $\varepsilon m = m - \alpha$, we see this is equivalent to

$$m \geqslant \alpha + \frac{1}{1 - \alpha\delta - o_n(1)} \ ,$$

which will be satisfied for large n as long as $m, \varepsilon, \alpha, \delta$ are constants and $m > \alpha + \frac{1}{1-\alpha\delta}$. $\qquad\square$

Thus, the above shows that obtaining two-source rank condensers meeting the probabilistic method bound of Proposition C.14 will yield dimension expanders essentially meeting the probabilistic method bound for dimension expanders (Proposition C.10).

In the definition of a two-source condenser, we say that it is bilinear if each coordinate of the output function is a bilinear form. It will sometimes be more convenient to consider all of the coordinates together. In this case, a bilinear condenser acts as a matrix E times the tensor product $\overline{v} \otimes \overline{w}$ as we now show.

Lemma 9.2. *Let $\mathbb{F}$ be a field and let $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$. Let $E \in \mathbb{F}^{t \times nm}$ and for $i \in [t]$ define $E_i := E_{i,\bullet} \in \mathbb{F}^{n \times m}$, the i-th row of E interpreted via the isomorphism $\mathbb{F}^{n \times m} = \mathbb{F}^{nm}$. Then f has $f(\overline{v}, \overline{w}) = (\overline{v}^{\mathrm{tr}} E_i \overline{w})_{i=1}^t$ for $E_i \in \mathbb{F}^{n \times m}$ iff $f(\overline{v}, \overline{w}) = E \cdot (\overline{v} \otimes \overline{w})$, where $\overline{v} \otimes \overline{w} \in \mathbb{F}^{nm}$ is the tensor product of $\overline{v}$ and $\overline{w}$.*

Proof:

$$\left(E \cdot (\overline{v} \otimes \overline{w})\right)_i = \sum_{j \in [nm]} E_{i,j} \cdot (\overline{v} \otimes \overline{w})_j$$

using the isomorphism between $\mathbb{F}^{nm}$ and $\mathbb{F}^{n \times m}$ so that we now index by $[n] \times [m]$,

$$= \sum_{j \in [n], k \in [m]} (E_i)_{j,k} \cdot (\overline{v} \otimes \overline{w})_{j,k}$$

factoring the tensor product,

$$= \sum_{j \in [n], k \in [m]} (E_i)_{j,k} \cdot v_j \cdot w_k$$
$$= \sum_{j \in [n], k \in [m]} v_j (E_i)_{j,k} w_k$$
$$= \overline{v}^{\mathrm{tr}} E_i \overline{w} \ . \qquad\square$$

As such, for bilinear rank condensers we can simply consider the matrix E to be the condenser and not discuss the function f.

Similar to Lemma 2.3 we can obtain that lossless two-source condensers automatically work for *all* smaller rank bounds.

Lemma 9.3. *Let $\mathbb{F}$ be a field and let $n \geqslant r \geqslant 1$ and $m \geqslant s \geqslant 1$. Let $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$. Then f is a lossless $(r, s, 0)$-two-source rank condenser iff f is a lossless $(\leqslant r, \leqslant s, 0)$-condenser.* $\qquad\square$

However, just as in Lemma 6.6 this is provably false for lossy condensers. However, while for single-source condensers we can expect to (and do) obtain "rank $\leqslant r$" results essentially for free from "natural" constructions obtaining results for "rank $= r$", this is provably not the case for two-source condensers. In particular, we now show that condensing all small enough sources can induce a linear output lower bound, showing that the linear dependence in the output given by Proposition C.14 for $(\leqslant r, s, \varepsilon)$-condensers is needed (in contrast to $(\leqslant r, s, \varepsilon)$-condensers which can do better).

Proposition 9.4. *Let $\mathbb{F}$ be a field and let $n \geqslant 1$ and $m \geqslant s \geqslant 1$. Let $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ where f is a bilinear $(1, s, \varepsilon)$-two-source rank condenser for $\varepsilon < 1$. Then $t \geqslant m - \varepsilon s$.*

Proof: The condenser is defined by a matrix $E \in \mathbb{F}^{t \times nm}$ such that for any $A \in \mathbb{F}^{n \times 1}$ of rank 1 and $B \in \mathbb{F}^{m \times s}$ of rank s we have that $\operatorname{rank} E(A \otimes B) \geqslant (1 - \varepsilon)1 \cdot s$. Thus, take $A = \{\overline{e}_1\}$ where $\overline{e}_1 \in \mathbb{F}^n$ is the first standard basis vector. Thus $E(A \otimes B) = E_1 B$, for some $E_1 \in \mathbb{F}^{t \times m}$ as once we fix $\overline{v}$, the map $\overline{w} \mapsto E(\overline{v} \otimes \overline{w})$ is linear.

Thus, we need that $E_1 \in \mathbb{F}^{t \times m}$ has that for any $B \in \mathbb{F}^{m \times s}$ that $\operatorname{rank} E_1 B \geqslant (1 - \varepsilon) \operatorname{rank} B$. If $\ker E_1$ has dimension $> \varepsilon s$ then we can find a subspace V with $\dim V = s$ and $\dim(V \cap \ker E_1) > \varepsilon s$, so $\dim V - \dim E_1 V > \varepsilon s$ and thus E is not such a condenser. Thus, it must be that $\dim \ker E_1 \leqslant \varepsilon s$. However, as $\dim \ker E_1 \geqslant m - t$ it follows that $t \geqslant m - \varepsilon s$. $\qquad\square$

9.1   Constructing Optimal Lossless Two-Source Condensers

We now turn to constructing two-source rank condensers. In this subsection, we consider the lossless case ($\varepsilon = 0$). Thus, we are given $A \in \mathbb{F}^{n \times r}$ of rank r and $B \in \mathbb{F}^{m \times s}$ and wish to obtain rank rs from them. One can trivially do this via the tensor product, so that $A \otimes B \in \mathbb{F}^{nm \times rs}$ has rank rs. However, we would like a condenser with a smaller output size. To achieve this, we observe that we can apply seeded rank condensers $\mathcal{E} : \mathbb{F}^n \to \mathbb{F}^r$ and $\mathcal{E}' : \mathbb{F}^m \to \mathbb{F}^s$ and after this dimension reduction apply the tensor product. This will certainly yield rank rs for *some* $E \in \mathcal{E}$ and $E' \in \mathcal{E}'$, so we must enumerate over all such choices. While this would naively yield an output size of $|\mathcal{E}||\mathcal{E}'|rs$, we can use that our lossless condensers ([Corollary 6.3](#)) only have a finite number of bad values so that a union bound yields an output size of $(|\mathcal{E}| + |\mathcal{E}'|)rs$, as stated in the following proposition.

Proposition 9.5. *Let $\mathbb{F}$ be a field. Let $n \geqslant r \geqslant 1$ and $m \geqslant s \geqslant 1$. Let $\omega \in \mathbb{F}$ be an element of multiplicative order $\geqslant n, m$. Identify $\mathbb{F}^n$ and $\mathbb{F}^m$ with the spaces of low-degree univariate polynomials, so that $\mathbb{F}^n = \mathbb{F}[x]^{<n}$ and $\mathbb{F}^m = \mathbb{F}[y]^{<m}$ so that $\mathbb{F}^n \otimes \mathbb{F}^m = \mathbb{F}^{nm} = \mathbb{F}[x,y]^{<n,<m}$ is the space of bivariate functions with the respective individual degree bounds. Let $S \subseteq \mathbb{F} \setminus \{0\}$ be a set where $|S| = r(n - r) + s(m - s) + 1$. Then $E : \mathbb{F}[x,y]^{<n,<m} \to \mathbb{F}^{rs \cdot |S|}$ defined by*

$$h(x,y) \mapsto (h(\omega^i \alpha, \omega^j \alpha))_{i \in [\![r]\!], j \in [\![s]\!], \alpha \in S}$$

is a bilinear $(r, s, 0)$-two-source rank condenser with output size $\leqslant rs(rn + sm)$.

Proof: By construction and [Lemma 9.2](#) we see that E is bilinear and has the desired output size (as $r, s \geqslant 1$), so that it suffices to show the condensing property.

Assume the setup of [Construction 2.7](#). In this construction, we see that $\mathrm{W}_r(\alpha)$ is a linear map $\mathrm{W}_r(\alpha) : \mathbb{F}[x]^{<n} \to \mathbb{F}^r$ given by $f(\alpha) \mapsto (f(\alpha), f(\omega\alpha), \ldots, f(\omega^{r-1}\alpha))$, and $\mathrm{W}_s(\alpha) : \mathbb{F}[y]^{<m} \to \mathbb{F}^s$ defined similarly (where we abuse the notation of [Construction 2.7](#) so that we allow $\mathrm{W}_s(\alpha)$ to act on $\mathbb{F}[y]^{<m}$). Thus, we see that $(\mathrm{W}_r(\alpha) \otimes \mathrm{W}_s(\alpha)) : \mathbb{F}[x,y]^{<n,<m} \to \mathbb{F}^{rs}$ is a linear map sending $h(x,y) \mapsto (h(\omega^i \alpha, \omega^j \alpha))_{i \in [\![r]\!], j \in [\![s]\!]}$.

Consider $A \in (\mathbb{F}[x]^{<n})^r$ of rank r and $B \in (\mathbb{F}[y]^{<m})^s$ of rank s. It suffices to show that $\operatorname{rank} E(A \otimes B) \geqslant rs$ and as E is simply the collection of the maps $(\mathrm{W}_r(\alpha) \otimes \mathrm{W}_s(\alpha))$ for $\alpha \in S$, it suffices to show that $\operatorname{rank}\big((\mathrm{W}_r(\alpha) \otimes \mathrm{W}_s(\alpha)) \cdot (A \otimes B)\big) = \operatorname{rank}((\mathrm{W}_r(\alpha)A) \otimes (\mathrm{W}_s(\alpha)B))$ has rank $\geqslant rs$ for some $\alpha \in S$.

It follows from [Proposition 6.2](#) that there are at most $r(n - r)$ values of α where $\operatorname{rank} \mathrm{W}_r(\alpha)A < \operatorname{rank} A$. Similarly, there are at most $s(m - s)$ values of α where $\operatorname{rank} \mathrm{W}_s(\alpha)B < \operatorname{rank} B$. Thus, as $|S| > r(n - r) + s(m - s)$ it follows there is some $\alpha_0 \in S$ where $\operatorname{rank} \mathrm{W}_r(\alpha_0)A = \operatorname{rank} A$ and $\operatorname{rank} \mathrm{W}_s(\alpha_0)B = \operatorname{rank} B$. Thus, it follows that $\operatorname{rank}((\mathrm{W}_r(\alpha_0)A) \otimes (\mathrm{W}_s(\alpha_0)B)) = rs$ as desired. $\qquad\square$

Note that in the balanced case of $n = m$ and $r = s$, this yields an output size of $\approx 2nr^3$ which while better than the trivial n^2 is still far from the probabilistic method bound of $2nr$ (Proposition C.13). However, we observe that the above map has a fair bit of *redundancy* which we can prune, and we can prune this redundancy because of the following lemma.

Lemma 9.6. *Let $\mathbb{F}$ be a field. Let $n, m \geqslant 1$, $1 \leqslant r \leqslant n$ and $1 \leqslant s \leqslant m$. Let $E \in \mathbb{F}^{t \times nm}$ be a bilinear (r, s, ε)-two-source rank condenser and suppose that $E' \in \mathbb{F}^{t' \times nm}$ has* row-span $E \subseteq$ row-span E'. *Then E' is a bilinear (r, s, ε)-two-source rank condenser.*

Proof: That row-span $E \subseteq$ row-span E' means that there is a $P \in \mathbb{F}^{t \times t'}$ so that $E = PE'$. Now consider some $A \in \mathbb{F}^{n \times r}$ of rank r and $B \in \mathbb{F}^{m \times s}$ of rank s. That E is such a condenser means that rank $E(A \otimes B) \geqslant (1 - \varepsilon)rs$. But then rank $E(A \otimes B) = \text{rank } PE'(A \otimes B) \leqslant \text{rank } E'(A \otimes B)$. Thus it follows that rank $E'(A \otimes B) \geqslant (1 - \varepsilon)rs$, so that E' is also such a condenser. $\qquad\square$

We now use the above lemma to prune the redundancies of Proposition 9.5 to obtain a near-optimal lossless two-source condenser.

Corollary 9.7. *Let $\mathbb{F}$ be a field. Let $n \geqslant r \geqslant 1$ and $m \geqslant s \geqslant 1$. Let $\omega \in \mathbb{F}$ be an element of multiplicative order $\geqslant n, m$. Identify $\mathbb{F}^n$ and $\mathbb{F}^m$ with the spaces of low-degree univariate polynomials, so that $\mathbb{F}^n = \mathbb{F}[x]^{<n}$ and $\mathbb{F}^m = \mathbb{F}[y]^{<m}$ so that $\mathbb{F}^n \otimes \mathbb{F}^m = \mathbb{F}^{nm} = \mathbb{F}[x, y]^{<n, <m}$ is the space of bivariate functions with the respective individual degree bounds. Let $T \subseteq \mathbb{F} \setminus \{0\}$ be a set where $|T| = n + m - 1$. Then $E' : \mathbb{F}[x, y]^{<n, <m} \to \mathbb{F}^{(r+s-1) \cdot |T|}$ defined by*

$$h'(x, y) \mapsto (h'(\beta, \omega^k \beta))_{-r < k < s, \beta \in T}$$

is a bilinear $(r, s, 0)$-two-source rank condenser with output size $(r + s - 1)(n + m - 1)$.

Proof: Note that the map of Proposition 9.5 consists of evaluating $h(x, y) \in \mathbb{F}[x, y]^{<n, <m}$ at the points $h(\omega^i \alpha, \omega^j \alpha)$ for $i \in [\![r]\!]$, $j \in [\![s]\!]$, and α in some set. Note that each such evaluation is the evaluation of the corresponding *univariate* polynomial $h(x, \omega^{j-i} x)$ at the point $\beta = \omega^i \alpha$. As the polynomial $h(x, \omega^{j-i} x)$ is of degree $\leqslant (n-1) + (m-1)$ and there are $(r-1) + (s-1) + 1$ values of $k = i - j$, we see that by polynomial interpolation that the evaluations of the map Proposition 9.5 are contained in the linear combinations of the evaluations

$$h'(x, y) \mapsto (h'(\beta, \omega^k \beta))_{-r < k < s, \beta \in T} \ ,$$

since $|T| \geqslant (n-1) + (m-1) + 1 = n + m - 1$. Appealing to Lemma 9.6 then yields the claim. $\qquad\square$

Thus, in the balanced case of $n = m$ and $r = s$ we obtain an output of size $\leqslant 4nr$, which matches the probabilistic method up to a factor of 2. While this seems to close the story on lossless two-source rank condensers, we note here that this construction is *exactly* the construction of Forbes and Shpilka [FS12] for an object called a *rank metric code*. While a priori it might seem that rank-metric codes are weaker objects (so that the above analysis would add something new), we show in Appendix B the rank metric codes are *equivalent* to lossless two-source rank condensers.

Proposition (Proposition B.4, Proposition B.5). *Let $\mathbb{F}$ be a field. Let $n, m \geqslant 1$ and $n, m \geqslant r \geqslant 0$. A matrix $E \in \mathbb{F}^{t \times nm}$ is a bilinear $(r, r, 0)$-two-source rank condenser iff $\mathcal{C} := \ker E \subseteq \mathbb{F}^{n \times m}$ is a distance $\geqslant r + 1$ rank-metric code.* $\qquad\square$

Thus, we can then leverage constructions of rank-metric codes to slightly improve upon the above condenser. In particular, Gabidulin [Gab85] rank-metric codes work over any fixed finite field, and the codes of Roth [Rot91] and Forbes-Shpilka [FS12] improve upon the above condensers by removing further redundancy. Perhaps more importantly, these codes are *optimal* for their respective regimes. In particular, we can translate the known bounds for rank-metric codes into the language of condensers.

Proposition ([Proposition B.7](#), [Proposition B.9](#)). *Let $\mathbb{F}$ be a field and $m \geqslant n \geqslant r \geqslant 1$. Let $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ be a bilinear $(r, r, 0)$-two source rank condenser. Then $t \geqslant rm$. Further, if $\mathbb{F}$ is algebraically closed and $n = m$ then $t \geqslant r(2n - r)$.* $\qquad\square$

We can then match these bounds with constructions of rank-metric codes, appealing to the above equivalence with condensers.

Proposition 9.8 ([Proposition B.8](#), [Proposition B.10](#)). *Let $\mathbb{F}$ be a field and $m \geqslant n \geqslant r \geqslant 1$.*

If $\mathbb{F} = \mathbb{F}_q$ is a finite field, and $\mathbb{F}_{q^m}$ is given explicitly as an extension of $\mathbb{F}_q$ then there is an explicit $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ which is a bilinear $(r, r, 0)$-two source rank condenser with $t = rm$.

If $|\mathbb{F}| \geqslant n$ then there is another such explicit f with $t = r(n + m - r)$. $\qquad\square$

We note that the last part of the above result can be proven by a tighter analysis of the construction used in [Corollary 9.7](#), but the cleanest exposition goes through rank-metric codes.

9.2   Constructing Lossy Condensers

We now turn to constructions of lossy two-source condensers, where our results are still far from optimal. Because of this distance to optimality, we restrict our attention to balanced sources, so that we seek (r, r, ε)-condensers $f : \mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F}^t$. We begin by discussing how the strategy in the lossless case (condense each source with a seeded condenser, tensor the results, then finally prune the output) does not seem to give any better results. Despite this, we show that this idea can serve well as a "outer condenser", so that if we had a good "inner condenser" for rank r subspaces of an r^3-dimensional space we would obtain near-optimal results.

Recall the strategy of the previous section. That is, we obtained lossless two-source condensers by using a seeded condenser based on the folded Wronskian ([Construction 2.7](#)), where we map $f(x) \mapsto (f(\alpha), f(\omega\alpha), \ldots, f(\omega^{t-1}\alpha))$ for various α. While using this approach yielded suboptimal results when applied naively ([Proposition 9.5](#)) we saw how to prune the output to yield an essentially optimal result ([Corollary 9.7](#)). The key idea to this pruning was that the induced map

$$f(x, y) \mapsto (f(\omega^i \alpha, \omega^j \alpha))_{i,j \in [\![t]\!], \alpha \in S}$$

evaluates f through evaluations of *univariate* polynomials $f(x, \omega^k x)$, and we can then restrict the number of evaluations to each univariate polynomial to be at most the degree bound of that polynomial. This helps as $k \cdot |S|$ was much larger than the degree of $f(x, \omega^k x)$.

One can thus implement the above strategy in the lossy case, where we now use the analysis of the folded Wronskian as a lossy rank condenser ([Proposition 6.8](#)). Thus, before any attempt at pruning we see that the output size is $t^2 \cdot \frac{n}{\varepsilon(t-r)}$, which is minimized at $\approx nr$ for $t = 2r$. While this improves upon the non-pruned lossless condenser of [Proposition 9.5](#) it does not even yield a smaller output size than the *lossless* constructions we give above.

Thus, to obtain better results we might try to prune this construction. However, we see that each univariate $f(x, \omega^k x)$ is evaluated only at $\frac{n}{\varepsilon(t-r)}$ points, which is *below* the degree bound of $\deg f(x, \omega^k x) = 2n$. Thus, there seems to be no pruning available to improve the above strategy.

Despite this, we observe that the above strategy has now reduced the problem to a smaller dimensional space, for which one could apply different methods. We now show that taking $t = r^3$ in this reduction can lead to near-optimal results.

Proposition 9.9. *Let $\mathbb{F}$ be a field and $n \geqslant r \geqslant 1$. Assume the setup of [Construction 2.7](). Suppose that $E \in \mathbb{F}^{t' \times t^2}$ is a bilinear $(\lceil (1-\varepsilon)r \rceil, \lceil (1-\varepsilon)r \rceil, \varepsilon)$-two source rank condenser. Let $S \subseteq \{(\omega^t)^j \mid j \geqslant 0\}$ with $|S| = \left\lceil \frac{2n}{\varepsilon(t-r+1)} \right\rceil$. Then $E' := (E \cdot (\mathrm{W}_t(\alpha) \otimes \mathrm{W}_t(\alpha)))_{\alpha \in S} \in \mathbb{F}^{t' \cdot |S| \times n^2}$ is $(r, r, 1-(1-\varepsilon)^3)$-condenser.*

In particular, taking $t = r^3$ and taking E to be a condenser meeting the bound of [Proposition C.13]() so that $t' \leqslant \frac{2r^3}{\varepsilon r} + (1-\varepsilon)r^2 + O(1)$, we obtain that E' is a $(r, r, 1-(1-\varepsilon)^3)$-condenser in $\mathbb{F}^{t'' \times n^2}$ for $t'' \leqslant O(n/\varepsilon^2 r)$.

Proof: Given $A, B \in \mathbb{F}^{n \times r}$ of rank r, we see that by [Proposition 6.8]() there are $< 2\frac{n}{\varepsilon(t-r+1)}$ values of $\alpha \in S$ such that $\mathrm{rank}\, \mathrm{W}_t(\alpha)A < (1-\varepsilon)\mathrm{rank}\, A$ or $\mathrm{rank}\, \mathrm{W}_t(\alpha)B < (1-\varepsilon)\mathrm{rank}\, B$. Thus, as $|S| \geqslant \frac{2n}{\varepsilon(t-r+1)}$ there is some α_0 where $\mathrm{rank}\, \mathrm{W}_t(\alpha_0)A, \mathrm{rank}\, \mathrm{W}_t(\alpha_0)B \geqslant (1-\varepsilon)r$. Thus, it follows that as E is a $(\lceil (1-\varepsilon)r \rceil, \lceil (1-\varepsilon)r \rceil, \varepsilon)$-condenser that $E((\mathrm{W}_t(\alpha_0)A) \otimes (\mathrm{W}_t(\alpha_0)B))$ has rank $\geqslant (1-\varepsilon) \cdot (1-\varepsilon)^2 r^2$. Thus, it follows that $(E \cdot (\mathrm{W}_t(\alpha) \otimes \mathrm{W}_t(\alpha)))_{\alpha \in S}$ applied to $A \otimes B$ has rank $\geqslant (1-\varepsilon)^3 r^2$, showing that this is indeed the desired condenser.

To obtain the bound on t', note that

$$
\begin{aligned}
t'' &= |S| \cdot t' \\
&\leqslant \left\lceil \frac{2n}{\varepsilon(r^3 - r + 1)} \right\rceil \cdot \left(\frac{2r^3}{\varepsilon r} + (1-\varepsilon)r^2 + O(1) \right) \\
&= O\left(\frac{n}{\varepsilon^2 r} \right) .
\end{aligned}
$$
$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad$ $\square$

In particular, one could find such a condenser E via brute force when $r = O(1)$.

10   Open Questions

This work leaves several directions for future work.

1. Can one obtain (r, ε)-lossy seeded rank *extractors*, where the output is $\approx (1-\varepsilon)r$? Our methods require the output to be $\geqslant r$.

2. Can one develop of theory of "code concatenation" to improve our results in [Section 8]() 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

References

[ASS13]   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.

[BISW04]   Boaz Barak, Russell Impagliazzo, Amir Shpilka, and Avi Wigderson. Personal Commutation to Dvir-Shpilka [DS11], 2004.

[Bou09]   Jean Bourgain. Expanders and dimensional expansion. *Comptes Rendus Mathematique*, 347(7-8):357–362, 2009.

[BS12]   Avraham Ben-Aroya and Igor Shinkar. A note on subspace evasive sets. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:95, 2012.

[BY13]   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 [Bou09].

[CKL13]   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)*.

[CRVW02]   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.

[Del78]   Philippe Delsarte. Bilinear forms over a finite field, with applications to coding theory. *J. Combin. Theory Ser. A*, 25(3):226–241, 1978.

[DGW09]   Zeev Dvir, Ariel Gabizon, and Avi Wigderson. Extractors and rank extractors for polynomial sources. *Computational Complexity*, 18(1):1–58, 2009. Preliminary version in the *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007)*.

[DL12]   Zeev Dvir and Shachar Lovett. Subspace evasive sets. In *Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC 2012)*, pages 351–358, 2012. Full version at arXiv:1110.5696.

[DS07]   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)*.

[DS11]   Zeev Dvir and Amir Shpilka. Towards dimension expanders over finite fields. *Combinatorica*, 31(3):305–320, 2011. Preliminary version in the *23rd Annual IEEE Conference on Computational Complexity (CCC 2008)*.

[DW10]   Zeev Dvir and Avi Wigderson. Monotone expanders: Constructions and applications. *Theory of Computing*, 6(12):291–308, 2010.

[For14]   Michael Forbes. *Polynomial Identity Testing of Read-Once Oblivious Algebraic Branching Programs*. PhD thesis, Massachusetts Institute of Technology, 2014.

[FS12]   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.

[FSS14]   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.

[Gab85]   Ernst M. Gabidulin. Theory of codes with maximum rank distance. *Probl. Inform. Transm.*, 21(1):1–12, 1985.

[GK13]   Venkatesan Guruswami and Swastik Kopparty. Explicit subspace designs. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2013)*, pages 608–617, 2013. Journal version to appear in *Combinatorica*.

[GR08a]   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)*.

[GR08b]   Venkatesan Guruswami and Atri Rudra. Explicit codes achieving list decoding capacity: Error-correction with optimal redundancy. *IEEE Transactions on Information Theory*, 54(1):135–150, 2008. Preliminary version in the *38th Annual ACM Symposium on Theory of Computing (STOC 2006)*.

[Gur11]   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 [GW13].

[GW13]   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).

[GW14]   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.

[GX12]   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.

[GX13]   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.

[Har08]   Aram W. Harrow. Quantum expanders from any classical cayley graph expander. *Quantum Information & Computation*, 8(8–9):715–721, 2008.

[KS11]   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)*.

[LMPS14]   Daniel Lokshtanov, Pranabendu Misra, Fahad Panolan, and Saket Saurabh. Deterministic truncation of linear matroids. *arXiv*, 1404.4506, 2014.

[LZ08]   Alexander Lubotzky and Efim Zelmanov. Dimension expanders. *J. Algebra*, 319(2):730–738, 2008.

[Mar09]   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)*.

[PR04]   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.

[Raz05]   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.

[Rot91]   Ron M. Roth. Maximum-rank array codes and their application to crisscross error correction. *IEEE Transactions on Information Theory*, 37(2):328–336, 1991.

[RS05]   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)*.

[RVW02]   Omer Reingold, Salil Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Annals of Mathematics*, 155(1):157–187, January 2002. Preliminary version in the *41st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2000)*.

[SY10]   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.

[Vad12]   Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012.

[Wig04]   Avi Wigderson. Expanders: Old and new applications and problems. Lecture at the Institute for Pure and Applied Mathematics (IPAM), February 2004.

A   Toward Iterative Constructions of Subspace Evasive Sets

We describe here a motivation for studying two-source rank condensers via a potential application to constructing another object called a *subspace evasive set*, as defined by Guruswami [Gur11].

Definition A.1 (Guruswami [Gur11]). *A set $S \subseteq \mathbb{F}^n$ is a (r,L)-subspace evasive set if for every $(\leqslant r)$-dimensional subspace $V \subseteq \mathbb{F}^n$, $|V \cap S| \leqslant L$. Equivalently, a set S is (r,L)-subspace evasive if for every subset $T \subseteq S$ with $|T| = L+1$, $\operatorname{rank} T \geqslant r+1$.* ◊

This notion was introduced as a method to prune the lists in list-decodable codes when using a linear algebraic approach that pins down candidate messages to a low-dimensional subspace [Vad12,Gur11,GW13]. To obtain a high rate code after this pruning we desire a *large* subspace evasive set. Such large sets are guaranteed by the probabilistic method, as given in the following lemma.

Lemma A.2 (Guruswami [Gur11]). *Let $\mathbb{F}_q$ be a finite field. Let $\varepsilon > 0$ and $n \geqslant r \geqslant 1$. There exists a $(r, 2r/\varepsilon)$-subspace evasive set in $\mathbb{F}_q^n$ of size $\lfloor q^{(1-\varepsilon)n} \rfloor$ when $r \leqslant \varepsilon n/2$.* □

The work of Ben-Aroya and Shinkar [BS12] showed that for $r \geqslant (1/\varepsilon)^{\Omega(1)}$ the above size-bound of "$2r/\varepsilon$" is asymptotically optimal up to constants.

Note that the most basic construction of a subspace evasive set is based on taking the set of moment curve vectors $S := \{(1, \alpha, \ldots, \alpha^{n-1}) \mid \alpha \in \mathbb{F}_q\}$. This set S is highly evasive (as it is $(n, n-1)$-subspace-evasive), but has size q, which is quite far from the above $\lfloor q^{(1-\varepsilon)n} \rfloor$. Thus far, the best explicit constructions are due to Dvir-Lovett [DL12] and Ben-Aroya and Shinkar [BS12].

Theorem (Dvir and Lovett [DL12]). *Let $\mathbb{F}_q$ be a finite field. Let $\varepsilon > 0$ and $n \geqslant r \geqslant 1$. Then there exists an explicit $(r, (r/\varepsilon)^r)$-subspace evasive set in $\mathbb{F}_q^n$ of size $> q^{(1-\varepsilon)n}$.* □

Theorem (Ben-Aroya and Shinkar [BS12]). *Let $\mathbb{F}_q$ be a finite field. Let $0 < \varepsilon < 1/2$ and $n \geqslant r \geqslant 1$ where $r, 1/\varepsilon \leqslant O(1)$. Then there exists an explicit $(r, (2/\varepsilon)^r)$-subspace evasive set in $\mathbb{F}_q^n$ of size $\geqslant q^{(1-\varepsilon)n}$.* □

Note that while the latter result only works for constant r and ε, this is still an interesting parameter regime for list-decoding applications.

Given that we are far from optimal explicit constructions for subspace evasive sets, we consider in this work an iterative approach to constructing subspace evasive sets along the lines of the zig-zag product of Reingold, Vadhan and Wigderson [RVW02] for constructing expanders. That is, suppose that $S \subseteq \mathbb{F}_q^n$ is (r, L)-subspace evasive. It follows that $S \otimes S = \{\overline{v} \otimes \overline{w} \mid \overline{v}, \overline{w} \in S\}$ is in a certain sense subspace-evasive. That is, for all $T, R \subseteq S$ with $|T|, |R| = L + 1$, the product set $T \otimes R \subseteq S \otimes S$ has dimension $\geqslant (r+1)^2$. Thus, large enough "product sets" among $S \otimes S$ have high rank. While having this property for all product sets is far from having it for all sets of similar size, and thus this is far from a iterative approach for boosting evasiveness, there is also the more basic problem that the *rate* of this construction is poor. That is, ideally we have that the rate $\frac{\log_q |S|}{n}$ is at least $\Omega(1)$. However, $\frac{\log_q |S \otimes S|}{n^2} = \frac{2}{n} \frac{\log_q |S|}{n}$ and thus this tensoring operation decreases the rate dramatically.

While the "product versus non-product set" problem seems fundamental to the above approach, we try to ameliorate the rate issue by "derandomizing" the tensor product using a two-source rank condenser. That is, if S is (r, L)-subspace evasive and $f : \mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F}^t$ is a $(r+1, r+1, 1/2)$-condenser, then we can define $S \otimes_f S := \{f(\overline{v}, \overline{w}) \mid \overline{v}, \overline{w} \in S\}$ and observe that for any subset $R, T \subseteq S$ with $|R|, |T| \geqslant L + 1$ we have that $R \otimes_f T \subseteq S \otimes_f S$ has rank $\geqslant (r+1)^2/2$. While this evades slightly smaller subspaces than using the standard tensor product, this operation still qualitatively squares the evasiveness as long as $r \geqslant \Omega(1)$ initially. Further, the rate has not dropped substantially if good enough condensers are used. That is, the probabilistic method (Proposition C.13) shows that we can take $t = \Theta(n/r + r^2)$. Thus, as long as $r \leqslant O(\sqrt[3]{n})$, we get that $t = \Theta(n/r)$ and that $\frac{\log_q |S \otimes_f S|}{t} = \Theta(2r \frac{\log_q |S|}{n})$. Thus, as long as $r \geqslant \Omega(1)$ the rate has actually *increased* under this derandomized tensor product. Thus, from the perspective of iteratively constructing subspace evasive sets such condensers allow one to focus on the "product versus non-product set" problem.

Unfortunately in this work we do not construct such condensers with $t \approx n/r$, nor are we even able to obtain $t = O(n)$ for any $r \geqslant \omega(1)$.

B Lossless Two-source Condensers versus Rank-Metric Codes

In this section we define the notion of a *rank metric code* and show that it is equivalent to our notion of a two-source condenser Definition 4.1. We then review constructions and limitations of rank-metric codes. Finally we derive the corresponding results in the language of two-source condensers.

B.1 Equivalence of Condensers and Rank-Metric Codes

We begin with the definition of rank metric codes.

Definition B.1. *Let $\mathbb{F}$ be a field and $n, m \geqslant r \geqslant 1$. A set $\mathcal{C} \subseteq \mathbb{F}^{n \times m}$ is a **rank metric code with distance** r if for all $A, B \in \mathcal{C}$ with $A \neq B$, $\operatorname{rank}(A - B) \geqslant r$. The code is **linear** if $\mathcal{C}$ is a linear space. The **parity checks** of a linear code $\mathcal{C}$ consist of a basis for the dual space $\mathcal{C}^{\perp}$. A linear rank metric code $\mathcal{C}$ is **explicit** if one can construct a basis of $\mathcal{C}$ in $\operatorname{poly}(n, m)$ operations in $\mathbb{F}$.* $\diamondsuit$

When the linear code $\mathcal{C}$ is simply the zero subspace we define the distance to be $\min\{n, m\} + 1$. We now observe that via standard coding theory arguments linear codes have their distance equal to the minimum weight of a non-zero element in the code (where "weight" here means "rank").

Lemma B.2. *Let $\mathbb{F}$ be a field and $n, m \geqslant 1$. Let $\mathcal{C} \subseteq \mathbb{F}^{n \times m}$ be a linear rank-metric code. Then the distance of $\mathcal{C}$ is $\min_{0 \neq M \in \mathcal{C}} \operatorname{rank} M$.* $\square$

We now give a lemma relating the inner product of matrices to the trace of their product, which will be helpful in the below results.

Lemma B.3. *Let $\mathbb{F}$ be a field and $n, m \geqslant 1$. Let $M, N \in \mathbb{F}^{n \times m}$. Then the inner product of M and N as vectors in $\mathbb{F}^{nm}$ can be expressed as $\langle M, N \rangle = \operatorname{tr}(M^{\mathrm{tr}} N) = \operatorname{tr}(N^{\mathrm{tr}} M)$.*

Proof: It suffices to prove the first part, as $\langle M, N \rangle = \langle N, M \rangle$.

$$\langle M, N \rangle = \sum_{i \in [n], j \in [m]} M_{i,j} N_{i,j} = \sum_{i \in [n], j \in [m]} (M^{\mathrm{tr}})_{j,i} N_{i,j} = \sum_{j \in [m]} ((M^{\mathrm{tr}}) \cdot N)_{j,j} = \operatorname{tr}(M^{\mathrm{tr}} N) .$$ $\square$

We now show that a matrix $E \in \mathbb{F}^{t \times nm}$ defining a bilinear lossless two-source condenser has a kernel which is a good rank-metric code. In particular, if E had a low-rank matrix $M \in \mathbb{F}^{n \times m}$ in its kernel, than E would not losslessly preserve the row-span and column-span of M.

Proposition B.4. *Let $\mathbb{F}$ be a field and let $n, m \geqslant 1$ and $n, m \geqslant r \geqslant 0$. Let $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ be a bilinear lossless $(r, r, 0)$-two-source rank condenser defined by $f(\overline{v}, \overline{w}) = (\overline{v}^{\mathrm{tr}} E_k \overline{w})_{k \in [t]}$ with $E_k \in \mathbb{F}^{n \times m}$. Then $\{E_k\}_k$ are the parity checks of a linear rank-metric code with distance $\geqslant r + 1$.*

Proof: We show the contrapositive for $r > 0$, as when $r = 0$ the claim is trivial as every rank-metric code has distance $\geqslant 1$. That $\{E_k\}_k$ are not such parity checks means that there is a non-zero matrix $M \in \mathbb{F}^{n \times m}$ with rank $s \leqslant r$ such that $(\langle E_k, M \rangle)_{k \in [t]} = \overline{0}$. Let $M = AB^{\mathrm{tr}}$ with $A \in \mathbb{F}^{n \times s}$ and $B \in \mathbb{F}^{m \times s}$.

As $(f(\overline{v}, \overline{w}))_k = \overline{v}^{\mathrm{tr}} E_k \overline{w}$, it follows that we can express the k-th output of $f(A, B)$ as $(f(A, B))_k = A^{\mathrm{tr}} E_k B \in \mathbb{F}^{s \times s}$. Now consider the inner product of these vectors with the $s \times s$ identity matrix $\mathrm{I}_s \in \mathbb{F}^{s \times s}$, and appealing to Lemma B.3 we have that $\langle \mathrm{I}_s, A^{\mathrm{tr}} E_k B \rangle = \operatorname{tr}(\mathrm{I}_s A^{\mathrm{tr}} E_k B) = \operatorname{tr}(B A^{\mathrm{tr}} E_k) = \operatorname{tr}(M^{\mathrm{tr}} E_k) = \langle E_k, M \rangle$, using the fact that the trace is invariant under cyclic permutations.

As $(\langle E_k, M \rangle)_{k \in [t]} = \overline{0}$ it follows that $\langle \mathrm{I}_s, A^{\mathrm{tr}} E_k B \rangle = 0$ for all k. That is, the $t \times s^2$ matrix which is the output $f(A, B)$ has a non-trivial nullspace as it contains the matrix I_s. Thus, $\operatorname{rank} f(A, B) < s^2 = \operatorname{rank} A \cdot \operatorname{rank} B$, so that f is not a lossless rank condenser. $\square$

Somewhat surprisingly, we also show the converse to the above, showing that the parity checks of any rank-metric code yield a bilinear lossless two-source rank condenser. This follows from the observation that if a matrix $E \in \mathbb{F}^{t \times nm}$ fails to condense the pair of subspaces $A \in \mathbb{F}^{n \times r}$ and $B \in \mathbb{F}^{m \times s}$ then $E(A \otimes B)$ must have rank $< rs$ so that it has a linear dependence $E(A \otimes B)C = 0$ where $C \in \mathbb{F}^{rs \times 1}$. However, using that $\mathbb{F}^{rs \times 1} = \mathbb{F}^{r \times s}$, we can see that $(A \otimes B)C$ can be interpreted as the matrix $ACB^{\mathrm{tr}} \in \mathbb{F}^{n \times m}$, which is of rank $\leqslant r, s$. Thus, E has a low-rank matrix in its kernel, and so the kernel does not define a good rank-metric code.

Proposition B.5. *Let $\mathbb{F}$ be a field and let $n, m \geqslant 1$ and $n, m \geqslant d \geqslant 0$. Let $E_1, \ldots, E_t \in \mathbb{F}^{n \times m}$ be the parity check matrix of a linear rank-metric code with distance $d + 1$. Define $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ by $(f(\overline{v}, \overline{w}))_k := \overline{v}^{\mathrm{tr}} E_k \overline{w}$. Then f is a bilinear lossless $(d, m, 0)$-two-source rank condenser, as well as a bilinear lossless $(n, d, 0)$-two-source rank condenser.*

Proof: We consider the contrapositive. Suppose that $A \in \mathbb{F}^{n \times r}$ with $\mathrm{rank}\, A = r$ and $B \in \mathbb{F}^{m \times s}$ with $\mathrm{rank}\, B = s$ is not losslessly condensed, so that $\mathrm{rank}\, f(A, B) < rs$. The output of $f(A, B)$ is a set of rs vectors, each in $\mathbb{F}^t$, which we can consider as a $t \times rs$ matrix. Viewing this matrix row-wise we can express the k-th output of $f(A, B)$ as $(f(A, B))_k = A^{\mathrm{tr}} E_k B$. Thus, that $\mathrm{rank}\, f(A, B) < rs$ means that there is some non-zero matrix $C \in \mathbb{F}^{r \times s}$ such that $\langle C, (f(A, B))_k \rangle = 0$ for all k. Thus, (appealing to Lemma B.3)

$$0 = \langle C, (f(A, B))_k \rangle = \mathrm{tr}(C^{\mathrm{tr}} A^{\mathrm{tr}} E_k B) = \mathrm{tr}(BC^{\mathrm{tr}} A^{\mathrm{tr}} E_k) = \langle E_k, ACB^{\mathrm{tr}} \rangle \,,$$

for all $k \in [t]$.

Note that we can decompose C into the basis $\{\overline{e}_i \overline{e}_j^{\mathrm{tr}}\}_{i \in [r], j \in [s]}$ for $\mathbb{F}^{r \times s}$ where $\overline{e}_i$ is the i-th standard basis vector. Thus,

$$ACB^{\mathrm{tr}} = A \left(\sum_{i \in [r], j \in [s]} C_{i,j} \overline{e}_i \overline{e}_j^{\mathrm{tr}} \right) B^{\mathrm{tr}} = \sum_{i \in [r], j \in [s]} C_{i,j} \cdot A(\overline{e}_i \overline{e}_j^{\mathrm{tr}}) B^{\mathrm{tr}} = \sum_{i \in [r], j \in [s]} C_{i,j} \cdot A_{\bullet,i} (B_{\bullet,j})^{\mathrm{tr}} \,,$$

where $A_{\bullet,i}$ is the i-th column of A and likewise for B. Now note that the outer-products $\{A_{\bullet,i} (B_{\bullet,j})^{\mathrm{tr}}\}_{i \in [r], j \in [s]} \subseteq \mathbb{F}^{n \times m}$ are linearly independent as this is simply saying that the tensor product multiplies rank. Thus, as $C \neq 0$, it follows that ACB^{tr} is a non-trivial linear combination of the $\{A_{\bullet,i} (B_{\bullet,j})^{\mathrm{tr}}\}_{i,j}$ and thus $ACB^{\mathrm{tr}} \neq 0$.

However, now note that $ACB^{\mathrm{tr}} \in \mathbb{F}^{n \times m}$ has $\mathrm{rank}\, ACB^{\mathrm{tr}} \leqslant \min\{\mathrm{rank}\, A, \mathrm{rank}\, B\}$ so that $\mathrm{rank}\, ACB^{\mathrm{tr}} \leqslant \min\{r, s\}$. Thus, if $\min\{r, s\} \leqslant d$ then this shows that the $\{E_k\}_k$ are not the parity checks of a distance $d + 1$ rank metric as ACB^{tr} is a non-zero matrix with rank $< d + 1$ where $(\langle E_k, ACB^{\mathrm{tr}} \rangle)_{k \in [t]} = \overline{0}$. Thus, the claim follows by taking first $r \leqslant d$ and $s \leqslant m$, then taking $r \leqslant n$ and $s \leqslant d$. $\qquad\square$

It thus follows that focusing on balanced sources in lossless two-source condensers is sufficient, using that lossless condensers also condense smaller sources (Lemma 9.3).

Corollary B.6. *Let $\mathbb{F}$ be a field and let $n, m \geqslant r \geqslant 1$. Let $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ be a bilinear lossless $(r, r, 0)$-two-source rank condenser. Then f is also a $(\leqslant r, \leqslant m, 0)$- and $(\leqslant n, \leqslant r, 0)$-condenser.* $\qquad\square$

B.2 Constructions of Rank-Metric Codes

We now review for completeness known constructions of rank-metric codes, which by the above equivalence yields constructions of bilinear lossless rank condensers. In this section, we will assume that $m \geqslant n$, so that we work with short and fat matrices. We begin with a basic limitation of rank-metric codes which is a direct generalization of the Singleton bound for codes in the Hamming metric.

Proposition B.7 (Gabidulin [Gab85], Delsarte [Del78], and Roth [Rot91])**.** *Let $\mathbb{F}$ be a field and $m \geqslant n \geqslant 1$ and $n \geqslant r \geqslant 0$. Let $\mathcal{C} \subseteq \mathbb{F}^{n \times m}$ be a rank-metric code with distance $\geqslant r+1$. If $\mathcal{C}$ is a linear code then $\dim \mathcal{C} \leqslant m(n-r)$. If $\mathbb{F}$ is finite (and $\mathcal{C}$ possibly non-linear) then $\log_{|\mathbb{F}|} |\mathcal{C}| \leqslant m(n-r)$.* $\qquad\square$

We now give a well-known construction of rank-metric codes, now called *Gabidulin codes*, which are a q-linearized version of Reed-Solomon codes. Correspondingly, these codes meet the above analogue of the Singleton bound.

Proposition B.8 (Gabidulin [Gab85], Delsarte [Del78], and Roth [Rot91])**.** *Let $\mathbb{F}_q$ be a finite field and $m \geqslant n \geqslant 1$ and $n \geqslant r \geqslant 0$. Given a presentation of $\mathbb{F}_{q^m}$ as a extension field $\mathbb{F}_q$, there is an explicit linear rank-metric code $\mathcal{C} \subseteq \mathbb{F}^{n \times m}$ with distance $r+1$ with $\dim \mathcal{C} = m(n-r)$.*

Proof: <u>$n = m$</u>: First, we call a polynomial $f(x) \in \mathbb{F}_{q^m}[x]$ to be *q-linearized* if it can written as $f(x) = \sum_{i=0}^{k} \beta_i x^{q^i}$. Now note that the evaluation map $f : \mathbb{F}_{q^m} \to \mathbb{F}_{q^m}$ is $\mathbb{F}_q$-linear by the Frobenius endomorphism, thus we can consider f as a linear map $M_f : \mathbb{F}_q^m \to \mathbb{F}_q^m$. Further, we can treat M_f as a matrix $M_f \in \mathbb{F}_q^{m \times m}$ by choosing some $\mathbb{F}_q$-basis $\alpha_1, \ldots, \alpha_m$ for $\mathbb{F}_{q^m}$ so that $\sum_i \gamma_i \alpha_i \to \sum_i \gamma_i f(\alpha_i)$ for $\gamma_i \in \mathbb{F}_q$.

Now note that the roots of a q-linearized polynomial, by the above linearity, form a $\mathbb{F}_q$-linear space within $\mathbb{F}_{q^m}$. In particular, if f is q-linearized and has $\deg f \leqslant q^k$ then it has $\leqslant q^k$ roots, so the map M_f has a kernel of $\mathbb{F}_q$-dimension $\leqslant k$, so that $\operatorname{rank} M_f \geqslant m - k$.

Finally, define $\mathcal{C}_k := \{M_f \mid f \in \mathbb{F}_{q^m}[x], f = \sum_{i=0}^{k} \beta_i x^{q^i}\} \subseteq \mathbb{F}_q^{m \times m}$ for $k < m$. Noting that the degree $\leqslant q^k$ q-linearized polynomials form a $\mathbb{F}_q$-vector-space, it follows that $\mathcal{C}_k$ is a linear space of dimension $\leqslant (k+1)m$. Further, $\mathcal{C}_k$ is a linear space of dimension $\geqslant (k+1)m$ as the map from polynomials to matrices is injective. That is, a degree $\leqslant q^k$ q-linearized polynomial that yields a zero matrix must be zero as the zero matrix has a kernel of size q^n but a non-zero such polynomial has at most q^k roots. Thus, $\mathcal{C}_k$ is a rank metric code of distance $m - k$ and dimension $(k+1)m$. Taking $k = m - (r+1) < m$ yields the claim (for $r = m$ we get $k = -1$ so that $\mathcal{C}_{-1}$ is just the zero polynomial).

<u>$n < m$</u>: Consider the above code $\mathcal{C}_{m-(r+1)} \subseteq \mathbb{F}^{m \times m}$ with distance $r+1$. Now consider the subcode $\mathcal{C}'$ of the above code where we insist that the last $m - n$ rows are all zero. It follows that $\mathcal{C}'$ is still a linear rank-metric code with distance $r+1$, and we can now embed it into $\mathbb{F}^{n \times m}$. The dimension is at least $\dim \mathcal{C}' \geqslant \dim \mathcal{C}_{m-(r+1)} - m(m-n) = m(m-r) - m(m-n) = m(n-r)$. By the above Singleton-type bound ([Proposition B.7](#)) it follows that this lower bound is met with equality.

<u>explicitness:</u> This is clear from construction as the presentation of $\mathbb{F}_{q^m}$ yields a $\mathbb{F}_q$-basis for $\mathbb{F}_{q^m}$ and a way to compute in $\mathbb{F}_{q^m}$. $\qquad\square$

We remark that these codes are also efficiently decodable, as shown by Gabidulin [Gab85], Delsarte [Del78], and Roth [Rot91].

While the above codes are thus optimal, they are only defined for *finite* fields, and in particular have no analogue over the complex numbers. In fact, there is provably no analogue as in algebraically closed fields there is the following upper bound on the dimension of rank metric codes that is lower than the Singleton-type bound.

Proposition B.9 (Roth [Rot91])**.** *Let $\mathbb{F}$ be an algebraically closed field and $n \geqslant 1$ and $n \geqslant r \geqslant 0$. Let $\mathcal{C} \subseteq \mathbb{F}^{n \times n}$ be a rank-metric code with distance $\geqslant r+1$. If $\mathcal{C}$ is a linear code then $\dim \mathcal{C} \leqslant (n-r)^2$.* $\qquad\square$

Thus, while Gabidulin codes are no longer applicable in this regime, a different construction was given by Roth [Rot91] that works over any large field and meets this bound.

Proposition B.10 (Roth [Rot91])**.** *Let $\mathbb{F}$ be a field and $m \geqslant n \geqslant 1$, $n \geqslant r \geqslant 0$ and $|\mathbb{F}| \geqslant n$. Then there is an explicit linear rank-metric code $\mathcal{R} \subseteq \mathbb{F}^{n \times m}$ with distance $r+1$ and $\dim \mathcal{R} = (n-r)(m-r)$.*

Proof: For a matrix $M \in \mathbb{F}^{[\![n]\!] \times [\![m]\!]}$ define the *k-diagonal* $M^{(k)}$ *of* M to be those entries $(M_{i,j})_{i+j=k}$ where we do the addition in the integers, recalling that this notation means we index M from 0. Thus, for $0 \leqslant k < n$ there are $k+1$ entries in the k-diagonal of M, for $n \leqslant k < m$ the diagonal $M^{(k)}$ has n entries, and for $m \leqslant k < n+m-1$ the diagonal $M^{(k)}$ has $(n+m-1) - k$ entries.

Now recall that as $|\mathbb{F}| \geqslant n$ there is an explicit[5] linear code $\mathcal{C}_\ell := [\ell, k, r+1]_{\mathbb{F}}$ with $k = \ell - (r+1) + 1$ for any $r \leqslant \ell \leqslant n$. For $1 \leqslant \ell \leqslant r$ define $\mathcal{C}_\ell$ be the single codeword $\mathcal{C}_\ell = \{0\}$ so that $\mathcal{C}_\ell$ has distance $\geqslant r+1$ (trivially).

Now take $\mathcal{R} \subseteq \mathbb{F}^{n \times m}$ to be

$$\mathcal{R} := \{M \in \mathbb{F}^{n \times m} \mid \forall 0 \leqslant k < n+m-1, M^{(k)} \in \mathcal{C}_{|M^{(k)}|}\} \,.$$

Thus, it follows that for any matrix $M \in \mathcal{R}$ that all non-zero diagonals $M^{(k)}$ have at least $r+1$ non-zero entries.

Note that $\dim \mathcal{C}_\ell = \max\{0, \ell - r\}$ for all $1 \leqslant \ell \leqslant n$. Thus, let us now count the dimension of $\mathcal{R}$ as measured from the full-dimension.

$$
\begin{aligned}
&nm - \dim \mathcal{R} \\
&= \sum_{0 \leqslant k < n+m-1} |M^{(k)}| - \dim \mathcal{C}_{|M^{(k)}|} \\
&= \sum_{0 \leqslant k < n+m-1} |M^{(k)}| - \max\{0, |M^{(k)}| - r\} \\
&= \sum_{0 \leqslant k < n+m-1} \min\{|M^{(k)}|, r\} \\
&= \sum_{0 \leqslant k < r} \min\{|M^{(k)}|, r\} + \sum_{r \leqslant k < (n+m-1)-r} \min\{|M^{(k)}|, r\} + \sum_{(n+m-1)-r \leqslant k < n+m-1} \min\{|M^{(k)}|, r\} \\
&= \sum_{0 \leqslant k < r} (k+1) + \sum_{r \leqslant k < (n+m-1)-r} r + \sum_{(n+m-1)-r \leqslant k < n+m-1} (n+m-1) - k \\
&= \binom{r+1}{2} + r(n+m-1-2r) + \binom{r+1}{2} \\
&= r(n+m-r) \,.
\end{aligned}
$$

Thus, $\dim \mathcal{R} = nm - r(n+m-r) = (n-r)(m-r)$ as desired.

Now consider a matrix M with rank $\leqslant r$. Consider the first k where the diagonal $M^{(k)}$ is non-zero. Note that the rows with non-zero entries in this diagonal $M^{(k)}$ must be linearly independent as they form a triangular system. Thus, $M^{(k)}$ is r-sparse, from which it follows that $M \notin \mathcal{R}$. Thus $\mathcal{R}$ is a distance $\geqslant r+1$ rank-metric code. That $\mathcal{R}$ is explicit is clear from construction. $\square$

While Roth [Rot91] did not provide an algorithm to decode the above codes, such an algorithm was provided Forbes and Shpilka [FS12]. Perhaps more relevant for this paper is that Forbes and Shpilka [FS12] realized that when the maximum distance separable (MDS) codes $\mathcal{C}_\ell$ used in the above proof are instantiated as dual Reed-Solomon codes, one can show that the above codes actually correspond to evaluation-based codes of bivariate polynomials.

---

[5]That is, taking $d = r+1$, if $\alpha_1, \ldots, \alpha_\ell \in \mathbb{F}$ are distinct then $H \in \mathbb{F}^{[\![d]\!] \times \ell}$ with $H_{i,j} := \alpha_j^i$ is a Vandermonde matrix with every $k \times k$ minor of full-rank. Thus, $\ker H \subseteq \mathbb{F}^\ell$ has no vectors of sparsity $\leqslant d-1$ so that $\text{row-span}(H)^\perp$ defines the desired code.

Proposition B.11 (Forbes and Shpilka [FS12])**.** *Let $\mathbb{F}$ be a field and $m \geqslant n \geqslant 1$ and $n \geqslant r \geqslant 0$. Let $\omega \in \mathbb{F}$ be an element with multiplicative order $\geqslant n$. Identifying matrices $\mathbb{F}^{n \times m}$ with low-degree bivariate polynomials $\mathbb{F}[x,y]^{<n,<m}$, define*

$$\mathcal{C} := \{f(x,y) \mid f \in \mathbb{F}[x,y]^{<n,<m}, \forall i \in [\![r]\!], f(x, \omega^i x) = 0\} \ .$$

Then $\mathcal{C}$ is an explicit linear rank-metric code with distance $r+1$ and dimension $(n-r)(m-r)$. $\qquad\square$

While one can prove the above by showing that they are an instantiation of Roth's [Rot91] code, Forbes and Shpilka [FS12] also gave a proof that we presented as the construction of a lossless two-source condenser in Corollary 9.7 (and actually that construction can be further pruned to match the above parameters).

C   Existential Arguments

In this section, we give arguments via the probabilistic method to show that the relevant pseudorandom objects exist. These arguments are completely standard, but we include them to demonstrate the quantitative gaps between our explicit constructions and the existential constructions.

C.1   Probabilistic Tools

We begin by establishing some lemmas that detail probabilistic estimates for when a matrix is low-rank. Some lemmas are standard (for which proofs can be found for example in the thesis of Forbes [For14]) and some of which we lack a good reference and thus include for completeness. The first lemma shows how random linear maps affect full rank matrices.

Lemma C.1. *Let $\mathbb{F}$ be a finite field and $M \in \mathbb{F}^{n \times r}$ be of rank r. Let E be a random variable uniformly distributed over matrices in $\mathbb{F}^{m \times n}$. Then $\mathsf{E} \cdot M$ is uniformly distributed over $\mathbb{F}^{m \times r}$.* $\qquad\square$

We now give an simple estimate for the number of low-dimensional subspaces (which is notably smaller than the number of low-rank matrices).

Lemma C.2. *Let $\mathbb{F}_q$ be a finite field and for $n \geqslant r$, the number of dimension r subspaces of $\mathbb{F}_q^n$ is $\leqslant \min\{q^{rn}, q^{(n-r)n}\}$. Further, this inequality is strict for $r \in (0,n)$.* $\qquad\square$

The above estimate is good enough for many purposes, but we will need a more refined version which we now develop. We first give an estimate for the probability that a random square matrix is full rank.

Lemma C.3 (see for example Forbes [For14, Lemma D.2.4])**.** *Let $\mathbb{F}_q$ be a finite field and $n \geqslant r \geqslant 1$. The probability a random matrix in $\mathbb{F}_q^{r \times r}$ has rank r is $> \left(1 - \frac{1}{q}\right)^{\frac{q}{q-1}}$.* $\qquad\square$

We now refine this to more amenable form using that $1 + x \leqslant e^x$.

Corollary C.4. *Let $\mathbb{F}_q$ be a finite field and $n \geqslant r \geqslant 1$. The probability a random matrix in $\mathbb{F}_q^{r \times r}$ has rank r is $> e^{-\frac{q}{(q-1)^2}}$.* $\qquad\square$

Turning this probabilistic estimate into a counting result gives the following.

Corollary C.5. *Let $\mathbb{F}_q$ be a finite field and $r \geqslant 1$. The number of matrices $\mathbb{F}_q^{r \times r}$ with rank r is $> e^{-\frac{q}{(q-1)^2}} q^{r^2}$.* $\qquad\square$

We now use this matrix counting result to count subspaces.

Lemma C.6. *Let $\mathbb{F}_q$ be a finite field and $n \geqslant r \geqslant 1$. The number of r-dimensional subspaces of $\mathbb{F}_q^n$ is $< e^{\frac{q}{(q-1)^2}} \cdot q^{r(n-r)}$.*

Proof: Consider the map from rank r matrices in $\mathbb{F}_q^{n \times r}$ to r-dimensional subspaces of $\mathbb{F}_q^n$ defined by $M \mapsto$ col-span(M). This map is well-defined and surjective. Consider then the pre-image of a rank r subspace $V \subseteq \mathbb{F}_q^n$. Each of these matrices is related by an invertible change of basis given by an invertible $r \times r$ matrix. Thus, the pre-images of this map are of size $> e^{-\frac{q}{(q-1)^2}} q^{r^2}$ by [Corollary C.5](#). The claim follows by applying double-counting using that there are q^{rn} matrices in $\mathbb{F}_q^{n \times r}$. $\qquad\square$

We remark that the number of subspaces is always at least $q^{r(n-r)}$, so the above bound is asymptotically correct for large q. We now bound the probability a random matrix is low-rank.

Lemma C.7. *Let $\mathbb{F}_q$ be a finite field and $n, m \geqslant r \geqslant 1$. Let M be a random variable uniformly distributed over matrices in $\mathbb{F}_q^{n \times m}$. Then,*

$$\Pr[\operatorname{rank} \mathsf{M} \leqslant r] < \min\{q^{-(nm - r(n+m))}, q^{-(n-m)(m-r)}, e^{\frac{q}{(q-1)^2}} q^{-(n-r)(m-r)}\} \,.$$

Proof: A matrix M is rank $\leqslant r$ iff there exists some $V \subseteq \mathbb{F}_q^m$ with $\dim V = r$ such that col-span $\mathsf{M} \subseteq V$. Equivalently, there exists some $V^\perp \subseteq \mathbb{F}_q^m$ with $\dim V^\perp = m - r$ such that $(\text{col-span } \mathsf{M})^\perp \supseteq V^\perp$. Taking $A \in \mathbb{F}^{m \times (m-r)}$ to be some basis for $V^\perp$, so that $\operatorname{rank} A = m - r$, this is equivalent to saying that $\mathsf{M}A = 0$. Thus, we have that

$$\Pr[\operatorname{rank} \mathsf{M} \leqslant r] = \Pr[\exists V \mid \text{col-span} A = V^\perp, \mathsf{M}A = 0]$$
$$\leqslant \sum_V \Pr[\mathsf{M}A = 0]$$

By [Lemma C.1](#) the random variable $\mathsf{M}A$ is uniformly distributed over $\mathbb{F}_q^{n \times (m-r)}$,

$$= \sum_V q^{-n(m-r)}$$

Counting such subspaces V by [Lemma C.2](#), using that $0 < r < m$,

$$< q^{m \cdot \min\{r, m-r\}} q^{-n(m-r)}$$
$$\leqslant \min\{q^{-(nm-r(n+m))}, q^{-(n-m)(m-r)}\} \,.$$

Alternatively, if we count subspaces with the more refined estimate of [Lemma C.6](#), we get that

$$\Pr[\operatorname{rank} \mathsf{M} \leqslant r] < e^{\frac{q}{(q-1)^2}} q^{r(m-r)} \cdot q^{-n(m-r)}$$
$$= e^{\frac{q}{(q-1)^2}} q^{-(m-r)(n-r)} \,. \qquad\square$$

Finally, as the quantity will come up more than once, we give the following approximation. It follows since for $q \geqslant 3$ we have that $\ln q \geqslant 1$ and for $q = 4$ that $\frac{2q}{(q-1)^2} = 8/9 < 1$.

Lemma C.8. *For $q \geqslant 4$,*

$$\frac{2q}{(q-1)^2 \ln q} \leqslant 1 \,. \qquad\square$$

C.2 Dimension Expanders

We now turn to non-constructive existence of good dimension expanders. We first study when random matrices expand subspaces of a given dimension r to a larger dimension t (dimension expanders must expand all subspaces of dimension *at most r*).

Proposition C.9. *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant t \geqslant r \geqslant 1$. Let $A_1, \ldots, A_d$ be random variables uniformly distributed over matrices in $\mathbb{F}_q^{n \times n}$. Then with probability $> 1 - 1/q^r$, for any subspace $V \subseteq \mathbb{F}_q^n$ of dimension r we have that*

$$\dim \sum_{i \in [d]} A_i(V) \geqslant t ,$$

assuming that

$$d \geqslant \frac{t-1}{r} + \frac{n-r+1}{n-t+1} + \frac{2q}{(q-1)^2 \ln q} .$$

In particular, if $q \geqslant 4$ then it suffices for

$$d \geqslant \frac{t-1}{r} + \frac{n-r+1}{n-t+1} + 1 .$$

Proof: Fix some subspace $V \subseteq \mathbb{F}_q^n$ of dimension r, and let $M \in \mathbb{F}_q^{n \times r}$ be a rank r matrix with col-span $M = V$. Then we see that $\dim \sum_i A_i(V) \geqslant t$ iff the $\mathbb{F}_q^{n \times rd}$ block matrix

$$N(V) := \left[A_1 M | \cdots | A_d M \right]$$

has rank $\geqslant t$. As M has full rank and the A_i are uniformly random, it follows from Lemma C.1 that $N(V)$ is uniformly random matrix over $\mathbb{F}^{n \times rd}$. Thus, by the refined version of Lemma C.7,

$$\Pr[\operatorname{rank} N(V) < t] = \Pr[\operatorname{rank} N(V) \leqslant t - 1]$$
$$< e^{\frac{q}{(q-1)^2}} q^{-(n-t+1)(rd-t+1)}$$

Thus, applying a union bound,

$$\Pr[\exists V \mid \operatorname{rank} N(V) < t] \leqslant \sum_V \Pr[\operatorname{rank} N(V) \leqslant t - 1]$$
$$< \sum_V e^{\frac{q}{(q-1)^2}} q^{-(n-t+1)(rd-t+1)}$$

counting subspaces by Lemma C.6,

$$\leqslant e^{\frac{q}{(q-1)^2}} q^{r(n-r)} \cdot e^{\frac{q}{(q-1)^2}} q^{-(n-t+1)(rd-t+1)}$$

This quantity is $\leqslant q^{-r}$ iff

$$(n-t+1)(rd-t+1) \geqslant r(n-r+1) + \frac{2q}{(q-1)^2 \ln q} .$$

Dividing by $r(n-t+1)$ on both sides, this is equivalent to

$$d \geqslant \frac{t-1}{r} + \frac{n-r+1}{n-t+1} + \frac{2q}{r(n-t+1)(q-1)^2 \ln q}$$

using that $r \geqslant 1$ and $n - t + 1 \geqslant 1$, it thus suffices for

$$d \geqslant \frac{t-1}{r} + \frac{n-r+1}{n-t+1} + \frac{2q}{(q-1)^2 \ln q} \,,$$

as desired. Appealing to Lemma C.8 yields the claim for $q \geqslant 4$. $\qquad\square$

We now apply a union bound to the above to existentially obtain dimension expanders.

Proposition C.10. *Let $\mathbb{F}_q$ be a finite field, $n \geqslant 1$, $\varepsilon > 0$ and $\alpha \in \mathbb{R}$ with $1 \leqslant \alpha < 1/\varepsilon$. 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*

$$d \geqslant \alpha + \frac{1}{1 - \alpha\varepsilon} + \frac{2q}{(q-1)^2 \ln q} \,.$$

In particular, if $q \geqslant 4$ then it suffices for

$$d \geqslant \alpha + \frac{1}{1 - \alpha\varepsilon} + 1 \,.$$

Proof: For $\mathcal{A}$ to be (ε, α) expander means that for subspaces $V \subseteq \mathbb{F}^n$ of dimension $r \leqslant \varepsilon n$ we need that $\dim \sum_i A_i(V) \geqslant \alpha \dim V$. That is, $\dim \sum_i A_i(V) \geqslant \lceil \alpha r \rceil$. For any fixed $r \leqslant \varepsilon n$, Proposition C.9 shows that a random collection $\mathcal{A}$ will have this property for all such V with probability $> 1 - 1/q^r$ as long as

$$d \geqslant \frac{\lceil \alpha r \rceil - 1}{r} + \frac{n - (r-1)}{n - (\lceil \alpha r \rceil - 1)} + \frac{2q}{(q-1)^2 \ln q}$$

As $\lceil \alpha r \rceil - 1 \leqslant \alpha r$ it follows that it is sufficient for

$$d \geqslant \alpha + \frac{n}{n - \alpha r} + \frac{2q}{(q-1)^2 \ln q}$$

and thus sufficient for

$$d \geqslant \alpha + \frac{1}{1 - \alpha\varepsilon} + \frac{2q}{(q-1)^2 \ln q} \,,$$

where this bound is independent of r.

Now, taking the union bound over all $1 \leqslant r \leqslant \varepsilon n$ we see that the failure probability is at $< \sum_{r=1}^{\infty} 1/q^r < 1$, so it follows that such a collection $\mathcal{A}$ exists that expands by a factor of α each subspace of dimension of at most $\leqslant \varepsilon n$. For $q \geqslant 4$ we can appeal to the latter part of Proposition C.9. $\qquad\square$

C.3 Lossy Rank Condensers

We now give an argument that good lossy rank condensers exist.

Proposition C.11. *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$, $\varepsilon \geqslant 0$ and $t > (1-\varepsilon)r$. Let $\mathsf{E}_1, \ldots, \mathsf{E}_k$ be random variables uniformly distributed over matrices in $\mathbb{F}_q^{t \times n}$. Then with probability $> 1 - 1/q^r$, $\mathcal{E} := \{\mathsf{E}_1, \ldots, \mathsf{E}_k\}$ is a (r, ε)-lossy rank condenser whenever*

$$k \geqslant \frac{rn + \frac{q}{(q-1)^2 \ln q}}{(t - (1-\varepsilon)r)(\lfloor \varepsilon r \rfloor + 1) - \frac{q}{(q-1)^2 \ln q}} \,,$$

whenever the denominator is positive. In particular, if $q \geqslant 4$ then it suffices for

$$k \geqslant \frac{rn+1}{(t-(1-\varepsilon)r)(\lfloor \varepsilon r \rfloor + 1) - 1} \, ,$$

whenever the denominator is positive.

Proof: We bound the probability that $\mathcal{E}$ is not such a condenser. Fix some matrix $M \in \mathbb{F}_q^{n \times r}$ of rank r. We then bound the probability

$$\Pr_{\mathsf{E}_i}[\forall i \in [k], \operatorname{rank} \mathsf{E}_i M < (1-\varepsilon) \operatorname{rank} M] = \Pr_{\mathsf{E}_i}[\forall i \in [k], \operatorname{rank} \mathsf{E}_i M \leqslant r - \lfloor \varepsilon r \rfloor - 1]$$

using independence of the E_i,

$$= \left(\Pr_{\mathsf{E}_1}[\operatorname{rank} \mathsf{E}_1 M \leqslant r - \lfloor \varepsilon r \rfloor - 1] \right)^k$$

as $\mathsf{E}_1 M$ is uniformly distributed over $\mathbb{F}_q^{t \times r}$ (by Lemma C.1), and invoking the finer form of Lemma C.7,

$$< \left(\mathrm{e}^{\frac{q}{(q-1)^2}} q^{-(t-(r-\lfloor \varepsilon r \rfloor - 1))(r-(r-\lfloor \varepsilon r \rfloor - 1))} \right)^k .$$

Union bounding over all such M (and using Lemma C.6 to count such M),

$$\Pr[\mathcal{E} \text{ not a } (r,\varepsilon)\text{-condenser}] = \Pr[\exists M \mid \forall i, \operatorname{rank} \mathsf{E}_i M < (1-\varepsilon) \operatorname{rank} M]$$

$$< \mathrm{e}^{\frac{q}{(q-1)^2}} q^{r(n-r)} \cdot \left(\mathrm{e}^{\frac{q}{(q-1)^2}} q^{-(t-(r-\lfloor \varepsilon r \rfloor - 1))(r-(r-\lfloor \varepsilon r \rfloor - 1))} \right)^k .$$

This quantity is at most q^{-r} iff

$$k \geqslant \frac{r(n-r+1) + \frac{q}{(q-1)^2 \ln q}}{(t-(r-\lfloor \varepsilon r \rfloor - 1))(\lfloor \varepsilon r \rfloor + 1) - \frac{q}{(q-1)^2 \ln q}}$$

so that it is sufficient (as $r \geqslant 1$, and $\lfloor \varepsilon r \rfloor + 1 \geqslant \varepsilon r$) that

$$k \geqslant \frac{rn + \frac{q}{(q-1)^2 \ln q}}{(t-(1-\varepsilon)r)(\lfloor \varepsilon r \rfloor + 1) - \frac{q}{(q-1)^2 \ln q}} \, ,$$

whenever this denominator is positive. Appealing to Lemma C.8 yields the claim for $q \geqslant 4$. $\qquad\square$

We note that the finer form of Lemma C.7 was not strictly needed in the above to obtain non-trivial results, but this finer form allows us to see that the explicit construction of Proposition 6.8 is slightly suboptimal as it requires $t \geqslant r$ (and this sub-optimality manifests in the constructions of dimension expanders as discussed after Corollary 7.3).

We now observe that by the above we can find a collection $\mathcal{E}$ that is a $(\leqslant r, \varepsilon)$-condenser. As any $(r,0)$-lossy condenser is also a $(s,0)$-condenser for all $s \leqslant r$ (Lemma 6.5), this is only non-trivial for $\varepsilon > 0$. Note that this is indeed non-trivial, as there are $(3n, 2/2)$-lossy condensers on $\mathbb{F}^{4n}$ that are not $(s, 2/3)$-condensers for any $s \leqslant n$ (see Lemma 6.6).

Proposition C.12. *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$, $\varepsilon > 0$ and $t > (1-\varepsilon)r$. Then there is a collection $\mathcal{E}$ of matrices $\mathcal{E} \subseteq \mathbb{F}_q^{t \times n}$ that is a $(\leqslant r, \varepsilon)$-lossy rank condenser whenever*

$$k \geqslant \frac{n + \frac{q}{(q-1)^2 \ln q}}{\varepsilon(t - (1-\varepsilon)r) - \frac{q}{(q-1)^2 \ln q}} ,$$

whenever this denominator is positive. In particular, if $q \geqslant 4$ then it suffices for

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

whenever this denominator is positive.

Proof: Take $\mathcal{E} := \{\mathsf{E}_1, \ldots, \mathsf{E}_k\}$ where E_i are uniformly and independently distributed over $\mathbb{F}_q^{t \times n}$. By Proposition C.11 we see that $\mathcal{E}$ is a (s, ε)-condenser with probability $> 1 - q^s$ as long as

$$k \geqslant \frac{sn + \frac{q}{(q-1)^2 \ln q}}{(t - (1-\varepsilon)s)(\lfloor \varepsilon s \rfloor + 1) - \frac{q}{(q-1)^2 \ln q}}$$

as $\lfloor \varepsilon s \rfloor + 1 \geqslant \varepsilon s$ and $s \leqslant r$, it is then sufficient for

$$k \geqslant \frac{sn + \frac{q}{(q-1)^2 \ln q}}{(t - (1-\varepsilon)r)\varepsilon s - \frac{q}{(q-1)^2 \ln q}}$$

whenever this denominator is positive. Simplifying further, we see that

$$k \geqslant \frac{n + \frac{q}{s(q-1)^2 \ln q}}{\varepsilon(t - (1-\varepsilon)r) - \frac{q}{s(q-1)^2 \ln q}}$$

and thus as $s \geqslant 1$, that

$$k \geqslant \frac{n + \frac{q}{(q-1)^2 \ln q}}{\varepsilon(t - (1-\varepsilon)r) - \frac{q}{(q-1)^2 \ln q}} ,$$

is sufficient whenever this denominator is positive, where this last bound is independent of s. Thus, with this value of k we see that $\mathcal{E}$ is a (s, ε)-condenser for all $1 \leqslant s \leqslant r$ except with probability $< \sum_{s=1}^{r} q^{-s} \leqslant \sum_{s=1}^{\infty} q^{-s} < 1$. Thus, such a condenser $\mathcal{E}$ exists, where for $q \geqslant 4$ we appeal to the latter half of Proposition C.11. $\quad\square$

C.4 Two-Source Rank Condensers

We now give an argument that good two-source rank condensers exist.

Proposition C.13. *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$, $m \geqslant s \geqslant 1$ and $\varepsilon \geqslant 0$. Let E be a random variable uniformly distributed over matrices in $\mathbb{F}_q^{t \times nm}$. Then $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ defined by $f(\overline{v}, \overline{w}) = E \cdot (\overline{v} \otimes \overline{w})$ is a bilinear (r, s, ε)-two-source rank condenser with probability $> 1 - 1/q^r$, assuming that*

$$t \geqslant \frac{n}{\varepsilon s} + \frac{m}{\varepsilon r} + (1-\varepsilon)rs + \frac{2q}{(q-1)^2 \ln q} .$$

for $\varepsilon > 0$. If $\varepsilon = 0$, then f is such a $(r, s, 0)$-condensers with probability $> 1 - \frac{1}{q^r}$, assuming that

$$t \geqslant rn + sm + rs + \frac{2q}{(q-1)^2 \ln q} - 1 .$$

Proof: Note that f is indeed a bilinear function as seen from the equivalence of the two definitions (Lemma 9.2). That f fails to be a (r,s,ε)-condenser means that there are full-rank matrices $A \in \mathbb{F}_q^{n \times r}$ and $B \in \mathbb{F}_q^{m \times s}$ such that $f(A,B) = \mathsf{E} \cdot (A \otimes B)$ has rank $f(A,B) < (1-\varepsilon)rs$, that is rank $f(A,B) \leqslant rs - (\lfloor \varepsilon rs \rfloor + 1)$. Thus, appealing to the above lemmas (and that $A \otimes B$ is rank rs),

$$\begin{aligned}
\Pr[f \text{ is not a } (r,s,\varepsilon)\text{-condenser}] &= \Pr[\exists A, B \mid \operatorname{rank} f(A,B) \leqslant rs - (\lfloor \varepsilon rs \rfloor + 1)] \\
&\leqslant \sum_{A,B} \Pr[\operatorname{rank} f(A,B) \leqslant rs - (\lfloor \varepsilon rs \rfloor + 1)] \\
&= \sum_{A,B} \Pr[\operatorname{rank}(\mathsf{E} \cdot (A \otimes B)) \leqslant rs - (\lfloor \varepsilon rs \rfloor + 1)]
\end{aligned}$$

appealing to Lemma C.1 to see that $\mathsf{E} \cdot (A \otimes B)$ is a random $t \times rs$ matrix, and then applying the third form of Lemma C.7, as well as Lemma C.2,

$$\begin{aligned}
&< \sum_{A,B} e^{\frac{q}{(q-1)^2}} q^{-(t-(rs-(\lfloor \varepsilon rs \rfloor+1)))(rs-(rs-(\lfloor \varepsilon rs \rfloor+1)))} \\
&< e^{\frac{q}{(q-1)^2}} q^{r(n-r)} \cdot q^{sm} \cdot e^{\frac{q}{(q-1)^2}} q^{-(t-(rs-(\lfloor \varepsilon rs \rfloor+1)))(\lfloor \varepsilon rs \rfloor+1)}
\end{aligned}$$

The above quantity is $\leqslant q^{-r}$ iff

$$t \geqslant \frac{r(n-r+1)+sm}{\lfloor \varepsilon rs \rfloor + 1} + (rs - \lfloor \varepsilon rs \rfloor + 1) + \frac{2q}{(\lfloor \varepsilon rs \rfloor + 1)(q-1)^2 \ln q}$$

using that $r, s \geqslant 1$, it suffices for

$$t \geqslant \frac{rn+sm}{\lfloor \varepsilon rs \rfloor + 1} + (rs - \lfloor \varepsilon rs \rfloor - 1) + \frac{2q}{(q-1)^2 \ln q}$$

which yields the result for $\varepsilon = 0$. For $\varepsilon > 0$, we use that $\varepsilon rs \leqslant \lfloor \varepsilon rs \rfloor + 1$, so see that it suffices for

$$t \geqslant \frac{rn+sm}{\varepsilon rs} + (1-\varepsilon)rs + \frac{2q}{(q-1)^2 \ln q} \ . \qquad \square$$

We now use the above to derive a lossy two-source condenser that works even when the first source is *small*. Note that as before this is only interesting when $\varepsilon > 0$, as when $\varepsilon = 0$ condensers automatically work for smaller ranks (Lemma 9.3).

Proposition C.14. *Let $\mathbb{F}_q$ be a finite field. Let $n \geqslant r \geqslant 1$, $m \geqslant s \geqslant 1$ and $\varepsilon > 0$. Then there exists a $f : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{F}^t$ which is a bilinear $(\leqslant r, s, \varepsilon)$-two-source rank condenser, assuming that*

$$t \geqslant \frac{n}{\varepsilon s} + \frac{m}{\varepsilon} + (1-\varepsilon)rs + \frac{2q}{(q-1)^2 \ln q} \ .$$

Proof: Let E be a random variable uniformly distributed over matrices in $\mathbb{F}_q^{t \times nm}$ and define f by $f(\overline{v}, \overline{w}) = \mathsf{E} \cdot (\overline{v} \otimes \overline{w})$. For r' with $1 \leqslant r' \leqslant r$, Proposition C.13 yields that f is a bilinear (r', s, ε)-two-source condenser with probability $> 1 - 1/q^{r'}$, assuming that

$$t \geqslant \frac{n}{\varepsilon s} + \frac{m}{\varepsilon r'} + (1-\varepsilon)r's + \frac{2q}{(q-1)^2 \ln q}$$

in particular, as $1 \leqslant r' \leqslant r$, if

$$t \geqslant \frac{n}{\varepsilon s} + \frac{m}{\varepsilon} + (1 - \varepsilon)rs + \frac{2q}{(q-1)^2 \ln q} \ ,$$

where this last bound is independent of r'. The probability of failure, union bounding over all $1 \leqslant r' \leqslant r$, is at most $\sum_{r'=1}^{\infty} q^{-r'} < 1$. Thus the desired f exists. $\qquad \square$