

Towards Dimension Expanders Over Finite Fields

Zeev Dvir

Department of Computer Science
Weizmann institute of science
Rehovot, Israel.
zeev.dvir@weizmann.ac.il

Amir Shpilka

Faculty of Computer Science
Technion - Israel Institute of Technology,
Haifa, Israel.
shpilka@cs.technion.ac.il

May 27, 2008

Abstract

In this paper we study the problem of explicitly constructing a *dimension expander* raised by [3]: Let \mathbb{F}^n be the n dimensional linear space over the field \mathbb{F} . Find a small (ideally constant) set of linear transformations from \mathbb{F}^n to itself $\{A_i\}_{i \in I}$ such that for every linear subspace $V \subset \mathbb{F}^n$ of dimension $\dim(V) < n/2$ we have

$$\dim \left(\sum_{i \in I} A_i(V) \right) \geq (1 + \alpha) \cdot \dim(V),$$

where $\alpha > 0$ is some constant. In other words, the dimension of the subspace spanned by $\{A_i(V)\}_{i \in I}$ should be at least $(1 + \alpha) \cdot \dim(V)$. For fields of characteristic zero Lubotzky and Zelmanov [8] completely solved the problem by exhibiting a set of matrices, of size independent of n , having the dimension expansion property. In this paper we consider the finite field version of the problem and obtain the following results.

1. We give a constant number of matrices that expand the dimension of every subspace of dimension $d < n/2$ by a factor of $(1 + 1/\log n)$.
2. We give a set of $O(\log n)$ matrices with expanding factor of $(1 + \alpha)$, for some constant $\alpha > 0$.

Our constructions are algebraic in nature and rely on expanding Cayley graphs for the group $\mathbb{Z}/\mathbb{Z}n$ and small-diameter Cayley graphs for the group $\text{SL}_2(p)$.

1 Introduction

Let \mathbb{F} be a field, and $V \subset \mathbb{F}^n$ a linear subspace of dimension $\dim(V) < n/2$. It is easy to verify that a random subspace $U \subset \mathbb{F}^n$ of dimension at most $n/2$ will have a small intersection, w.h.p.,¹ with

¹When $|\mathbb{F}|$ grows, V and U are likely to intersect only at the $\vec{0}$ vector.

V . Similarly, if we pick a linear transformation A from \mathbb{F}^n to itself, at random, then w.h.p V and $A(V)$ will have a small intersection. Stated differently this is equivalent to having $\dim(V + A(V))$ significantly larger than $\dim(V)$. It is not difficult to show that if we pick a small number (say 100) linear transformations $\{A_i\}_{i=1}^{100}$ at random then w.h.p. for every such V we will have that

$$\dim\left(\sum_{i=1}^{100} A_i(V)\right) \geq \frac{11}{10} \cdot \dim(V).$$

It is thus a natural problem (indeed it was raised in the discussions of [3] and was posed as an open problem by Wigderson [12]) to find an explicit construction of such set $\{A_i\}$. We refer to a set of linear transformations having the dimension expansion property as a *dimension expander*.

In this paper we study the problem of explicitly constructing dimension expanders, of a constant size, over finite fields and obtain some partial results. We start by giving the formal definition and statement of the problem. Then we state our results and discuss the context of the problem.

1.1 Our results

Before stating our result we shall need the following formal definition of a dimension expander.

Definition 1.1 (Dimension Expander). *Let \mathbb{F} be a field and let $A_1, \dots, A_k : \mathbb{F}^n \rightarrow \mathbb{F}^n$ be linear mappings. The set $\mathcal{A} = \{A_i\}_{i=1}^k$ is a (d, α) -dimension expander if for every subspace $V \subset \mathbb{F}^n$ of dimension at most d we have*

$$\dim\left(\sum_{i=1}^k A_i(V)\right) \geq (1 + \alpha) \cdot \dim(V). \quad (1)$$

We say that \mathcal{A} is explicit if there exists a $\text{poly}(n)$ -time algorithm that, on input n , outputs \mathcal{A} .

Problem 1. *Construct an explicit (d, α) -dimension expander $\mathcal{A} = \{A_i\}_{i=1}^k$, with $d = \Omega(n)$, $\alpha = \Omega(1)$ and $k = O(1)$.*

We give two constructions of dimension expanders. The first gives a set of $\log n$ linear transformations that have a constant expansion factor. The second gives a constant number of linear transformation with an expansion factor of $1 + 1/\log n$.

Theorem 1. *Let \mathbb{F} be a field. There exists a constant $\alpha > 0$ such that for every n there exists a set $\mathcal{A}(n)$ of $O(\log(n))$ linear mappings from \mathbb{F}^n to \mathbb{F}^n that is an $(\Omega(n), \alpha)$ -dimension expander. Moreover, the construction of $\mathcal{A}(n)$ is explicit and independent of the field \mathbb{F} .*

Theorem 2. *Let \mathbb{F} be a field. There exists a constant $k_0 > 0$ such that for every n there exists a set $\mathcal{A}(n)$ of k_0 linear mappings from \mathbb{F}^n to \mathbb{F}^n that is an $(\Omega(n), \Omega(1/\log(n)))$ -dimension expander. Moreover, the construction of $\mathcal{A}(n)$ is explicit and independent of the field \mathbb{F} .*

1.2 Background

The question of explicitly constructing a dimension expander is a very natural derandomization problem. It is very easy to show that a random choice of linear transformations yields an $(\Omega(n), \alpha)$ -dimension expander (for some constant $\alpha > 0$), and we wish to find, deterministically, an explicit construction.

Giving an explicit construction of an object whose existence is easily proved by probabilistic arguments has been a very active and fruitful field of research. Two combinatorial objects that are closely related to dimension expanders are affine extractors and bounded degree expander graphs.

An *affine extractor* is a function $E : \mathbb{F}^n \rightarrow \mathbb{F}^k$ that, when restricted to any affine subspace V of dimension $k' = O(k)$, satisfy that $E(V)$ is (almost) uniformly distributed. Clearly such extractors seem very related to dimension expanders. The main difference between the two concepts is that extractors do not need to act linearly on the input. In fact, it is not difficult to show that it is impossible for E to be linear. The problem of constructing affine extractors was solved almost completely by Gabizon and Raz [5] over large (polynomial in n) fields and partially by Bourgain [4] over \mathbb{F}_2 (Bourgain's extractors work when $k = \Omega(n)$). For subspaces of small dimension in \mathbb{F}_2 , and for small fields other than \mathbb{F}_2 , the problem is still open. It may seem surprising, but the affine extractors of [5] already found an application in the work of [6] regarding the derandomization of polynomial identity testing for depth three circuits. We hope that dimension expanders will prove to be useful objects for other derandomization problems.

Another related combinatorial object is constant degree expander graphs. Dimension expanders can be thought of as constant degree expanders that need to expand (in dimension) any linear space. Indeed, a different formulation of Problem 1 is: We wish to construct a k -regular graph whose nodes are the elements of \mathbb{F}^n , such that each element v is connected to $\{A_i v\}_{i=1}^k$, that has the following property: the neighborhood of any d dimensional vector space is not contained in any $(1 + \alpha) \cdot d$ dimensional space, for $d = O(n)$. We note that a usual edge/vertex expander is not necessarily a dimension expander as even if the size of the neighborhood of a subspace V is of size $k \cdot |V|$, it does not mean that it is not contained in a linear space of dimension, say, $\dim(V) + \log k$. It is interesting to note though that both our constructions and the construction of [8] are based on expanding Cayley graphs (see Section 1.3 for definition of a Cayley graph).

Beside being a natural derandomization problem, the question of constructing dimension expanders is related to an interesting problem in the theory of representations of finite groups over finite fields. Specifically it is related to an attempt to generalize the notion of property T of unitary representations of finite groups (over \mathbb{C}) to representations over finite fields.

1.3 Property T

Below we give the formal definition of property T . For more on this topic we refer the reader to the excellent books [7, 9] where some applications of property T are also discussed. In the following $\|\cdot\|$ denotes the ℓ_2 norm.

Definition 1.2 (Property T). *Let G be a finite group, $S \subset G$ a set of generators for G and $\rho : G \rightarrow U(\mathbb{C}^n)$ an irreducible unitary representation of G . The Kazhdan constant of G and S with respect to ρ is defined as:*

$$\kappa_G^S(\rho) = \min_{0 \neq v \in \mathbb{C}^n} \max_{s \in S} \frac{\|\rho(s)v - v\|}{\|v\|}. \quad (2)$$

The Kazhdan constant of G with respect to S is defined as $\kappa_G(S) = \inf_{\rho \in \mathcal{U}} \kappa_G^S(\rho)$, where \mathcal{U} is the set of nontrivial irreducible unitary representations of G . We say that G has property T with respect to S if $\kappa_G(S) > 0$.

It is not hard to see that G has property T with respect to S if the Cayley graph $\text{Cay}(G, S)$ is an expander. Recall that the Cayley graph $\text{Cay}(G, S)$ is the graph on the elements of G , in which g and g' are connected by an edge if and only if $g \cdot g'^{-1} \in S \cup S^{-1}$. Moreover, if we replace in Equation (2) $\max_{s \in S} \frac{\|\rho(s)v - v\|}{\|v\|}$ with the average $\frac{\|\frac{1}{|S|} \sum_{s \in S} \rho(s)v - v\|}{\|v\|}$ then we get (following the notations of [11])

$$\tilde{\kappa}_G(S) = \inf_{\rho \in \mathcal{U}} \min_{0 \neq v \in \mathbb{C}^n} \frac{\|\frac{1}{|S|} \sum_{s \in S} \rho(s)v - v\|}{\|v\|} \quad (3)$$

the ‘‘averaged’’ Kazhdan constant. Meshulam and Wigderson [11] observed that

$$\lambda(\text{Cay}(G, S)) = 1 - \tilde{\kappa}_G(S),$$

where $\lambda(\text{Cay}(G, S))$ is the second largest eigenvalue of the normalized adjacency matrix of $\text{Cay}(G, S)$. This highlights the tight connection between property T and expansion of the corresponding Cayley graph.

The connection between property T and dimension expanders was demonstrated in the work of Lubotzky and Zelmanov [8] that proved that if G has property T with respect to S , with Kazhdan constant $\kappa = \kappa_G(S)$, and $\rho : G \rightarrow \mathbb{C}^n$ is any nontrivial irreducible unitary representation then the set $\{\rho(s)\}_{s \in S}$ is an $(O(n), \alpha)$ dimension expander over \mathbb{C} for $\alpha \geq \frac{\kappa^2}{4}$. Interestingly, Lubotzky and Zelmanov were not motivated by the question of constructing a dimension expander but rather with the problem of generalizing the notion of property T to finite fields as well. Indeed, property T is defined through unitary representations, a concept that does not exist for finite fields. As property T proved to be such a useful mathematical concept it is very desirable to have a finite field analog of it. The work of [8] suggests that a possible generalization of property T for representations over finite fields is to ask that $\{\rho(s)\}_{s \in S}$ is a dimension expander. Namely, G will have property T if there exists a set of generators S such that for every non trivial representation ρ , $\{\rho(s)\}_{s \in S}$ is a dimension expander. Currently however little is known on the relation between dimension expanders and Property T. In particular Lubotzky and Zelmanov [8] ask the following question:³

Question 1. *Let G be a group generated by a finite set S . If $\rho : G \rightarrow \mathbb{C}^n$ is an irreducible unitary representation such that $\{\rho(s)\}_{s \in S}$ is a dimension expander, then is it true that $\kappa_G^S(\rho) > 0$?*

² $U(\mathbb{C}^n)$ is the group of unitary transformations from \mathbb{C}^n to itself.

³The question of [8] is slightly different and concerns the Kazhdan constant of the representation $\text{adj}(\rho)$ that is derived from ρ , but for the purpose of this exposition we give a slightly stronger question.

Our results also give the feeling that there is a strong connection between expanding Cayley graphs and dimension expanders (over finite fields), and so it hints that dimension expanders are the “correct” analog of property T , but we do not have any formal theorem of that sort.

1.4 Organization

In Section 2 we give some definitions and prove some basic properties of dimension expanders. We also discuss expanding Cayley graphs there. In Section 3 we prove Theorem 1 and in Section 4 we prove Theorem 2.

2 Preliminaries

2.1 Subspaces and their set of degrees

Let $v = (v_1, \dots, v_n) \in \mathbb{F}^n$ be a non-zero vector. We denote by $\deg(v)$ (the degree of v) the largest index $i \in [n]$ such that $v_i \neq 0$. Let V be a subspace of dimension k in \mathbb{F}^n and let

$$D_V = \{\deg(v) \mid v \in V, v \neq 0\}$$

be the set of degrees of all vectors in V . It is clear that $|D_V| = k$, since vectors with distinct degrees are always linearly independent and since we can always find a basis of V such that the degrees of the basis vectors are distinct. The following claim is trivial, but we make it explicit since it will be used many times in our construction.

Claim 2.1. *Let $V \subset \mathbb{F}^n$ be a k -dimensional subspace and let D_V be its set of degrees. Let $A : \mathbb{F}^n \rightarrow \mathbb{F}^n$ be a linear mapping and let*

$$D_{A(V)} = \{\deg(A(v)) \mid v \in V, A(v) \neq 0\}.$$

Then,

$$\dim(V + A(V)) \geq |D_V \cup D_{A(V)}|$$

2.2 Expanding generators for $(\mathbb{Z}_n, +)$

Let $G = (V, E)$ be an undirected graph on n vertices. For a set $S \subset V$ we denote by $\Gamma(S)$ the set of neighbors of S in G . We say that G is an (s, β) -expander if for every set $S \subset V$ such that $|S| \leq s$ we have $|\Gamma(S)| \geq (1 + \beta)|S|$. We will be interested in the case in which the set of vertices of G is a group. Let H be a finite group (possibly non-abelian). Let $M \subset H$ be a set of generators for H . The *Cayley graph* induced by M on H is the (undirected) graph with vertex set H and such that two vertices v, u have an edge between them iff there exists $m \in M$ such that $u = mv$ or $v = mu$. We denote this graph by $\text{Cay}(H, M)$.

An important ingredient of our construction will be a set of integers $J = \{j_1, \dots, j_d\} \subset [n]$ such that the Cayley graph they induce on $(\mathbb{Z}_n, +)$ (the group of integers modulo n with the

operation of addition) is an $(\gamma n, \beta)$ -expander for constants $0 < \beta, \gamma$ independent of n . Since $(\mathbb{Z}_n, +)$ is an abelian group, we know that the size of J must be at least logarithmic in n . It is also known that $O(\log(n))$ generator are sufficient to give an expanding Cayley graph [1]. A result of Wigderson and Xiao [13] allows us to compute such a set J in polynomial time (the result in [13] is more general and regards *any* group H and not just \mathbb{Z}_n).

Theorem 2.2 (Special case of [13]). *There exist constants $\beta, \gamma > 0$ and an algorithm T such that on input n , the algorithm runs in $\text{poly}(n)$ time and returns a set $J \subset [n]$ of size $O(\log(n))$ such that J generates $(\mathbb{Z}_n, +)$ and the graph $\text{Cay}(\mathbb{Z}_n, J)$ is a $(\gamma n, \beta)$ -expander.*

2.3 Generating cyclic shifts with small diameter

For an integer n we denote by $s_1, \dots, s_n : \mathbb{F}^n \rightarrow \mathbb{F}^n$ the n right cyclic shifts of the coordinates of \mathbb{F}^n . That is, for a vector $v = (v_1, \dots, v_n) \in \mathbb{F}^n$ we have

$$s_j(v) = (v_{n-j+1}, \dots, v_n, v_1, v_2, \dots, v_{n-j}).$$

In order to prove Theorem 2 we will need to find a way to express all n cyclic shifts as words of length $O(\log(n))$ using some small (constant) set of permutations on $[n]$. The next lemma allows us to do something almost as good, which will be sufficient for our needs. One restriction which we will later need to overcome is that $n = p + 1$ for a prime p .

Lemma 2.3. *Let $n = p + 1$, for a prime p . Let S_{p+1} denote the set of permutations on $\{1, \dots, p + 1\}$. Let $s_1, \dots, s_p \in S_{p+1}$ denote the p right cyclic shifts on the set $\{1, \dots, p\}$ (as defined above), and such that $s_j(p + 1) = p + 1$ for every j . Then, there exists a set $M \subset S_{p+1}$ of size $|M| \leq 7$ such that for every $j \in [p]$, the permutation s_j can be written as a word of length $O(\log(p))$ using elements from $M \cup M^{-1}$ (elements from M and their inverses). Moreover, this set M can be generated in time polynomial in n .*

Proof. Let $G = SL(2, p)$ be the group of 2×2 matrices over \mathbb{F}_p with determinant one. It was shown in [2] that there exists a set $M' \subset G$ of size at most seven such that the diameter of $\text{Cay}(G, M')$ is $O(\log |G|) = O(\log(p))$ (the result of [2] holds for any finite simple group).

Let us identify the set $[n] = [p + 1]$ with the projective line $\mathbb{P}^1 = \mathbb{Z}_p \cup \{\infty\}$ via the bijection $\xi : [n] \rightarrow \mathbb{P}^1$ that sends $\xi(j) = j - 1$ for $1 \leq j \leq p$ and $\xi(p + 1) = \infty$. The group G acts on \mathbb{P}^1 in a natural way via a group homomorphism $\phi : G \rightarrow S_{\mathbb{P}^1} \simeq S_{p+1}$ that is defined as follows: for a matrix $A = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in G$ we let $\phi(A)$ be the permutation on \mathbb{P}^1 defined by

$$\phi(A)(x) = \frac{ax + b}{cx + d},$$

where we use the convention that $\frac{j}{0} = \infty$ for $j \neq 0$ and that $\phi(A)(\infty) = \frac{a}{c}$ (notice that we never have to deal with the case $\frac{0}{0}$ since $\det(A) = 1$). It is easy to verify that ϕ is indeed a group homomorphism. The reason ϕ is useful is that the image of ϕ contains all p cyclic shifts on \mathbb{Z}_p that fix ∞ (these are the images of the matrices $\begin{pmatrix} 1 & j \\ 0 & 1 \end{pmatrix}$) and these are precisely the permutations (now viewed as elements of S_{p+1}) that we wish to write as short words.

All is left now is to observe that $M = \phi(M')$ satisfies the requirement of the lemma. For $j \in \mathbb{Z}_p$ let $A_j = \begin{pmatrix} 1 & j \\ 0 & 1 \end{pmatrix}$. Let $s_j = \phi(A_j)$ be a cyclic shift on \mathbb{Z}_p that fixes ∞ . Using the result of [2] we can write $A_j = M'_1 M'_2 \dots M'_t$, where for each $i \in [t]$, $M'_i \in M' \cup M'^{-1}$ and $t \leq O(\log(p))$. Applying ϕ to this equality we get that

$$s_j = \phi(M'_1) \circ \phi(M'_2) \circ \dots \circ \phi(M'_t)$$

and we are done. □

2.4 Expansion of a composition of linear maps

In the proof of Theorem 2 we will need to use the following easy lemma. What the lemma says is that for a fixed subspace the expansion of a composition of linear maps is at most the sum of their individual expansion w.r.t that particular subspace. Because of its simplicity and usefulness we give the lemma here and not as part of the proof of Theorem 2.

Lemma 2.4. *Let \mathbb{F} be a field. Let $V \subset \mathbb{F}^n$ be a subspace of dimension k . Let $A_1, \dots, A_t : \mathbb{F}^n \rightarrow \mathbb{F}^n$ be linear maps such that for every $i \in [t]$ we have*

$$\dim(V + A_i(V)) \leq (1 + \alpha) \cdot k,$$

where α is some positive number (α can depend on n). Let $A = A_1 \circ \dots \circ A_t$, then

$$\dim(V + A(V)) \leq (1 + t\alpha) \cdot k.$$

Proof. The proof is by induction on t . If $t = 1$ then the claim is trivial. Suppose $t > 1$. Set $A' = A_2 \circ \dots \circ A_t$. We have that $A = A_1 \circ A'$. From the inductive hypothesis we know that

$$\dim(V + A'(V)) \leq (1 + (t-1)\alpha) \cdot k.$$

This means that $A'(V)$ contains at most $(t-1)\alpha \cdot k$ linearly independent vectors that do not belong to V . Applying A_1 to $A'(V)$ can add to these at most an additional $\alpha \cdot k$ independent vectors from $\mathbb{F}^n \setminus V$ and so the total number of linearly independent vectors in $A(V)$ that are not in V is at most $t\alpha \cdot k$. This means that $\dim(V + A(V)) \leq (1 + t\alpha) \cdot k$. □

3 Constant expansion using $O(\log(n))$ maps

In this section we give a construction of $O(\log(n))$ linear mappings that expand the dimension of any (not too large) subspace by a constant factor. Throughout this section we will assume w.l.o.g that n is even. If n is odd we can fix one coordinate in \mathbb{F}^n to zero. This can reduce the expansion by an additive term of at most one, that can be safely ignored. Hence, we write $n = 2m$.

Recall that $s_1, \dots, s_n : \mathbb{F}^n \rightarrow \mathbb{F}^n$ are the n right cyclic shifts of the coordinates of \mathbb{F}^n . We construct our dimension expander by first taking $O(\log(n))$ cyclic shifts $\{s_j\}_{j \in J}$ where $J \subset [m]$ is

such that $\text{Cay}(\mathbb{Z}_m, J)$ is an expander (see Section 2.2). Notice that the mappings s_j are cyclic shifts of \mathbb{F}^n and not \mathbb{F}^m , even though we use an expander for \mathbb{Z}_m . This set of linear maps alone cannot be a dimension expander since it is known that for certain values of n there are low dimensional subspaces of \mathbb{F}^n that are invariant to *all* cyclic shifts (see [10]). To fix this, we include in our dimension expander two additional maps we denote by $P_L, P_R : \mathbb{F}^n \rightarrow \mathbb{F}^n$, which are defined as follows: Let $v = (v', v'') \in \mathbb{F}^n$ be a vector such that v' denotes the first m coordinates of v and v'' denotes the last m coordinates of v . We define $P_L(v', v'') = (v'', \bar{0})$ and $P_R(v', v'') = (\bar{0}, v')$. That is, P_L is a composition of a projection on the last m coordinates and a left shift of m places and P_R is a composition of a projection on the *first* m coordinates and a *right* shift of m places⁴.

The intuition for the proof is the following: Consider the set of degrees of a k -dimensional vector space V (as defined in Section 2.1). We can split this set into two sets: one containing the degrees smaller than m and the other the degrees larger than m . Call these sets D_L and D_R (L and R for Left and Right). If P_L, P_R do *not* expand V than it means that D_L is approximately equal to D_R when considering the sets modulo m . Now, for a vector v such that $\deg(v) \in D_L$ and for a shift $s_j, j \in J$, we have that $\deg(s_j(v)) = \deg(v) + j$. This means that the shifts we use act on the set D_L as if it was a set in \mathbb{Z}_m (the fact that $D_L \sim D_R$ is what makes this work). Using the expansion properties of $\text{Cay}(\mathbb{Z}_m, J)$ we get that applying all the shifts in J gives us many new degrees that did not appear in $D_L \cup D_R$ and, using Claim 2.1, the proof is completed.

The next theorem, which directly implies Theorem 1, describes this construction more formally.

Theorem 3.1. *Let \mathbb{F} be a field and let $n = 2m$ be an even integer. Let $\beta, \gamma > 0$ and $J \subset [m]$ be given by Theorem 2.2. That is, the graph $\text{Cay}(\mathbb{Z}_m, J)$ is a $(\gamma n, \beta)$ -expander and $|J| \leq O(\log(m))$. Let $\mathcal{A}(n) = \{s_j\}_{j \in J} \cup \{P_L, P_R\}$, where s_1, \dots, s_n and P_L, P_R are defined as above. Then, $\mathcal{A}(n)$ is a $(\gamma' n, \beta')$ -dimension expander, where $\gamma', \beta' > 0$ depend only on γ and β respectively.*

Proof. Let $\alpha > 0$ be some small constant to be determined later (α will be a function of β). Let $V \subset \mathbb{F}^n$ be a subspace of dimension k . We will assume that P_L, P_R do *not* expand V by a factor of α and will conclude that the set of shifts $\{s_j\}_{j \in J}$ expand V by some constant factor depending on β .

Let D be the set of degrees of V (see Section 2.1) and let $D_L = D \cap \{1, \dots, m\}$ and $D_R = D \setminus D_L$. By definition we have $|D| = |D_L| + |D_R| = k$. We denote $(D_R - m) = \{i - m \mid i \in D_R\}$ and similarly for other sets (and shifts).

Claim 3.2. *Suppose that $\dim(V + P_L(V) + P_R(V)) \leq (1 + \alpha) \cdot k$. Then*

1. $|D_L \cup (D_R - m)| \leq (1 + \alpha) \cdot \frac{k}{2}$.
2. $|D_L \cap (D_R - m)| \geq (1 - \alpha) \cdot \frac{k}{2}$.

Proof. It is easy to see that the set of degrees of the subspace $V + P_L(V) + P_R(V)$ contains the four sets $D_L, D_R, (D_L + m)$ and $(D_R - m)$. If the first item of the claim is false, that is:

⁴The reader familiar with cyclic codes can easily verify that P_L alone expands any cyclic subspace. However, we do not use this fact directly in our proof.

$|D_L \cup (D_R - m)| > (1 + \alpha) \cdot \frac{k}{2}$, then it also holds that $|D_R \cup (D_L + m)| > (1 + \alpha) \cdot \frac{k}{2}$. This implies, using Claim 2.1, that the dimension of $V + P_L(V) + P_R(V)$ is larger than $2 \cdot (1 + \alpha) \cdot \frac{k}{2} = (1 + \alpha) \cdot k$ - a contradiction.

The second item of the claim follows directly from the first item and the fact that $|D_L| + |D_R - m| = k$, without using any special properties of the sets. \square

We now proceed under the assumption (and conclusions) of Claim 3.2. Our goal is to show (again, using Claim 2.1) that the set of shifts $\{s_j\}_{j \in J}$ expand V by a constant factor. In order to do so we define the set

$$R = \bigcup_{j \in J} (D_L + j).$$

Observe that since both D_L and J are subsets of $\{1, \dots, m\}$ it holds that for a vector $v \in V$ such that $\deg(v) \in D_L$ we have $\deg(s_j(v)) = \deg(v) + j$. Therefore, the set of degrees of the subspace $\sum_{j \in J} s_j(V)$ contains the set R . Hence, showing that $|R \cup D| \geq (1 + \beta') \cdot k$ for some constant β' will complete the proof.

Let $R' = (R \bmod m) = \{i \bmod m \mid i \in R\}$. The set R' corresponds to the set of neighbors of D_L in the graph $\text{Cay}(\mathbb{Z}_m, J)$. We would like to use the expansion properties of $\text{Cay}(\mathbb{Z}_m, J)$ to show that R' is larger than D_L by a constant factor. In order to do so we need to make sure that D_L is not too large. Taking $\gamma' = \gamma/2$ and observing that $|D_L| \leq k \leq \gamma' n = \gamma m$ we can indeed use the above expander to conclude that $|R'| \geq (1 + \beta) \cdot |D_L|$. Using this fact and both parts of Claim 3.2 we can derive the following inequality

$$\begin{aligned} |R'| &\geq (1 + \beta) \cdot |D_L| \\ &\geq (1 + \beta) \cdot (1 - \alpha) \cdot \frac{k}{2} \\ &= (1 + \beta) \cdot \frac{1 - \alpha}{1 + \alpha} \cdot (1 + \alpha) \cdot \frac{k}{2} \\ &\geq (1 + \beta) \cdot \frac{1 - \alpha}{1 + \alpha} \cdot |D_L \cup (D_R - m)| \\ &\geq (1 + \beta/2) \cdot |D_L \cup (D_R - m)|, \end{aligned} \tag{4}$$

where the last inequality holds for small enough $\alpha = \alpha(\beta)$.

From Eq. 4 it follows that R' contains at least

$$(\beta/2) \cdot |D_L \cup (D_R - m)| \geq (\beta/2) \cdot (k/2) = (\beta/4) \cdot k$$

elements *not* in $D_L \cup (D_R - m)$. This implies that R contains the same number of elements not in $D_L \cup D_R$. From this we conclude that $|R \cup D| \geq (1 + \beta/4) \cdot k$. Combining all of the above we get that $\mathcal{A}(n)$ expands V by $\beta' = \min\{\alpha, \beta/4\}$ as was required. \square

4 Inverse logarithmic expansion using $O(1)$ maps

We now turn to proving Theorem 2. We will use ideas similar to the ones appearing in Section 3 together with some new observations. The main new ingredient will be an application of Lemma 2.3.

Since Lemma 2.3 requires n to be $p + 1$ for a prime p we first give a construction for this case. Later we will show how to deal with general n .

4.1 The case $n = p + 1$

In this section we describe a construction of a dimension expander as in Theorem 2, for the case $n = p + 1$, p prime. Let p be an odd prime and let $n = p + 1$. Our dimension expander will include the set of coordinate permutations $M \subset S_{p+1}$ given by Lemma 2.3 together with their inverses and one additional mapping $P : \mathbb{F}^n \rightarrow \mathbb{F}^n$ defined as follows: Let $v = (v_1, \dots, v_n) \in \mathbb{F}^n$. We define $P(v) = (v_{(p+3)/2}, \dots, v_{p+1}, 0, \dots, 0)$. That is, P is defined the same way as P_L from Section 3 with $n = p + 1$ (notice that n is even).

In view of the discussion in Section 3, the intuition for the proof is quite clear: If P does not expand a subspace V then the set of ‘small’ degrees D_L must be of size at least, say, $k/3$. From this fact, it follows quite easily that there *exists* a cyclic shift of the first p coordinates that expand V by a constant (consider the expected number of ‘new’ degrees in a randomly chosen shift of D_L). Writing this shift as a composition of $O(\log(p))$ permutations from $M \cup M^{-1}$ and using Lemma 2.4 we deduce that there must be a mapping in $M \cup M^{-1}$ that expand V by at least $\Omega(1/\log(p))$.

The following theorem describes this construction more formally.

Theorem 4.1. *Let \mathbb{F} be a field and let $n = p + 1$ where p is an odd prime. Let M be the set of coordinate permutations of \mathbb{F}^n given by Lemma 2.3 and let P be the mapping defined above. Recall that $|M| \leq 7$. Let $\mathcal{A}(n) = M \cup M^{-1} \cup \{P\}$. Then, $\mathcal{A}(n)$ is an $(n/5, \Omega(1/\log(n)))$ -dimension expander.*

Proof. Let V be a subspace of \mathbb{F}^n of dimension $k \leq n/5$. Let $D \subset [n]$ denote its set of degrees (see Section 2.1) and let $D_L = D \cap \{1, \dots, n/2\}$ and $D_R = D \setminus D_L$. If $|D_L| < k/3$ then we have $\dim(V + P(V)) \geq (1 + 1/3) \cdot k$. To see this observe that in this case $|D_R| \geq (2/3) \cdot k$ and that the subspace $V + P(V)$ contains degrees both in D_R and in $D_R - (n/2)$, which are disjoint from one another.

We now proceed under the assumption that $|D_L| \geq k/3$. We denote by $s_1, \dots, s_p : \mathbb{F}^n \rightarrow \mathbb{F}^n$ the p right cyclic shifts of the first p coordinates of \mathbb{F}^n (the last coordinate remains fixed). The next claim shows that there exists j such that s_j expands V by a constant.

Claim 4.2. *There exists $j \in [p]$ such that*

$$\dim(V + s_j(V)) \geq (1 + 1/12) \cdot k.$$

Proof. Observe that it is enough to show that there exists $j \leq p/2$ for which

$$|D \cup (D_L + j)| \geq (1 + 1/12) \cdot k.$$

This is because for $j \leq p/2$ and for a vector $v \in V$ for which $\deg(v) \in D_L$ we have $\deg(s_j(v)) = \deg(v) + j$. This will follow if we show that there exists $j \leq p/2$ such that $|D \cap (D_L + j)| \leq (1/4) \cdot k$. To show this, we will consider the expectation $\mu = \mathbb{E}_j[|D \cap (D_L + j)|]$ when j is chosen uniformly

from the set $\{1, \dots, (p-1)/2\}$. We can write μ as a sum $\mu = \sum_{i \in D, j \in D_L} \chi_{i,j}$, where the $\chi_{i,j}$'s are indicator variables for the event $i \in (D_L + j)$. We note that

$$\begin{aligned} \mu &= \frac{1}{p/2} \cdot \sum_{i \in D, j \in D_L} \chi_{i,j} \\ &= \frac{2}{p} \cdot \sum_{j < i: j \in D_L, i \in D} \chi_{i,j} \leq \frac{2}{p} \cdot \binom{|D|}{2} \\ &< \frac{|D|}{p} \cdot |D|. \end{aligned}$$

Using the fact that $|D| = k \leq n/5 \leq p/4$ we get that $\mu \leq (1/4)|D| = k/4$. Therefore, there must exist $j \leq p/2$ for which $|D \cap (D_L + j)| \leq k/12$. \square

Let s_{j^*} be such that $\dim(V + s_{j^*}(V)) \geq (1 + 1/12) \cdot k$. Using Lemma 2.3 we can write

$$s_{j^*} = M_1 \circ \dots \circ M_t,$$

where for each $i \in [t]$, $M_i \in M \cup M^{-1}$ and such that $t \leq O(\log(p)) = O(\log(n))$. Now, using Lemma 2.4, we get that there exists $M_{i^*} \in M \cup M^{-1} \subset \mathcal{A}(n)$ that expands V by at least $\Omega(1/\log(n))$. \square

4.2 The case of general n

Having constructed a dimension expander as in Theorem 2 for the case $n = p + 1$, p prime, we now wish to reduce the case of general n to this case. The reduction will be rather easy and will require us to add only one more mapping on top of the mappings in the dimension expander we previously constructed.

Let n be some integer. We can find (in time polynomial in n) a prime $n/2 < p < n$. The fact that such a prime exists is known as Bertrand's Postulate and finding it in polynomial time can be done using a trivial search. The idea of the construction is to apply the dimension expander $\mathcal{A}(p+1)$ given by Theorem 4.1 on the first $p+1$ coordinates and to include one more mapping that will ensure that these coordinates contain a non-negligible part of the subspace. This additional mapping $Q_p : \mathbb{F}^n \rightarrow \mathbb{F}^n$ is defined by $Q_p(v_1, \dots, v_n) = (v_{p+2}, \dots, v_n, 0, \dots, 0)$. The next theorem, which directly implies Theorem 2, describes the above construction more formally.

Theorem 4.3. *Let \mathbb{F} be a field and let n be an integer. Let p be a prime such that $n/2 < p < n$. Let $\mathcal{A}(p+1)$ be the dimension expander given by Theorem 4.1. We treat the mappings in $\mathcal{A}(p+1)$ as acting on \mathbb{F}^n by applying them only on the first $p+1$ coordinates and leaving the last $n-p-1$ coordinates untouched. Let $\mathcal{A}'(n) = \mathcal{A}(p+1) \cup \{Q_p\}$, where Q_p is defined as above. Then, $\mathcal{A}'(n)$ is an $(n/10, \Omega(1/\log(n))$ -dimension expander.*

Proof. Let V be a subspace of \mathbb{F}^n of dimension $k \leq n/10$. Let $D \subset [n]$ be its set of degrees (see Section 2.1), let $D_L = D \cap \{1, \dots, p+1\}$ and $D_R = D \setminus D_L$. If $|D_L| \leq k/3$ then, as in the

proof of Theorem 4.1, we have that $\dim(V + Q_p(V)) \geq (1 + 1/3) \cdot k$. We can thus assume that $|D_L| \geq k/3$. Let V_L be the projection of V onto the first $p + 1$ coordinates, so that D_L is equal to the set of degrees of V_L . Using Theorem 4.1 we have that $\mathcal{A}(p + 1)$ expands V_L by a factor of $\Omega(1/\log(n))$ (here we use the fact that $k \leq n/10 \leq (p + 1)/5$). This means that the image of V under the mappings in $\mathcal{A}(p + 1)$ contains at least $\Omega(|D_L|/\log(n)) = \Omega(k/\log(n))$ linearly independent vectors that are not in V and such that these vectors are zero in the last $n - p - 1$ coordinates. This means that the image of V under $\mathcal{A}(p + 1)$ has dimension at least

$$|D_L| \cdot (1 + \Omega(1/\log(n))) + |D_R| \geq (1 + \Omega(1/\log(n))) \cdot k.$$

□

5 Acknowledgements

We thank Boaz Barak, Russel Impagliazzo and Avi Wigderson for discussions that led to the question of constructing dimension expanders and for helpful conversations. We also thank Prahladh Harsha, Alex Lubotzky, Anup Rao and Chris Umans and for helpful discussions. A. Shpilka was supported by the Israel Science Foundation (grant number 439/06).

References

- [1] N. Alon and Y. Roichman. Random cayley graphs and expanders. *Random Structures and Algorithms*, 5(2):271–285, 1994.
- [2] L. Babai, W. M. Kantor, and A. Lubotzky. Small-diameter cayley graphs for finite simple groups. *Europ. J. Combinatorics*, 10(6):507–522, 1989.
- [3] B. Barak, R. Impagliazzo, A. Shpilka, and A. Wigderson, 2004. Private communication.
- [4] J. Bourgain. On the construction of affine extractors. *Geometric And Functional Analysis*, 17(1):33–57, 2007.
- [5] A. Gabizon and R. Raz. Deterministic extractors for affine sources over large fields. In *46th Annual FOCS*, pages 407–418, 2005.
- [6] Z. Karnin and A. Shpilka. Black box polynomial identity testing of generalized depth-3 arithmetic circuits with bounded top fan-in. In *23rd Annual CCC*, 2008.
- [7] A. Lubotzky. *Discrete Groups, Expanding Graphs and Invariant Measures*. Progress in Mathematics. Birkhauser, 1994.
- [8] A. Lubotzky and Y. Zelmanov. Dimension expanders. *Journal of Algebra*, 319(2):730–738, 2008.
- [9] A. Lubotzky and A. Žuk. On property (τ) . In preperation. <http://www.ma.huji.ac.il/~alexlub/BOOKS/On%20property/On%20property.pdf>.
- [10] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes, Part II*. North-Holland, 1977.
- [11] R. Meshulam and A. Wigderson. Expanders in group algebras. *Combinatorica*, 24(4):659–680, 2004.
- [12] A. Wigderson. A lecture at IPAM, UCLA, February 2004.
- [13] A. Wigderson and D. Xiao. Derandomizing the aw matrix-valued chernoff bound using pessimistic estimators and applications, 2006.