THEORY OF COMPUTING, Volume 20 (5), 2024, pp. 1–22
www.theoryofcomputing.org

RESEARCH EXPOSITION

On the Elementary Construction of High-Dimensional Expanders by Kaufman and Oppenheim

Prahladh Harsha* Ramprasad Saptharishi*

Received June 9, 2020; Revised November 26, 2022; Published October 26, 2024

Abstract. At STOC 2018, Kaufman and Oppenheim presented an elementary construction of high-dimensional spectral expanders using elementary matrices. We give a short, self-contained elementary proof of the correctness of their construction. As a bonus, this also yields a simple construction and analysis of standard expanders of bounded degree.

ACM Classification: F.2.2, G.2.2

AMS Classification: 68R10

Key words and phrases: high-dimensional expanders, groups, cosets

1 Introduction

In the last few years, there has been a surge of activity related to high-dimensional expanders (HDXs). High-dimensional expanders are high-dimensional generalizations of classical graph

This work was done while the authors were visiting the Simons Institute for the Theory of Computing, Berkeley to participate in the summer cluster on Error-Correcting Codes and High Dimensional Expanders, 2019.

*Research of the authors supported by the Department of Atomic Energy, Government of India, under project no. 12-R&D-TFR-5.01-0500. Research of the first and second authors supported in part by the Swarnajayanti and Ramanujan Fellowships, respectively.

 DOI: 10.4086/toc.2024.v020a005

expanders. Depending on which definition of graph expansion is generalized, there are several different (and unfortunately, often inequivalent) definitions of HDXs. For the purposes of this note, we will restrict ourselves to the spectral definition of HDXs (see Definition 2.4). *Ramanujan graphs* are families of expander graphs which have optimal spectral expansion (see Definition 2.3). Lubotzky, Phillips and Sarnak [5] and Margulis [8] independently gave explicit constructions of Ramanujan graphs. *Ramanujan complexes* are high-dimensional generalizations of Ramanujan graphs. Lubotzky, Samuels and Vishne [6, 7] and Sarveniazi [12] gave explicit constructions of Ramanujan complexes, which are the high-dimensional analgoues of the construction of Lubotzky, Phillips and Sarnak [5]. These constructions were the first construction of constant-degree spectral HDXs. The Ramanujan graphs constructed by Lubotzky, Phillips and Sarnak [5] and Margulis [8] have the nice property that they are simple to describe, however the proof of the optimality of their expansion is involved. The Ramanujan complexes constructed by Lubotzky, Samuels and Vishne [6, 7] and Sarveniazi [12], on the other hand, are non-trivial to describe and it is difficult to prove their high-dimensional expansion property. Subsequently Kaufman and Oppenheim [4, See STOC version] gave an elegant elementary construction of bounded-degree spectral HDXs using elementary matrices. While their HDXs are not Ramanujan, their construction gives rise to new families of expander graphs whose spectral gap is close to the optimal Ramanujan bound.

Despite the Kaufman–Oppenheim construction being elementary and simple, the proof of expansion is not elementary. The purpose of this exposition is to give an elementary proof of the expansion of the Kaufman–Oppenheim HDX construction. In particular, we obtain the same eigenvalue bound as their proof.

The 1-skeleton (see Def. 2.1) of a HDX (even of a one-sided spectral HDX) is a two-sided spectral expander (see Def. 2.3). Thus, this construction has the added advantage that it yields an elementary construction (accompanied by a simple proof) of a standard two-sided spectral expander (though not an optimal one).

2 Preliminaries

We begin by recalling what a simplicial complex is.

Definition 2.1 (Simplicial complex)**.** A *simplicial complex* X over a finite set U is a collection of subsets of U with the property that if $S \in X$ then any $T \subseteq S$ is also in X.

- For all $i \geq -1$, define $X(i) := \{S \in X : |S| = i+1\}$. Thus, if X is non-empty, then $X(-1) = \{\emptyset\}$.
- The elements of X are called *simplices* or *faces*. The elements of $X(0)$, $X(1)$ and $X(2)$ are usually referred to as *vertices*, *edges* and *triangles*, respectively.
- For any $1 \leq k \leq d$, the *k-skeleton* of the complex X is the subcomplex $X(-1) \cup X(0) \cup X(1) \cup \cdots \cup X(k)$. We identify the *1-skeleton* with the graph defined by $X(0)$ and $X(1)$.
- The *dimension* of the simplicial complex X is defined as the largest d such that $X(d)$ (which consists of faces of size $d+1$) is non-empty.

- The simplicial complex is said to be *pure* if every face is contained in some face in $X(d)$, where $d = \dim(X)$.
- For a face $S \in X$, the *link* of S, denoted by X_S, is the simplicial complex defined as

$$X_S := \{T \setminus S \ : \ T \in X \ , \ S \subseteq T\} .$$

Thus, a graph $G = (V, E)$ is just a simplicial complex G of dimension one with $G(0) = V$ and $G(1) = E$. We will deal with *weighted* pure simplicial complexes where the weight function satisfies a certain *balance* condition.

Definition 2.2 (weighted pure simplicial complexes)**.** Given a d-dimensional pure simplicial complex X and an associated weight function $w\colon X \to \mathbb{R}_{\geq 0}$, we say the weight function is *balanced* if the following two conditions are satisfied.

$$\sum_{S \in X(d)} w(S) = 1 \ ; \qquad w(S) = \frac{1}{i+2} \sum_{T \in X(i+1), T \supset S} w(T), \ \text{ for all } i < d \text{ and } S \in X(i). \tag{2.1}$$

A *weighted simplicial complex* (X, w) is a pure simplical complex accompanied with a balanced weight function w. If no weight function is specified, then we work with the balanced weight function w induced by the uniform distribution on the set $X(d)$ of maximal faces.

For a face $S \in X$, the balanced weight function w_S associated with the link X_S is the restricted weight function, suitably normalized, more precisely $w_S := {}^{w|_{X_S}}/_{w(S)}$.

Condition (2.1) states that the weight function can be interpreted as a family of joint distributions $(w|_{X(-1)}, \ldots, w|_{X(d)})$ where $w|_{X(i)}$ is a probability distribution on $X(i)$. The distribution $w|_{X(d)}$ is specified by the first condition in (2.1) while the second condition implies that the weight distribution $w|_{X(i)}$ is the distribution on $X(i)$ obtained by picking a random $\tau \in X(d)$ according to $w|_{X(d)}$ and then removing $(d - i)$ elements uniformly at random.

We now recall the classical definition of what it means for a graph to be a spectral expander. We will be exclusively discussing only **undirected graphs.**

Definition 2.3 (spectral expander)**.** Given a weighted graph $G = (V, E, w)$ on n vertices, let A_G be its normalized adjacency matrix given as follows:

$$A_G(u, v) := \begin{cases} \frac{w(u,v)}{w(u)} & \text{if } \{u, v\} \in E, \\ 0 & \text{otherwise.} \end{cases}$$

Let $1 = \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n \geq -1$ be the n eigenvalues[1] of A_G. We denote the second largest eigenvalue of G by $\lambda(G)$.

G is said to be a λ*-spectral expander* if $\max\{\lambda_2, |\lambda_n|\} \leq \lambda$. This is sometimes also referred to as a λ-two-sided spectral expander.

G is said to be a λ*-one-sided spectral expander* if $\lambda_2 \leq \lambda$.

[1]By the balance condition, w satisfies $w(v) = \sum_{\{u,v\} \in E} w(u, v)$. The matrix A_G is self-adjoint with respect to the inner product $\langle f, g \rangle_w := \mathbb{E}_{v \sim w}[f(v)g(v)]$ since $\langle f, Ag \rangle_w = \langle Af, g \rangle_w = \mathbb{E}_{\{u,v\} \sim w}[f(u)g(v)]$. Hence, A_G has n real eigenvalues which can be obtained using the Courant–Fischer Theorem (Theorem A.1).

This spectral definition of expanders is generalized to higher-dimensional simplicial complexes as follows.

Definition 2.4 (λ-spectral HDX)**.** A weighted simplicial complex (X, w) of dimension $d \geq 1$ is said to be a *λ-spectral HDX* (or a *λ-two-sided spectral HDX*)[2] if for every $-1 \leq i \leq d-2$ and $s \in X(i)$, the weighted 1-skeleton of the link (X_s, w_s) is a λ-spectral expander.

A weighted simplicial complex (X, w) of dimension $d \geq 1$ is said to be a *λ-one-sided spectral HDX* if for every $-1 \leq i \leq d-2$ and $s \in X(i)$, the weighted 1-skeleton of the link (X_s, w_s) is a λ-one-sided spectral expander.

Using Garland's technique [2], Oppenheim [10] showed that if the 1-skeletons of all the links are connected, then a spectral gap at dimension $(d-2)$ descends to all lower levels.

Descent Theorem 2.5 (Oppenheim, 2018)**.** *Let (X, w) be a d-dimensional weighted simplicial complex with the following properties.*

- *For all $s \in X(d-2)$, the link (X_s, w_s) is a λ-one-sided spectral expander for some $\lambda < \frac{1}{d-1}$.*
- *The 1-skeleton of every link is connected.*

Then, (X, w) is a $\left(\frac{\lambda}{1-(d-1)\lambda}\right)$-one-sided spectral HDX.

Thus to prove that a given simplicial complex is a spectral HDX, it suffices to show that the 1-skeleton of every link is connected and to show a spectral gap at the top level. For the sake of completeness, we give a proof of the Descent Theorem in Appendix A which includes a descent theorem for the least eigenvalue as well.

3 Coset complexes

The HDX construction of Kaufmann and Oppenheim is a particular instantiation of a certain type of simplicial complex called a *coset complex* based on a group and its subgroups. In this section, we give an exposition of these objects. For a basic primer on group theory, see Section B.

Definition 3.1 (coset complex)**.** Let G be a group and let $K_1, \ldots, K_d$ be d subgroups of G. The coset complex $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ is a $(d-1)$-dimensional simplicial complex defined as follows:

- The set of *vertices*, $\mathcal{X}(0)$, consists of the left cosets of $K_1, \ldots, K_d$ and we shall say the left cosets of K_i are of *type* i.
- The set of *maximal faces*, $\mathcal{X}(d-1)$, consists of the d-sets of cosets of different types with a non-empty intersection. That is,

$$\{g_1 K_1, \ldots, g_d K_d\} \in \mathcal{X}(d-1) \iff g_1 K_1 \cap \cdots \cap g_d K_d \neq \emptyset.$$

[2] These are sometimes also referred to as λ-link HDXs or λ-local-expanders to distinguish from an alternative global definition of high-dimensional expansion.

An equivalent way of stating this is that $\{g_1K_1, \ldots, g_dK_d\} \in \mathcal{X}(d-1)$ if and only if there is some $g \in G$ such that $g_iK_i = gK_i$ for all i, since

$$g_iK_i = gK_i \iff K_i = g_i^{-1}gK_i \iff g_i^{-1}g \in K_i \iff g \in g_iK_i.$$

- The lower-dimensional faces are obtained by *down-closing* the maximal faces. Hence, for $0 \le r \le d$, $\{g_{i_1}K_{i_1}, \ldots, g_{i_r}K_{i_r}\} \in \mathcal{X}(r-1)$ if and only if $i_j \neq i_k$ for all $j \neq k$ and

  $$g_{i_1}K_{i_1} \cap \cdots \cap g_{i_r}K_{i_r} \neq \emptyset.$$

  We shall call the set $\{i_1, \ldots, i_r\}$ the *type* of this face.

- The dimension of this complex is $d-1$.

- The weight function we will use is the one induced by the uniform distribution on the set $\mathcal{X}(d-1)$ of maximal faces.

A simplicial complex constructed this way is *partite* in the sense that each maximal face consists of vertices of distinct types.

It follows from the definition, that $\mathcal{X}(i)$ is precisely the set of cosets of the form gK_S where $K_S = \bigcap_{j\in S} K_j$ for sets $S \subseteq [d]$ of size exactly $i+1$. In particular, $\mathcal{X}(d-1)$, the set of maximal faces, is in 1-1 correspondence with the group G if $\bigcap_{j\in[d]} K_j = \{\mathrm{id}\}$ where "id" is the identity element of the group G.

Connectivity

Observation 3.2. $g_1K_1 \cap g_2K_2 \neq \emptyset$ if and only if $g_1^{-1}g_2 \in K_1K_2$.

Proof. ($\Rightarrow$) Say $x = g_1k_1 = g_2k_2$ for $k_1 \in K_1$ and $k_2 \in K_2$. Then $g_1^{-1}g_2 = g_1^{-1}x \cdot x^{-1}g_2 = k_1k_2^{-1} \in K_1K_2$.

($\Leftarrow$) If $g_1^{-1}g_2 = k_1k_2$ for some $k_1 \in K_1$ and $k_2 \in K_2$, then $g_1k_1 = g_2k_2^{-1} \in g_1K_1 \cap g_2K_2$. □

Lemma 3.3 (Criterion for connected 1-skeletons)**.** *The 1-skeleton of $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ is connected if and only if $G = \langle K_1, \ldots, K_d\rangle$.*

Proof. ($\Leftarrow$) Since there is always an edge between gK_i and gK_j for $i \neq j$, it suffices to show that K_1 is connected to gK_1 for an arbitrary $g \in G$. Suppose, for an arbitrary element $g \in G$, we have $g = g_1 \ldots g_r$ where $g_j \in K_{i_j}$ and $i_j \neq i_{j+1}$ for each j. We might, without loss of generality, assume that (a) $g_1 \in K_1$ (otherwise set $g = 1 \cdot g_1 \cdots g_r$) and (b) if $r \ge 2$, then $i_r \neq 1$ (since otherwise we might then have worked with $g' = g_1g_2 \ldots g_{r-1}$ as $gK_1 = g'g_rK_1 = g'K_1$).

Then, we get the following path connecting K_1 and gK_{i_r}

$$K_1 = g_1K_{i_1} \to (g_1g_2)K_{i_2} \to (g_1g_2g_3)K_{i_3} \to \ldots \to (g_1 \cdots g_r)K_{i_r} = gK_{i_r}.$$

Note that, due to Observation 3.2, each successive pair of cosets are connected by an edge in the simplicial complex. Now, since gK_{i_r} is adjacent to gK_1 (as $i_r \neq 1$), we have that K_1 is connected to gK_1.

($\Rightarrow$) For an arbitrary $g \in G$, since the 1-skeleton is connected we have a path

$$K_1 = g_0 K_{i_0} \to g_1 K_{i_1} \to \cdots \to g_r K_{i_r} = gK_1.$$

By Observation 3.2, for every $j = 0, \ldots, r-1$, we have $g_j^{-1} g_{j+1} \in K_{i_j} K_{i_{j+1}} \in \langle K_1, \ldots, K_d \rangle$. Therefore,

$$g = (g_0^{-1} g_1) \cdot (g_1^{-1} g_2) \cdots (g_{r-1}^{-1} g_r) \in \langle K_1, \ldots, K_d \rangle . \qquad \square$$

Structure of links of the coset complex

For any set $S \subseteq [d]$, define the group $K_S := \bigcap_{i \in S} K_i$; let $K_\emptyset := \langle K_1, \ldots, K_d \rangle$. The following lemma shows that the links of a coset complex are themselves coset complexes.

Lemma 3.4. *For any $v \in \mathcal{X}(k)$ of type $S \subseteq [d]$, the link X_v is isomorphic to the simplicial complex defined by $\mathcal{X}(K_S, \{K_S \cap K_i \ : \ i \notin S\})$.*

Proof. It suffices to prove this lemma for $v \in \mathcal{X}(0)$ as links of higher levels can be obtained by inductive applications of this case.

Observe that if g is any element of G, then $(g_{i_1} K_{i_1}, \ldots, g_{i_r} K_{i_r}) \in \mathcal{X}(r-1)$ if and only if $(g g_{i_1} K_{i_1}, \ldots, g g_{i_r} K_{i_r}) \in \mathcal{X}(r-1)$. Therefore, the link of a coset gK_i is isomorphic to the link of the coset K_i. Thus, it suffices to prove the lemma for links of the type X_{K_i} for some $i \in [d]$.

Let v be the coset K_1, without loss of generality. The *vertices* of the link, $X_v(0)$, are cosets of $K_2, \ldots, K_d$ that have a non-empty intersection with K_1. Note that any non-empty intersection $g_j K_j \cap K_1$ of a coset with K_1 is itself a coset $\tilde{g}_j(K_j \cap K_1)$ of the intersection subgroup $K_j \cap K_1$ in K_1. Indeed, suppose that $g_j h_j \in K_1$ for some $h_j \in K_j$. Then, $g_j h_j K_j = g_j K_j$ and $g_j h_j K_1 = K_1$ and hence

$$g_j K_j \cap K_1 = g_j h_j K_j \cap g_j h_j K_1 = g_j h_j (K_j \cap K_1).$$

Therefore, the vertices of the link $X_v(0)$ are in bijective correspondence with the cosets of the subgroups $\left\{K_j \cap K_1 \ : \ j \in \{2, \ldots, d\}\right\}$.

The maximal faces in X that contain the coset K_1 are precisely the d-sets $\{K_1, g_2 K_2, \ldots, g_d K_d\}$ of cosets with a non-empty intersection and hence

$$\emptyset \neq K_1 \cap g_2 K_2 \cap \cdots \cap g_d K_d = (g_2 K_2 \cap K_1) \cap \cdots \cap (g_d K_d \cap K_1) = \tilde{g}_2 (K_2 \cap K_1) \cap \cdots \cap \tilde{g}_d (K_d \cap K_1),$$

which are precisely the maximal faces of the coset complex $\mathcal{X}\left(K_1, \left\{K_j \cap K_1 \ : \ j \in \{2, \ldots, d\}\right\}\right)$. This establishes the isomorphism between X_v and $\mathcal{X}\left(K_1, \left\{K_j \cap K_1 \ : \ j \in \{2, \ldots, d\}\right\}\right)$. $\square$

4 A concrete instantiation

The simplicial complex of Kaufman and Oppenheim [4, See STOC version] is a specific instantiation of the above *coset complex* construction. This section is devoted to an exposition of this instantiation of Kaufman and Oppenheim. We will need some notation to describe their group.

Notation

- Let p be a prime power and consider the ring $\mathbb{F}_p[t]$ of polynomials over the finite field $\mathbb{F}_p$. Let R denote the quotient ring $\mathbb{F}_p[t]/\langle t^s \rangle$ where $\langle t^s \rangle$ denotes the ideal generated by t^s. This is a ring whose elements can be identified with polynomials in $\mathbb{F}_p[t]$ of degree less than s (where addition and multiplication are performed modulo t^s), with t being a formal variable and s a positive integer[3]. By increasing the value of s, the construction will provide a family of complexes on a growing number of vertices.

- For any $d \geq 3$, and $1 \leq i, j \leq d$ with $i \neq j$ and an element $r \in R$, we define $e_{i,j}(r)$ to be the $d \times d$ elementary matrix with 1's on the diagonal and r on the (i, j)-th entry.

  For the sake of notational convenience, we shall extend this notation and write $e_{k,\ell}(r) := e_{i,j}(r)$ for all $k, \ell \in \mathbb{Z}$ such that $k \equiv i \pmod{d}$ and $\ell \equiv j \pmod{d}$ $(1 \leq i, j \leq d)$. For example, $e_{d,d+1}(r)$ refers to $e_{d,1}(r)$.

  We will extend further and use $\{i, i+1, \ldots, j-1\}_{\text{mod } d}$ to denote the set $i, i+1, \cdots j-1$ when $i < j$, and the set $\{i, i+1, \ldots d, 1, 2, \ldots, j-1\}$ when $j \leq i$.

We are now ready to describe the groups in the construction.

$$\text{For } i \in \{1, \ldots, d\}, \quad K_i = \left\langle e_{j,j+1}(at+b) \ : \ a, b \in \mathbb{F}_p \, , \, j \in [d] \setminus \{i\} \right\rangle .$$
$$G = \langle K_1, \ldots, K_d \rangle$$

Each K_i is generated by elementary matrices that have 1's on the diagonal and an arbitrary linear polynomial in one entry of the generalised diagonal $\{(i, j) \ : \ i+1 \equiv j \pmod{d}\}$.

It so happens that the group G generated by the subgroups $K_1, \ldots, K_d$ is $\mathrm{SL}_d(R)$, the group of $d \times d$ matrices with entries in R whose determinant is 1 (in R). This is a non-trivial fact (see [3, Theorem 4.3.9]). All we will need is the simpler fact that $|G|$ grows exponentially with s (for fixed p and d) while the size of the groups K_i are functions of p and d (and independent of s). This will follow from the sequence of observations and lemmas developed in the following section.

Given the above definition, there are two "different" subgroups we can define.

$$K_S = \bigcap_{i \in S} K_i ,$$
$$\widetilde{K_S} := \left\langle e_{i,i+1}(at+b) \ : \ a, b \in \mathbb{F}_p \, , \, i \notin S \right\rangle .$$

That is, K_S is the intersection of the groups $\{K_i \ : \ i \in S\}$, and $\widetilde{K_S}$ is the group generated by the intersection of the sets of generators of the K_i. Thus, clearly, $\widetilde{K_S} \subseteq K_S$. The following lemma shows that in fact the two groups are identical.

[3]In the subsequent journal version [4], the authors generalise the ring R to any unital commutative ring. For the purposes of this exposition, we focus exclusively on finite rings of the form described above.

Lemma 4.1 (Intersections of K_i's)**.** *For any $S \subseteq [d]$,*

$$\widetilde{K_S} = \left\langle e_{i,i+1}(at+b) \;:\; a, b \in \mathbb{F}_p \,,\, i \notin S \right\rangle = \bigcap_{i \in S} K_i = K_S$$

In other words, the group generated by the intersection of generators equals the group intersection.

We will prove this lemma in the following section by giving an explicit description of the groups that makes the above lemma evident. An immediate consequence of this lemma is that $K_{[d]} = \{\mathrm{id}\}$ and hence $\mathcal{X}(d-1)$, the set of maximal faces, is in 1-1 correspondence with the group G.

4.1 Explicit description of the groups

The following is an easy consequence of the definition of $e_{i,j}(r)$. Note that $e_{i,j}(r)$ is defined only if $i \neq j$.

Observation 4.2. (a) *Sum:* $e_{i,j}(r_1) \cdot e_{i,j}(r_2) = e_{i,j}(r_1 + r_2)$.

As a corollary, $e_{i,j}(r)^{-1} = e_{i,j}(-r)$.

(b) *Product:* If $i \neq \ell$, the commutator[4] $[e_{i,j}(r_1), e_{k,\ell}(r_2)]$ behaves as follows.

$$[e_{i,j}(r_1), e_{k,\ell}(r_2)] = \begin{cases} e_{i,\ell}(r_1 r_2) & \text{if } j = k, \\ \mathrm{id} & \text{if } j \neq k. \end{cases}$$

Proof. Let $\mu_{i,j}$ denote the matrix that has a 1 at the (i,j)-th entry, and 0 everywhere. Then, (a) follows as

$$\begin{aligned} e_{i,j}(r_1) e_{i,j}(r_2) &= (I + \mu_{i,j} r_1) \cdot (I + \mu_{i,j}(r_2)) \\ &= I + \mu_{i,j} \cdot (r_1 + r_2) \qquad (\text{since } \mu_{i,j}^2 = 0 \text{ when } i \neq j). \end{aligned}$$

As for (b), we follow along a similar calculation. Note that

$$\mu_{i,j} \cdot \mu_{k,\ell} = \begin{cases} 0 & \text{if } j \neq k \\ \mu_{i,\ell} & \text{if } j = k. \end{cases}$$

Therefore,

$$[e_{i,j}(r_1), e_{k,\ell}(r_2)] = (I - \mu_{i,j} r_1) \cdot (I - \mu_{k,\ell} r_2) \cdot (I + \mu_{i,j} r_1) \cdot (I + \mu_{k,\ell} r_2).$$

When $j = k$ (along with the assumption that $i \neq j$, $k \neq \ell$), this simplifies to

$$\begin{aligned} [e_{i,j}(r_1), e_{k,\ell}(r_2)] &= (I - \mu_{i,j} r_1 - \mu_{k,\ell} r_2 + \mu_{i,\ell} \cdot r_1 r_2) \cdot (I + \mu_{i,j} r_1 + \mu_{k,\ell} r_2 + \mu_{i,\ell} \cdot r_1 r_2) \\ &= I + \mu_{i,j}(r_1 - r_1) + \mu_{k,\ell}(r_2 - r_2) + \mu_{i,\ell}(r_1 r_2 + r_1 r_2 - r_1 r_2) \\ &= I + r_1 r_2 \mu_{i,\ell}. \end{aligned}$$

[4]The commutator of two elements g, h, denoted by $[g, h]$, is defined as $g^{-1} h^{-1} g h$. (Definition B.3)

If $j \neq k$, then we get

$$\begin{aligned}[e_{i,j}(r_1), e_{k,\ell}(r_2)] &= (I - \mu_{i,j} r_1) \cdot (I - \mu_{k,\ell} r_2) \cdot (I + \mu_{i,j} r_1) \cdot (I + \mu_{k,\ell} r_2). \\ &= I + \mu_{i,j}(r_1 - r_1) + \mu_{k,\ell}(r_2 - r_2) = I\end{aligned}$$

□

Therefore, for distinct $i, j, k \in [d]$ (which exist when $d \geq 3$), we have

$$\left[e_{i,j}(r_1), \left[e_{j,i}(r_2), e_{i,k}(r_3)\right]\right] = e_{i,k}(r_1 r_2 r_3)$$

Thus, for all $d \geq 3$, using the above observation along with Observation 4.2(a), we get that $e_{i,j}(r)$ for any $r \in R$ can be generated by $\left\{e_{k,\ell}(at+b) \ : \ k, \ell \in [d]\ ,\ a, b \in \mathbb{F}_p \right\}$. This in particular implies that $|G|$ is at least p^s. On the other hand, the size of K_i depends only on d, p and is independent of s. The lemma below describes K_d; the other K_i are just rearrangements of rows and columns in K_d.

Lemma 4.3 (Explicit description of K_d)**.** *The group $K_d = \left\langle e_{i,i+1}(at+b) \ : \ a, b \in \mathbb{F}_p\ ,\ i \neq d \right\rangle$ consists of the matrices $A = (A_{i,j})$ of the following form:*

$$A_{i,j} = \begin{cases} 1 & \textit{if } i = j, \\ \textit{a polynomial of degree} \leq j - i & \textit{if } i < j, \\ 0 & \textit{if } i > j. \end{cases}$$

Stating the above differently, for any $n \in [d]$, the group $K_n = \left\langle e_{i,i+1}(at+b) \ : \ a, b \in \mathbb{F}_p\ ,\ i \neq n \right\rangle$ consists of the matrices $A = (A_{i,j})$ of the following form:

$$A_{i,j} = \begin{cases} 1 & \textit{if } i = j, \\ \textit{a polynomial of degree } (j - i) \bmod d & \textit{if } j \neq i \textit{ and } n \notin \{i, i+1, \ldots, j-1\}_{\bmod d}\,, \\ 0 & \textit{otherwise}. \end{cases}$$

Proof. Follows by repeated application of Observation 4.2. □

Therefore, we can obtain a crude bound of $|K_i| \leq p^{O(d^3)}$ for any i. Also, the above lemma also gives an explicit description of the groups K_S.

Corollary 4.4 (Explicit description of K_S)**.** *For any subset $S \subseteq [d]$, the group $K_S = \bigcap_{i \in S} K_i$ consists of matrices $A = (A_{i,j})$ of the following form:*

$$A_{i,j} = \begin{cases} 1 & \textit{if } i = j, \\ \textit{a polynomial of degree } (j - i) \bmod d & \textit{if } j \neq i \textit{ and } \{i, i+1, \ldots, j-1\}_{\bmod d} \cap S = \emptyset, \\ 0 & \textit{otherwise}. \end{cases}$$

Recall the *other* set of subgroups defined for each $S \subseteq [d]$:

$$\widetilde{K_S} := \left\langle e_{i,i+1}(at+b) \ : \ a, b \in \mathbb{F}_p\ ,\ i \notin S \right\rangle.$$

These groups can also be explicitly described.

Lemma 4.5 (Explicit description of $\widetilde{K_S}$)**.** *For any $\emptyset \neq S \subseteq [d]$, the group $\widetilde{K_S}$ is the set of all $d \times d$ matrices $A = (a_{ij})$ of the form*

$$a_{i,j} = \begin{cases} 1 & \text{if } i = j, \\ \text{a polynomial of degree} \leq j - i & \text{if } j \neq i \text{ and } \{i, i+1, \ldots, j-1\}_{\bmod d} \cap S = \emptyset, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Any $A \in \widetilde{K_S}$ can be expressed as $A = B_1 \cdots B_m$ where each $B_r = e_{i_r, i_r+1}(\ell_r)$, for some linear polynomial ℓ_r, with $i_r \notin S$. Then,

$$A_{i,j} = \sum_{\substack{i_1, \ldots, i_{m+1} \\ i_1 = i\ ,\ i_{m+1} = j}} (B_1)_{i_1, i_2} (B_2)_{i_2, i_3} \cdots (B_m)_{i_m, i_{m+1}}.$$

From the structure of each B_r, any nonzero contribution from the RHS must involve either $i_{r+1} = i_r$, or $i_{r+1} = i_r + 1$ if $r \notin S$. This forces that the only entries of A that are nonzero, besides the diagonal, are at (i, j) with none of $\{i, i+1, \ldots, j-1\}$ in S.

In the case when $\{i, i+1, \ldots, j-1\} \cap S = \emptyset$, the above argument also shows that the entry $A_{i,j}$ has degree at most $j - i$. Furthermore, using Observation 4.2, we can easily see that $e_{i,j}(f) \in \widetilde{K_S}$ for an arbitrary polynomial $f(t)$ of degree at most $j - i$. From this, we can deduce that the structure of $\widetilde{K_S}$ is exactly as claimed. ◻

Proof of Lemma 4.1. Follows immediately from Corollary 4.4 and Lemma 4.5. ◻

From this point on, since the groups K_S and $\widetilde{K_S}$ are identical, we drop the tilde notation and use K_S for $\widetilde{K_S}$.

4.2 Connectivity of the coset complex

Lemma 4.6. *Let $S \subset [d]$ with $|S| \leq d - 2$. Then,*

$$K_S = \langle K_S \cap K_i \colon i \in [d] \setminus S \rangle \,.$$

Proof. It is clear that K_S is a superset of the RHS. It only remains to show that the other containment also holds. To see this, consider an arbitrary generator $e_{j,j+1}(r)$ of K_S. Since $j \notin S$ and $|S| \leq d - 2$, there is some $i \in [d] \setminus (S \cup \{j\})$. Therefore, $e_{j,j+1}(r) \in K_S \cap K_i$ and hence is generated by the RHS. ◻

Combining the above lemma with Lemma 3.3 and Lemma 3.4, we have the following corollary.

Corollary 4.7. *For the coset complex $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ defined by the above groups, the 1-skeleton of every link is connected.*

5 Spectral expansion of the complex

In this section we prove that the coset complex $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ is a good spectral HDX. The Descent Theorem (Theorem 2.5) states that it suffices to show that the 1-dimensional links of faces in $\mathcal{X}(d-3)$ are good spectral expanders.

5.1 Structure of 1-dimensional links

One-dimensional links of the coset complex constructed are links of $v \in \mathcal{X}(G, \{K_1, \ldots, K_d\})$ of size exactly $d-2$ (which are elements of $\mathcal{X}(d-3)$). Any such v can be written as $\{gK_1, \ldots, gK_d\} \setminus \left\{gK_i, gK_j\right\}$ for $i, j \in [d]$ with $i \neq j$ and $g \in G$. Since the link of v is isomorphic to the link of $\{K_1, \ldots, K_d\} \setminus \left\{K_i, K_j\right\}$, we might as well assume that $g = \mathrm{id}$. These happen to be of two types depending on whether i and j are consecutive or not.

Observation 5.1. Consider $v = \{K_1, \ldots, K_d\} \setminus \left\{K_i, K_j\right\}$ where i and j are *not* consecutive (i. e., $(i-j) \neq \pm 1 \bmod d$). Then the 1-dimensional link of v is a complete bipartite graph.

Proof. Note that since $j \neq i \pm 1$, we have $[e_{i,i+1}(r_1), e_{j,j+1}(r_2)] = \mathrm{id}$ by Observation 4.2. Hence, these two elements commute.

The link of v corresponds to the coset complex $\mathcal{X}(H, \{H_1, H_2\})$ where

$$\begin{aligned}
H &= K_{[d]\setminus\{i,j\}} = \left\langle e_{i,i+1}(at+b), e_{j,j+1}(at+b) \ : \ a, b \in \mathbb{F}_p \right\rangle , \\
H_1 &= K_{[d]\setminus\{i\}} = \left\langle e_{i,i+1}(at+b) \ : \ a, b \in \mathbb{F}_p \right\rangle , \\
H_2 &= K_{[d]\setminus\{j\}} = \left\langle e_{j,j+1}(at+b) \ : \ a, b \in \mathbb{F}_p \right\rangle .
\end{aligned}$$

Thus, the groups H_1 and H_2 commute with each other and hence any element of $h \in H$ can be written as $h = g_1 \cdot g_2$ where $g_1 \in H_1$ and $g_2 \in H_2$. Observation 3.2 implies that the resulting graph is the complete bipartite graph. ◻

The interesting case is when $v = \{K_1, \ldots, K_d\} \setminus \{K_i, K_{i+1}\}$. Without loss of generality, we may focus on the link of $v = \{K_3, K_4, \ldots, K_d\}$. This corresponds to the coset complex $\mathcal{X}(H, \{H_1, H_2\})$ where

$$\begin{aligned}
H &= K_{3,4,\ldots,d} = \left\langle e_{1,2}(at+b), e_{2,3}(at+b) \ : \ a, b \in \mathbb{F}_p \right\rangle , \\
H_1 &= K_{2,3,4,\ldots,d} = \left\langle e_{1,2}(at+b) \ : \ a, b \in \mathbb{F}_p \right\rangle , \\
H_2 &= K_{1,3,4,\ldots,d} = \left\langle e_{2,3}(at+b) \ : \ a, b \in \mathbb{F}_p \right\rangle .
\end{aligned}$$

Hence, it suffices to focus on the first three rows and columns of these matrices as the rest of them are constant. Written down explicitly,

$$H = \left\{ \begin{bmatrix} 1 & \ell_1 & Q \\ 0 & 1 & \ell_2 \\ 0 & 0 & 1 \end{bmatrix} \;:\; \begin{matrix} \ell_1, \ell_2 \text{ are linear polynomials in } \mathbb{F}_p[t] \\ \text{and } Q \text{ is a quadratic polynomial in } \mathbb{F}_p[t] \end{matrix} \right\},$$

$$H_1 = \left\{ \begin{bmatrix} 1 & \ell & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \;:\; \ell \text{ is a linear polynomial in } \mathbb{F}_p[t] \right\},$$

$$H_2 = \left\{ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & \ell \\ 0 & 0 & 1 \end{bmatrix} \;:\; \ell \text{ is a linear polynomial in } \mathbb{F}_p[t] \right\}.$$

Multiplication of an arbitrary element of H with an arbitrary element of H_1 is of the form

$$\begin{bmatrix} 1 & \ell_1 & Q \\ 0 & 1 & \ell_2 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & \ell & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & \ell_1 + \ell & Q \\ 0 & 1 & \ell_2 \\ 0 & 0 & 1 \end{bmatrix}.$$

Note that the a unique choice of ℓ that makes the $(1,2)$-th entry of the RHS zero is $\ell = -\ell_1$. Thus, each coset of H_1 in H has a unique representative of the form $M_1(\ell, Q)$ described below, and similarly, each coset of H_2 in H has a unique representative of the form $M_2(\ell, Q)$:

$$M_1(\ell, Q) := \begin{bmatrix} 1 & 0 & Q \\ 0 & 1 & \ell \\ 0 & 0 & 1 \end{bmatrix} \quad , \quad M_2(\ell, Q) := \begin{bmatrix} 1 & \ell & Q \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

respectively, where ℓ is a linear polynomial and Q is a quadratic polynomial in $\mathbb{F}_p[t]$. This is because any element of H can be uniquely written as

$$\begin{aligned} \begin{bmatrix} 1 & \ell_1 & Q \\ 0 & 1 & \ell_2 \\ 0 & 0 & 1 \end{bmatrix} &= \begin{bmatrix} 1 & \ell_1 & Q - \ell_1\ell_2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & \ell_2 \\ 0 & 0 & 1 \end{bmatrix} \\ &= \begin{bmatrix} 1 & 0 & Q - \ell_1\ell_2 \\ 0 & 1 & \ell_2 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & \ell_1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \end{aligned}$$

Lemma 5.2. *For linear polynomials $\ell_1, \ell_2 \in \mathbb{F}_p[t]$ and quadratic polynomials $Q_1, Q_2 \in \mathbb{F}_p[t]$, we have that*

$$M_1(\ell_1, Q_1)H_1 \cap M_2(\ell_2, Q_2)H_2 \neq \emptyset \quad \Longleftrightarrow \quad \ell_1\ell_2 = Q_1 - Q_2.$$

Proof. Note that matrices in H_1H_2 are of the form

$$\begin{bmatrix} 1 & \ell_1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & \ell_2 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & \ell_1 & \ell_1\ell_2 \\ 0 & 1 & \ell_2 \\ 0 & 0 & 1 \end{bmatrix}.$$

By Observation 3.2, the cosets have a non-empty intersection if and only if

$$H_1H_2 \ni M_1(\ell_1, Q_1)^{-1}M_2(\ell_2, Q_2) = \begin{bmatrix} 1 & 0 & -Q_1 \\ 0 & 1 & -\ell_1 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & \ell_2 & Q_2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & \ell_2 & Q_2 - Q_1 \\ 0 & 1 & -\ell_1 \\ 0 & 0 & 1 \end{bmatrix}$$

which happens if and only if $(\ell_2)(-\ell_1) = Q_2 - Q_1$ which is the same as $Q_1 - Q_2 = \ell_1\ell_2$. □

Therefore, the 1-dimensional link is the bipartite graph $A = (U, V, E)$ with left and right vertices identified by pairs (ℓ, Q) where ℓ and Q are linear and quadratic polynomials in $\mathbb{F}_p[t]$, respectively, with $(\ell_1, Q_1) \sim (\ell_2, Q_2) \Leftrightarrow \ell_1\ell_2 = Q_1 + Q_2$ (by associating $M_1(\ell, Q)$ with the tuple (ℓ, Q) on the left, and $M_2(\ell, Q)$ with the tuple $(\ell, -Q)$ on the right).

Note that A is a p^2-regular bipartite graph with p^5 vertices on each side. It suffices to show that A is a good expander.

Kaufman and Oppenheim [4] prove the expansion properties of this graph using the representation theory of the associated groups, while we directly analyse the spectral gap of the adjacency matrix associated with this graph. O'Donnell and Pratt [9, Case 2 in the Proof of Theorem 3.23] give yet another proof of the spectral gap using the Polynomial Identity Lemma (also referred to as the Schwartz–Zippel lemma).

5.2 A related graph

The following graph is the "lines-points" or the "affine plane" graph used by Reingold, Vadhan and Wigderson [11] (as the *base graph* in their construction of constant-degree expanders, using the zig-zag product). Let $\mathbb{F}_q$ be a finite field. Consider the bipartite graph $B_q = (U', V', E')$ defined as follows:

$$U' = V' = \mathbb{F}_q \times \mathbb{F}_q \qquad\qquad E' = \{((a, b), (c, d)) \ : \ ac = b + d\}\,.$$

Note that the graph B_q is q-regular as for any $a, b, c \in \mathbb{F}_q$, there is a unique $d \in \mathbb{F}_q$ such that $ac = b + d$ and thus the vertex (a, b) has exactly q neighbours in B_q.

Lemma 5.3 ([11, Lemma 5.1])**.** *The q-regular bipartite graph B_q is a $\frac{1}{\sqrt{q}}$-one-sided spectral expander.*

Proof. Let B_q^2 denote the multigraph whose adjacency matrix is the square of the adjacency matrix of B_q. Note that B_q^2 is a multigraph as each edge in B_q^2 corresponds to a length-two path in B_q and there may be more than one such path between a pair of vertices. It is easy to see that

$$\text{Number of edges between } (a, b) \text{ and } (c, d) \text{ in } B_q^2 = \begin{cases} 1 & \text{if } a \neq c, \\ q & \text{if } a = c \text{ and } b = d, \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, the adjacency matrix of B_q^2 can be written[5] (under a suitable order of listing vertices) as

$$qI_{q^2} + (J_q - I_q) \otimes J_q \quad \text{(where } J_q \text{ is the } q \times q \text{ matrix of 1s).}$$

By observing that J_q has eigenvalue of q with multiplicity 1, and eigenvalue 0 with multiplicity $(q-1)$, a simple calculation shows that B_q^2 has eigenvalue of q^2 with multiplicity 1, eigenvalue q with multiplicity $q(q-1)$ and eigenvalue 0 with multiplicity $q-1$. Hence the unnormalized second largest eigenvalue of B_q^2 is q and hence we have that the normalized second largest eigenvalue of B_q is $1/\sqrt{q}$. □

5.3 Relating the graph B_q with A

Set $q = p^3$ so that $\mathbb{F}_q = \mathbb{F}_p[y]/\langle \mu(y) \rangle$ for some irreducible polynomial $\mu(y)$ of degree 3. Therefore, each element in $\mathbb{F}_q$ is expressible as $a_0 + a_1 y + a_2 y^2$ for some $a_0, a_1, a_2 \in \mathbb{F}_p$. Thus, the graph $B_q = (U', V', E')$ defined above, for $q = p^3$, is a p^3-regular bipartite graph with p^6 vertices on either side.

Let $U'' = V'' = \left\{(a_0 + a_1 y, b_0 + b_1 y + b_2 y^2) \ : \ a_0, a_1, b_0, b_1, b_2 \in \mathbb{F}_p\right\}$, which is a subset of U' and V', respectively, of size p^5 each.

Observation 5.4. The induced subgraph of B_q on U'', V'' is exactly the graph $A = (U, V, E)$ described earlier.

Proof. Note that $((\ell_1(y), Q_1(y)), (\ell_2(y), Q_2(y))) \in E'$ if and only if

$$\ell_1(y) \cdot \ell_2(y) = Q_1(y) + Q_2(y) \bmod \mu(y).$$

However, since the above equation has degree at most 2, we have

$$\ell_1(y) \cdot \ell_2(y) = Q_1(y) + Q_2(y) \quad \Leftrightarrow \quad \ell_1(y) \cdot \ell_2(y) = Q_1(y) + Q_2(y) \pmod{\mu(y)},$$

and the first equation is exactly the adjacency condition of the graph A. Hence, the induced subgraph of B_q on U'', V'' is indeed the graph A. □

Normally, induced subgraphs of expanders need not even be connected. However, the following lemma shows that there are some instances where we may be able to give non-trivial bounds on λ.

Lemma 5.5. *Let X be a d-regular graph that is an induced subgraph of a D-regular graph Y. Then,*

$$\lambda(X) \le \frac{D\lambda(Y)}{d}.$$

[5] In the equation, the notation $\otimes$ refers to the Kronecker product, or tensor product of matrices. It is well known that, for square matrices A and B, the multiset of eigenvalues of $A \otimes B$ is all products of the form $\lambda_i \cdot \nu_j$ where λ_i is an eigenvalue of A and ν_j is an eigenvalue of B.

Proof. The Courant–Fischer characterization of the second largest eigenvalue (Theorem A.1) tells us that $\lambda(X) = \max_{a \perp \mathbb{1}_{|X|}} \frac{a^T G a}{d \cdot a^T a}$ where G is the adjacency matrix of X. Consider an arbitrary $a \in \mathbb{R}^{|X|}$ such that $a \perp \mathbb{1}_{|X|} = 0$. Since X is an induced subgraph of Y, the vector a can be padded with zeroes to obtain a vector $b_a \in \mathbb{R}^{|Y|}$ such that $b_a \perp \mathbb{1}_{|Y|}$. Therefore, if A_X and A_Y are the normalised adjacency matrices of X and Y, we have

$$\lambda(X) = \max_{a \perp \mathbb{1}_{|X|}} \frac{a^T A_X a}{a^T a} = \frac{D}{d} \cdot \max_{a \perp \mathbb{1}_{|X|}} \frac{b_a^T A_Y b_a}{b_a^T b_a} \le \frac{D}{d} \cdot \max_{b \perp \mathbb{1}_{|Y|}} \frac{b^T A_Y b}{b^T b} = \frac{D\lambda(Y)}{d} \quad . \qquad \square$$

Corollary 5.6. *The graph $A(U, V, E)$ corresponding to the 1-dimensional links of $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ is a $\frac{1}{\sqrt{p}}$-one-sided spectral expander.*

Proof. The graph B_{p^3} is a bipartite, p^3-regular graph with $\lambda(B_{p^3}) \le \frac{1}{p^{3/2}}$ and $A(U, V, E)$ is a p^2-regular graph that is an induced subgraph of B_{p^3}. Hence, by Lemma 5.5,

$$\lambda(A) \le \frac{p^3 \cdot (1/p^{3/2})}{p^2} = \frac{1}{\sqrt{p}}. \qquad \square$$

The final expansion bounds

From the corollary above, we obtain the following theorem of Kaufman and Oppenheim.

Theorem 5.7 ([4])**.** *For $p > (d-2)^2$, the $(d-1)$-dimensional coset complex $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ is a $\frac{1}{\sqrt{p}-(d-2)}$-one-sided spectral HDX.*

Proof. It follows directly from Theorem 2.5 that $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ is a γ-one-sided spectral HDX for

$$\gamma \le \frac{1/\sqrt{p}}{1-(d-2)(1/\sqrt{p})} = \frac{1}{\sqrt{p}-(d-2)}. \qquad \square$$

Constructing two-sided spectral HDXs and standard expanders: The $(d-1)$-dimensional coset complex $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ is not a two-sided spectral HDX as the 1-skeletons of the links of the faces in $\mathcal{X}(d-3)$ are bipartite. However, if we restrict attention to the k-skeleton of $\mathcal{X}$ for some $k < d-1$ then we can bound the least eigenvalue by applying the Descent Theorem to the least eigenvalue (Theorem A.2(2)). This is summarized in the following corollary.

Corollary 5.8. *For $p > (d-2)^2$ and any $1 \le k < d$ the k-skeleton of the $(d-1)$-dimensional coset complex $\mathcal{X}(G, \{K_1, \ldots, K_d\})$ is a $\max\left\{\frac{1}{\sqrt{p}-(d-2)}, \frac{1}{d-k}\right\}$-two-sided spectral HDX.*

In particular, if we set $k = 1$ in the above corollary, we get a standard $\max\left\{\frac{1}{\sqrt{p}-(d-2)}, \frac{1}{d-1}\right\}$-two-sided spectral expander. This graph is a d-partite graph and hence its least eigenvalue is at most $-1/(d-1)$, while the above argument shows that it is least (and hence equal to) $-1/(d-1)$.

Thus, this not only yields an elementary construction and proof of one-sided spectral HDXs (Theorem 5.7), but also one of standard spectral expander (Corollary 5.8).

A Proof of the Descent Theorem

For the sake of completeness, we present the proof of Theorem 2.5 that asserts that proving spectral expansion for the maximal faces is sufficient to obtain expansion of any link. This exposition is essentially from the lecture notes by Dikstein [1].

Let (X, w) be a weighted d-dimensional simplicial complex. Let $\mu_d = w|_{X(d)}$ be the distribution on the set $X(d)$ of $(d+1)$-sized faces. This distribution induces distributions μ_i on $X(i)$ in the natural way.

For two functions $f, g \colon X(0) \to \mathbb{R}$, define their *inner product* $\langle f, g\rangle_X = \mathbb{E}_{u\sim\mu_0}[f(u)g(u)]$. We will drop the subscript X if it is clear from context. Note that, by the definition of μ_1, sampling u according to μ_0 can be equivalently achieved by sampling an edge (u, v) according to μ_1 and returning one of the points uniformly at random. Therefore,

$$\langle f, g\rangle_X = \mathbb{E}_{u\sim\mu_0}[f(u)g(u)] = \mathbb{E}_{\{u,v\}\sim\mu_1}[f(u)g(u)] = \mathbb{E}_{v\sim\mu_0}\mathbb{E}_{u\sim X_v(0)}[f(u)g(u)] = \mathbb{E}_{v\sim\mu_0}[\langle f_v, g_v\rangle_{X_v}], \tag{A.1}$$

where $f_v, g_v \colon X_v(0) \to \mathbb{R}$ are the restrictions to the link of v.

Define the *adjacency operator* A that, on a function $f \colon X(0) \to \mathbb{R}$ on vertices returns another function Af on vertices defined via

$$Af(v) = \mathbb{E}_{u\sim v}[f(u)],$$

where $u \sim v$ refers to a random neighbour of v according to the distribution $u \sim \mu_0(X_v)$. In other words, A *averages* f over neighbours. Furthermore, A is self-adjoint with respect to the above inner product, i.e, $\langle Af, g\rangle = \langle f, Ag\rangle$. Hence, it has n real eigenvalues and an orthonormal set of eigenvectors. Clearly $A\mathbb{1} = \mathbb{1}$; the constant 1 function is an eigenvector for this operator (in fact, it is an eigenvector corresponding to the largest eigenvalue 1). The remaining eigenvalues are characterized by the Courant–Fischer Theorem, stated below.

Theorem A.1 (Courant–Fischer)**.** *Let $A \in \mathbb{R}^{n\times n}$ be an $n \times n$ matrix over the reals that is self-adjoint with respect to some inner product $\langle\cdot,\cdot\rangle : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$. Then A has n real eigenvalues $\lambda_1 \geq \cdots \geq \lambda_n$ which have the following characterization.*

$$\lambda_i = \max_{V\colon \dim V = i} \min_{0\neq x\in V} \frac{\langle x, Ax\rangle}{\langle x, x\rangle} = \min_{V\colon \dim V = n-i+1} \max_{0\neq x\in V} \frac{\langle x, Ax\rangle}{\langle x, x\rangle}.$$

Similar to (A.1), we have

$$\begin{aligned} \langle Af, g\rangle_X &= \mathbb{E}_{\{u,w\}\sim\mu_1}[f(u)g(w)] = \mathbb{E}_{\{u,v,w\}\sim\mu_2}[f(u)g(w)] \\ &= \mathbb{E}_{v\sim\mu_0}\left[\mathbb{E}_{\{u,w\}\sim\mu_1(X_v)}[f(u)g(w)]\right] \\ &= \mathbb{E}_{v\sim\mu_0}\left[\langle A_v f_v, g_v\rangle_{X_v}\right] \end{aligned} \tag{A.2}$$

where A_v denotes the adjacency operator restricted to the link X_v.

With the above notation, we can now state the theorem we wish to prove. It suffices to prove the theorem in the case of $d = 3$ as we can obtain Theorem 2.5 by induction.

Theorem A.2. *Suppose (X, w) is weighted 2-dimensional simplicial complex. Then, we have the following two implications:*

1. *Suppose the 1-skeleton of X is connected and for every vertex $v \in X(0)$, $\langle A_v f, f \rangle \leq \lambda \langle f, f \rangle$ for all $f: X_v(0) \to \mathbb{R}$ with $f \perp \mathbb{1}_{X_v}$ for some $\lambda \in [0, 1)$. Then, for any $g: X(0) \to \mathbb{R}$ with $g \perp \mathbb{1}_X$, we have $\langle Ag, g \rangle \leq \gamma \langle g, g \rangle$ where $\gamma \leq \frac{\lambda}{1-\lambda}$.*

2. *Suppose the 1-skeleton of X is non-empty and for every vertex $v \in X(0)$, we have $\langle A_v f, f \rangle \geq \eta \langle f, f \rangle$ for all $f: X_v(0) \to \mathbb{R}$ for some $\eta \in [-1, 1)$. Then, for any $g: X(0) \to \mathbb{R}$, we have $\langle Ag, g \rangle \geq \gamma \langle g, g \rangle$ where $\gamma \geq \frac{\eta}{1-\eta}$.*

Before we see a proof of this, let us see how Theorem 2.5 follows from this.

Descent Theorem A.3 (Theorem 2.5 restated)**.** *Suppose (X, w) is a non-empty d-dimensional weighted simplicial complex with the following properties.*

- *The 1-skeleton of every link is connected.*
- *For all $v \in X(d-2)$, the link (X_v, w_v) is a λ-one-sided spectral expander for some $\lambda < \frac{1}{d-1}$. I. e., there is a $\lambda > 0$ such that, for every $v \in X(d-2)$ and every $g: X_v(0) \to \mathbb{R}$ with $g \perp \mathbb{1}$, we have*

$$\langle A_v g, g \rangle \leq \lambda \langle g, g \rangle .$$

Then, (X, w) is a γ-one-sided spectral HDX for $\gamma \leq \frac{\lambda}{1-(d-1)\lambda}$. That is, for any $v \in X(-1) \cup \cdots \cup X(d-2)$ and every $g: X_v(0) \to \mathbb{R}$ with $g \perp \mathbb{1}$, we have $\langle A_v g, g \rangle \leq \gamma \langle g, g \rangle$.

Furthermore, suppose we also know that there is a $\eta \in [-1, 0)$ such that, for every $v \in X(d-2)$ and every $g: X_v(0) \to \mathbb{R}$, we have $\langle A_v g, g \rangle \geq \eta \langle g, g \rangle$. Then, X is a γ-two-sided spectral HDX with

$$\gamma \leq \max\left(\frac{\lambda}{1-(d-1)\lambda}, \left|\frac{\eta}{1-(d-1)\eta}\right|\right).$$

That is, for every $g: X_v(0) \to \mathbb{R}$ with $g \perp \mathbb{1}$, we have $|\langle A_v g, g \rangle| \leq \gamma \langle g, g \rangle$.

Proof. For any $i \leq d-2$, let

$$\lambda_i = \min_{v \in X(i)} \max_{\substack{g: X_v(0) \to \mathbb{R} \\ g \perp \mathbb{1}}} \frac{\langle A_v g, g \rangle}{\langle g, g \rangle},$$

the smallest one-sided spectral expansion with respect to $X(i)$. From repeated applications of Theorem A.2,

$$\lambda_{-1} \leq \frac{\lambda_0}{1-\lambda_0} \leq \frac{\lambda_1/(1-\lambda_1)}{1-(\lambda_1/(1-\lambda_1))} = \frac{\lambda_1}{1-2\lambda_1} \leq \cdots \leq \frac{\lambda_{d-2}}{1-(d-1)\lambda_{d-2}}$$

which eventually completes the proof for one-sided spectral expansion.

For two-sided spectral expansion, we also have to show that all the eigenvalues are bounded away from −1. One again, let η_i be such that

$$\eta_i = \max_{v \in X(i)} \min_{\substack{g : X_v(0) \to \mathbb{R} \\ g \perp \mathbb{1}}} \frac{\langle A_v g, g \rangle}{\langle g, g \rangle}.$$

By repeated applications of Theorem A.2 (2), we obtain

$$\eta_{-1} \geq \frac{\eta_0}{1-\eta_0} \geq \frac{\eta_1/(1-\eta_1)}{1-(\eta_1/(1-\eta_1))} = \frac{\eta_1}{1-2\eta_1} \geq \cdots \geq \frac{\eta_{d-2}}{1-(d-1)\eta_{d-2}}$$

Together, we have that X is a γ-two-sided spectral HDX for

$$\gamma = \max\left(\frac{\lambda}{1-(d-1)\lambda}, \left|\frac{\eta}{1-(d-1)\eta}\right|\right). \qquad \square$$

Proof of Theorem A.2. Let g be an eigenvector that satisfies $\langle g, g \rangle = 1$ and $g \perp \mathbb{1}_X$ that maximises (or minimises) $\langle Ag, g \rangle$, and $\gamma = \langle Ag, g \rangle$ be the extremal value. In particular, $Ag = \gamma \cdot g$. From (A.2) we have $\gamma = \langle Ag, g \rangle = \mathbb{E}_v \left[\langle A_v g_v, g_v \rangle\right]$.

Even though $g \perp \mathbb{1}_X$, the *local* component g_v need not be perpendicular to $\mathbb{1}_{X_v}$. Hence, let us write $g_v = \alpha_v \mathbb{1}_{X_v} + g_v^\perp$ where $g_v^\perp \perp \mathbb{1}_{X_v}$; we shall drop the subscript from $\mathbb{1}_{X_v}$ for the sake of brevity as the length of the vector will be clear from context. Note that $\alpha_v = \langle g_v, \mathbb{1} \rangle = \mathbb{E}_{w \in X_v(0)}[g_v] = Ag(v)$. Therefore, $\mathbb{E}_v[\alpha_v^2] = \langle Ag, Ag \rangle = \gamma^2$. Hence,

$$\gamma = \langle Ag, g \rangle = \mathbb{E}_v \left[\langle A_v g_v, g_v \rangle\right] = \mathbb{E}_v \left[\alpha_v^2 + \left\langle A_v g_v^\perp, g_v^\perp \right\rangle\right] \tag{A.3}$$

We shall now focus on the proof of Theorem A.2 (1). The other direction is exactly identical with the inequality flipped.

In the case of Theorem A.2 (1), where we are given $\langle A_v g_v^\perp, g_v^\perp \rangle \leq \lambda \langle g_v^\perp, g_v^\perp \rangle$ for all $v \in X(0)$, we have

$$\begin{aligned} \gamma &= \mathbb{E}_v \left[\alpha_v^2 + \left\langle A_v g_v^\perp, g_v^\perp \right\rangle\right] \leq \mathbb{E}_v \left[\alpha_v^2 + \lambda \left\langle g_v^\perp, g_v^\perp \right\rangle\right] \\ &= \mathbb{E}_v \left[(1-\lambda)\alpha_v^2 + \lambda \langle g_v, g_v \rangle\right] \\ &= (1-\lambda)\gamma^2 + \lambda. \\ \implies \gamma(1-\gamma) &\leq \lambda(1-\gamma^2) \\ \implies \gamma &\leq \lambda(1+\gamma) \qquad \text{(connected, thus } \gamma < 1) \\ \implies \gamma &\leq \frac{\lambda}{1-\lambda}. \end{aligned}$$

In the case of Theorem A.2 (2), where we are given $\langle A_v g_v^\perp, g_v^\perp \rangle \geq \eta \langle g_v^\perp, g_v^\perp \rangle$ for all $v \in X(0)$, the same argument yields

$$\begin{aligned} \gamma = \mathbb{E}_v \left[\alpha_v^2 + \left\langle A_v g_v^\perp, g_v^\perp \right\rangle\right] \geq \mathbb{E}_v \left[\alpha_v^2 + \eta \left\langle g_v^\perp, g_v^\perp \right\rangle\right] = (1-\eta)\gamma^2 + \eta \\ \implies \gamma \geq \frac{\eta}{1-\eta} \end{aligned} \qquad \square$$

B Primer on group theory

In this section, for completeness, we shall note the basic definitions and properties of groups that are used in this exposition.

Definition B.1 (Groups and subgroups)**.** A set of elements G equipped with a binary operation $\star : G \times G \to G$ is said to be a *group* if it satisfies the following properties:

Associativity: For all $g_1, g_2, g_3 \in G$, we have $(g_1 \star g_2) \star g_3 = g_1 \star (g_2 \star g_3)$.

Identity: There exists an *identity* element $\mathrm{id} \in G$ such that, for all $g \in G$, we have $g \star \mathrm{id} = \mathrm{id} \star g = g$.

Inverses: For every element $g \in G$, there is an element $g^{-1} \in G$ such that $g \star g^{-1} = g^{-1} \star g = \mathrm{id}$.

A subset $H \subseteq G$ is said to be a *subgroup* of G if H the binary operation $\star$ restricted to H satisfies the above three properties (including the fact that $h_1 \star h_2 \in H$ for all $h_1, h_2 \in H$).

Often the binary operation $\star$ is omitted and products just expressed as concatenation of elements.

Definition B.2 (Cosets)**.** Given a subgroup H of a group G, if $x \in G$ is an arbitrary element, the *left coset* of H containing x, denoted by xH, is defined as the set

$$xH = \{xh \ : \ h \in H\}.$$

Two cosets xH and yH are identical if and only if $x^{-1}y \in H$. Hence, any element $x' \in xH$ is also referred to as a *coset representative* of xH as $x'H = xH$.

Right cosets are defined similarly. A subgroup H is said to be *normal* if the right cosets and the left cosets agree for all x, i. e., $xH = Hx, \forall x \in G$.

Since two cosets of a subgroup H of G are either identical or disjoint, the set of distinct cosets of a subgroup H of G partition the elements of G. If a subgroup H is normal, this set of cosets forms a group G/H, called the *quotient group* of H in G.

For subgroups H, K of G, we will often consider the product HK (or $H \star K$) which refers to the set $\{hk \ : \ h \in H \ , \ k \in K\}$. It is worth stressing that HK *need not* be a subgroup of G and the above just refers to a set of elements that can be expressed as an (ordered) product of an element in H and an element in K.

For an arbitrary subset S of G, we will define $\langle S \rangle$ as the smallest subgroup of G that contains the set S. This is also referred to as the *group generated* by S.

In general, the binary operation $\star$ is order dependent. Groups where $g_1 g_2 = g_2 g_1$ for all $g_1, g_2 \in G$ are said to be *commutative* or *abelian* groups. The following notion of *commutators* (and *commutator subgroups*) is a way to measure *how non-commutative* a group G is.

Definition B.3 (Commutators)**.** For a pair of elements $g, h \in G$, we shall define the *commutator* of g, h (denoted by $[g, h]$) as

$$[g, h] := g^{-1}h^{-1}gh.$$

The *commutator subgroup* of G, denoted by $[G, G]$ is the *group generated by all commutators*. That is,

$$[G, G] := \langle \{[g, h] \ : \ g, h \in G\} \rangle .$$

Note that if G is abelian, then $[G, G] = \{\mathrm{id}\}$. As mentioned earlier, the commutator subgroup can be thought of as a way of describing how non-abelian a group is. In fact, the commutator subgroup of G is the smallest *normal* subgroup H of G such that the *quotient* G/H is abelian (although these are concepts that are not necessary to follow this exposition).

References

[1] YOTAM DIKSTEIN: Oppenheim's Trickling Down Theorem, 2019. Lecture notes from the Error-Correcting Codes and High-Dimensional Expansion Boot Camp at Simons Institute for the Theory of Computing, Available at https://simons.berkeley.edu/sites/default/files/docs/14119/ecctalkii.pdf. 16

[2] HOWARD GARLAND: p-adic curvature and the cohomology of discrete subgroups of p-adic groups. *Ann. Math.*, 97(3):375–423, 1973. [doi:10.2307/1970829] 4

[3] ALEXANDER J. HAHN AND O. TIMOTHY O'MEARA: *The Classical Groups and K-Theory*. Springer, 1989. [doi:10.1007/978-3-662-13152-7] 7

[4] TALI KAUFMAN AND IZHAR OPPENHEIM: High dimensional expanders and coset geometries. *Europ. J. Combinat.*, 111(103696), 2023. Preliminary version in STOC'18. [doi:10.1016/j.ejc.2023.103696, arXiv:1710.05304] 2, 6, 7, 13, 15

[5] ALEXANDER LUBOTZKY, RALPH PHILLIPS, AND PETER SARNAK: Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. Preliminary version in STOC'86. [doi:10.1007/BF02126799] 2

[6] ALEXANDER LUBOTZKY, BETH SAMUELS, AND UZI VISHNE: Explicit constructions of Ramanujan complexes of type $\tilde{A}_d$. *Europ. J. Combinat.*, 26(6):965–993, 2005. [doi:10.1016/j.ejc.2004.06.007, arXiv:math/0406217] 2

[7] ALEXANDER LUBOTZKY, BETH SAMUELS, AND UZI VISHNE: Ramanujan complexes of type $\tilde{A}_d$. *Israel J. Math.*, 149(1):267–299, 2005. [doi:10.1007/BF02772543, arXiv:math/0406208] 2

[8] GRIGORII ALEKSANDROVICH MARGULIS (Григорий Александрович Маргулис): Явные теоретико-групповые конструкции комбинаторных схем и их применения в построении расширителей и концентраторов (Russian) [Explicit group-theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators]. *Problemy Peredachi Informatsii*, 24(1):51–60, 1988. Available at https://www.mathnet.ru/eng/ppi686, (English translation in *Problems Inform. Transmission*, 24(1):39–46, 1988). 2

[9] RYAN O'DONNELL AND KEVIN PRATT: High-dimensional expanders from Chevalley groups. In *Proc. 37th Comput. Complexity Conf. (CCC'22)*, pp. 18:1–26. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2022. [doi:10.4230/LIPIcs.CCC.2022.18, arXiv:2203.03705] 13

[10] IZHAR OPPENHEIM: Local spectral expansion approach to high dimensional expanders Part I: Descent of spectral gaps. *Discr. Comput. Geom.*, 59(2):293–330, 2018. [doi:10.1007/s00454-017-9948-x, arXiv:1709.04431] 4

[11] OMER REINGOLD, SALIL VADHAN, AND AVI WIGDERSON: Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Ann. Math.*, 155(1):157–187, 2005. Preliminary version in FOCS'00. [doi:10.2307/3062153, arXiv:math.CO/0406038, ECCC:TR01-018] 13

[12] ALIREZA SARVENIAZI: Explicit construction of a Ramanujan $(n_1, n_2, \ldots, n_{d-1})$ -regular hypergraph. *Duke Math. J.*, 136(1):141–171, 2007. [doi:10.1215/S0012-7094-07-13913-9] 2

AUTHORS

Prahladh Harsha
School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, India
prahladh@tifr.res.in
http://www.tifr.res.in/~prahladh/

Ramprasad Saptharishi
School of Technology and Computer Science
Tata Institute of Fundamental Research
Mumbai, India
ramprasad@tifr.res.in
http://www.tifr.res.in/~ramprasad.saptharishi/

ABOUT THE AUTHORS

PRAHLADH HARSHA is a theoretical computer scientist at the Tata Institute of Fundamental Research (TIFR). He received his B. Tech. degree in Computer Science and Engineering from the IIT Madras in 1998. He then obtained his S. M. and Ph. D. degrees in Computer Science from MIT in 2000 and 2004, respectively, advised by Madhu Sudan. Prior to joining TIFR in 2010, he was at Microsoft Research, Silicon Valley and at the Toyota Technological Institute at Chicago. His areas of interest include computational complexity, hardness of approximation, coding theory and information theory. Prahladh credits his mother for his interest in mathematics and dance. He is also deeply indebted to U Koteswara Rao, his high school mentor, for exposing him to both the beauty and rigour in mathematics.

RAMPRASAD SAPTHARISHI is a theoretical computer scientist at the School of Technology and Computer Science at the Tata Institute of Fundamental Research, Mumbai, and is generally interested in most questions of an algebraic nature. He did his Ph. D. at Chennai Mathematical Institute, advised by Manindra Agrawal, and spent a couple of years at Microsoft Research, India and Tel Aviv University for his postdoc. His other interests include board games, reading ridiculously long web-series and writing software code. However, writing reasonable biographies continues to remain outside his area of expertise.