CAYLEY GRAPHS WITHOUT A BOUNDED EIGENBASIS

ASHWIN SAH, MEHTAAB SAWHNEY, AND YUFEI ZHAO

Abstract. Does every n-vertex Cayley graph have an orthonormal eigenbasis all of whose coordinates are $O(1/\sqrt{n})$? While the answer is yes for abelian groups, we show that it is no in general.

On the other hand, we show that every n-vertex Cayley graph (and more generally, vertex-transitive graph) has an orthonormal basis whose coordinates are all $O(\sqrt{\log n/n})$, and that this bound is nearly best possible.

Our investigation is motivated by a question of Assaf Naor, who proved that random abelian Cayley graphs are small-set expanders, extending a classic result of Alon–Roichman. His proof relies on the existence of a bounded eigenbasis for abelian Cayley graphs, which we now know cannot hold for general groups. On the other hand, we navigate around this obstruction and extend Naor's result to nonabelian groups.

1. Introduction

1.1. **Bounded eigenbasis.** It is a fundamental problem to understand the spectral decomposition of a Cayley graph. Since every vertex in a Cayley graph has the same degree, it does not matter whether we are talking about the adjacency matrix or the Laplacian matrix, but we will stick with the adjacency matrix for concreteness. Enormous attention has been given to the eigenvalues of Cayley graphs, especially the spectral gap, due to an intimate connection with the expansion properties of the graph [12, 15]. In this paper, we study the complementary question of what can arise as eigenvectors of Cayley graphs.

We adopt the following normalization[1]. Given a finite set S and a function $x\colon S \to \mathbb{C}$ (sometimes viewed as a vector $x \in \mathbb{C}^S$), we denote its L^p norm by

$$\|x\|_{L^p(S)} := (\mathbb{E}_{s\in S} |x(s)|^p)^{1/p}.$$

The Hermitian inner product is defined by $\langle x, y\rangle = \mathbb{E}_{s\in S}\overline{x(s)}y(s)$. We say that x is *C-bounded* if $\|x\|_{L^\infty(S)} \le C$, and we say that a set of functions is C-bounded if all of its elements are C-bounded. In a unitary or orthonormal eigenbasis, each eigenfunction x is normalized as $\|x\|_{L^2(S)} = 1$.

Below is the main question that we study. Here our Cayley graphs are unweighted and undirected. A Cayley graph on a finite group G with symmetric generator $S = S^{-1}$ (not containing the identity) has edges of form (g, sg) ranging over all $g \in G$ and $s \in S$.

Question 1.1. What is the minimum $C(n)$ so that every n-vertex Cayley graph has a $C(n)$-bounded unitary (or orthonormal) eigenbasis?

Every abelian Cayley graph has a 1-bounded unitary eigenbasis. Indeed, given an abelian group G, the basis of Fourier characters of G forms an eigenbasis for every Cayley graph on G. In particular, all coordinates of such a Fourier basis are are roots of unity. The existence of a bounded eigenbasis is useful in certain applications. In fact, the initial motivation for this work is a result of Naor [16] that proves a certain small-set expansion property of random abelian Cayley graphs, extending a classic result of Alon–Roichman [1] that random Cayley graphs on arbitrary groups are

Zhao was supported by NSF Award DMS-1764176, a Sloan Research Fellowship, and the MIT Solomon Buchsbaum Fund.

[1]This normalization, viewing x as a function on a set or group equipped with the averaging measure, is different from the normalization used in the abstract, where we use the usual Euclidean distance in $\mathbb{R}^n$.

expanders. Naor's argument uses that every abelian Cayley graph has a 1-bounded eigenbasis. He asks whether his results also hold for nonabelian groups. Here we show that general Cayley graphs do not always have a bounded eigenbasis, therefore exhibiting an obstruction to Naor's argument for nonabelian groups. On the other hand, we provide an alternative argument showing that Naor's theorem indeed extends to general groups. See Theorem 1.7 below for a precise statement.

Our first result below implies that Cayley graphs do not always have a bounded unitary eigenbasis. It gives a lower bound $C(n) \gtrsim \sqrt{\log n}/\log\log n$ for Question 1.1 for infinitely many n. (Notation: we write $A \lesssim B$ and $A = O(B)$ to mean that $A \le CB$ for some constant $C > 0$.)

Theorem 1.2. *There exist infinitely many Cayley graphs G whose adjacency matrix has an eigenspace all of whose eigenfunctions $x\colon G \to \mathbb{C}$ satisfy $\|x\|_{L^\infty(G)} \ge c\,\|x\|_{L^2(G)}\sqrt{\log n}/\log\log n$, where n is the number of vertices and $c > 0$ is some absolute constant.*

The next result gives a nearly matching upper bound of $C(n) \lesssim \sqrt{\log n}$ for Question 1.1.

Theorem 1.3. *Every Cayley graph has an orthonormal $C\sqrt{\log n}$-bounded eigenbasis, where n is the number of vertices and C is some absolute constant.*

More generally, the same upper bound holds for vertex-transitive graphs.

Theorem 1.4. *Every vertex-transitive graph has an orthonormal $C\sqrt{\log n}$-bounded eigenbasis, where n is the number of vertices and C is some absolute constant.*

It remains an intriguing open problem to close the gap between the upper and lower bounds. This problem appears to be related to a recent deep and difficult result of Green [8], who showed that the maximum possible width of a finite transitive subset of the unit sphere in $\mathbb{R}^d$ is on the order of $1/\sqrt{\log d}$ (in sharp constrast to infinite subsets, e.g., the entire sphere has width 1). Green's theorem answers a question of the third author, which was in turn prompted by [5] and this work. Green's proof relies on the classification of finite simple groups.

Let us mention a few directions worth further investigation. First, our construction proving Theorem 1.2 uses graphs of increasing degree. Can one also find bounded degree Cayley graphs without a bounded eigenbasis?

Conjecture 1.5. *There exists some d such that for every C there exists a d-regular Cayley graph without an orthonormal C-bounded eigenbasis.*

Another direction worth exploring further is to understand what families of groups always have bounded eigenbasis. Extending the example of abelian groups, it is not hard to show using non-abelian Fourier analysis that in a group where every irreducible representation has dimension at most d, every Cayley graph has a $\sqrt{d}$-bounded unitary eigenbasis. Given these examples, a natural question is if for more natural classes of "nearly abelian" groups, every Cayley graph has a bounded eigenbasis.

Question 1.6. Do Cayley graphs on nilpotent groups of bounded step always have bounded eigenbasis? What about affine groups?

The general problem of characterizing groups with the bounded eigenbasis property is somewhat reminiscent of the characterization of approximate groups by Breuillard, Green, and Tao [4], which unifies classic theorems of Freiman on sets of bounded doubling [7] and Gromov on groups of polynomial growth [10].

Let us mention that another instance where a bounded eigenbasis came in handy was in studying the relationship between discrepancy and eigenvalues of Cayley graphs. Kohayakawa, Rödl, and Schacht [13] showed that for abelian Cayley graphs, having small discrepancy is equivalent to having small second eigenvalues, with a spectral proof suggested by Gowers. The proof relies on the bounded eigenbasis of abelian Cayley graphs. The abelian hypothesis was later removed by Conlon and Zhao [5] via an application of Grothendieck's inequality.

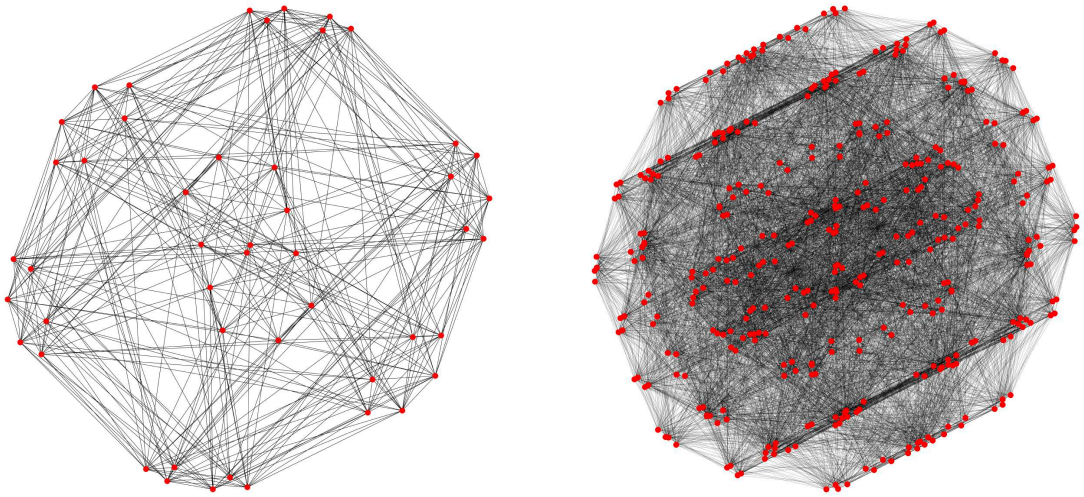

(A) $G = S_3 \ltimes (\mathbb{Z}/2\mathbb{Z})^3$ (B) $G = S_4 \ltimes (\mathbb{Z}/2\mathbb{Z})^4$

FIGURE 1. Some spectral drawings of Cayley graphs used in the proof of Theorem 1.2, corresponding to the group $G = S_d \ltimes (\mathbb{Z}/2\mathbb{Z})^d$, drawn for $d \in \{3, 4\}$. Such a spectral drawing requires a canvas of side-lengths on the order of $\sqrt{d/\log d}$, where vertices have uncorrelated x and y coordinates each having variance 1. See the remark at the end of Section 3.3 on how these figures were generated.

The boundedness of eigenfunctions has an appealing interpretation for *spectral graph drawings*. Hall's spectral drawing of a graph [11] (also see Spielman's survey [18], which contains some nice figures) places each vertex v at $(x(v), y(v)) \in \mathbb{R}^2$, where x and y are eigenfunctions corresponding to the second and third eigenvalues of the graph Laplacian (here x and y are assumed orthogonal and properly scaled). This drawing has the property that it minimizes the sum of squared edge-lengths among all drawings of the graph in $\mathbb{R}^2$ with the vertices in isotropic position (so that x and y coordinates each have variance 1 and are uncorrelated). Every abelian Cayley graph has a spectral drawing where all the coordinates are bounded. On the other hand, Theorem 1.2 gives us an example of a Cayley graph where no spectral drawing can fit inside a disk of radius $c\sqrt{\log n}/\log\log n$ (provided that the eigenspace in the theorem corresponds to the second and third eigenvalues, which can be achieved; see the end of Section 3 for further details). Some examples of spectral drawings of Cayley graphs used in the proof of Theorem 1.2 are shown in Figure 1.

1.2. **Random Cayley graphs are small-set expanders.** A classic result due to Alon and Roichman [1] shows that in a random Cayley graph of a group G generated by $k > C\epsilon^{-2}\log|G|$ independent and uniform random group elements, all eigenvalues other than the top one has absolute value at most $k\epsilon$. In particular, via the expander mixing lemma, it implies that for every $\emptyset \neq X \subsetneq G$,

$$\left| \frac{e(X, G \setminus X)}{\frac{2k}{|G|}\,|X|\,|G \setminus X|} - 1 \right| \leq \epsilon,$$

where $e(A, B)$ counts the number of edges with one endpoint in A and the other in B.

Naor [16] developed a new Banach space-valued Azuma inequality and proved more refined small-set isoperimetry inequalities in random Cayley graphs of abelian groups, and he asked whether his result can be extended to all groups. Here we answer his question affirmatively. In the following theorem, by "the Cayley (multi)graph associated k independent uniformly chosen random group elements" we mean the following: select random $g_1, \dots, g_k$ and take the Cayley graph generated by $g_1, g_1^{-1}, \dots, g_k, g_k^{-1}$, taken with multiplicity. Allowing multiplicities makes the result a bit easier to state, and is a technicality that one should feel free to ignore (in many parameter ranges multiplicities are unlikely to occur).

Theorem 1.7. *There exists a universal constant $C > 0$ such that for every positive integer k and every group G, with probability at least $1/2$, the Cayley (multi)graph associated to k independent uniformly chosen random group elements has the property that for every $X \subseteq G$ with $1 < |X| \leq |G|/2$, the number of edges $e(X, G \setminus X)$ between X and $G \setminus X$ satisfies*

$$\left| \frac{e(X, G \setminus X)}{\frac{2k}{|G|} |X| |G \setminus X|} - 1 \right| \leq C \sqrt{\frac{\log |X|}{k}}.$$

Naor proved Theorem 1.7 for abelian groups. His proof relies on a bounded eigenbasis of abelian Cayley graphs. In Section 5 we explain how to bypass this obstacle in order to prove the result for nonabelian groups.

Acknowledgments. Zhao thanks Assaf Naor for discussions and for encouraging him to work on this problem. We thank Shengtong Zhang for pointing out some typographical errors.

2. Preliminaries

2.1. **Nonabelian Fourier transform.** We begin by summarizing some standard facts on nonabelian Fourier analysis (e.g., [9]). Given a finite group G, let $\widehat{G}$ denote the set of irreducible unitary representations of G. For each representation $\rho \in \widehat{G}$, call its dimension d_ρ, and call the space that it acts on $W_\rho \cong \mathbb{C}^{d_\rho}$. For any $f : G \to \mathbb{C}$ and $\rho \in \widehat{G}$, its Fourier transform evaluated at ρ is given by

$$\widehat{f}(\rho) = \mathbb{E}_{g \in G} f(g) \rho(g),$$

which is an endomorphism of W_ρ (i.e., $\widehat{f}(\rho) \in \operatorname{End} W_\rho$). There is an inversion formula, namely

$$f(g) = \sum_{\rho \in \widehat{G}} d_\rho \langle \rho(g), \widehat{f}(\rho) \rangle_{\mathrm{HS}}$$

where $\langle A, B \rangle_{\mathrm{HS}} = \operatorname{Tr}(A^\dagger B)$ is the Hilbert–Schmidt inner product, which is just the entry-wise Hermitian product of matrices. The Hilbert–Schmidt norm is written as $\|A\|_{\mathrm{HS}} = \sqrt{\operatorname{Tr}(A^\dagger A)}$. We have Parseval's identity

$$\langle f_1, f_2 \rangle_{L^2(G)} = \mathbb{E}_{g \in G} \overline{f_1(g)} f_2(g) = \sum_{\rho \in \widehat{G}} d_\rho \langle \widehat{f_1}(\rho), \widehat{f_2}(\rho) \rangle_{\mathrm{HS}},$$

and in particular,

$$\mathbb{E}_{g \in G} |f(g)|^2 = \sum_{\rho \in \widehat{G}} d_\rho \|\widehat{f}(\rho)\|_{\mathrm{HS}}^2.$$

Finally, we define a convolution of two functions $f_1, f_2 : G \to \mathbb{C}$ via

$$(f_1 * f_2)(g) = \mathbb{E}_{h \in G} f_1(gh^{-1}) f_2(h).$$

The Fourier transform turns convolution into matrix multiplication:

$$\widehat{f_1 * f_2}(\rho) = \widehat{f_1}(\rho) \widehat{f_2}(\rho)$$

for all $\rho \in \widehat{G}$.

2.2. **Eigendecomposition.** Given a function $f : G \to \mathbb{C}$, we consider the operator M_f acting on $\mathbb{C}^G$, the space of functions $G \to \mathbb{C}$, via

$$M_f x = f * x,$$

i.e., $(M_f x)(g) = \mathbb{E}_{h\in G} f(gh^{-1})x(h)$ for all $x\colon G \to \mathbb{C}$ and $g \in G$. Equivalently, one can also view M_f as a matrix with rows and columns indexed by G, whose entry in position $(g, h) \in G \times G$ is $f(gh^{-1})/|G|$. Then, viewing $x \in \mathbb{C}^G$ as a vector, the matrix product $M_f x$ agrees with the definition above. The matrix can be thought of as the adjacency matrix (after suitable normalization) of a Cayley graph. Let us explain how to analyze the eigendata of M_f using the Fourier transform.

Assume from now on that $f(g^{-1}) = \overline{f(g)}$ for every $g \in G$. Then

$$\widehat{f}(\rho) = \mathbb{E}_{g\in G} f(g)\rho(g)$$

is Hermitian. For any $x : G \to \mathbb{C}$, applying the Fourier transform, we see that x is an eigenfunction of M_f with eigenvalue λ (i.e., $f * x = \lambda x$) if and only if

$$\widehat{f}(\rho)\widehat{x}(\rho) = \lambda\widehat{x}(\rho) \quad \text{for all } \rho \in \widehat{G},$$

i.e., all columns of $\widehat{x}(\rho)$ (when viewed as a $d_\rho \times d_\rho$ matrix) lie in the eigenspace of $\widehat{f}(\rho)$ corresponding to the eigenvalue λ.

Let V_ρ be the subspace of functions whose Fourier transform is supported on ρ:

$$\begin{aligned} V_\rho &= \{x \in L^2(G) : \operatorname{supp} \widehat{x} \subseteq \{\rho\}\} \\ &= \{x \in L^2(G) : x(g) = d_\rho \langle \rho(g), A\rangle \text{ for some } A \in \operatorname{End} W_\rho\}. \end{aligned} \tag{2.1}$$

For any column vector $\mathbf{v} \in W_\rho$, we define

$$\begin{aligned} V_{\rho,\mathbf{v}} &= \{x \in V_\rho : \text{every column of } \widehat{x}(\rho) \text{ is a multiple of } \mathbf{v}\} \\ &= \{x \in L^2(G) : x(g) = d_\rho \langle \rho(g), \mathbf{v}\mathbf{w}^\dagger\rangle_{\mathrm{HS}} \text{ for some } \mathbf{w} \in W_\rho\} \subseteq V_\rho. \end{aligned} \tag{2.2}$$

In particular, if $\widehat{f}(\rho)\mathbf{v} = \lambda\mathbf{v}$ for some $\lambda \in \mathbb{R}$, then $M_f x = \lambda x$ for all $x \in V_{\rho,\mathbf{v}}$ (as can be seen by taking the Fourier transform). Furthermore, if $\mathbf{v}, \mathbf{v}' \in W_\rho$ with $\mathbf{v}^\dagger\mathbf{v}' = 0$, then $\langle x, x'\rangle = 0$ for all $x \in V_{\rho,\mathbf{v}}$ and $x' \in V_{\rho,\mathbf{v}'}$.

To summarize, we have an orthogonal decomposition (the orthogonality is easy to check via the Fourier transform)

$$\mathbb{C}^G = \bigoplus_\rho V_\rho.$$

For each $\rho \in \widehat{G}$, let $\mathbf{v}_1^\rho, \dots, \mathbf{v}_{d_\rho}^\rho \in W_\rho$ be an eigenbasis of $\widehat{f}(\rho) \in \operatorname{End} W_\rho$, and call the corresponding eigenvalues $\lambda_{\rho,1}, \dots, \lambda_{\rho,d_\rho}$. We have an orthogonal decomposition

$$V_\rho = \bigoplus_{j=1}^{d_\rho} V_{\rho,\mathbf{v}_j^\rho}$$

and $M_f x = \lambda_j^\rho x$ for each $x \in V_{\rho,\mathbf{v}_j^\rho}$. Thus the eigenvalues of M_f consists of $\lambda_{\rho,j}$ with multiplicity d_ρ, ranging over all $\rho \in \widehat{G}$ and $j \in [d_\rho]$. The eigenspace of M_f corresponding to an eigenvalue λ is the direct sum of all $V_{\rho,\mathbf{v}_j^\rho}$ ranging over all $\rho \in \widehat{G}$ and $j \in [d_\rho]$ with $\lambda_{\rho,j} = \lambda$.

2.3. **Schatten norms.** The Schatten p-norm $\|A\|_{S_p}$ of a matrix $A \in \mathbb{C}^{n\times n}$ is defined via

$$\|A\|_{S_p}^p = \sum_{i=1}^{n} \sigma_i(A)^p,$$

where $\sigma_1(A), \dots, \sigma_n(A)$ are the singular values of A.

The Schatten p-norm satisfies a noncommutative Hölder's inequality (e.g., [3, Corollary IV.2.6]): for $1 \le p \le q \le \infty$ with $1/p + 1/q = 1$, we have

$$\langle A, B\rangle_{\mathrm{HS}} \le \|A\|_{S_p}\|B\|_{S_q}. \tag{2.3}$$

Given a function $f\colon G \to \mathbb{C}$ on a finite group G, we define its Schatten p-norm $\|f\|_{S_p}$ to be the Schatten p-norm of its associated matrix M_f (giving the linear map $x \mapsto f * x$ on $\mathbb{C}^G$):

$$\|f\|_{S_p} = \left(\sum_i \sigma_i(M_f)^p\right)^{1/p} = \left(\sum_\rho d_\rho \|\widehat{f}(\rho)\|_{S_p}^p\right)^{1/p}. \tag{2.4}$$

3. Construction

In this section, we prove Theorem 1.2 by constructing a Cayley graph on a group G with an eigenspace all of whose eigenfunctions satisfy $\|x\|_{L^\infty(G)} \gtrsim \|x\|_{L^2(G)} \sqrt{\log|G|}/\log\log|G|$.

To motivate our construction, we first explain in Section 3.1 what happens for the unitary group $G = U(d)$, which is simpler to analyze although it is not finite. Then, in Section 3.2, we explain how to construct an edge-weighted Cayley graph on a certain finite subgroup of $U(d)$. We then explain in Section 3.3 how to convert the edge-weighted construction to an unweighted construction via sampling, and show that eigenvectors maintain their desired properties. Only Section 3.3 is required for the proof of Theorem 1.2, and the earlier subsections are solely for motivation, but we hope that they are helpful to the readers.

3.1. **Unitary group.** Let $G = U(d)$. Let ρ denote the standard representation of G on $\mathbb{C}^d$, which is irreducible. Let V_ρ denote the subspace of $L^2(G)$ consisting of all $x \in L^2(G)$ of the form $x(g) = d\langle \rho(g), A\rangle_{\mathrm{HS}}$ for some $A \in \mathbb{C}^{d\times d}$, i.e., the Fourier transform $\widehat{x}$ is supported at ρ and $\widehat{x}(\rho) = A$. Note that this definition of V_ρ agrees with our earlier definition in (2.1) for finite groups.

Claim 3.1. *For any $x \in V_\rho$, we have $\|x\|_{L^\infty(G)} \ge \sqrt{d}\|x\|_{L^2(G)}$.*

Proof. Let $A \in \mathbb{C}^{d\times d}$ be such that $x(g) = d\langle \rho(g), A\rangle_{\mathrm{HS}}$. By Fourier inversion and Parseval, $\|x\|_{L^2(G)} = \sqrt{d}\|A\|_{\mathrm{HS}}$. Thus we have

$$\|x\|_{L^\infty(G)} = d \sup_{U\in U(d)} \langle U, A\rangle_{\mathrm{HS}} = d\|A\|_{S_1} \ge d\|A\|_{S_2} = d\|A\|_{\mathrm{HS}} = \sqrt{d}\|x\|_{L^2(G)}$$

by definition of ρ and duality of the Schatten norms. □

For a sufficiently generic $f\colon G \to \mathbb{C}$ (assuming no unwanted eigenvalue collisions), the subspace V_ρ is a direct sum of eigenspaces of the operator M_f, there by giving a continuous analog of Theorem 1.2. The actual construction proving Theorem 1.2 will involve a discretization of this construction.

3.2. **Weighted construction.** In this section we prove a weighted analogue of Theorem 1.2 which serves as a stepping stone towards the entire proof. Recall from earlier that for a function $f : G \to \mathbb{C}$ we have $M_f(g, h) = f(gh^{-1})/|G|$.

Theorem 3.2. *There exist some constant $c > 0$ and infinitely many groups G and functions $f : G \to \mathbb{R}$ such that M_f has an eigenspace all of whose elements x satisfy*

$$\|x\|_{L^\infty(G)} \ge c\frac{\sqrt{\log|G|}}{\log\log|G|}\|x\|_{L^2(G)}.$$

Let $G = S_d \ltimes (\mathbb{Z}/2\mathbb{Z})^d$, where S_d acts on $(\mathbb{Z}/2\mathbb{Z})^d$ by permuting coordinates. The group has a natural d-dimensional representation on $\mathbb{C}^d$, which we call ρ, where S_d permutes the coordinates and $(\mathbb{Z}/2\mathbb{Z})^d$ flips the signs of the coordinates. It is easy to check that ρ is irreducible.

We need a G-orbit on the unit sphere in $\mathbb{C}^d$ with large width in every direction. The next lemma serves as a finitary analogue of Claim 3.1. This lemma also appears in [8] and we include its proof here for the convenience of the reader.

Lemma 3.3. *Let* $\mathbf{a}$ *be the unit vector in the direction of* $(1, 1/\sqrt{2}, \dots, 1/\sqrt{d})$. *Then for any* $\mathbf{v} \in \mathbb{C}^d$ *we have*

$$\sup_{g\in G} |\langle \mathbf{v}, \rho(g)\mathbf{a}\rangle| \gtrsim \frac{|\mathbf{v}|}{\sqrt{\log d}}.$$

Remark. A difficult recent result of Green [8], confirming a conjecture of Zhao, showed that for every *finite* subgroup G of $U(d)$ and every unit vector $\mathbf{a} \in \mathbb{C}^d$, there is some unit vector $\mathbf{v}$ such that $\sup_{g\in G} |\langle \mathbf{v}, \rho(g)\mathbf{a}\rangle| \lesssim 1/\sqrt{\log d}$ (i.e., a tight upper bound on the width of every finite transitive subset of a sphere). In contrast, the width of an infinite transitive subset of the sphere can be as large as 1 (e.g., the entire sphere). It is initially quite counterintuitive that a finiteness assumption implies such a dramatic reduction in the width of an orbit.

Proof. Let us first assume that $\mathbf{v} \in \mathbb{R}^d$, so

$$\sup_{g\in G} |\langle \mathbf{v}, \rho(g)\mathbf{a}\rangle| = \sup_{g\in G} |\langle \rho(g)\mathbf{v}, \mathbf{a}\rangle| \geq \langle \mathbf{w}, \mathbf{a}\rangle,$$

where $\mathbf{w}$ is the vector obtained by making the coordinates of $\mathbf{v}$ nonnegative and then rearranging them in nonincreasing order. Let $\mathbf{w} = (w_1, \dots, w_d)$ with $w_1 \geq \cdots \geq w_d \geq 0$. Then

$$\begin{aligned}
\left(\sum_{j=1}^d \frac{1}{j}\right) \langle \mathbf{w}, \mathbf{a}\rangle^2 = \left(\sum_{j=1}^d \frac{w_j}{\sqrt{j}}\right)^2 \\
&\gtrsim w_1^2 + \frac{w_2}{\sqrt{2}}\left(w_1 + \frac{w_2}{\sqrt{2}}\right) + \frac{w_3}{\sqrt{3}}\left(w_1 + \frac{w_2}{\sqrt{2}} + \frac{w_3}{\sqrt{3}}\right)\cdots \\
&\gtrsim w_1^2 + \frac{w_2^2}{\sqrt{2}}\left(1 + \frac{1}{\sqrt{2}}\right) + \frac{w_3^2}{\sqrt{3}}\left(1 + \frac{1}{\sqrt{2}} + \frac{1}{\sqrt{3}}\right) + \cdots \\
&\gtrsim w_1^2 + \cdots + w_d^2 = |\mathbf{w}|^2 = |\mathbf{v}|^2.
\end{aligned}$$

Thus

$$\sup_{g\in G} |\langle \mathbf{v}, \rho(g)\mathbf{a}\rangle| \gtrsim \frac{|\mathbf{v}|}{\sqrt{\log d}}.$$

Finally, when $\mathbf{v} \in \mathbb{C}^d$, we can reduce to the real case. Assume without loss of generality (since ρ is real) that $|\operatorname{Re}\mathbf{v}| \geq |\operatorname{Im}\mathbf{v}|$. Then

$$\sup_{g\in G} |\langle \mathbf{v}, \rho(g)\mathbf{a}\rangle| \geq \sup_{g\in G} |\langle \operatorname{Re}\mathbf{v}, \rho(g)\mathbf{a}\rangle| \gtrsim \frac{|\operatorname{Re}\mathbf{v}|}{\sqrt{\log d}} \geq \frac{|\mathbf{v}|}{\sqrt{2\log d}}. \qquad \square$$

Proof of Theorem 3.2. Let $\mathbf{a}$ be as in Lemma 3.3 (viewed as a column vector). Define $f\colon G \to \mathbb{R}$ via

$$f(g) = d\langle \rho(g), \mathbf{a}\mathbf{a}^\intercal\rangle_{\mathrm{HS}}.$$

Therefore, by Fourier inversion, $\widehat{f}$ is supported at ρ and $\widehat{f}(\rho) = \mathbf{a}\mathbf{a}^\intercal$. That is, $f \in V_\rho$ (as in (2.1)).

As in (2.2), let $V_{\rho,\mathbf{a}}$ denote the subspace of $L^2(G)$ consisting of all $x \in L^2(G)$ of the form $x(g) = d\langle \rho(g), \mathbf{a}\mathbf{v}^\dagger\rangle_{\mathrm{HS}}$ for some $\mathbf{v} \in \mathbb{C}^d$. From the discussions in Section 2.2, we see that M_f has exactly one nonzero eigenvalue, namely 1, and its eigenspace is $V_{\rho,\mathbf{a}}$.

We claim that

$$\|x\|_{L^\infty(G)} \gtrsim \sqrt{\frac{d}{\log d}}\|x\|_{L^2(G)}$$

for all $x \in V_{\rho,\mathbf{a}}$. Letting $x(g) = d\langle \rho(g), \mathbf{a}\mathbf{v}^\dagger\rangle_{\mathrm{HS}}$, we have

$$\|x\|_{L^\infty(G)} = d\sup_{g\in G}|\langle \rho(g), \mathbf{a}\mathbf{v}^\dagger\rangle_{\mathrm{HS}}| = d\sup_{g\in G}|\langle \mathbf{v}, \rho(g)\mathbf{a}\rangle| \gtrsim \frac{d}{\sqrt{\log d}}|\mathbf{v}|,$$

where the last inequality is Lemma 3.3. Furthermore, by Parseval,

$$\|x\|_{L^2(G)} = \sqrt{d}\,\|\widehat{x}(\rho)\|_{\mathrm{HS}} = \sqrt{d}\|\mathbf{a}\mathbf{v}^\dagger\|_{\mathrm{HS}} = \sqrt{d}|\mathbf{v}|.$$

Since $|G| = 2^d d!$, we have $d = (1+o(1))\log|G|/\log\log|G|$, which completes the proof. □

3.3. **Unweighted construction.** Let us first explain the setup for this entire subsection. Let G be any finite group (later on we will specialize to $G = S_d \ltimes (\mathbb{Z}/2\mathbb{Z})^d$). Let $\rho \in \widehat{G}$ be a real irreducible representation of G of dimension $d = d_\rho$, i.e., $\rho\colon G \to O(d)$ is a homomorphism. Finally, we assume that $\sqrt{|G|/\log|G|} > 15d$.

Fix a unit vector $\mathbf{a} \in \mathbb{R}^d$ and let

$$f(g) = \frac{1 - \mathbf{a}^\intercal\rho(g)\mathbf{a}}{2}. \tag{3.1}$$

By Cauchy–Schwarz, $f(g) \in [0,1]$. Furthermore, $f(g^{-1}) = f(g)$ since $\rho(g)$ has real entries.

Note that f is quite similar to the example given in the proof of Theorem 3.2, but shifted and scaled so that its values lie in $[0,1]$. The idea is to sample a random Cayley graph from f. Then its eigenvalues will be close to the original. Furthermore, we will show that the top nontrivial eigenspace (which corresponds to ρ) does not change much, so the estimate Lemma 3.3 will remain valid.

We now sample a random function h based on f. Let G' be the set of elements $g \in G$ with $g = g^{-1}$ and G'' be a subset of $G \setminus G'$ containing exactly one element of each set $\{g, g^{-1}\} \in \binom{G}{2}$ such that $g \neq g^{-1}$. For $g \in G'$, let $h(g)$ be 1 with probability $f(g)$ and zero otherwise. For $g \in G''$, let $h(g) = h(g^{-1}) = 1$ with probability $f(g)$ and 0 otherwise. The choices are independent across $G' \cup G''$. Note that $f(e) = 0$, so $h(e) = 0$.

For each $g \in G$, let P_g denote the matrix with columns and columns indexed by $G \times G$ with entry $1/|G|$ at position (gx, x) for each $x \in G$ and zero elsewhere. Viewing M_f and M_h as matrices (as described in Section 2.2), we have

$$\begin{aligned} M_h &= \sum_{g\in G'} h(g)P_g + \sum_{g\in G''} h(g)(P_g + P_{g^{-1}}) \\ &= \sum_{g\in G'} (h(g) - f(g))P_g + \sum_{g\in G''} (h(g) - f(g))(P_g + P_{g^{-1}}) + \sum_{g\in G} f(g)P_g. \end{aligned}$$

We first compute the spectrum of $\mathbb{E}[M_h] = M_f$.

Lemma 3.4. *Let h be as above. The spectrum of $\mathbb{E}[M_h]$ is $1/2$ with multiplicity 1, $-1/(2d)$ with multiplicity d, and 0 with multiplicity $|G| - d - 1$.*

Proof. By linearity of expectation, $\mathbb{E}[M_h] = M_f$. Furthermore, since

$$f(g) = \frac{1 - \mathbf{a}^\intercal\rho(g)\mathbf{a}}{2} = \frac{1}{2} - \frac{1}{2d}(d\langle \rho(g), \mathbf{a}\mathbf{a}^\intercal\rangle_{\mathrm{HS}}),$$

we see that $\widehat{f}(\mathrm{triv}_G) = 1/2$ and $\widehat{f}(\rho) = -\mathbf{a}\mathbf{a}^\intercal/(2d)$ by Fourier inversion, with $\widehat{f}$ only supported at these two representations (here triv_G is the trivial representation of G).

The analysis in Section 2.2 therefore shows that M_f has eigenvalues $1/2$ with multiplicity 1, $-1/(2d)$ with multiplicity d, and 0 for the rest. □

In order to establish quantitative concentration bounds regarding our sampling, we use the matrix Hoeffding inequality by Tropp [19]. Recall that for self-adjoint A, B, the notation $A \preceq B$ means that $B - A$ is positive semidefinite.

Theorem 3.5 ([19, Theorem 1.3])**.** *Consider a finite sequence $\{\mathbf{X}_k\}$ of independent, random, self-adjoint matrices with dimension d, and let $\{\mathbf{A}_k\}$ be a sequence of fixed self-adjoint matrices. Assume that each random matrix satisfies*

$$\mathbb{E}\mathbf{X}_k = \mathbf{0} \quad \text{and} \quad \mathbf{X}_k^2 \preceq \mathbf{A}_k^2 \quad \text{almost surely.}$$

Then, for all $t \geq 0$,

$$\mathbb{P}\left(\lambda_{\max}\left(\sum_k \mathbf{X}_k\right) \geq t\right) \leq d \cdot e^{-t^2/(8\sigma^2)} \quad \text{where} \quad \sigma^2 := \left\|\sum_k \mathbf{A}_k^2\right\|_{\mathrm{op}}.$$

Remark. The constant 8 can be replaced by 2 if $\mathbf{X}_k$ and $\mathbf{A}_k$ commute almost surely, which will hold true in our application. See [19, Remark 7.4].

Lemma 3.6. *Let h be as above. Then*

$$\mathbb{P}\left[\|M_h - \mathbb{E}[M_h]\|_{\mathrm{op}} \leq 4\sqrt{\frac{\log(6\,|G|)}{|G|}}\right] \geq \frac{2}{3}.$$

Proof. Note that when $g = g^{-1}$ we have that $P_g^2 = I/\,|G|^2$. Otherwise note that

$$(P_g + P_{g^{-1}})^2 = 2I/\,|G|^2 + P_{g^2} + P_{g^{-2}} \preceq 4I/\,|G|^2$$

as

$$-2I/\,|G|^2 + P_{g^2} + P_{g^{-2}} = (P_g - P_{g^{-1}})^2 \preceq \mathbf{0}.$$

Here we are using that $P_g - P_{g^{-1}}$ is antisymmetric.

Using that $|h(g) - f(g)| \leq 1$ almost surely, and applying Theorem 3.5 to

$$M_h - \mathbb{E}[M_h] = \sum_{g \in G'} (h(g) - f(g))P_g + \sum_{g \in G''} (h(g) - f(g))(P_g + P_{g^{-1}})$$

we find that

$$\mathbb{P}\Big(\lambda_{\max}\big(M_h - \mathbb{E}[M_h]\big) \geq t\Big) \leq |G| \exp\left(\frac{-t^2\,|G|}{16}\right)$$

for all $t \geq 0$. Applying the same inequality to $-M_h$ yields

$$\mathbb{P}\Big(\,\|M_h - \mathbb{E}[M_h]\|_{\mathrm{op}} \geq t\Big) \leq 2\,|G| \exp\left(\frac{-t^2\,|G|}{16}\right).$$

Setting $t = 4\,|G|^{-1/2}\,(\log 6\,|G|)^{1/2}$ yields the lemma. □

This allows us to control the spectrum of M_h.

Corollary 3.7. *With h as above, we have that with probability at least $2/3$ the number of eigenvalues of M_h in $[-1/d, -1/(3d)]$ is exactly d.*

Proof. This is an immediate consequence of Weyl's inequality on deviation of eigenvalues along with Lemmas 3.4 and 3.6.

More specifically, we have that M_h and $\mathbb{E}[M_h]$ are self-adjoint and $\|M_h - \mathbb{E}[M_h]\|_{\mathrm{op}} \leq 4\sqrt{\log(6\,|G|)/\,|G|}$ with probability at least $2/3$. Thus, writing $\nu_1 \geq \cdots \geq \nu_{|G|}$ for the eigenvalues of $\mathbb{E}[M_h]$ (which we know to be $1/2, 0, \ldots, 0, -1/(2d), \ldots, -1/(2d)$ with d copies of $-1/(2d)$ by Lemma 3.4) and $\mu_1 \geq \cdots \geq \mu_{|G|}$ for the eigenvalues of M_h, we have

$$|\mu_j - \nu_j| \leq 4\sqrt{\frac{\log(6\,|G|)}{|G|}} \leq \frac{1}{6d}$$

for all $1 \le j \le |G|$ by Weyl's inequality. The final inequality uses the assumption that $\sqrt{|G| / \log |G|} > 15d$. □

We now show that $\widehat{h}(\rho)$ and $\widehat{f}(\rho) = \mathbb{E}[\widehat{h}(\rho)]$ are close.

Lemma 3.8. *Let h and ρ be as above. Then*

$$\mathbb{P}\left[\|\widehat{h}(\rho) - \mathbb{E}[\widehat{h}(\rho)]\|_{\mathrm{op}} \le 4\sqrt{\frac{\log(6d)}{|G|}}\right] \ge \frac{2}{3}.$$

Proof. The proof is essentially identical to the proof of Lemma 3.6. Note that if $g = g^{-1}$ then $\rho(g)^2 = I_d$. Otherwise

$$(\rho(g) + \rho(g^{-1}))^2 = 2I_d + \rho(g^2) + \rho(g^{-2}) \preceq 4I_d$$

as

$$(\rho(g) - \rho(g^{-1}))^2 \preceq \mathbf{0}.$$

Here we are using that $\rho(g) - \rho(g^{-1}) = \rho(g) - \rho(g)^\intercal$ is antisymmetric.

Then, using the matrix Hoeffding bound Theorem 3.5, it immediately follows that

$$\mathbb{P}\left(\left\|\widehat{h}(\rho) - \mathbb{E}[\widehat{h}(\rho)]\right\|_{\mathrm{op}} \ge t\right) \le 2d \exp\left(\frac{-t^2 |G|}{16}\right).$$

Setting $t = 4 |G|^{-1/2} (\log 6 |G|)^{1/2}$ yields the lemma. □

We now show that the top eigenvector of $\widehat{h}(\rho)$ and $\widehat{f}(\rho) = \mathbb{E}[\widehat{h}(\rho)]$ are close. This is a special case of the Davis–Kahan Theorem [6]. We include a proof (adapted from [17, Theorem 5.9]) for completeness. Recall that $\widehat{h}(\rho)$ is real so its eigenvectors are real.

Lemma 3.9. *Let $h, \rho, \mathbf{a}$ be as above. Let $\mathbf{b}$ be a real unit eigenvector of the top eigenvalue of $\widehat{h}(\rho)$. With probability at least $2/3$ we have*

$$\min\{|\mathbf{a} + \mathbf{b}|, |\mathbf{a} - \mathbf{b}|\} \le 16\sqrt{2}d\sqrt{\frac{\log(6d)}{|G|}}.$$

Proof. Let

$$\Sigma = \mathbb{E}[\widehat{h}(\rho)] = \widehat{f}(\rho) = -\frac{\mathbf{a}\mathbf{a}^\intercal}{2d}$$

and

$$\Sigma' = \widehat{h}(\rho).$$

Then we have that

$$\begin{aligned}
\mathbf{a}^\intercal\Sigma\mathbf{a} - \mathbf{b}^\intercal\Sigma\mathbf{b} &= \mathbf{a}^\intercal\Sigma'\mathbf{a} - \mathbf{b}^\intercal\Sigma\mathbf{b} - \mathbf{a}^\intercal(\Sigma' - \Sigma)\mathbf{a} \\
&\le \mathbf{b}^\intercal\Sigma'\mathbf{b} - \mathbf{b}^\intercal\Sigma\mathbf{b} - \mathbf{a}^\intercal(\Sigma' - \Sigma)\mathbf{a} \\
&= \langle \Sigma - \Sigma', \mathbf{a}\mathbf{a}^\intercal - \mathbf{b}\mathbf{b}^\intercal\rangle_{\mathrm{HS}} \\
&\le \left\|\Sigma' - \Sigma\right\|_{\mathrm{op}} \cdot ||\mathbf{a}\mathbf{a}^\intercal - \mathbf{b}\mathbf{b}^\intercal||_{S_1} \\
&\le \sqrt{2}\left\|\Sigma' - \Sigma\right\|_{\mathrm{op}} \cdot \|\mathbf{a}\mathbf{a}^\intercal - \mathbf{b}\mathbf{b}^\intercal\|_{\mathrm{HS}}.
\end{aligned}$$

The second inequality is an application of (2.3), noting that $\|\cdot\|_{\mathrm{op}} = \|\cdot\|_{\mathrm{S}_\infty}$. The last step uses that for a matrix M of rank at most 2, one has $\|M\|_{S_1} \le \sqrt{2}\|M\|_{S_2} = \sqrt{2}\|M\|_{\mathrm{HS}}$.

Furthermore we have

$$\mathbf{a}^\intercal\Sigma\mathbf{a} - \mathbf{b}^\intercal\Sigma\mathbf{b} = \frac{1 - (\mathbf{b} \cdot \mathbf{a})^2}{2d}$$

and

$$\|\mathbf{a}\mathbf{a}^\intercal - \mathbf{b}\mathbf{b}^\intercal\|_{\mathrm{HS}}^2 = \mathrm{Tr}((\mathbf{a}\mathbf{a}^\intercal - \mathbf{b}\mathbf{b}^\intercal)^2) = 2 - 2(\mathbf{b} \cdot \mathbf{a})^2.$$

Therefore we have that

$$(1-(\mathbf{b}\cdot\mathbf{a})^2)^{1/2} \le 4d\left\|\Sigma'-\Sigma\right\|_{\mathrm{op}}.$$

Then

$$(\mathbf{a}\cdot\mathbf{b})^2 \ge 1-16d^2\left\|\Sigma'-\Sigma\right\|_{\mathrm{op}}^2.$$

Negating $\mathbf{b}$ if necessary so that $\mathbf{a}\cdot\mathbf{b}\ge 0$, we have

$$|\mathbf{a}-\mathbf{b}|^2 = 2-2(\mathbf{a}\cdot\mathbf{b}) \le 2\left(1-\sqrt{1-16d^2\left\|\Sigma'-\Sigma\right\|_{\mathrm{op}}^2}\right) \le 32d^2\left\|\Sigma'-\Sigma\right\|_{\mathrm{op}}^2.$$

By Lemma 3.8, with probability at least $3/4$, one has $\|\Sigma'-\Sigma\|_{\mathrm{op}}^2 \le 16\log(6d)/|G|$, and the lemma follows. □

We combine the concentration results derived so far.

Proposition 3.10. *Let $h, \mathbf{a}$ be defined as above. For $|G|$ sufficiently large, we have with probability at least $1/3$ that all of the following hold:*

- *M_h has exactly d eigenvalues in the interval $[-1/d, -1/(3d)]$.*
- *$\widehat{h}(\rho)$ has exactly one eigenvalue λ in $[-1/d, -1/(3d)]$.*
- *There is a real unit eigenvector $\mathbf{b}$ of $\widehat{h}(\rho)$ of eigenvalue λ with*

$$|\mathbf{a}-\mathbf{b}| \le 16\sqrt{2}d\sqrt{\frac{\log(6d)}{|G|}}.$$

Proof. This is an immediate application of Corollary 3.7 and Lemmas 3.8 and 3.9. Note that although we are union-bounding over the failures of three statements (with failure rate at most $1/3$ each), the event used in Lemma 3.9 is precisely that of Lemma 3.8. □

We are now in position to prove Theorem 1.2. The proof will mimic that of Theorem 3.2.

Proof of Theorem 1.2. As in Section 3.2, let $G = S_d \ltimes (\mathbb{Z}/2\mathbb{Z})^d$ and ρ be its standard representation on $\mathbb{C}^d$ (permuting and negating coordinates), which is easily seen to be real. Furthermore let $\mathbf{a}$ be the unit vector in the direction of $(1, 1/\sqrt{2}, \dots, 1/\sqrt{d})$, viewed as a column.

We sample h as in the beginning of this subsection (the assumption $\sqrt{|G|/\log|G|} > 15d$ holds for sufficiently large d). Let our graph be the Cayley graph with adjacency matrix $|G|\,M_h$. As we care only about scale-invariant properties of eigenspaces, we restrict attention to M_h, which acts on $V = \mathbb{C}^G$.

By Proposition 3.10, for $|G|$ large enough, with probability at least $2/3$ there are exactly d eigenvalues of M_h in $[-1/d, -1/(3d)]$, and $\widehat{h}(\rho)$ has one eigenvalue λ in this range with a unit eigenvector $\mathbf{b}$ satisfying

$$|\mathbf{a}-\mathbf{b}| \le 16\sqrt{2}d\sqrt{\frac{\log(6d)}{|G|}}. \tag{3.2}$$

By the characterization of eigenspaces of Cayley graphs in Section 2.2, we see that $\widehat{h}(\rho)$ contributes a d-dimensional eigenspace to M_h for each of its eigenvalues. Therefore we see that the d eigenvalues of M_h in $[-1/d, -1/(3d)]$ are precisely d copies of this eigenvalue λ.

In particular, M_h has an eigenvalue λ which has eigenspace precisely $V_{\rho,\mathbf{b}}$, which recall from (2.2) is

$$V_{\rho,\mathbf{b}} = \{x \in L^2(G) : x(g) = d\langle \rho(g), \mathbf{b}\mathbf{w}^\dagger\rangle_{\mathrm{HS}} \text{ for some } \mathbf{w} \in \mathbb{C}^d\}.$$

Now to show the construction satisfies the conclusion of the theorem. Let $x \in V_{\rho,\mathbf{b}}$. We wish to show that (recall $|G| = 2^d d!$ so that $d = (1+o(1))\log|G|/\log\log|G|$)

$$\|x\|_{L^\infty(G)} \gtrsim \sqrt{\frac{d}{\log d}}\,\|x\|_{L^2(G)} = (1+o(1))\frac{\sqrt{\log|G|}}{\log\log|G|}\,\|x\|_{L^2(G)}.$$

If $x(g) = d\left\langle \rho(g), \mathbf{b}\mathbf{v}^\dagger \right\rangle_{\mathrm{HS}} = d\mathbf{v}^\dagger \rho(g)\mathbf{b}$ for some $\mathbf{v} \in \mathbb{C}^d$, then by Parseval,

$$\|x\|_{L^2(G)} = \sqrt{d}\|\mathbf{b}\mathbf{v}^\dagger\|_{\mathrm{HS}} = \sqrt{d}|\mathbf{v}|.$$

Furthermore,

$$\begin{aligned}
\|x\|_{L^\infty(G)} &= \sup_{g\in G} |d\langle \rho(g), \mathbf{b}\mathbf{v}^\dagger\rangle_{\mathrm{HS}}| \\
&\geq \sup_{g\in G} |d\langle \rho(g), \mathbf{a}\mathbf{v}^\dagger\rangle_{\mathrm{HS}}| - \sup_{g\in G} |d\langle \rho(g), (\mathbf{a}-\mathbf{b})\mathbf{v}^\dagger\rangle_{\mathrm{HS}}| \\
&\gtrsim \frac{d|\mathbf{v}|}{\sqrt{\log d}} - d|\mathbf{a}-\mathbf{b}||\mathbf{v}| \\
&\gtrsim \sqrt{\frac{d}{\log d}}\|x\|_{L^2(G)}.
\end{aligned}$$

for sufficiently large d, by Lemma 3.3 and (3.2). □

Remark. To produce the graphs in Figure 1, we produce a Cayley graph on $G = S_d \ltimes (\mathbb{Z}/2\mathbb{Z})^d$ with each possible generator $g \neq e$ included with probability $C(1 - f(g))$. ($C = 2/3$ for the first figure and $C = 1/7$ for the second; these constants merely serve to sparsify the graphs for aesthetic purposes.)

In the proof of Theorem 1.2 above, we sample a graph via a similar procedure and deduce that with positive probability it has an unbounded eigenspace of dimension d. In fact, we can further deduce that this unbounded eigenspace has negative eigenvalue, and is the largest eigenvalue in magnitude after the trivial eigenvalue. Thus, a similar proof shows that the graph we sampled above, with positive probability, has an unbounded eigenspace of multiplicity $d \geq 2$ which contains the second and third eigenvalues. Therefore any possible spectral drawing of such a graph will have width at least $c\sqrt{\log|G|}/\log\log|G|$, as required.

Finally, in practice, we only sampled the small values $d = 3$ and $d = 4$. In this situation there is a decent probability of sampling a graph not satisfying the desired properties, namely of having the second and third eigenvalues come from the standard representation ρ of G, and that their eigenspace is precisely d-dimensional. To produce Figure 1, we check for these properties and resample until they hold.

4. Upper Bound

In this section we prove Theorems 1.3 and 1.4, showing that all Cayley graphs (Section 4.1) and transitive graphs (Section 4.2) on n vertices have an $O(\sqrt{\log n})$-bounded orthonormal eigenbasis.

4.1. **Cayley graphs.**

Lemma 4.1. *Given a set S of n unit vectors in $\mathbb{R}^d$ (resp. $\mathbb{C}^d$) we can find an orthonormal (resp. unitary) basis L of $\mathbb{R}^d$ (resp. $\mathbb{C}^d$) such that*

$$\max_{\mathbf{w}\in L, \mathbf{v}\in S} |\langle \mathbf{w}, \mathbf{v}\rangle| \lesssim \sqrt{\frac{\log(dn)}{d}}.$$

Furthermore, when $S \subseteq \mathbb{C}^d$ we can choose L to have all real vectors.

Proof. Let us first do the real case. Recall the following standard bound on the volume of spherical caps in high dimensions (e.g., [2, Lemma 2.2]): for a uniformly random unit vector $\mathbf{w} \in \mathbb{R}^d$ and fixed unit vector $\mathbf{v} \in \mathbb{R}^d$, one has

$$\mathbb{P}(\langle \mathbf{w}, \mathbf{v}\rangle \geq \epsilon) = \mathbb{P}(|\mathbf{w} - \mathbf{v}|^2 \leq 2 - 2\epsilon) \leq e^{-d\epsilon^2/2}. \tag{4.1}$$

Let $\epsilon = \sqrt{2\log(4dn)/d}$. Then applying union bound with (4.1), we find that an orthonormal basis $L = \{\mathbf{w}_1, \dots, \mathbf{w}_d\}$ uniformly at random satisfies

$$\mathbb{P}\left(\max_{\mathbf{w}\in L, \mathbf{w}\in S} |\langle \mathbf{w}, \mathbf{v}\rangle| \geq \epsilon\right) \leq 2dne^{-d\epsilon^2/2} \leq \frac{1}{2}.$$

So there is some L such that $\langle \mathbf{w}, \mathbf{v}\rangle < \epsilon \lesssim \sqrt{\log(dn)/d}$ for all $\mathbf{w} \in L$ and $\mathbf{v} \in S$.

For the complex case, for each $v \in S$ write $\mathbf{v} = \mathbf{v}_1 + i\mathbf{v}_2$ for $\mathbf{v}_1, \mathbf{v}_2 \in \mathbb{R}^n$. We apply the real case of this lemma to the set

$$S' = \left\{\frac{\mathbf{v}_1}{\|\mathbf{v}_1\|_2} : \mathbf{v} \in S\right\} \cup \left\{\frac{\mathbf{v}_2}{\|\mathbf{v}_2\|_2} : \mathbf{v} \in S\right\}$$

to obtain a basis L of $\mathbb{R}^d \subseteq \mathbb{C}^d$. Then for each $\mathbf{w} \in L$ and $\mathbf{v} \in S$, we have

$$|\langle \mathbf{w}, \mathbf{v}\rangle| \leq |\langle \mathbf{w}, \mathbf{v}_1\rangle| + |\langle \mathbf{w}, \mathbf{v}_2\rangle| \leq \left|\left\langle \mathbf{w}, \frac{\mathbf{v}_1}{\|\mathbf{v}_1\|_2}\right\rangle\right| + \left|\left\langle \mathbf{w}, \frac{\mathbf{v}_2}{\|\mathbf{v}_2\|_2}\right\rangle\right| \lesssim \sqrt{\frac{\log(dn)}{d}}. \qquad \square$$

Now we are ready to prove Theorem 1.3. In essence our argument amounts to choosing a random unitary basis (via Lemma 4.1) for each eigenspace coming from the representation theory of the group G of the Cayley graph. However, in order to choose a *real* orthonormal eigenbasis, we essentially pair up conjugate irreducible representations. This technicality is unnecessary if we only wish to find a unitary eigenbasis.

Proof of Theorem 1.3. Let S be the symmetric generating set of the Cayley graph and $\mathbb{1}_S$ be the corresponding indicator function. Recall, from Section 2.2, the orthogonal decomposition of $V = \mathbb{C}^G$ as

$$V = \bigoplus_{\rho\in\widehat{G}} V_\rho = \bigoplus_{\rho\in\widehat{G}} \bigoplus_{j=1}^{d_\rho} V_{\rho,\mathbf{v}_j^\rho}. \tag{4.2}$$

See Section 2.2 for the definitions of V_ρ and $V_{\rho,\mathbf{v}}$. Here the vectors $\mathbf{v}_j^\rho$, $j \in [d_\rho]$, form a unitary eigenbasis of $\widehat{\mathbb{1}_S}(\rho)$ with respective eigenvalues $\lambda_{\rho,j}$. The eigenspace of $M_{\mathbb{1}_S}$ corresponding to some eigenvalue λ is a direct sum of all components $V_{\rho,\mathbf{v}_j^\rho}$ with $\lambda_{\rho,j} = \lambda$. Using this decomposition, we shall construct a real-valued unitary eigenbasis for the operator $M_{\mathbb{1}_S}$ on V as follows:

- For each conjugate pair $(\rho, \overline{\rho})$ and eigenvector $\mathbf{b}$ of $\widehat{\mathbb{1}_S}(\rho)$, we will find a real-valued unitary basis of $V_{\rho,\mathbf{b}} \oplus V_{\overline{\rho},\overline{\mathbf{b}}}$.
- For each irreducible representation $\rho \in \widehat{G}$ such that ρ and $\overline{\rho}$ are isomorphic (such representations are called *self-dual*), we will find a special eigenbasis of $\widehat{\mathbb{1}_S}(\rho)$, and a specific construction giving a unitary basis of V_ρ that will depend on whether the matrix Q satisfying $\rho Q = Q\overline{\rho}$ is symmetric or antisymmetric.

Note that the second case includes ρ which can be realized as a real representation, but not all self-dual representations are of this form (e.g. the two-dimensional irreducible representation of the quaternion group Q_8). Furthermore, we will ensure that all the functions x chosen as basis elements above satisfy

$$\|x\|_{L^\infty(G)} \lesssim \sqrt{\log|G|}\, \|x\|_{L^2(G)}.$$

As a model case we consider $\rho \in \widehat{G}$ a real irreducible representation of dimension $d = d_\rho$, acting on $\mathbb{R}^d$. We will not, strictly speaking, need this analysis in the final argument as our treatment of self-dual representations is strictly more general. Let $\mathbf{b} \in \mathbb{R}^d$ be a real eigenvector of $\widehat{\mathbb{1}_S}(\rho)$. Recall

$$V_{\rho,\mathbf{b}} = \{x \in L^2(G) : x(g) = d\langle \rho(g), \mathbf{b}\mathbf{w}^\dagger\rangle_{\mathrm{HS}} \text{ for some } \mathbf{w} \in \mathbb{C}^d\}.$$

Similar to in the proof of Theorem 1.2, writing $x_{\mathbf{v}} : G \to \mathbb{C}$ for the function $x_{\mathbf{v}}(g) = d\langle \rho(g), \mathbf{b}\mathbf{v}^\dagger\rangle_{\mathrm{HS}} = d\langle \rho(g)\mathbf{v}, \mathbf{b}\rangle$, we have

$$\|x_{\mathbf{v}}\|_{L^\infty(G)} = d \cdot \sup_{g\in G} |\langle \mathbf{v}, \rho(g)\mathbf{b}\rangle|$$

and

$$\|x_{\mathbf{v}}\|_{L^2(G)} = \sqrt{d}|\mathbf{b}||\mathbf{v}| = \sqrt{d}|\mathbf{v}|.$$

By the real version of Lemma 4.1 applied to $S = \{\rho(g)\mathbf{b} : g \in G\}$, there is an orthonormal basis L of $\mathbb{R}^d$ such that (note that $d^2 \le |G|$)

$$\sup_{\mathbf{v}\in L, g\in G} |\langle \mathbf{v}, \rho(g)\mathbf{b}\rangle| \lesssim \sqrt{\frac{\log(d\,|G|)}{d}} \lesssim \sqrt{\frac{\log|G|}{d}}.$$

Then for each $\mathbf{v} \in L$ we have

$$\|x_{\mathbf{v}}\|_{L^\infty(G)} \lesssim \sqrt{\log|G|}\|x_{\mathbf{v}}\|_{L^2(G)}.$$

By Parseval's identity, we see that $\{x_{\mathbf{v}}/\sqrt{d} : \mathbf{v} \in L\}$ forms a unitary basis of $V_{\rho,\mathbf{b}}$. Also note that $x_{\mathbf{v}}$ is real-valued for each $\mathbf{v} \in \mathbb{R}^d$. This completes the case of real ρ.

Next, let $(\rho, \overline{\rho})$ be a conjugate pair of irreducible representations with ρ and $\overline{\rho}$ not isomorphic to each other. Again let $d = d_\rho$. For each eigenvector $\mathbf{b} \in \mathbb{C}^d$ of $\widehat{\mathbb{1}_S}(\rho)$, by the complex version of Lemma 4.1, we find a unitary basis L of $\mathbb{C}^d$ so that $\sup_{\mathbf{v}\in L, g\in G} |\langle \mathbf{v}, \rho(g)\mathbf{b}\rangle| \lesssim \sqrt{(\log|G|)/d}$. Again writing $x_{\mathbf{v}}(g) = d\langle \rho(g), \mathbf{b}\mathbf{v}^\dagger\rangle$, we find that $\{x_{\mathbf{v}}/\sqrt{d} : \mathbf{v} \in L\}$ is a unitary basis of $V_{\rho,\mathbf{b}}$ and

$$\|x_{\mathbf{v}}\|_{L^\infty(G)} \lesssim \sqrt{\log|G|}\|x_{\mathbf{v}}\|_{L^2(G)}.$$

Likewise, taking conjugates, we see that $\{\overline{x_{\mathbf{v}}}/\sqrt{d} : \mathbf{v} \in L\}$ is a unitary basis of $V_{\overline{\rho},\overline{\mathbf{b}}}$. Recall that if $M_{\mathbb{1}_S} x_{\mathbf{v}} = \lambda x_{\mathbf{v}}$ then $M_{\mathbb{1}_S}\overline{x_{\mathbf{v}}} = \lambda\overline{x_{\mathbf{v}}}$ as $M_{\mathbb{1}_S}$ is symmetric and hence λ is real.

The collection of $2d$ vectors

$$y^0_{\mathbf{v}} = \frac{x_{\mathbf{v}} + \overline{x_{\mathbf{v}}}}{\sqrt{2d}} \quad \text{and} \quad y^1_{\mathbf{v}} = \frac{x_{\mathbf{v}} - \overline{x_{\mathbf{v}}}}{i\sqrt{2d}},$$

as $\mathbf{v}$ ranges over L, forms a real-valued unitary basis of $V_{\rho,\mathbf{b}} \oplus V_{\overline{\rho},\overline{\mathbf{b}}}$. Furthermore,

$$\|y^0_{\mathbf{v}}\|_{L^\infty(G)} \le \sqrt{\frac{2}{d}}\,\|x_{\mathbf{v}}\|_{L^\infty(G)} \lesssim \sqrt{\frac{\log|G|}{d}}\,\|x_{\mathbf{v}}\|_{L^2(G)} = \sqrt{\log|G|}\,\|y^0_{\mathbf{v}}\|_{L^2(G)}$$

and similarly for $y^1_{\mathbf{v}}$. This completes the case of non-self-dual complex irreducible representations.

Finally, let ρ be a self-dual irreducible representation. Again let $d = d_\rho$. Note that $\widehat{\mathbb{1}_S}(\rho)$ is Hermitian, hence we can choose coordinates on the representation so that it is a real diagonal matrix. Having done so, we now note that $g \mapsto \rho(g)$ and $g \mapsto \overline{\rho(g)}$ are isomorphic representations on the same space, since ρ is self-dual (where complex conjugation is done in the natural way with respect to the coordinates chosen on the space). Hence there is a unitary operator Q so that

$$\rho(g)Q = Q\overline{\rho(g)}$$

for all $g \in G$. Thus

$$\rho(g)Q\overline{Q} = Q\overline{\rho(g)Q} = Q\overline{Q}\rho(g)$$

for all $g \in G$. By Schur's lemma, we deduce that

$$Q\overline{Q} = \omega I$$

for some $\omega \in \mathbb{C}$. Since Q is unitary, $QQ^\dagger = I$ (we use † to denote Hermitian transpose and $^\intercal$ for transpose), which yields

$$Q = \omega Q^\intercal = \omega^2 Q.$$

Since Q is invertible, we deduce $\omega^2 = 1$, and hence $\omega \in \{\pm 1\}$.

From $\widehat{\mathbb{1}_S}(\rho) = \sum_{g\in S}\rho(g)/|G|$ and $\rho(g)Q = Q\overline{\rho(g)}$ we obtain

$$Q^\dagger\widehat{\mathbb{1}_S}(\rho)Q = \overline{\widehat{\mathbb{1}_S}(\rho)} = \widehat{\mathbb{1}_S}(\rho),$$

since $\widehat{\mathbb{1}_S}(\rho)$ is a real diagonal matrix, as noted earlier. Therefore Q and $\widehat{\mathbb{1}_S}(\rho)$ commute. Now choose a unitary simultaneous eigenbasis of Q and $\widehat{\mathbb{1}_S}(\rho)$, which can be done by the spectral theorem as both operators are normal.

We will actually take a specific basis with more structure. First, note that if $\mathbf{b}$ is an eigenvector of Q with eigenvalue λ then

$$Q^{-1}\overline{\mathbf{b}} = Q^\dagger\overline{\mathbf{b}} = \overline{Q^T\mathbf{b}} = \omega^{-1}\overline{Q\mathbf{b}} = \omega^{-1}\overline{\lambda\mathbf{b}}.$$

Thus $\overline{\mathbf{b}}$ is an eigenvector of Q with eigenvalue $(\omega^{-1}\overline{\lambda})^{-1} = \omega\lambda$, using $|\lambda| = 1$ since Q is unitary. Now we break into sub-cases depending on the value of $\omega \in \{\pm 1\}$.

If $\omega = -1$, then we see that our unitary simultaneous eigenbasis of Q and $\widehat{\mathbb{1}_S}(\rho)$ can be chosen so that if $\mathbf{b}$ is in it, then so is $\overline{\mathbf{b}}$, since $\mathbf{b}$ and $\overline{\mathbf{b}}$ lie in distinct orthogonal eigenspaces of Q. For such an eigenvector $\mathbf{b}$, as earlier we can apply Lemma 4.1 to obtain a unitary basis L of $\mathbb{C}^d$ so that $\sup_{\mathbf{v}\in L, g\in G}|\langle\mathbf{v}, \rho(g)\mathbf{b}\rangle| \lesssim \sqrt{(\log|G|)/d}$. Again writing $x_{\mathbf{v}}(g) = d\langle\rho(g), \mathbf{b}\mathbf{v}^\dagger\rangle$, we find that $\{x_{\mathbf{v}}/\sqrt{d} : \mathbf{v}\in L\}$ is a unitary basis of $V_{\rho,\mathbf{b}}$ and $\|x_{\mathbf{v}}\|_{L^\infty(G)} \lesssim \sqrt{\log|G|}\|x_{\mathbf{v}}\|_{L^2(G)}$.

Note that $\overline{x_{\mathbf{v}}} \in V_{\rho,\overline{\mathbf{b}}}$ since

$$\overline{x_{\mathbf{v}}}(g) = d\overline{\langle\rho(g)\mathbf{v}, \mathbf{b}\rangle} = d\langle\overline{\rho(g)}\mathbf{v}, \overline{\mathbf{b}}\rangle = d\langle Q^\dagger\rho(g)Q\overline{\mathbf{v}}, \overline{\mathbf{b}}\rangle = d\langle\rho(g)Q\overline{\mathbf{v}}, Q\overline{\mathbf{b}}\rangle = d\langle\rho(g)Q\overline{\mathbf{v}}, \omega\lambda\overline{\mathbf{b}}\rangle.$$

This shows that $\overline{V_{\rho,\mathbf{b}}} = V_{\rho,\overline{\mathbf{b}}}$. Furthermore, as $M_{\mathbb{1}_S}x_{\mathbf{v}} = \lambda x_{\mathbf{v}}$, we have $M_{\mathbb{1}_S}\overline{x_{\mathbf{v}}} = \lambda\overline{x_{\mathbf{v}}}$ since λ is real. Then, as $\mathbf{v}$ varies over L, the functions

$$y^0_{\mathbf{v}} = \frac{x_{\mathbf{v}} + \overline{x_{\mathbf{v}}}}{\sqrt{2d}} \quad\text{and}\quad y^1_{\mathbf{v}} = \frac{x_{\mathbf{v}} - \overline{x_{\mathbf{v}}}}{i\sqrt{2d}}$$

form a real-valued unitary basis of $V_{\rho,\mathbf{b}} \oplus V_{\rho,\overline{\mathbf{b}}}$. This completes the proof of the case $\omega = -1$.

If $\omega = 1$, then for every eigenvector $\mathbf{b}$ of Q with eigenvalue λ, $\overline{\mathbf{b}}$ is another eigenvector of Q also with the same eigenvalue λ. Thus every eigenspace U of Q satisfies $U = \overline{U}$. A $\mathbb{C}$-vector space U satisfying $U = \overline{U}$ is always the $\mathbb{C}$-extension of the $\mathbb{R}$-vector space $\operatorname{Re}U = \{(\operatorname{Re}u_1, \dots, \operatorname{Re}u_d) : (u_1, \dots, u_d) \in U\}$ (since every $\mathbf{v}\in U$ can be written as $\mathbf{x}+i\mathbf{y}$ with $\mathbf{x} = (\mathbf{v}+\overline{\mathbf{v}})/2$ and $\mathbf{y} = (\mathbf{v}-\overline{\mathbf{v}})/(2i)$ both having real coordinates). Thus we can choose an orthonormal basis of $\mathbb{R}^d$ consisting of real-valued eigenvectors $\mathbf{b}$ of Q.

Fix such a real eigenvector $\mathbf{b}$ of Q. For any $\mathbf{v}\in\mathbb{C}^d$, setting $x_{\mathbf{v}}(g) = d\langle\rho(g), \mathbf{b}\mathbf{v}^\dagger\rangle$ as before, we have $\overline{x_{\mathbf{v}}} \in V_{\rho,\mathbf{b}}$ since

$$\overline{x_{\mathbf{v}}}(g) = d\overline{\langle\rho(g)\mathbf{v}, \mathbf{b}\rangle} = d\langle\overline{\rho(g)}\mathbf{v}, \mathbf{b}\rangle = d\langle Q^\dagger\rho(g)Q\overline{\mathbf{v}}, \mathbf{b}\rangle = d\langle\rho(g)Q\overline{\mathbf{v}}, Q\mathbf{b}\rangle = d\langle\rho(g)Q\overline{\mathbf{v}}, \lambda\mathbf{b}\rangle.$$

Thus $\overline{V_{\rho,\mathbf{b}}} = V_{\rho,\mathbf{b}}$, which then must be the $\mathbb{C}$-extension of the d-dimensional $\mathbb{R}$-vector space $\operatorname{Re}V_{\rho,\mathbf{b}}$.

We now apply Lemma 4.1 to find a unitary basis L of the d-dimensional $\mathbb{R}$-vector space $\{\mathbf{v}\in\mathbb{C}^d : x_{\mathbf{v}} \in \operatorname{Re}V_{\rho,\mathbf{b}}\}$ satisfying $\sup_{\mathbf{v}\in L, g\in G}|\langle\mathbf{v}, \rho(g)\mathbf{b}\rangle| \lesssim \sqrt{(\log|G|)/d}$. Then $\{x_{\mathbf{v}}/\sqrt{d} : \mathbf{v}\in L\}$ is a real-valued unitary basis of $V_{\rho,\mathbf{b}}$ with $\|x_{\mathbf{v}}\|_{L^\infty(G)} \lesssim \sqrt{\log|G|}\|x_{\mathbf{v}}\|_{L^2(G)}$. □

4.2. **Vertex-transitive graphs.** We now extend Theorem 1.3 to vertex-transitive graphs; the idea is the same as before, except we first lift to a Cayley graph on the automorphism group G of the original. This trick is closely related to the proof of [5, Theorem 2.2].

Proof of Theorem 1.4. Let G denote the automorphism group of the given vertex-transitive graph, acting on the vertex-set from the right. Fix a vertex as the root of the graph. Let H denote the stabilizer of the root. Then the vertices of the graph are given by right cosets Hg, $g \in H\backslash G$, with the root corresponding to the trivial coset H. Thus $|H\backslash G| = n$.

Let $f\colon H\backslash G \to \mathbb{C}$ denote the edge-weights from the root to other vertices of the graph. Since G induces automorphisms on the graph, the edgeweight of (H, Hg) equals that of $(Hh, Hgh) = (H, Hgh)$ for all $h \in H$. Hence $f(gh) = f(g)$ for all $g \in G$ and $h \in H$. So we can view f as a function $f\colon G \to \mathbb{C}$ that is H-invariant from both the left and right.

A function on the vertex set is represented as $x\colon H\backslash G \to \mathbb{C}$, which we will view as a left-H-invariant $x\colon G \to \mathbb{C}$, i.e., $x(hg) = x(g)$ for all $g \in G$ and $h \in H$.

A function $x\colon G \to \mathbb{C}$ satisfies $x(hg) = x(g)$ for all $g \in G$ and $h \in H$ (i.e., it is left-H-invariant) if and only if $\widehat{x}(\rho) = \rho(h)\widehat{x}(\rho)$ for all $\rho \in \widehat{G}$ and $h \in H$. The latter condition is equivalent to saying that the column-space of $\widehat{x}(\rho)$ lies in U_ρ, the 1-eigenspace of $\rho|_H$:

$$U_\rho := \{v \in W_\rho : \rho(h)v = v \text{ for all } h \in H\}.$$

The forward implication follows from the Fourier transform formula $\widehat{x}(\rho) = \mathbb{E}_{g\in G} x(g)\rho(g)$, while the reverse implication follows from the inversion formula $x(g) = \sum_\rho d_\rho \operatorname{tr}(\rho(g)^\dagger \widehat{x}(\rho))$.

Let $m_\rho = \dim U_\rho$. By counting the dimension of the space of all left-H-invariant functions, we obtain

$$\sum_{\rho\in\widehat{G}} d_\rho m_\rho = |H\backslash G| = n. \tag{4.3}$$

Indeed, the condition that $\widehat{x}(\rho)$ has column-space contained within U_ρ restricts $\widehat{x}(\rho)$ to a $d_\rho m_\rho$-dimensional subspace of $\operatorname{End} W_\rho$.

Since $f\colon G \to \mathbb{C}$ is both left- and right-H-invariant, $\widehat{f}(\rho) = \rho(h)\widehat{f}(\rho)\rho(h')$ for all $h, h' \in H$. So $\widehat{f}(\rho)$ leaves U_ρ invariant. Let $v_1^\rho, \dots, v_{m_\rho}^\rho \in U_\rho$ be an eigenbasis of the action of $\widehat{f}(\rho)$ on U_ρ.

For each $\rho \in \widehat{G}$, choose a unitary basis $a_1^\rho, \dots, a_{d_\rho}^\rho$ of W_ρ, and for each $j \in [d_\rho]$ and $k \in [m_\rho]$, define $x_{j,k}^\rho\colon G \to \mathbb{C}$ by setting $\widehat{x_{j,k}^\rho}(\rho) = v_k^\rho (a_j^\rho)^\dagger/\sqrt{d_\rho}$ and $\widehat{x}(\rho') = 0$ for all $\rho' \neq \rho$.

The functions $x_{j,k}^\rho$, with $\rho \in \widehat{G}$, $j \in [d_\rho]$, $k \in [m_\rho]$ satisfy the following properties.

(1) $x_{j,k}^\rho$ is left-H-invariant, i.e., are functions $H\backslash G \to \mathbb{C}$ (since the columns of $\widehat{x_{j,k}^\rho}$ are in U_ρ),
(2) $x_{j,k}^\rho$ is an eigenfunctions of M_f (since $\widehat{f}(\rho)\widehat{x_{j,k}^\rho}(\rho)$ is a scalar multiple of $\widehat{x_{j,k}^\rho}(\rho)$)
(3) The functions $x_{j,k}^\rho$ are pairwise orthogonal and $\|x_{j,k}^\rho\|_2 = 1$. Indeed, by Parseval, one has $\langle x_{j,k}^\rho, x_{j',k'}^{\rho'}\rangle_{\mathrm{HS}} = 0$ if $\rho \neq \rho'$, and

$$\langle x_{j,k}^\rho, x_{j',k'}^\rho\rangle = d_\rho \langle \widehat{x_{j,k}^\rho}(\rho), \widehat{x_{j',k'}^\rho}(\rho)\rangle_{\mathrm{HS}} = \langle a_{j'}^\rho, a_j^\rho\rangle\langle v_k^\rho, v_{k'}^\rho\rangle = 1_{j=j'}1_{k=k'}.$$

(4) They form a basis of all functions $H\backslash G \to \mathbb{C}$ (by orthogonality and dimension counting (4.3)).

Furthermore, we have for each ρ, $j \in [d_\rho]$, and $k \in [m_\rho]$ that

$$x_{j,k}^\rho(g) = d_\rho \left\langle \rho(g), \widehat{x_{j,k}^\rho}(\rho)\right\rangle_{\mathrm{HS}} = \sqrt{d_\rho}\left\langle \rho(g), v_k^\rho (a_j^\rho)^\dagger\right\rangle_{\mathrm{HS}} = \sqrt{d_\rho}\left\langle \rho(g)a_j^\rho, v_k^\rho\right\rangle.$$

For each fixed $\rho \in \widehat{G}$, set $S_\rho = \{\rho(g)v_k^\rho : g \in G, k \in [m_\rho]\}$. Now since $\rho(h)v_k^\rho = v_k^\rho$ for all $k \in [m_\rho]$ (because M_ρ is the 1-eigenspace of $\rho|_H$), we see that $|S| \le |G/H|m_\rho \le n^2$. By the complex version of Lemma 4.1, there exists a choice of the unitary basis $a_1^\rho, \dots, a_{d_\rho}^\rho$ in the definition of $x_{j,k}^\rho(g)$ earlier so that

$$\sup_{j\in[d_\rho],k\in[m_\rho]} \sup_{g\in G} |x_{j,k}^\rho(g)| = \sup_{j\in[d_\rho],k\in[m_\rho]} \sqrt{d_\rho}\left\langle \rho(g)a_j^\rho, v_k^\rho\right\rangle \lesssim \sqrt{\log n}.$$

Thus the functions $x_{j,k}^\rho$, ranging over all irreducible representations ρ and indices $j \in [d_\rho]$ and $k \in [m_\rho]$, form a unitary $\sqrt{\log n}$-bounded eigenbasis.

To obtain a real orthonormal eigenbasis, we can repeat the technique in the proof of Theorem 1.3 in the previous subsection. We omit the details. □

5. Small-set expansion in random Cayley graphs

In this section we prove Theorem 1.7, extending Naor's theorem [16] on small-set expansion for random Cayley graphs to nonabelian groups. Recall the definition of the Schatten p-norm $\|\cdot\|_{S_p}$ from Section 2.3. We state below a Hausdorff–Young inequality for groups, which is standard though we include its short proof (see [14] for a proof for locally compact unimodular groups).

Lemma 5.1 (Hausdorff—Young inequality for groups)**.** *Let G be a finite group and $f\colon G \to \mathbb{C}$. For any $1 \le p \le 2 \le q \le \infty$ with $1/p + 1/q = 1$, one has*

$$\|f\|_{S_q} \le \|f\|_{L^p(G)}\,.$$

Proof. By the Riesz—Thorin interpolation theorem, it suffices to check the inequality for $(p,q) = (2,2)$ and $(1,\infty)$. For $(p,q) = (2,2)$, we have $\|f\|_{S_2} = \|f\|_2$ by Parseval. For $(p,q) = (1,\infty)$ we have

$$\|f\|_{S_\infty} = \max_{\rho\in\widehat{G}} \|\widehat{f}(\rho)\|_{\mathrm{op}} = \max_{\rho\in\widehat{G}} \|\mathbb{E}_g f(g)\rho(g)\|_{\mathrm{op}} \le \mathbb{E}_{g\in G}\,|f(g)| = \|f\|_1$$

as $\|\rho(g)\|_{\mathrm{op}} = 1$ for all $g \in G$. □

Lemma 5.2. *Let G be a finite group. For functions $f, x\colon G \to \mathbb{C}$ and real $p \ge 1$, one has*

$$|\langle x, f * x\rangle| \le \|f\|_{S_p}\,\|x\|^2_{2p/(p+1)}\,.$$

Proof. For each $\rho \in \widehat{G}$ we have

$$\begin{aligned}\left|\left\langle \widehat{x}(\rho), \widehat{f}(\rho)\widehat{x}(\rho)\right\rangle_{\mathrm{HS}}\right| &= \left|\mathrm{Tr}\left(\widehat{x}(\rho)^\dagger \widehat{f}(\rho)\widehat{x}(\rho)\right)\right| = \left|\mathrm{Tr}\left(\widehat{x}(\rho)\widehat{x}(\rho)^\dagger \widehat{f}(\rho)\right)\right| \\ &\le \|\widehat{f}(\rho)\|_{S_p}\|\widehat{x}(\rho)\widehat{x}(\rho)^\dagger\|_{S_{p/(p-1)}} = \|\widehat{f}(\rho)\|_{S_p}\|\widehat{x}(\rho)\|^2_{S_{2p/(p-1)}}. \qquad (5.1)\end{aligned}$$

Here the inequality step uses the matrix Hölder inequality (2.3): $\mathrm{Tr}(AB) \le \|A\|_{S_p}\,\|B\|_{S_{p/(p-1)}}$ for all $p \in [1,\infty]$. The last step uses that the singular values of a matrix A are the square roots of the singular values of $AA^\dagger$.

Thus, applying the convolution and Parseval identities for the nonabelian Fourier transform (Section 2.1), we have

$$\begin{aligned}|\langle x, f * x\rangle| &= \left|\sum_{\rho\in\widehat{G}} d_\rho \left\langle \widehat{x}(\rho), \widehat{f}(\rho)\widehat{x}(\rho)\right\rangle_{\mathrm{HS}}\right| && \text{[Convolution \& Parseval]}\\ &\le \sum_{\rho\in\widehat{G}} d_\rho \|\widehat{f}(\rho)\|_{S_p}\|\widehat{x}(\rho)\|^2_{S_{2p/(p-1)}} && \text{[by (5.1)]}\\ &\le \left(\sum_\rho d_\rho \|\widehat{f}(\rho)\|^p_{S_p}\right)^{1/p}\left(\sum_\rho d_\rho \|\widehat{x}(\rho)\|^{2p/(p-1)}_{S_{2p/(p-1)}}\right)^{(p-1)/p} && \text{[Hölder's inequality]}\\ &= \|f\|_{S_p}\,\|x\|^2_{S_{2p/(p-1)}} && \text{[by (2.4)]}\\ &\le \|f\|_{S_p}\,\|x\|^2_{2p/(p+1)}\,. && \text{[by Lemma 5.1]}\end{aligned}$$

This proves the desired inequality. □

Naor proved the following uniform bound on the Schatten norms of random Cayley graphs via a novel Azuma-type concentration inequality in uniformly smooth normed spaces.

Lemma 5.3 ([16, Lemma 4.1])**.** *There exists a universal constant $C > 0$ with the following property. For any positive integer k and any finite group G, if $g_1, \dots, g_k \in G$ are chosen independently and*

uniformly at random, then, with probability at least $1/2$*, the function* $f\colon G \to \mathbb{R}$ *given by*

$$f = \mathbb{1}_{\{g_1\}} + \mathbb{1}_{\{g_1^{-1}\}} + \cdots + \mathbb{1}_{\{g_k\}} + \mathbb{1}_{\{g_k^{-1}\}} - \frac{2k}{|G|} \tag{5.2}$$

satisfies

$$\|f\|_{S_p} \le C\,|G|^{-1+1/p}\sqrt{pk}$$

simultaneously for every integer $p \ge 2$*.*

Using this concentration lemma, we can now prove Theorem 1.7. Unlike the proof in [16] for abelian groups, we do not need to rely on a bounded eigenbasis.

Proof of Theorem 1.7. Let $g_1, \dots, g_k \in G$ be the random group elements generating the Cayley (multi)graph. Define $f\colon G \to \mathbb{R}$ as in (5.2).

Let $X \subseteq G$ with $1 < |X| \le |G|/2$. Define a function $x\colon G \to \mathbb{R}$ by

$$x = |G \setminus X|\,\mathbb{1}_X - |X|\,\mathbb{1}_{G\setminus X}.$$

It is straightforward to check that

$$\langle x, f * x\rangle = \frac{2k}{|G|}\,|X|\,|G\setminus X| - e(X, G\setminus X).$$

We also have

$$\|x\|_{2p/(p+1)} = \left(\frac{|X|\,|G\setminus X|^{\frac{2p}{p+1}} + |X|^{\frac{2p}{p+1}}\,|G\setminus X|}{|G|}\right)^{\frac{p+1}{2p}}.$$

Applying the inequality $|\langle f * x, x\rangle| \le \|f\|_{S_p}\,\|x\|_{2p/(p+1)}^2$ from Lemma 5.2, and with the upper bound $\|f\|_{S_p} \lesssim |G|^{-1+1/p}\sqrt{pk}$ from Lemma 5.3, we obtain that with probability at least $1/2$, one has

$$\left|\frac{2k}{|G|}\,|X|\,|G\setminus X| - e(X, G\setminus X)\right| \lesssim |G|^{-1+1/p}\sqrt{pk}\left(\frac{|X|\,|G\setminus X|^{\frac{2p}{p+1}} + |X|^{\frac{2p}{p+1}}\,|G\setminus X|}{|G|}\right)^{\frac{p+1}{p}}$$

simultaneously for all positive integers p. Dividing both sides by $2k\,|X|\,|G\setminus X|\,/\,|G|$, we obtain

$$\begin{aligned}\left|\frac{e(X, G\setminus X)}{\frac{2k}{|G|}\,|X|\,|G\setminus X|} - 1\right| &\lesssim |G|^{1/p}\sqrt{\frac{p}{k}}\left(\frac{|X|^{\frac{1}{p+1}}\,|G\setminus X|^{\frac{p}{p+1}} + |X|^{\frac{p}{p+1}}\,|G\setminus X|^{\frac{1}{p+1}}}{|G|}\right)^{\frac{p+1}{p}}\\ &\lesssim |X|^{1/p}\sqrt{\frac{p}{k}},\end{aligned}$$

where in the last step we apply the inequality $x^t(1-x)^{1-t} + x^{1-t}(1-x)^t \le x^t + x^{1-t} \le 2x^t$ for $x = |X|\,/\,|G| \le 1/2$ and $t = 1/(p+1) \in [0, 1/2]$. Finally, setting $p = \lceil \log |X| \rceil$, we see the final expression has an upper bound of $O(\sqrt{(\log |X|)/k})$. □

References

[1] Noga Alon and Yuval Roichman, *Random Cayley graphs and expanders*, Random Structures Algorithms **5** (1994), 271–284.

[2] Keith Ball, *An elementary introduction to modern convex geometry*, Flavors of geometry, Math. Sci. Res. Inst. Publ., vol. 31, Cambridge Univ. Press, Cambridge, 1997, pp. 1–58.

[3] Rajendra Bhatia, *Matrix analysis*, Graduate Texts in Mathematics, vol. 169, Springer-Verlag, New York, 1997.

[4] Emmanuel Breuillard, Ben Green, and Terence Tao, *The structure of approximate groups*, Publ. Math. Inst. Hautes Études Sci. **116** (2012), 115–221.

[5] David Conlon and Yufei Zhao, *Quasirandom Cayley graphs*, Discrete Anal. (2017), Paper No. 6, 14pp.

[6] Chandler Davis and W. M. Kahan, *The rotation of eigenvectors by a perturbation. III*, SIAM J. Numer. Anal. **7** (1970), 1–46.

[7] G. A. Freĭman, *Foundations of a structural theory of set addition*, American Mathematical Society, Providence, R. I., 1973, Translated from the Russian, Translations of Mathematical Monographs, Vol 37.
[8] Ben Green, *On the width of transitive sets: Bounds on matrix coefficients of finite groups*, Duke Math. J. **169** (2020), 551–578.
[9] Ben Green and Avi Wigderson, *Lecture notes for the 22nd McGill invitational workshop on computational complexity*, 2010, `https://www.math.ias.edu/~avi/TALKS/Green_Wigderson_lecture.pdf`.
[10] Mikhael Gromov, *Groups of polynomial growth and expanding maps*, Inst. Hautes Études Sci. Publ. Math. (1981), 53–73.
[11] Kenneth M. Hall, *An r-dimensional quadratic placement algorithm*, Management Science **17** (1970), 219–229.
[12] Shlomo Hoory, Nathan Linial, and Avi Wigderson, *Expander graphs and their applications*, Bull. Amer. Math. Soc. (N.S.) **43** (2006), 439–561.
[13] Yoshiharu Kohayakawa, Vojtěch Rödl, and Mathias Schacht, *Discrepancy and eigenvalues of Cayley graphs*, Czechoslovak Math. J. **66** (2016), 941–954.
[14] R. A. Kunze, L_p *Fourier transforms on locally compact unimodular groups*, Trans. Amer. Math. Soc. **89** (1958), 519–540.
[15] Alexander Lubotzky, *Expander graphs in pure and applied mathematics*, Bull. Amer. Math. Soc. (N.S.) **49** (2012), 113–162.
[16] Assaf Naor, *On the Banach-space-valued Azuma inequality and small-set isoperimetry of Alon-Roichman graphs*, Combin. Probab. Comput. **21** (2012), 623–634.
[17] Philippe Rigollet and Jan-Christian Hütter, *High dimensional statistics (lecture notes)*, 2019, draft version November 5, 2019, available at `http://www-math.mit.edu/~rigollet/PDFs/RigNotes17.pdf`.
[18] Daniel Spielman, *Spectral graph theory*, Combinatorial scientific computing, Chapman & Hall/CRC Comput. Sci. Ser., CRC Press, Boca Raton, FL, 2012, pp. 495–524.
[19] Joel A. Tropp, *User-friendly tail bounds for sums of random matrices*, Found. Comput. Math. **12** (2012), 389–434.

Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
Email address: `{asah,msawhney,yufeiz}@mit.edu`