

QUASIRANDOM CAYLEY GRAPHS

DAVID CONLON AND YUFEI ZHAO

ABSTRACT. We prove that the properties of having small discrepancy and having small second eigenvalue are equivalent in Cayley graphs, extending a result of Kohayakawa, Rödl, and Schacht, who treated the abelian case. The proof relies on Grothendieck's inequality. As a corollary, we also prove that a similar result holds in all vertex-transitive graphs.

1. INTRODUCTION

A fundamental result of Chung, Graham, and Wilson [6], building on earlier work of Thomason [19, 20], states that for a sequence of graphs of density p , where p is a fixed positive constant, a number of seemingly distinct notions of quasirandomness are equivalent. The following theorem details some of these equivalences.

Theorem 1.1 (Chung–Graham–Wilson). *For any fixed $0 < p < 1$ and any sequence of graphs $(\Gamma_n)_{n \in \mathbb{N}}$ with $|V(\Gamma_n)| = n$, the following properties are equivalent:*

- P_1 : for all subsets $S, T \subseteq V(\Gamma_n)$, $e(S, T) = p|S||T| + o(pn^2)$;
- P_2 : $e(\Gamma_n) \geq (1 + o(1))p\binom{n}{2}$, $\lambda_1(\Gamma_n) = (1 + o(1))pn$ and $|\lambda_2(\Gamma_n)| = o(pn)$, where $\lambda_i(\Gamma_n)$ is the i th largest eigenvalue, in absolute value, of the adjacency matrix of Γ_n ;
- P_3 : for all graphs H , the number of labeled copies of H in Γ_n is $(1 + o(1))p^{e(H)}n^{v(H)}$.

Here, the adjacency matrix $A(\Gamma)$ of an n -vertex graph Γ is the $n \times n$ matrix $(a_{st})_{s, t \in [n]}$ where $a_{st} = 1$ if s and t are adjacent in Γ and 0 otherwise.

In this paper, we will be concerned with studying the extent to which this theorem, and particularly the equivalence between the first two properties, extends to sparse graphs, that is, graphs where the density p tends to zero. We will focus our attention on regular graphs. We then say that an n -vertex d -regular graph $\Gamma = (V, E)$ is ϵ -uniform if, for all $S, T \subseteq V$,

$$\left| e(S, T) - \frac{d}{n}|S||T| \right| \leq \epsilon dn.$$

The equivalence between properties P_1 and P_2 of Theorem 1.1 now implies that if p is fixed and Γ_n is a sequence of graphs with $|V(\Gamma_n)| = n$ and Γ_n regular of degree $d_n = pn$, then the sequence Γ_n is $o(1)$ -uniform if and only if $|\lambda_2(\Gamma_n)| = o(d_n)$.

One direction of this equivalence follows from the famous expander mixing lemma. This says that if $\Gamma = (V, E)$ is an (n, d, λ) -graph, that is, an n -vertex d -regular graph such that all eigenvalues of the adjacency matrix $A(\Gamma)$, except the largest, are bounded above in absolute value by λ , then

$$\left| e(S, T) - \frac{d}{n}|S||T| \right| \leq \lambda \sqrt{|S||T|}$$

for all $S, T \subseteq V$. Thus, if $|\lambda_2(\Gamma)| \leq \epsilon d$, Γ is ϵ -uniform. Note that this statement does not rely on knowing that d is large, applying just as well when it is constant as when it is on the order of n . This observation is of fundamental importance in the study of expander graphs (see [10]).

Date: February 25, 2017.

The first author was supported by a Royal Society University Research Fellowship and ERC Starting Grant 676632. The second author was supported by an Esmée Fairbairn Junior Research Fellowship at New College, Oxford.

Answering a question of Chung and Graham [5], Krivelevich and Sudakov [13] showed that the converse does not hold. That is, there are sequences of sparse graphs for which the discrepancy condition P_1 of Theorem 1.1 does not imply the eigenvalue condition P_2 . However, their example is not regular (and a later example due to Bollobás and Nikiforov [3] is not connected), so we give an alternative example below.

Example 1.2. Take three disjoint vertex sets U , V , and W , with $|U| = |V| = t$ and $|W| = n - 2t$. Choose a d -regular graph Γ on vertex set $U \cup V \cup W$ such that every vertex in U has no neighbours in U and exactly λ neighbours in V , and vice versa for V , while each vertex in W has the same number of neighbours in U and in V . By taking the vector $y = (y_1, \dots, y_n)$ with $y_j = 1$ if $j \in U$, -1 if $j \in V$, and 0 otherwise, we see that $-\lambda$ is an eigenvalue of the adjacency matrix $A(\Gamma)$. More concretely, suppose d is even, let $t = \lambda = d/2$, and let the bipartite graph between U and V be complete. We split the remaining set of vertices W into two subsets W_0 and W_1 , where $|W_1| = d^2/4$ and each vertex in W_1 is joined to both endpoints of a unique associated edge in $\Gamma[U, V]$, while no vertex in W_0 is joined to $U \cup V$. In W , we place a random graph such that every vertex in W_0 has degree d and every vertex in W_1 has degree $d - 2$. The resulting graph is d -regular. Moreover, provided $d \rightarrow \infty$ and $d = o(n)$, a simple calculation (similar to the proof of Proposition 12 in [3]) implies that G is $o(1)$ -uniform, despite having $-d/2$ as an eigenvalue.

Given such examples, a recent result of Kohayakawa, Rödl, and Schacht [12] comes as something of a surprise. Suppose that G is a finite group and S is a subset of G . The (directed) *Cayley graph* $\text{Cay}(G, S)$ is formed by taking the elements of G as the vertex set and $\{(sg, g) : g \in G, s \in S\}$ as the set of (directed) edges. If S is symmetric, that is, $S = S^{-1}$, the graph can be viewed as undirected. Unless stated otherwise, we will always use the term Cayley graph to refer to such an undirected graph with a symmetric generating set. We note that many of the standard examples of quasirandom graphs, including Paley graphs and the Ramanujan graphs of Lubotzky, Phillips, and Sarnak [15] and Margulis [16], are Cayley graphs.

Theorem 1.3 (Kohayakawa–Rödl–Schacht). *Let G be an abelian group. Then every ϵ -uniform Cayley graph $\text{Cay}(G, S)$ is an (n, d, λ) -graph with $n = |G|$, $d = |S|$, and $\lambda \leq C\epsilon d$ for some absolute constant C .*

The main result of this paper generalizes Theorem 1.3 to any finite group.

Theorem 1.4. *Every ϵ -uniform Cayley graph $\text{Cay}(G, S)$ is an (n, d, λ) -graph with $n = |G|$, $d = |S|$, and $\lambda \leq 8\epsilon d$.*

It is not hard to lift this result to all vertex-transitive graphs, where a graph is said to be *vertex transitive* if the automorphism group of the graph acts transitively on the vertices. Note that Cayley graphs are always vertex transitive, while the example of the Petersen graph shows that the converse is not true.

Corollary 1.5. *Every n -vertex d -regular ϵ -uniform vertex-transitive graph is an (n, d, λ) -graph with $\lambda \leq 8\epsilon d$.*

Suppose now that Γ is an n -vertex d -regular vertex-transitive graph. If Γ is ϵ -uniform, this means that

$$\left| e(S, T) - \frac{d}{n} |S| |T| \right| \leq \epsilon dn \quad (1)$$

for all $S, T \subseteq V(\Gamma)$. Corollary 1.5 then says that Γ is an (n, d, λ) -graph with $\lambda \leq 8\epsilon d$. In turn, the expander mixing lemma shows that

$$\left| e(S, T) - \frac{d}{n} |S| |T| \right| \leq 8\epsilon d \sqrt{|S| |T|} \quad (2)$$

for all $S, T \subseteq V(\Gamma)$. Thus, Corollary 1.5 may be seen as saying that in vertex-transitive graphs the discrepancy condition (1) bootstraps itself to the stronger condition (2).

For comparison, we note also a result of Bilu and Linial [2] which says that if Γ is an n -vertex d -regular graph satisfying the stronger discrepancy condition (2), then it is an (n, d, λ) -graph with $\lambda = O(\epsilon d \log(2/\epsilon))$. Our result shows that in Cayley graphs, one can derive the stronger conclusion $\lambda = O(\epsilon d)$ from the weaker condition (1).

We begin the process of proving Theorem 1.4 and Corollary 1.5 in the next section by introducing two matrix norms, the cut norm and the spectral norm, which capture the properties of having small discrepancy and having small second eigenvalue. The rough idea then will be to show that if A is the adjacency matrix of a vertex-transitive graph Γ , the spectral norm of A agrees with a semidefinite relaxation of the cut norm.

We note that some similar ideas have been used in this context before. In particular, a semidefinite relaxation was used by Alon et al. [1] to prove the related result that if a graph G has small discrepancy, then one can remove a small fraction of the vertices so that in the remaining graph all eigenvalues, except the largest, are small.

2. NORMS AND RELAXATIONS

For any vector $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, we define its p -norm, $1 \leq p \leq \infty$, by

$$|x|_p := (|x_1|^p + \dots + |x_n|^p)^{1/p}$$

and

$$|x|_\infty := \max\{|x_1|, \dots, |x_n|\}.$$

In particular, we write $|x| := |x|_2$ for the Euclidean norm. The inner product on \mathbb{R}^n is

$$\langle x, y \rangle := x^*y = x_1y_1 + \dots + x_ny_n.$$

Let $A = (a_{st})_{s \in [m], t \in [n]} \in \mathbb{R}^{m \times n}$ be a matrix (we write $[n] := \{1, \dots, n\}$). We define a number of matrix norms. The most familiar one is the *spectral norm*:

$$\|A\| := \sup_{\substack{x \in \mathbb{R}^n \\ |x| \leq 1}} |Ax| = \sup_{\substack{x \in \mathbb{R}^m, y \in \mathbb{R}^n \\ |x|, |y| \leq 1}} |x^*Ay|, \tag{3}$$

where the second equality above follows from an application of the Cauchy–Schwarz inequality. Equivalently, $\|A\|$ equals the largest singular value of A . If A is symmetric, then $\|A\|$ is also the maximum of the absolute values of the eigenvalues of A .

The notion of ϵ -uniformity is captured by the *cut norm*:

$$\|A\|_{\text{cut}} := \sup_{S \subseteq [m], T \subseteq [n]} \left| \sum_{s \in S, t \in T} a_{st} \right|.$$

It is not hard to see that the cut norm can be expressed as a linear relaxation:

$$\|A\|_{\text{cut}} = \sup_{x_1, \dots, x_m, y_1, \dots, y_n \in [0, 1]} \left| \sum_{s \in [m], t \in [n]} a_{st} x_s y_t \right| = \sup_{x \in [0, 1]^m, y \in [0, 1]^n} |x^t Ay|. \tag{4}$$

To see the validity of this reformulation, note that the expression inside the absolute value is linear individually in each of $x_1, \dots, x_m, y_1, \dots, y_n$, and hence its extremum is attained when all these variables are $\{0, 1\}$ -valued, which is equivalent to the earlier formulation with subsets S and T .

We digress for a moment to relate these norms to graphs. Let Γ be an n -vertex d -regular graph and let $A(\Gamma)$ be its adjacency matrix (recall that this is the $n \times n$ matrix $(a_{st})_{s, t \in [n]}$ where $a_{st} = 1$ if s and t are adjacent in Γ and 0 otherwise). Then Γ is an (n, d, λ) -graph if and only if

$$\|A(\Gamma) - \frac{d}{n}J\| \leq \lambda,$$

where J is the $n \times n$ all-1's matrix, and Γ is ϵ -uniform if and only if

$$\|A(\Gamma) - \frac{d}{n}J\|_{\text{cut}} \leq \epsilon dn.$$

Returning to the discussion of matrix norms, we further relax the cut norm by allowing each x_s and y_t to be numbers in $[-1, 1]$:

$$\begin{aligned} \|A\|_{\infty \rightarrow 1} &:= \sup_{\substack{x \in \mathbb{R}^n \\ |x|_\infty \leq 1}} |Ax|_1 = \sup_{\substack{x \in \mathbb{R}^m, y \in \mathbb{R}^n \\ |x|_\infty, |y|_\infty \leq 1}} |x^*Ay| \\ &= \sup_{x_1, \dots, x_m, y_1, \dots, y_n \in [-1, 1]} \left| \sum_{s \in [m], t \in [n]} a_{st} x_s y_t \right|. \end{aligned}$$

This norm is equivalent to the cut norm:

$$\|A\|_{\text{cut}} \leq \|A\|_{\infty \rightarrow 1} \leq 4\|A\|_{\text{cut}}. \quad (5)$$

Indeed, write $x = x_+ - x_-$ and $y = y_+ - y_-$, where $x_+, x_- \in [0, 1]^m$ and $y_+, y_- \in [0, 1]^n$, and then apply the triangle inequality and (4).

Finally, we consider a semidefinite relaxation, which we shall refer to as the *Grothendieck norm*:

$$\|A\|_{\mathbb{G}} := \sup_{x_1, \dots, x_m, y_1, \dots, y_n \in B(\mathbb{H})} \left| \sum_{s \in [m], t \in [n]} a_{st} \langle x_s, y_t \rangle \right|,$$

where $(\mathbb{H}, \langle \cdot, \cdot \rangle)$ is any Hilbert space¹ and $B(\mathbb{H})$ is the unit ball in \mathbb{H} , containing all points in \mathbb{H} of norm at most 1. In fact, since we are working with $m + n$ vectors, one can assume that $\mathbb{H} = \mathbb{R}^{m+n}$ in the definition above.

A key fact about the Grothendieck norm is that it is equivalent to the $\infty \rightarrow 1$ norm. Though this result was proved by Grothendieck [8], it was first stated in this form by Lindenstrauss and Pełczyński [14]. We refer the interested reader to the surveys [11, 17] for further information.

Theorem 2.1 (Grothendieck's inequality). *There exists a constant $K_{\mathbb{G}}$ such that for all real-valued matrices A ,*

$$\|A\|_{\infty \rightarrow 1} \leq \|A\|_{\mathbb{G}} \leq K_{\mathbb{G}} \|A\|_{\infty \rightarrow 1}.$$

We use the symbol $K_{\mathbb{G}}$ (the real Grothendieck constant) to denote the optimal constant in the above inequality.

The best current upper bound on the real Grothendieck constant is $K_{\mathbb{G}} < \frac{\pi}{2 \log(1+\sqrt{2})} = 1.78\dots$, due to Braverman, Makarychev, Makarychev, and Naor [4] (see Haagerup [9] for the best constant in the complex case). Combining (5) with Grothendieck's inequality (and using $4K_{\mathbb{G}} \leq 8$), we have

$$\|A\|_{\text{cut}} \leq \|A\|_{\mathbb{G}} \leq 8\|A\|_{\text{cut}}, \quad (6)$$

that is, the cut norm and the Grothendieck norm are equivalent.

The spectral norm provides an upper bound on the Grothendieck norm: for any $A \in \mathbb{R}^{m \times n}$,

$$\|A\|_{\mathbb{G}} \leq \sqrt{mn} \|A\|. \quad (7)$$

The proof is by a straightforward application of the Cauchy–Schwarz inequality.²

¹In general, one should take a complex Hilbert space, but since we only consider real-valued matrices A , it is equivalent to use real Hilbert spaces.

²We may assume that $\mathbb{H} = \mathbb{R}^k$ for some $k \leq m + n$ and we have $x_1, \dots, x_m, y_1, \dots, y_n \in B(\mathbb{H})$ such that $\|A\|_{\mathbb{G}} = \left| \sum_{s \in [m], t \in [n]} a_{st} \langle x_s, y_t \rangle \right| \leq \sum_{j=1}^k \left| \sum_{s \in [m], t \in [n]} a_{st} x_{sj} y_{tj} \right| \leq \sum_{j=1}^k \|A\| \sqrt{\sum_{s \in [m]} |x_{sj}|^2} \sqrt{\sum_{t \in [n]} |y_{tj}|^2} \leq \|A\| \sqrt{\sum_{s \in [m]} |x_s|^2} \sqrt{\sum_{t \in [n]} |y_t|^2} \leq \sqrt{mn} \|A\|$ by the triangle inequality, the definition of the spectral norm, and the Cauchy–Schwarz inequality.

In general, the spectral norm and the Grothendieck norm are inequivalent, i.e., $\sqrt{mn}\|A\|/\|A\|_{\mathbb{G}}$ can be arbitrarily large, as shown by Example 1.2. Our main theorem below shows, somewhat surprisingly, that the two norms agree for matrices associated to vertex-transitive graphs.

From now on we consider only square matrices. Let $A = (a_{s,t})_{1 \leq s,t \leq n}$. We say that a permutation σ of $[n]$ is an *automorphism* if $A_{\sigma(s),\sigma(t)} = A_{s,t}$ for all $s,t \in [n]$. We call a subgroup G of the permutations of $[n]$ *transitive* if for every $s,t \in [n]$ there is some $\sigma \in G$ such that $\sigma(s) = t$. We say that A is *vertex-transitive* if its group of automorphisms is transitive on $[n]$.

Theorem 2.2. $\|A\|_{\mathbb{G}} = n\|A\|$ for all $n \times n$ vertex-transitive real-valued matrices A .

Using (6), we obtain the following corollary relating the cut norm and the spectral norm, from which Corollary 1.5 clearly follows.

Corollary 2.3. $\|A\|_{\text{cut}} \leq n\|A\| \leq 8\|A\|_{\text{cut}}$ for all $n \times n$ vertex-transitive real-valued matrices A .

As discussed in the introduction, the proof will proceed through an analysis of Cayley graphs, which form the subject of the next section.

3. CAYLEY GRAPHS

We consider a weighted version of Cayley graphs. Let $f: G \rightarrow \mathbb{R}$. Define a weighted directed graph $\Gamma = \text{Cay}(G, f)$ with vertex set G and edge weights $\Gamma: G \times G \rightarrow \mathbb{R}$ given by $\Gamma(g, h) := f(gh^{-1})$ for all $g, h \in G$. Let $A(f)$ denote the associated matrix, whose rows and columns are indexed by G , with $a_{g,h} = \Gamma(g, h) = f(gh^{-1})$ for all $g, h \in G$. The natural analogue of the assumption that S is symmetric (giving undirected Cayley graphs) would be $f(g) = f(g^{-1})$, implying that $A(f)$ is symmetric, though we will not need to make such an assumption here.

We define various norms for functions $f: G \rightarrow \mathbb{R}$ corresponding to the matrix norms defined in Section 2. We use the averaging measure on G . Define the p -norms for functions on G by

$$\|f\|_p = (\mathbb{E}_{g \in G} |f(g)|^p)^{1/p}$$

and $\|f\|_{\infty} = \sup_{g \in G} |f(g)|$. Define the convolution $f_1 * f_2$ of two functions $f_1, f_2: G \rightarrow \mathbb{R}$ by

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

The spectral norm is defined as follows, agreeing with the matrix version (up to normalization):

$$\|f\| = \sup_{\substack{x: G \rightarrow \mathbb{R} \\ \|x\|_2 \leq 1}} \|f * x\|_2 = \sup_{\substack{x, y: G \rightarrow \mathbb{R} \\ \|x\|_2, \|y\|_2 \leq 1}} |\mathbb{E}_{g, h \in G} f(gh^{-1})x(g)y(h)| = |G|^{-1} \|A(f)\|. \quad (8)$$

The Grothendieck norm is defined by

$$\|f\|_{\mathbb{G}} := \sup_{x, y: G \rightarrow B(\mathbb{H})} \left| \mathbb{E}_{g, h \in G} f(gh^{-1}) \langle x(g), y(h) \rangle \right| = |G|^{-2} \|A(f)\|_{\mathbb{G}}.$$

Our main result now is that the spectral norm agrees with the Grothendieck norm for functions on a group.

Theorem 3.1. $\|f\| = \|f\|_{\mathbb{G}}$ for every function $f: G \rightarrow \mathbb{R}$ on a finite group G .

We thank Assaf Naor and Prasad Raghavendra for independently suggesting the following short proof of Theorem 3.1 after Y. Z. presented an initial version of this work at the Simons Symposium on Analysis of Boolean Functions in April 2016. We include their proof here with permission. Our original proof using representation theory can be found in the appendix.

Proof. Let $f: G \rightarrow \mathbb{R}$. Let $x, y: G \rightarrow \mathbb{R}$ with $\|x\|_2 \leq 1$ and $\|y\|_2 \leq 1$ be such that

$$\|f\| = |\mathbb{E}_{g, h \in G} f(gh^{-1})x(g)y(h)|.$$

Define $x_g(h) = x(gh)$ and $y_g(h) = y(gh)$ for all $g, h \in G$. We view x_g and y_h as vectors in the unit ball in $L^2(G)$ equipped with inner product $\langle x, y \rangle = \mathbb{E}_{g \in G} x(g)y(g)$ for all $x, y \in L^2(G)$. We have

$$\begin{aligned} \|f\|_{\mathbb{G}} &\geq |\mathbb{E}_{g,h \in G} f(gh^{-1}) \langle x_g, y_h \rangle| = |\mathbb{E}_{g,h,a \in G} f(gh^{-1}) x(ga) y(ha)| \\ &= |\mathbb{E}_{g,h,a \in G} f((ga)(ha)^{-1}) x(ga) y(ha)| = |\mathbb{E}_{g,h \in G} f(gh^{-1}) x(g) y(h)| = \|f\|. \end{aligned}$$

Combining with $\|f\|_{\mathbb{G}} \leq \|f\|$ from (7), we obtain $\|f\|_{\mathbb{G}} = \|f\|$. \square

4. TRANSITIVE GRAPHS

Now we show that Theorem 3.1 on functions $f: G \rightarrow \mathbb{R}$ implies Theorem 2.2 on vertex-transitive matrices. The idea is that we can lift the edges of a transitive graph to a Cayley graph over its automorphism group and this operation preserves all the norms that we are interested in.

Recall that $A = (a_{s,t})_{s,t \in [n]}$ is an $n \times n$ vertex-transitive matrix if the group of automorphisms of A is transitive on $[n]$. Let G be any transitive subgroup of automorphisms (it could be the whole automorphism group or any transitive subgroup). Here every element $g \in G \leq S_n$ is some permutation of $[n]$, i.e., $g: [n] \rightarrow [n]$, and the product of two elements $g, h \in S_n$ is defined by $(gh)(s) = h(g(s))$ for all $s \in [n]$. We have $a_{g(s),g(t)} = a_{s,t}$ for all $s, t \in [n]$ and $g \in G$.

Lemma 4.1. *Let $A = (a_{s,t})_{s,t \in [n]}$ be an $n \times n$ vertex-transitive real-valued matrix and let G be any transitive subgroup of automorphisms of A . Define $f: G \rightarrow \mathbb{R}$ by $f(g) = a_{g(1),1}$. Then*

$$n\|f\| = \|A\| \quad \text{and} \quad n^2\|f\|_{\mathbb{G}} = \|A\|_{\mathbb{G}}.$$

Proof. The proofs of these two identities are essentially the same. Let \mathbb{H} denote either an arbitrary Hilbert space (for the Grothendieck norm) or simply \mathbb{R} (for the spectral norm). For any $x_1, \dots, x_n, y_1, \dots, y_n \in \mathbb{H}$, define $x, y: G \rightarrow \mathbb{H}$ by $x(g) = x_{g(1)}$ and $y(g) = y_{g(1)}$ for all $g \in G$. Then

$$\begin{aligned} \mathbb{E}_{g,h \in G} [f(gh^{-1}) \langle x(g), y(h) \rangle] &= \mathbb{E}_{g,h \in G} [a_{(gh^{-1})(1),1} \langle x(g), y(h) \rangle] \\ &= \mathbb{E}_{g,h \in G} [a_{g(1),h(1)} \langle x_{g(1)}, y_{h(1)} \rangle] = \frac{1}{n^2} \sum_{s,t \in [n]} a_{s,t} \langle x_s, y_t \rangle. \end{aligned} \quad (9)$$

Furthermore, we have $n^{-1} \sum_{s \in [n]} |x_s|^2 = \mathbb{E}_{g \in G} |x(g)|^2$, and similarly with y . This shows $n\|f\| \geq \|A\|$ and $n^2\|f\|_{\mathbb{G}} \geq \|A\|_{\mathbb{G}}$.

Conversely, given $x, y: G \rightarrow \mathbb{H}$, we can set $x_1, \dots, x_n, y_1, \dots, y_n \in \mathbb{H}$ to be $x_s = \mathbb{E}_{g \in G: g(1)=s} x(g)$ and $y_s = \mathbb{E}_{g \in G: g(1)=s} y(g)$ for all $s \in [n]$. Then (9) still holds. We also have $n^{-1} \sum_{s \in [n]} |x_s|^2 \leq \mathbb{E}_{g \in G} |x(g)|^2$ by convexity, and similarly with y . This shows that $n\|f\| \leq \|A\|$ and $n^2\|f\|_{\mathbb{G}} \leq \|A\|_{\mathbb{G}}$. \square

5. BIPARTITE ANALOGUES

There are also variants of our results in the bipartite setting. Given a group G of order n and a function $f: G \rightarrow \mathbb{R}$, the *bipartite Cayley graph* $\Gamma = \text{Bip}(G, f)$ is the bipartite graph between two copies of G with edge weights $f: G \times G \rightarrow \mathbb{R}$ given by $\Gamma(g, h) = f(gh^{-1})$. We write $B(f)$ for the associated matrix, whose rows and columns are both indexed by G , with $b_{g,h} = \Gamma(g, h) = f(gh^{-1})$. Since we did not require that f be symmetric in Theorem 3.1, we can apply it directly to prove that $\|B(f)\|_{\mathbb{G}} = n\|B(f)\|$ for any $f: G \rightarrow \mathbb{R}$. An analogue of Theorem 1.4 then follows provided we replace the second largest eigenvalue with the largest singular value of $B(\Gamma) - \frac{d}{n}J$, where J is the all-1's matrix.

We also have an analogue of Theorem 1.5. We say that a bipartite graph Γ between sets U and V of orders m and n , respectively, is *bitransitive* if the group of automorphisms of Γ is transitive on each part. For example, the graph between any two consecutive layers of the Boolean cube satisfies this condition. A slight amendment to the proof of Lemma 4.1 then allows us to show that

if B is an $m \times n$ bitransitive real-valued matrix, then $\|B\|_G = \sqrt{mn}\|B\|$. The claimed analogue of Theorem 1.5 then follows easily, with the second largest eigenvalue now replaced by the largest singular value of $B(\Gamma) - pJ$, where p is the density of Γ .

Acknowledgments. We would like to thank Sean Eberhard for helpful discussions. We would also like to thank Noga Alon and Mathias Schacht for bringing the paper [1] to our attention. We are greatly indebted to Assaf Naor and Prasad Raghavendra for their contribution and Y. Z. would like to thank the Simons Foundation and the organisers of the Simons Symposium on Analysis of Boolean Functions for facilitating these conversations. Finally, we thank the referees for suggesting changes that greatly improved the presentation of the paper.

REFERENCES

- [1] N. Alon, A. Coja-Oghlan, H. Hàn, M. Kang, V. Rödl, and M. Schacht, Quasi-randomness and algorithmic regularity for graphs with general degree distributions, *SIAM J. Comput.* **39** (2010), 2336–2362.
- [2] Y. Bilu and N. Linial, Lifts, discrepancy and nearly optimal spectral gap, *Combinatorica* **26** (2006), 495–519.
- [3] B. Bollobás and V. Nikiforov, Hermitian matrices and graphs: singular values and discrepancy, *Discrete Math.* **285** (2004), 17–32.
- [4] M. Braverman, K. Makarychev, Y. Makarychev, and A. Naor, The Grothendieck constant is strictly smaller than Krivine’s bound, *Forum Math. Pi* **1** (2013), 42 pp.
- [5] F. R. K. Chung and R. L. Graham, Sparse quasi-random graphs, *Combinatorica* **22** (2002), 217–244.
- [6] F. R. K. Chung, R. L. Graham, and R. M. Wilson, Quasi-random graphs, *Combinatorica* **9** (1989), 345–362.
- [7] W. T. Gowers, Quasirandom groups, *Combin. Probab. Comput.* **17** (2008), 363–387.
- [8] A. Grothendieck, Résumé de la théorie métrique des produits tensoriels topologiques, *Bol. Soc. Mat. São Paulo* **8** (1953), 1–79.
- [9] U. Haagerup, A new upper bound for the complex Grothendieck constant, *Israel J. Math.* **60** (1987), 199–224.
- [10] S. Hoory, N. Linial, and A. Wigderson, Expander graphs and their applications, *Bull. Amer. Math. Soc.* **43** (2006), 439–561.
- [11] S. Khot and A. Naor, Grothendieck-type inequalities in combinatorial optimization, *Comm. Pure Appl. Math.* **65** (2012), 992–1035.
- [12] Y. Kohayakawa, V. Rödl, and M. Schacht, Discrepancy and eigenvalues of Cayley graphs, *Czechoslovak Math. J.* **66** (2016), 941–954.
- [13] M. Krivelevich and B. Sudakov, Pseudo-random graphs, in *More sets, graphs and numbers*, Bolyai Soc. Math. Stud., Vol. 15, 199–262, Springer, Berlin, 2006.
- [14] J. Lindenstrauss and A. Pełczyński, Absolutely summing operators in L_p -spaces and their applications, *Studia Math.* **29** (1968), 275–326.
- [15] A. Lubotzky, R. Phillips, and P. Sarnak, Ramanujan graphs, *Combinatorica* **8** (1988), 261–277.
- [16] G. A. Margulis, Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators (in Russian), *Problemy Peredachi Informatsii* **24** (1988), 51–60; translation in *Problems Inform. Transmission* **24** (1988), 39–46.
- [17] G. Pisier, Grothendieck’s theorem, past and present, *Bull. Amer. Math. Soc.* **49** (2012), 237–323.
- [18] J.-P. Serre, *Linear representations of finite groups*, Graduate Texts in Mathematics, Vol. 42, Springer-Verlag, New York-Heidelberg, 1977.
- [19] A. G. Thomason, Pseudorandom graphs, in *Random graphs ’85* (Poznań, 1985), North-Holland Math. Stud., Vol. 144, 307–331, North-Holland, Amsterdam, 1987.
- [20] A. G. Thomason, Random graphs, strongly regular graphs and pseudorandom graphs, in *Surveys in Combinatorics 1987*, London Math. Soc. Lecture Note Ser., Vol. 123, 173–195, Cambridge University Press, Cambridge, 1987.
- [21] A. Wigderson, Representation theory of finite groups, and applications, lecture notes available at <http://www.math.ias.edu/~avi/TALKS/additive-lectures-v2.pdf>.

APPENDIX A. REPRESENTATION-THEORETIC PROOF

We include in this appendix our original proof of Theorem 3.1, that $\|f\| = \|f\|_G$ for $f: G \rightarrow \mathbb{R}$, using representation theory. The proof follows the spirit of Gowers’ work on quasirandom groups [7] and is a natural generalization of the proof for the abelian case [12].

A.1. Abelian groups. As a warm up, we give a short proof in the abelian setting, which is significantly easier. This proof is instructive, though not strictly required for the general setting.

Here and throughout the appendix, we will work with functions $f: G \rightarrow \mathbb{C}$ and with complex analogues of our various norms. For the most part, the definitions of these norms are the same as in the real case or the same after some minor modification. For example, letting $\mathbb{D} := \{z \in \mathbb{C} : |z| \leq 1\}$ denote the unit disk in the complex plane, the complex $\infty \rightarrow 1$ norm for functions $f: G \rightarrow \mathbb{C}$ is defined by

$$\|f\|_{\infty \rightarrow 1}^{\mathbb{C}} := \sup_{x, y: G \rightarrow \mathbb{D}} \left| \mathbb{E}_{g, h \in G} f(gh^{-1}) \overline{x(g)} y(h) \right|.$$

The essence of the proof below can be found in Kohayakawa–Rödl–Schacht [12] (where the proof is attributed to Gowers). Here we give a more streamlined presentation.

Theorem A.1. $\|f\| = \|f\|_{\infty \rightarrow 1}^{\mathbb{C}}$ for every function $f: G \rightarrow \mathbb{C}$ on a finite abelian group G .

Proof. It is easy to see that $\|f\|_{\infty \rightarrow 1}^{\mathbb{C}} \leq \|f\|$ from the definitions, as $\|x\|_2 \leq 1$ for all $x: G \rightarrow \mathbb{D}$. It is well known that $\|f\| = |\mathbb{E}_{g \in G} f(g) \chi(g)|$ for some multiplicative character $\chi: G \rightarrow \mathbb{C}$. On the other hand, since $\|\chi\|_{\infty} = 1$, we have

$$\|f\|_{\infty \rightarrow 1}^{\mathbb{C}} \geq |\mathbb{E}_{g, h \in G} f(gh^{-1}) \overline{\chi(g)} \chi(h)| = |\mathbb{E}_{g \in G} f(g) \overline{\chi(g)}| = \|f\|.$$

Therefore, $\|f\| = \|f\|_{\infty \rightarrow 1}^{\mathbb{C}}$. □

Similarly to (5), it can be shown that the cut norm is equivalent to the complex $\infty \rightarrow 1$ norm for complex-valued matrices. Together with Theorem A.1, this implies that the cut norm and the spectral norm are equivalent for abelian Cayley graphs.

Remark. For non-abelian groups G , it is possible to have $\|f\| > \|f\|_{\infty \rightarrow 1}^{\mathbb{C}}$. If an undirected weighted Cayley graph on G admits a coordinate-wise uniformly bounded eigenbasis, then one has $\|f\| = O(\|f\|_{\infty \rightarrow 1}^{\mathbb{C}})$ by the above argument.³ However, as will be addressed in future work, not all Cayley graphs admit such eigenbases. Consequently, the above proof does not extend immediately to the non-abelian case. The proof given below instead generalizes the above proof by replacing the multiplicative characters with group representations.

In the next three subsections, we recall some standard mathematical tools used in our proof. The reader familiar with these tools should feel free to skip ahead to Section A.5.

A.2. Group representation theory. We recall some basic facts from group representation theory. They can be found in any standard textbook on the subject, e.g., [18].

Let G be a finite group, not necessarily abelian. We write \widehat{G} for the set of distinct irreducible representations of G . An irreducible representation $\rho \in \widehat{G}$ of dimension d_{ρ} is a homomorphism $\rho: G \rightarrow U(d_{\rho})$, where $U(n)$ is the group of unitary $n \times n$ matrices (so that $\rho(g^{-1}) = \rho(g)^*$). Note that by a standard averaging trick, we need to only consider unitary representations.

We will need the following orthogonality result known as Schur's lemma.

Theorem A.2 (Schur's lemma). *Let G be a finite group and let $\rho, \sigma \in \widehat{G}$. Let M be any $d_{\rho} \times d_{\sigma}$ matrix with complex entries. Then*

$$\mathbb{E}_{g \in G} [\rho(g) M \sigma(g^{-1})] = \begin{cases} \frac{\text{Tr } M}{d_{\rho}} I & \text{if } \rho = \sigma, \\ 0 & \text{otherwise.} \end{cases}$$

³We thank Assaf Naor for this remark.

A.3. Non-abelian Fourier analysis. Let us first recall the Fourier transform for abelian groups. Let G be a finite abelian group and \widehat{G} the group of multiplicative characters $\chi: G \rightarrow \mathbb{C}$. For any function $f: G \rightarrow \mathbb{C}$, we can define its Fourier transform $\widehat{f}: \widehat{G} \rightarrow \mathbb{C}$ by $\widehat{f}(\chi) = \mathbb{E}_{g \in G} f(g) \chi(g)$. We then have the inversion formula $f(g) = \sum_{\chi \in \widehat{G}} \widehat{f}(\chi) \overline{\chi(g)}$.

From now on, let G be any finite group, not necessarily abelian, and \widehat{G} the set of distinct irreducible unitary representations $\rho: G \rightarrow U(d_\rho)$ of G . For any $f: G \rightarrow \mathbb{C}$ and any irreducible representation $\rho \in \widehat{G}$, define the Fourier transform of f at ρ by

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

Note that $\widehat{f}(\rho)$ is a $d_\rho \times d_\rho$ matrix, so its dimension varies with ρ .

For any two complex-valued matrices $A = (a_{st})_{s \in [m], t \in [n]}$ and $B = (b_{st})_{s \in [m], t \in [n]}$ with the same dimensions, define the Hilbert–Schmidt inner product

$$\langle A, B \rangle_{\text{HS}} = \text{Tr}(AB^*) = \sum_{s \in [m], t \in [n]} a_{st} \overline{b_{st}}$$

and the Hilbert–Schmidt norm

$$\|A\|_{\text{HS}} = \sqrt{\langle A, A \rangle_{\text{HS}}} = \sqrt{\sum_{s \in [m], t \in [n]} |a_{st}|^2}.$$

We have a Fourier inversion formula:

$$f(g) = \sum_{\rho \in \widehat{G}} d_\rho \langle \widehat{f}(\rho), \rho(g) \rangle_{\text{HS}} = \sum_{\rho \in \widehat{G}} d_\rho \text{Tr}(\widehat{f}(\rho) \rho(g)^*). \quad (10)$$

A proof of this formula can be found in [21, Section 2.10]. We also have Plancherel’s identity: for all $f_1, f_2: G \rightarrow \mathbb{C}$,

$$\mathbb{E}_{g \in G} f_1(g) \overline{f_2(g)} = \sum_{\rho \in \widehat{G}} d_\rho \langle \widehat{f}_1(\rho), \widehat{f}_2(\rho) \rangle_{\text{HS}}.$$

In particular, setting $f_1 = f_2 = f$, we have that for all $f: G \rightarrow \mathbb{C}$,

$$\|f\|_2^2 = \mathbb{E}_{g \in G} [|f(g)|^2] = \sum_{\rho \in \widehat{G}} d_\rho \|\widehat{f}(\rho)\|_{\text{HS}}^2. \quad (11)$$

The Fourier transform converts convolution into multiplication. For all $f_1, f_2: G \rightarrow \mathbb{C}$, writing

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

for convolution, we have, for all $\rho \in \widehat{G}$,

$$\widehat{f_1 * f_2}(\rho) = \mathbb{E}_{g, h \in G} f_1(g) f_2(h) \rho(gh) = (\mathbb{E}_{g \in G} f_1(g) \rho(g)) (\mathbb{E}_{h \in G} f_2(h) \rho(h)) = \widehat{f}_1(\rho) \widehat{f}_2(\rho). \quad (12)$$

Or, more succinctly, $\widehat{f_1 * f_2} = \widehat{f}_1 \widehat{f}_2$.

Conceptually, the Fourier transform is a simultaneous block diagonalization of all $f: G \rightarrow \mathbb{C}$ (viewed as linear maps on the space of functions $x: G \rightarrow \mathbb{C}$ via the convolution $x \mapsto f * x$). For each $\rho \in \widehat{G}$, the $d_\rho \times d_\rho$ matrix $\widehat{f}(\rho)$ appears exactly d_ρ times in the block diagonalization (the multiplicity d_ρ is the multiplicity of ρ in the regular representation of G). As a result, the spectral norm of f must be the maximum of the spectral norms of the components of the block diagonalization:

$$\|f\| = \max_{\rho \in \widehat{G}} \|\widehat{f}(\rho)\|, \quad (13)$$

where on the left-hand side $\|f\|$ is the spectral norm (8) of a function $f: G \rightarrow \mathbb{C}$ and on the right-hand side $\|\widehat{f}(\rho)\|$ is the spectral norm (3) of a matrix. Let us also give a direct proof of (13) for the convenience of the reader. Using (12) and Plancherel's identity (11), we have

$$\|f\|^2 = \sup_{\substack{x: G \rightarrow \mathbb{C} \\ \|x\|_2 \leq 1}} \|f * x\|_2^2 = \sup_{\substack{x: G \rightarrow \mathbb{C} \\ \|x\|_2 \leq 1}} \sum_{\rho \in \widehat{G}} d_\rho \|\widehat{f}(\rho) \widehat{x}(\rho)\|_{\text{HS}}^2. \quad (14)$$

If A and B are $d \times d$ matrices and b_1, \dots, b_d are the columns of B , then $\|AB\|_{\text{HS}}^2 = |Ab_1|^2 + \dots + |Ab_d|^2 \leq \|A\|^2(|b_1|^2 + \dots + |b_d|^2) = \|A\|^2\|B\|_{\text{HS}}^2$. Therefore, for any $x: G \rightarrow \mathbb{C}$, we have

$$\begin{aligned} \sum_{\rho \in \widehat{G}} d_\rho \|\widehat{f}(\rho) \widehat{x}(\rho)\|_{\text{HS}}^2 &\leq \sum_{\rho \in \widehat{G}} d_\rho \|\widehat{f}(\rho)\|^2 \|\widehat{x}(\rho)\|_{\text{HS}}^2 \\ &\leq \left(\max_{\rho \in \widehat{G}} \|\widehat{f}(\rho)\|\right)^2 \sum_{\rho \in \widehat{G}} d_\rho \|\widehat{x}(\rho)\|_{\text{HS}}^2 = \left(\max_{\rho \in \widehat{G}} \|\widehat{f}(\rho)\|\right)^2 \|x\|_2^2, \end{aligned}$$

where the final step uses Plancherel (11) again. Combining with (14), this shows that $\|f\| \leq \max_{\rho \in \widehat{G}} \|\widehat{f}(\rho)\|$. Conversely, for a fixed ρ , by setting the columns of $\widehat{x}(\rho)$ to be the top right singular vector of $\widehat{f}(\rho)$ so that $\|\widehat{f}(\rho) \widehat{x}(\rho)\|_{\text{HS}} = \|\widehat{f}(\rho)\| \|\widehat{x}(\rho)\|_{\text{HS}}$, and setting $\widehat{x}(\sigma) = 0$ for all $\sigma \neq \rho$, we obtain $\|f\| \geq \|\widehat{f}(\rho)\|$ for each $\rho \in \widehat{G}$. This proves (13).

A.4. Singular value decompositions. Every $m \times n$ complex matrix M of rank r has a singular value decomposition (SVD):

$$M = \lambda_1 u_1 v_1^* + \dots + \lambda_r u_r v_r^*,$$

where $\lambda_1 \geq \dots \geq \lambda_r > 0$ (known as the singular values), u_1, \dots, u_r are mutually orthogonal unit vectors in \mathbb{C}^m , and v_1, \dots, v_r are mutually orthogonal unit vectors in \mathbb{C}^n .

Let $f: G \rightarrow \mathbb{C}$. For each $\rho \in \widehat{G}$, take any singular value decomposition of $\widehat{f}(\rho)$:

$$\widehat{f}(\rho) = \lambda_1^\rho u_1^\rho v_1^{\rho*} + \dots + \lambda_{d_\rho}^\rho u_{d_\rho}^\rho v_{d_\rho}^{\rho*},$$

where $\lambda_1^\rho \geq \dots \geq \lambda_{d_\rho}^\rho \geq 0$ (we allow some of the λ_j^ρ 's to be zero), and $\{u_1^\rho, \dots, u_{d_\rho}^\rho\}$ and $\{v_1^\rho, \dots, v_{d_\rho}^\rho\}$ are both orthonormal bases of \mathbb{C}^{d_ρ} . For any $\lambda \geq 0$, $u, v \in \mathbb{C}^{d_\rho}$, we have $\text{Tr}(\lambda u v^* \rho(g)^*) = \lambda v^* \rho(g)^* u$. Therefore, by the Fourier inversion formula (10), we have the following SVD for any $f: G \rightarrow \mathbb{C}$:

$$f(g) = \sum_{\rho \in \widehat{G}} d_\rho \sum_{k=1}^{d_\rho} \lambda_k^\rho v_k^{\rho*} \rho(g)^* u_k^\rho. \quad (15)$$

A.5. Proof of Theorem 3.1. Now we give the representation-theoretic proof that $\|f\| = \|f\|_{\mathbb{G}}$ for any $f: G \rightarrow \mathbb{C}$.

Consider the SVD (15) of f . From (13), we see that $\|f\|$ is the maximum of the top singular value λ_1^ρ ranging over all $\rho \in \widehat{G}$. Let $\sigma \in \widehat{G}$ be such that $\|f\| = \lambda_1^\sigma$.

We already know from (7) that $\|f\|_{\mathbb{G}} \leq \|f\|$. To show $\|f\|_{\mathbb{G}} \geq \|f\|$, it suffices to exhibit $x, y: G \rightarrow B(\mathbb{H})$ so that $\mathbb{E}_{g,h} f(gh^{-1}) \langle x(g), y(h) \rangle = \|f\|$ (here $\langle x, y \rangle := \overline{x_1} y_1 + \dots + \overline{x_n} y_n$ for $x, y \in \mathbb{C}^n$). Set $\mathbb{H} = \mathbb{C}^{d_\sigma}$ and, for all $g, h \in G$,

$$x(g) = \overline{\sigma(g^{-1})} v_1^\sigma \quad \text{and} \quad y(h) = \overline{\sigma(h^{-1})} u_1^\sigma,$$

where u_1^σ and v_1^σ are singular vectors from the SVD (15). Let $\rho \in \widehat{G}$ and $1 \leq k \leq d_\rho$. Writing $u = u_k^\rho$ and $v = v_k^\rho$ to reduce clutter in notation, we have

$$\begin{aligned} \mathbb{E}_{g,h} d_\rho v^* \rho(gh^{-1}) u \langle x(g), y(h) \rangle &= \mathbb{E}_{g,h} d_\rho v^* \rho(gh^{-1}) u \overline{\langle y(h), x(g) \rangle} \\ &= \mathbb{E}_{g,h} d_\rho v^* \rho(gh^{-1}) u \langle \sigma(h^{-1}) u_1^\sigma, \sigma(g^{-1}) v_1^\sigma \rangle \\ &= \mathbb{E}_{g,h} d_\rho v^* \rho(gh^{-1}) u u_1^{\sigma*} \sigma(hg^{-1}) v_1^\sigma \\ &= \mathbb{E}_g d_\rho v^* \rho(g) u u_1^{\sigma*} \sigma(g^{-1}) v_1^\sigma. \end{aligned}$$

By Schur's lemma (Theorem A.2), if $\rho \neq \sigma$, then this is zero. If $\rho = \sigma$, then it equals

$$v^* \text{Tr}(u u_1^{\sigma*}) v_1^\sigma = \langle u_1^\sigma, u \rangle \langle v, v_1^\sigma \rangle,$$

which is 1 if $k = 1$, i.e., $u = u_1^\sigma$ and $v = v_1^\sigma$, and zero otherwise due to the orthogonality of the singular vectors. It follows, by the SVD (15) and the above calculation, that

$$\mathbb{E}_{g,h} f(gh^{-1}) \langle x(g), y(h) \rangle = \lambda_1^\sigma = \|f\|.$$

This shows that $\|f\|_{\mathbb{G}} \geq \|f\|$, completing the proof. \square

Remark. The above proof shows that in defining $\|f\|_{\mathbb{G}}$ for $f: G \rightarrow \mathbb{C}$, one only needs to take \mathbb{H} (in the definition of the Grothendieck norm) to have dimension at most $\max_{\rho \in \widehat{G}} d_\rho$, the maximum dimension of an irreducible representation of G . In other words, if we define intermediate Grothendieck-type norms

$$\|f\|_{\mathbb{G},k} := \sup_{x,y: G \rightarrow B(\mathbb{C}^k)} |\mathbb{E}_{g,h \in G} f(gh^{-1}) \langle x(g), y(h) \rangle|$$

for every positive integer k , so that

$$\|f\|_{\infty \rightarrow 1} = \|f\|_{\mathbb{G},1} \leq \|f\|_{\mathbb{G},2} \leq \cdots \leq \|f\|_{\mathbb{G},m+n} = \|f\|_{\mathbb{G}},$$

then $\|f\|_{\mathbb{G}} = \|f\|_{\mathbb{G},k}$ whenever all irreducible representations of G have dimension at most k . In particular, if G is abelian, then all irreducible representations have dimension 1, so $\|f\| = \|f\|_{\mathbb{G},1} = \|f\|_{\infty \rightarrow 1}^{\mathbb{C}}$, as shown earlier in Theorem A.1.

MATHEMATICAL INSTITUTE, OXFORD OX2 6GG, UNITED KINGDOM
E-mail address: david.conlon@maths.ox.ac.uk

MATHEMATICAL INSTITUTE, OXFORD OX2 6GG, UNITED KINGDOM
E-mail address: yufei.zhao@maths.ox.ac.uk