Random Cayley graphs are expanders: a simple proof of the Alon–Roichman theorem

ZEPH LANDAU
Department of Mathematics
City College of New York
landau@sci.ccny.cuny.edu

ALEXANDER RUSSELL*
Department of Computer Science and Engineering
University of Connecticut
acr@cse.uconn.edu

Submitted: Jul 13, 2004; Accepted: Aug 26, 2004; Published: Sep 13, 2004
Mathematics Subject Classification: 05C25, 05C80, 20C15, 20F69

Abstract

We give a simple proof of the Alon–Roichman theorem, which asserts that the Cayley graph obtained by selecting $c_\varepsilon \log |G|$ elements, independently and uniformly at random, from a finite group G has expected second eigenvalue no more than ε; here c_ε is a constant that depends only on ε. In particular, such a graph is an expander with constant probability. Our new proof has three advantages over the original proof: (i.) it is extremely simple, relying only on the decomposition of the group algebra and tail bounds for operator-valued random variables, (ii.) it shows that the $\log |G|$ term may be replaced with $\log D$, where $D \leq |G|$ is the sum of the dimensions of the irreducible representations of G, and (iii.) it establishes the result above with a smaller constant c_ε.

1 Introduction

A beautiful theorem of N. Alon and Y. Roichman [4] asserts that *random Cayley graphs are expanders*. Specifically, they study the spectrum of the Cayley graphs obtained by selecting k elements, independently and uniformly at random, from a finite group G. They show that for every $\varepsilon > 0$ there is a constant c_ε so that the expected second eigenvalue of the normalized

*Supported, in part, by National Science Foundation grants CCR-0093065, CCR-0121277, CCR-0220264, CCR-0311368, and EIA-0218443.

adjacency matrix of the graph is less than ε as long as $k \geq (c_\varepsilon + o(1)) \log |G|$. Their proof involves a clever combinatorial argument that controls the behavior of random walks taken on the (random) graph. Invoking established relationships between graph expansion and the second eigenvalue, this implies bounds on the expected expansion of the Cayley graph formed from $k = O(\log |G|)$ random elements.

In this article, we give a simple proof of the result based on tail bounds for sums of independent operator-valued random variables established by R. Ahlswede and A. Winter [1]. Our proof yields a stronger relationship between c_ε and ε: we show that $k = (2 \ln 2/\varepsilon + o(1))^2 \log |G|$ elements suffice, whereas the original proof requires $(4e/\varepsilon^2 + o(1)) \log |G|$ elements. Moreover, using some elementary group representation theory, we show that the $\log |G|$ term may be replaced with the term $\log D$, where D is the sum of the dimensions of the irreducible representations of G. The improvement from $\log |G|$ to $\log D$ was independently discovered by L. Schulman and P. Loh [6].

We remark that the theorem is tight, up to the constant appearing before the logarithm, in the sense that there exist groups, $(\mathbb{Z}_2)^n$ for example, that cannot even be generated with fewer than $n = \log_2 2^n$ elements.

We begin, in Section 2, with a brief discussion of expander graphs, the representation theory of finite groups, and tail bounds for positive operator-valued random variables. The main theorem is proved in Section 3.

2 Background

We outline, below, the elements of graph theory, representation theory, and probability theory required for the statement of the theorem and the subsequent proof. The exposition here is primarily for the purposes of setting down notation; we refer the reader to the more complete accounts appearing in Alon and Spencer [5], Serre [7], and Ahlswede and Winter [1] for greater detail and discussion.

Let G be a finite group and $S \subset G$ a set of generators for G. The *Cayley* graph $X(G, S)$ is the graph obtained by taking the elements of G as vertices and including the edge (α, β) if $\alpha^{-1}\beta \in S \cup S^{-1}$, where $S^{-1} = \{s^{-1} \mid s \in S\}$. As the set $S \cup S^{-1}$ is closed under inverse, $\alpha^{-1}\beta \in S \cup S^{-1} \Leftrightarrow \beta^{-1}\alpha \in S \cup S^{-1}$ so that we may naturally treat $X(G, S)$ as an undirected graph. We overload the symbol 1, letting it denote the identity element of a group G.

Graph expansion and spectral gap. An undirected graph $G = (V, E)$ is an (n, d, ϵ)-*expander graph* if G has n vertices, every vertex has degree d or less, and for all subsets X of vertices with $|X| \leq |V|/2$, $|\Gamma(X) \setminus X| \geq \epsilon |X|$, where

$$\Gamma(X) = \{v \mid \exists u \in X, (u, v) \in E\} \ .$$

A *family of linear expanders* is a family of graphs $\{G_i \mid i > 0\}$, where G_i is a (n_i, d, ϵ)-expander, ϵ and d are constants independent of n_i, and the n_i tend to infinity in i. Graphs with these properties are the principal combinatorial elements featured in many pseudorandom constructions.

Graph expansion has a propitious relationship with the *spectral* properties of the graph G. Focusing, as we will, on regular graphs, define the *normalized adjacency matrix* $\mathcal{A}(G)$ of the d-regular graph G so that

$$\mathcal{A}(G)_{uv} = \begin{cases} \frac{1}{d} & \text{if } (u,v) \in E, \\ 0 & \text{otherwise.} \end{cases}$$

As $\mathcal{A}(G)$ is a symmetric, real matrix, its eigenvalues are real and it is easy to see that all eigenvalues lie in the interval $[-1, 1]$. $\mathcal{A}$ always possesses the eigenvalue 1 which, when G is connected, has multiplicity one; the corresponding eigenvectors are those with uniform entries (taking the same value at each $g \in G$). For a regular graph G, we let $\lambda_2(G)$ denote the second largest element of the multiset of absolute values of eigenvalues of $\mathcal{A}(G)$. As mentioned above, a strong relationship between $\lambda_2(\cdot)$ and expansion has been achieved. In particular, if G is a d-regular graph with n vertices and $\lambda_2(G) \leq \lambda$ then G is an (n, d, ϵ)-expander with $\epsilon \geq 2(1-\lambda)/(3-2\lambda)$ (see Alon and Milman [3]). Conversely, if G is an (n, d, ϵ)-expander then $\lambda_2(G) \leq 1 - \epsilon^2/[2d(2+\epsilon^2)]$ (see Alon [2]).

The representation theory of finite groups. Let G be a finite group. A *representation* ρ of G is a homomorphism $\rho : G \to \mathsf{U}(V)$, where V is a finite dimensional Hilbert space and $\mathsf{U}(V)$ is the group of unitary operators on V. The *dimension* of ρ, denoted d_ρ, is the dimension of the vector space V. By choosing a basis for V, then, each $\rho(g)$ is associated with a unitary matrix $[\rho(g)]$ so that for every $g, h \in G$, $[\rho(gh)] = [\rho(g)] \cdot [\rho(h)]$, where $\cdot$ denotes matrix multiplication.

Fixing a representation $\rho : G \to \mathsf{U}(V)$, we say that a subspace $W \subset V$ is *invariant* if $\rho(g)W \subset W$ for all $g \in G$; observe that in this case the restriction $\rho_W : G \to \mathsf{U}(W)$ given by restricting each $\rho(g)$ to W is also a representation. When ρ has no invariant subspace other than the trivial space $\{\mathbf{0}\}$ and V, ρ is said to be *irreducible*. In the case when ρ is *not* irreducible, then, there is a nontrivial invariant subspace $W \subset V$ and, as the inner product $\langle \cdot, \cdot \rangle$ is invariant under each of the unitary maps $\rho(g)$, it is immediate that the subspace

$$W^\perp = \{\mathbf{u} \mid \forall \mathbf{w} \in W, \langle \mathbf{u}, \mathbf{w} \rangle = 0\}$$

is also invariant. Associated with the decomposition $V = W \oplus W^\perp$ is the natural decomposition of the operators $\rho(g) = \rho_W(g) \oplus \rho_{W^\perp}(g)$. By repeating this process, any representation $\rho : G \to \mathsf{U}(V)$ may be decomposed into irreducible representations: we write $\rho = \sigma_1 \oplus \cdots \oplus \sigma_k$.

If two representations ρ and σ are the same up to an isometric change of basis, we say that they are *equivalent*. It is a fact that any finite group G has a finite number of distinct irreducible representations up to equivalence and, for a group G, we let $\widehat{G}$ denote a set of representations containing exactly one from each equivalence class.

Two representations play a special role in the following analysis. The first is the *trivial* representation $\mathbf{1}$, the one-dimensional representation that maps all elements of G to the identity operator on $\mathbb{C}$. The second is the *regular* representation R, given by the permutation action of G on itself. Specifically, let $\mathbb{C}[G]$ be the $|G|$-dimensional vector space of formal sums

$$\left\{ \sum_g \alpha_g \cdot g \mid \alpha_g \in \mathbb{C} \right\}$$

equipped with the unique inner product for which $\langle g, h\rangle$ is equal to one when $g = h$ and zero otherwise. Then R is the representation $R : G \to \mathsf{U}(\mathbb{C}[G])$ given by linearly extending the rule $R(g)[h] = gh$. While the trivial representation $\mathbf{1}$ is irreducible, R is not: in fact, every irreducible representation $\rho \in \widehat{G}$ appears in R with multiplicity equal to its dimension:

$$R = \bigoplus_{\rho \in \widehat{G}} \underbrace{\rho \oplus \cdots \oplus \rho}_{d_\rho} \quad . \tag{1}$$

By counting dimensions on each side of this equation, we have $|G| = \sum_\rho d_\rho^2$.

Tail bounds for operator-valued random variables. Our principal technical tool will be a tail bound for (positive) operator-valued random variables. The bound below was proved in [1], where it is modestly attributed to H. Chernoff.

Let $\mathsf{A}(V)$ denote the collection of self-adjoint linear operators on the finite dimensional Hilbert space V. For $A \in \mathsf{A}(V)$, we let $\|A\|$ denote the operator norm of A equal to the largest absolute value obtained by an eigenvalue of A. The cone of *positive* operators

$$\mathsf{P}(V) = \{A \in \mathsf{A}(V) \mid \forall \mathbf{v}, \langle A\mathbf{v}, \mathbf{v}\rangle \geq 0\}$$

gives rise to a natural partial order on $\mathsf{A}(V)$ by defining $A \geq B$ iff $A - B \in \mathsf{P}(V)$. We shall write $B \in [A, A']$ for $A \leq B \leq A'$.

Proposition 1 ([1]). *Let V be a Hilbert space of dimension d and let $A_1, \ldots, A_k$ be independent, identically distributed random variables taking values in $\mathsf{P}(V)$ with expected value $\mathbb{E}[A_i] = M \geq \mu\mathbb{1}$ and $A_i \leq \mathbb{1}$. Then for all $\varepsilon \in [0, 1/2]$,*

$$\Pr\left[\frac{1}{k}\sum_{i=1}^{k} A_i \notin \Big[(1-\varepsilon)M, (1+\varepsilon)M\Big]\right] \leq 2d \cdot e^{-\frac{\varepsilon^2 \mu k}{2\ln 2}} \quad .$$

3 A proof of the Alon-Roichman theorem

We shall demonstrate tail bounds on the distribution of $\lambda_2(X(G, S))$ and conclude from these a strengthened version (Corollary 3) of the following theorem of Alon and Roichman.

Theorem 1 ([4]). *For every $\varepsilon > 0$ there is a function $k = k(|G|) = \left[\frac{4e}{\varepsilon^2} + o(1)\right] \log |G|$ so that for all finite groups G,*

$$\mathbb{E}\big[\lambda_2(X(G, S))\big] \leq \varepsilon \quad ,$$

where $s_1, \ldots, s_k$ are independent random variables, uniformly distributed in G, and S is the set $\{s_1, \ldots, s_k\}$.

We begin with the development of tail bounds for the variable $\lambda_2(X(G, S))$; Corollary 3 will follow.

Theorem 2. *Let G be a finite group, $\varepsilon > 0$, $D = \sum_{\rho \in \widehat{G}} d_\rho$, and $k = (2\ln 2/\varepsilon)^2[\log D + b + 1]$. Then*

$$\Pr[\lambda_2(X(G,S)) > \varepsilon] \leq 2^{-b} \ ,$$

where $s_1, \ldots, s_k$ are independent random variables, uniformly distributed in G, and S is the set $\{s_1, \ldots, s_k\}$.

Proof. For an element $a = \sum_g a_g \cdot g \in \mathbb{C}[G]$ and a representation ρ, let $\widehat{a}(\rho) = \sum_g a_g \rho(g)$. Defining s to be the formal sum $1/(2k) \cdot \sum_{i=1}^k (s_i + s_i^{-1}) \in \mathbb{C}[G]$, observe that the normalized adjacency matrix $\mathcal{A}$ of the graph $X(G,S)$ is precisely the operator $\widehat{s}(R)$ expressed in the basis $\{1 \cdot g \mid g \in G\}$ of $\mathbb{C}[G]$. We consider the decomposition of R into irreducible representations given by Equation (1); as discussed above, this corresponds to an orthogonal direct sum decomposition of $\mathbb{C}[G]$ into spaces invariant under each $R(g)$. Observe that the eigenvalue 1 corresponds to the appearance of the trivial representation in $\mathbb{C}[G]$. It suffices, then, to bound the spectrum of $\widehat{s}(R)$ when restricted to the nontrivial representations appearing in the decomposition: specifically, $\lambda_2(X(G,S)) = \max_{\rho \neq \mathbf{1}} \|\widehat{s}(\rho)\|$, this maximum extended over all nontrivial irreducible representations of G. Let ρ be a nontrivial irreducible representation of G and define $a_i = 1/2 \cdot (s_i + s_i^{-1}) \in \mathbb{C}[G]$; then $s = 1/k \cdot \sum a_i$ and each $\widehat{a_i}(\rho) = 1/2 \cdot [\rho(s_i) + \rho(s_i)^{-1}]$ is self-adjoint as $\rho(s_i^{-1}) = \rho(s_i)^{-1} = \rho(s_i)^*$. Since $\|\widehat{a_i}(\rho)\| \leq 1$, define $p_i = 1/2 \cdot (1 + a_i)$ and observe that $\widehat{p_i}(\rho)$ is a positive operator satisfying $\widehat{p_i}(\rho) \leq \mathbb{1}$.

Recalling that R contains a single copy of the trivial representation and observing that the operator $\sum_g R(g)$ has rank 1 (indeed, in the basis above, each entry in the corresponding matrix is a 1), we conclude that $\mathbb{E}_{g \in G}[\rho(g)] = 0 \cdot \mathbb{1}$ for nontrivial ρ. Hence $\mathbb{E}[\widehat{p_i}(\rho)] = \frac{1}{2}\mathbb{1}$ and, by Proposition 1,

$$\begin{aligned} \Pr\left[\frac{1}{k}\sum_i \widehat{p_i}(\rho) \notin \left[\frac{1-\varepsilon}{2}\mathbb{1}, \frac{1+\varepsilon}{2}\mathbb{1}\right]\right] &\leq 2d_\rho \exp\left(-\frac{k\varepsilon^2}{4\ln 2}\right) \\ &= 2d_\rho \exp\bigl(-\ln(2)[\log D + b + 1]\bigr) = \frac{d_\rho}{D} 2^{-b} \ . \end{aligned}$$

Finally, $\Pr[\lambda_2(X(G,S)) > \varepsilon] = \Pr\bigl[\exists \rho \in \widehat{G} \setminus \{\mathbf{1}\}, \frac{1}{k}\sum_{i=1}^k \widehat{a_i}(\rho) \notin [-\varepsilon\mathbb{1}, \varepsilon\mathbb{1}]\bigr]$ so that

$$\begin{aligned} \Pr[\lambda_2(X(G,S)) > \varepsilon] &= \Pr\left[\exists \rho \in \widehat{G} \setminus \{\mathbf{1}\}, \frac{1}{k}\sum_{i=1}^k \widehat{p_i}(\rho) \notin \left[\frac{1-\varepsilon}{2}\mathbb{1}, \frac{1+\varepsilon}{2}\mathbb{1}\right]\right] \\ &\leq \sum_{\rho \in \widehat{G}} \frac{d_\rho}{D} 2^{-b} = 2^{-b} \ . \end{aligned}$$

□

Remark. An even simpler proof, relying on no representation theory, can be given by writing $\mathbb{C}[G] = T \oplus N$, where T is the one-dimensional eigenspace spanned by the uniform vector $\sum_g g$ and N is the orthogonal complement of T. By the reasoning above, the average of the operators $R(g)$ on the space N is zero and the proof may proceed by applying the tail bound (Proposition 1) over N. This results in the bound $k = ((2\ln 2)/\varepsilon)^2(\log|G| + b + 1)$.

Observe that if X is a random variable taking values in the interval $[0,1]$ for which $\Pr[X > \varepsilon] \leq \delta$, then $\mathbb{E}[X] \leq (1-\delta)\varepsilon + \delta \leq \varepsilon + \delta$. In particular, selecting a function $\delta(D)$ tending to zero for which $\log(\delta^{-1}) = o(\log D)$ and applying the bound above with $\varepsilon' = \varepsilon(1-\delta)$ and $k' = [(2\ln 2)/\varepsilon']^2(\log D - \log(\varepsilon\delta) + 1)$, we obtain the following corollary that implies Theorem 1 above.

Corollary 3. *For every $\varepsilon > 0$ there is a function $k = k(D) = \left[\frac{2\ln 2}{\varepsilon} + o(1)\right]^2 \log D$ so that for all finite groups G,*

$$\mathbb{E}\big[\lambda_2(X(G,S))\big] \leq \varepsilon \ ,$$

where $s_1, \ldots, s_k$ are independent random variables, uniformly distributed in G, S is the set $\{s_1, \ldots, s_k\}$, and $D = \sum_{\rho\in\widehat{G}} d_\rho$.

Remark. This improves upon Theorem 1 both by reducing the leading constant (from $4e/\varepsilon^2 \approx 10.87/\varepsilon^2$ to $(2\ln 2/\varepsilon)^2 \approx 1.93/\varepsilon^2$) and replacing the $\log|G|$ term with $\log D$. Recall that $\sum_\rho d_\rho^2 = |G|$, whence $D = \sum_\rho d_\rho \leq |G|$, and that for groups with large irreducible representations D can grow as slowly as $O(\sqrt{|G|})$ (e.g., the affine groups $\mathbb{Z}_p^* \ltimes \mathbb{Z}_p$).

4 Acknowledgements

We are grateful to Avi Wigderson and Leonard Schulman for helpful discussions.

References

[1] Rudolf Ahlswede and Andreas Winter. Strong converse for identification via quantum channels. Technical Report quant-ph/0012127v2, quant-ph e-Print archive, 2001.

[2] Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6:83–96, 1986.

[3] Noga Alon and Vitali D. Milman. Eigenvalues, expanders and superconcentrators (extended abstract). In *25th Annual Symposium on Foundations of Computer Science*, pages 320–322, Singer Island, Florida, 24–26 October 1984. IEEE.

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

[5] Noga Alon and Joel H. Spencer. *The Probabilistic Method*. John Wiley & Sons, Inc., 1992.

[6] Po-Shen Loh and Leonard Schulman. Improved expansion of random cayley graphs. 2004.

[7] Jean-Pierre Serre. *Linear Representations of Finite Groups*. Springer-Verlag, 1977.