On Non-localization of Eigenvectors of High Girth Graphs

Shirshendu Ganguly[*]
UC Berkeley

Nikhil Srivastava[†]
UC Berkeley

January 17, 2018

Abstract

We prove improved bounds how localized an eigenvector of a high girth regular graph can be, and present examples showing that these bounds are close to sharp. This study was initiated by Brooks and Lindenstrauss [BL13] who relied on the observation that certain suitably normalized averaging operators on high girth graphs are hyper-contractive and can be used to approximate projectors onto the eigenspaces of such graphs. Informally, their delocalization result in the contrapositive states that for any $\varepsilon \in (0, 1)$ and positive integer k, if a $(d + 1)-$regular graph has an eigenvector which supports ε fraction of the ℓ_2^2 mass on a subset of k vertices, then the graph must have a cycle of size $\tilde{O}(\log_d(k)/\varepsilon^2)$, suppressing logarithmic terms in $1/\epsilon$. In this paper, we improve the upper bound to $\tilde{O}(\log_d(k)/\varepsilon)$ and present a construction showing a lower bound of $\Omega(\log_d(k)/\varepsilon)$. Our construction is probabilistic and involves gluing together a pair of trees while maintaining high girth as well as control on the eigenvectors and could be of independent interest.

1 Introduction

Spectral graph theory studies graphs via associated linear operators such as the Laplacian and adjacency matrix. While the extreme eigenvectors of these operators are relatively well-understood and correspond to sparse cuts and colorings, much less is known about the combinatorial meaning of the interior eigenvectors. Most of the literature about them falls into two categories:

1. **Analysis of eigenvectors of random graphs.** For example, Dekel, Lee, Linial [DLL11] prove that any eigenvector of a dense random graph has a bounded number of nodal domains i.e., connected components where the eigenvector does not change sign. Following a sequence of results by various authors, in a recent breakthrough work Bauerschmidt, Huang, Yau [BHY], among various other things, show that with high probability, any 'bulk' eigenvector v of a random regular graph with large enough but fixed degree, is ℓ_∞ delocalized in the following sense:

$$\|v\|_\infty \le \frac{\log^C(n)}{\sqrt{n}}\|v\|_2,$$

[*]SG is supported by a Miller Research Fellowship.
[†]NS is supported by NSF Grant CCF-1553751 and a Sloan Research Fellowship.

where $\|\cdot\|_2$, and $\|\cdot\|_\infty$ denote the usual ℓ_2 and ℓ_∞ norms respectively and C is a constant. For a more precise statement see Theorem 1.2 in [BHY]. In another line of work, Backhausz and Szegedy [BS16] establish Gaussian behavior of the entry distribution of eigenvectors of random regular graphs by studying factors of i.i.d. processes on the regular infinite tree.

In all of these works the randomness of the model is used heavily, and weaker notions of delocalization are also considered (see e.g. [Gei13]). We refer the reader to [OVW16] for a survey of recent developments on delocalization of eigenvectors of random matrices.

2. A parallel story based on **asymptotic analysis of sequences of deterministic graphs**. The driving force for this is the so called Quantum Unique Ergodicity (QUE) conjecture by Rudnick and Sarnak [RS94]. The QUE conjecture states that on any compact negatively curved manifold all high energy eigenfunctions of the Laplacian equi-distribute. The conjecture is still widely open having been verified in only a few cases; perhaps most notably for the Hecke orthonormal basis on an arithmetic surface by Lindenstrauss [Lin06]. Brooks-Lindenstrauss [BL13] initiated the study of graph-theoretic analogues of this conjecture. The analogue of negatively curved manifolds are high girth regular graphs — the girth is defined as the length of the shortest cycle in a graph. Subsequently, Anantharaman and Le-Masson [ALM15] proved an asymptotic version of quantum ergodicity for regular expanders which converge (in the Benjamini-Schramm local topology) to the infinite $d-$regular tree.

The starting point of this paper is the beautiful result of [BL13]. Since the statement is a bit technical and could be hard to parse at first read we first explain the content informally in words. The theorem roughly says that if a graph does not have many short cycles, then eigenvectors cannot localize on small sets: for any eigenvector, any subset of the vertices representing a fraction of the ℓ_2^2 mass must have size n^δ for some δ depending on the fraction. The condition of not having many cycles is articulated as hyper-contractivity (i.e., control of $\|\cdot\|_{p\to q}$ norms for some $p<q$) of certain spherical mean operators on the graph.

Theorem 1.1 ([BL13]). *Suppose $G=(V,E)$ is a $(d+1)-$regular graph with adjacency matrix A. Let*

$$S_n(f)(x) := \frac{1}{d^{n/2}} \sum_{\mathrm{dist}(x,y)=n} f(y)$$

and suppose

$$\|S_n\|_{p\to q} \le C d^{-\alpha n} \tag{1}$$

($\|\cdot\|_{p\to q}$ denotes the norm of the naturally associated operator from ℓ_p to ℓ_q) for all $n\le N$, for some $1\le p\le q\le\infty$ and $\alpha\in[0,1]$. Then for any normalized eigenvector $v=(v_x)_{x\in V}$, of A and $S\subset V$ with $\|v_S\|_2^2 := \sum_{x\in S} v_x^2 \ge \varepsilon$,

$$|S| \ge \Omega_{d,C,\alpha,\varepsilon}(d^{\delta N}),$$

where $\delta = 2^{-7}\frac{\alpha p}{2-p}\varepsilon^2$.

In particular, the condition (1) is satisfied with $p=1, q=\infty, \alpha=1/2, C=d$ and $N=\lceil g/2\rceil-1$ for a graph of girth g. Viewed in the contrapositive, the theorem therefore says that the existence of an eigenvector of A with ε fraction of its mass on $k=|S|$ coordinates implies that the graph must contain a cycle of length $O(\log_d(k/\epsilon)/\varepsilon^2)$. In fact, a close examination of the proof reveals that it gives an upper bound which varies between $O(\log_d(k/\varepsilon)/\varepsilon)$ and $O(\log_d(k/\epsilon)/\varepsilon^2)$ depending on the diophantine properties of the eigenvalue being considered.

In this paper, we contribute to the understanding of this phenomenon in two ways.. First, we improve the above bound to $O(\log_d(k/\epsilon)/\varepsilon)$ for all eigenvalues of $d+1$-regular[1] graphs, irrespective of the number theoretic properties of the eigenvalue. The proof involves replacing the approximation-theoretic component of their proof by a simpler and more efficient method. Specifically, we prove the following theorem in Section 2.

Theorem 1.2. *Suppose G is a $(d+1)$-regular graph of girth g and v is a normalized eigenvector of the adjacency matrix of G. Then any subset S with $\|v_S\|_2^2 = \varepsilon$ must have*

$$|S| \geq \frac{d^{\varepsilon g/4}\varepsilon}{2d^2}.$$

The contrapositive of the above theorem implies that if there exists ε and k and S such that $|S| = k$ and $\|v_S\|_2^2 = \varepsilon$, then

$$g \leq \frac{4\log_d(k/\varepsilon) + O(1)}{\varepsilon}.$$

Remark 1.3 (Choice of Hypercontractive Norms)**.** *The paper [BL13] works with general $p \to q$ norms, but in this paper we will work solely with the $1 \to \infty$ since it is reveals all of the ideas and is easier to interpret combinatorially. Our proof of Theorem 1.2 can easily be modified to work with $p \to q$ norms, if desired.*

Remark 1.4 (Entropy Bounds from Delocalization)**.** *As already observed in [Bro09, Corollary 1], it is quite straightforward to obtain a lower bound on the entropy of an eigenvector v from a delocalization result such as Theorem 1.2, where the entropy of v is $-\sum_{x\in V} v_x^2 \log_d v_x^2$.*

Remark 1.5 (Tempered and Untempered Eigenvalues)**.** *Eigenvalues of A in the interval $[-2\sqrt{d}, 2\sqrt{d}]$ are referred to as* tempered *(indicating wave-like behavior) and those outside are called* untempered *(indicating exponential growth) in the QUE literature. It is known that a much stronger delocalization result, with dependence roughly $g = O(\log_d(k/\epsilon))$, can be proven for untempered eigenvalues using elementary arguments — see e.g. [Bro09, Page 59] or the arguments of [Kah92]. Note that any sequence of graphs with girth going to infinity must have a vanishingly small fraction of untempered eigenvalues. We will present bounds for arbitrary eigenvalues in this paper, without focusing on the distinction between tempered and untempered.*

Second, for every $d \geq 2$, sufficiently large k, and $\varepsilon \in (0,1)$, we exhibit a $(d+1)-$regular graph with a localized eigenvector which has girth at least $\Omega(\log_d(k)/\varepsilon)$, showing that our improved bound is sharp up to an additive $\log(1/\varepsilon)$ factor in the numerator, which is negligible whenever $k = \Omega(1/\varepsilon^c)$ for any c. We are able to construct such eigenvectors for a dense subset of eigenvalues in $(-2\sqrt{d}, 2\sqrt{d})$. The proof is probabilistic, and involves gluing together two trees without introducing any short cycles and while controlling their eigenvectors, which may be of independent interest.

Theorem 1.6. *For every $d \geq 2$, sufficiently large k and all $\varepsilon > 0$, there is a finite $(d+1)-$regular graph G with the following properties.*

1. *A_G has a normalized eigenvector v with eigenvalue $\lambda \in (-2\sqrt{d}, 2\sqrt{d})$ and*

$$\|v_S\|_2^2 = \Omega_\lambda(\varepsilon)$$

for a set S of size k, where the implicit constant depends on λ and is bounded away from zero on any subinterval of $(-2\sqrt{d}, 2\sqrt{d})$.

[1] We work with $(d+1)$-regular rather than $d-$regular graphs to avoid repeatedly writing $d-1$.

2. *G has girth at least*

$$\Omega\left(\frac{\log_d(k)}{\varepsilon}\right).$$

For every fixed ϵ (or for every fixed, sufficiently large k), the set of eigenvalues attained by the above graphs is dense in $(-2\sqrt{d}, 2\sqrt{d})$.

The proof of Theorem 1.6 appears in Section 3. Notice that the above theorem does not provide any bound as the eigenvalue λ approaches one of the edges $\pm 2\sqrt{d}$, which is consistent with Remark 1.5.

Remark 1.7 (Other Notions of Delocalization). *Various other notions of delocalization for eigenvectors have been studied — ℓ_∞ delocalization as mentioned above, lower bounds on the ℓ_1 norm, and "no-gaps" delocalization [RV15, RV16, ERS17] (see the surveys [OVW16, Rud17] for details). The examples we construct in Theorem 1.6 do* not *satisfy any of these notions, whereas they are all satisfied by random regular graphs. This establishes that the delocalization properties of high girth graphs are weaker than those of random graphs.*

1.1 Connection between localization and low girth : $\varepsilon = 1$ case.

Before proceeding to the proofs of these theorems, we give a quick proof of the upper bound in the extreme case $\varepsilon = 1$, to give some intuition about why a localized eigenvector implies a short cycle. Assume G is a $(d+1)$-regular graph with adjacency matrix A and $Av = \lambda v$ for a vector v with exactly k nonzero entries. Let H be the induced subgraph of G supported on the nonzero vertices. Observe that for the eigenvector equation to hold for any vertex $s \notin H$, we must have

$$\sum_{t \in E} A(s,t) = 0,$$

so in particular any such s must have at least two neighbors (of opposite signs) inside H. Thus, for every edge ts with $t \in H$ leaving H, there must be some $t' \in H$ such that tst' is a path of length 2 in G. Replace all such paths by new edges tt' to obtain a graph H' on the vertices of H (possibly creating multiedges), and observe that every vertex in this graph has degree at least $(d+1)$. Now, if H' has girth g, then any ball of radius $g/2 - 1$ does not contain cycles. Growing a ball from any vertex, we find that

$$d^{g/2-1} \leq |H'| \leq k,$$

which implies that $g \leq 2\log_d(k) + 2$. Since every edge in H' corresponds to a path of length at most 2 in G, G must contain a cycle of length at most $4\log_d(k) + 4$.

Theorem 1.2 shows that this continues to happen even when $\varepsilon = o(1)$. Note that since the girth of a $(d+1)$-regular graph on n vertices is at most $O(\log_d(n))$ by a similar argument, the only interesting regime is when $\varepsilon = \Omega(1/\log_d(n))$.

2 Improved Upper Bound

In this section we prove Theorem 1.2, at a high level following the approach of [BL13]. The main ingredient is the following hypercontractivity estimate. Let T_m be the Chebyshev polynomials of the first kind, i.e., $T_m(\cos\theta) = \cos(m\theta)$.

Lemma 2.1 (Hypercontractivity of Chebyshev Polynomials, [BL13]). *If A is a $d+1$-regular graph with girth g, then for all even $m < g/2$,*

$$\left\| T_m\left(A/(2\sqrt{d})\right)\right\|_{1\to\infty} = \frac{d-1}{2d^{m/2}}.$$

The proof appearing in [BL13] is based on a spectral decomposition in terms of spherical functions on trees. For completeness we give a quick proof of the above using connections to non-backtracking walks instead.

Proof. Let $U_m(\cdot)$ defined by $U_m(\cos\theta) = \frac{\sin((m+1)\theta)}{\sin\theta}$ be the Chebyshev polynomials of the second kind. It is well known that for any m, (see for e.g. Section 2 in [ABLS07])

$$B^{(m)} = d^{m/2}\left(U_m(A/(2\sqrt{d})) - \frac{1}{d}U_{m-2}(A/(2\sqrt{d}))\right),$$

where for any pair of vertices $u, v \in V$, the entry $B^{(m)}(u, v)$, is the number of non-backtracking walks of length m between u and v. At this point we use the following well known relation between the Chebyshev Polynomials of the first and second kind: $T_m = \frac{1}{2}\left(U_m - U_{m-2}\right)$. Putting the above together we get

$$T_m(A/(2\sqrt{d})) = \frac{1}{d}\left(T_{m-2}(A/(2\sqrt{d})) + \frac{1}{2}\Big(\frac{B^{(m)}}{d^{m/2}} - \frac{B^{(m-2)}}{d^{m/2-1}}\Big)\right). \tag{2}$$

Now note that $\left\| T_m\left(A/(2\sqrt{d})\right)\right\|_{1\to\infty}$ is nothing but the maximum entry of the corresponding matrix. Since $m < g/2$ by hypothesis, for all $j \le m$ and for all $u, v \in V$ we have $B^{(j)}(u, v) = \delta_{j,\mathsf{dist}(u,v)}$ where $\mathsf{dist}(u, v)$ is the graph distance between u and v. Summing (2) from 1 to m and using the last observation completes the proof. □

Using the above lemma, the next approximation result is at the heart of the proof of Theorem 1.2. As will be clear soon, given any eigenvalue λ_0 of $A/(2\sqrt{d})$, the proof of Theorem 1.2 demands the existence of a polynomial f, with the following two properties:

1. $f(A/(2\sqrt{d}))$ is hyper-contractive.
2. $f(\lambda_0)$ is large, and $f(\lambda)$ is not too negative for any other eigenvalue λ.

Since $f(\lambda_0)$ then approximates the projector onto the λ_0-eigenspace of $A/2\sqrt{d}$, the corresponding eigenvector cannot be localized. The following lemma states that such a polynomial exists. It is in the proof of this lemma that we achieve the required estimates needed to improve the bounds in [BL13].

Lemma 2.2 (Hypercontractive Polynomial Approximation). *If A has girth g, then for all positive integers r, m such that r is even, $mr < g/2$, and $\lambda \in \mathbb{R}$ there exists a polynomial f such that:*

1. $f(\lambda) \ge \frac{m}{4} - 1$.
2. $f(x) \ge -1$ *on* $\mathbb{R}$.
3. $\|f(A/2\sqrt{d})\|_{1\to\infty} \le \frac{2(d-1)}{d^{r/2}}$.

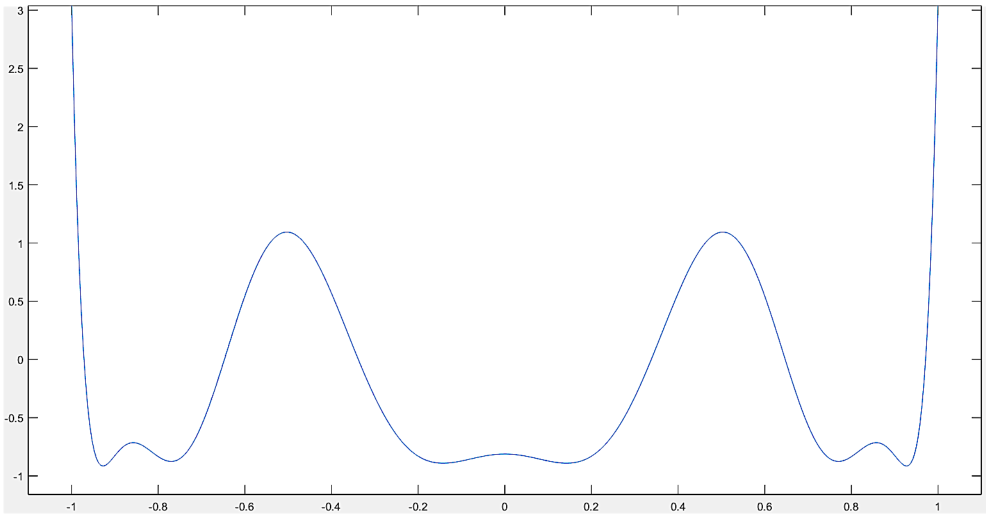

Figure 1: An example of the polynomial f in Lemma 2.2, with parameters $m = 8, r = 2, \phi = \pi/3$

Proof. Assume first that $\lambda \in [-1, 1]$. Let

$$F_m(\theta) := \sum_{j=-m}^{m} (1 - |j|/m)e^{ij\theta} = 1 + 2\sum_{j=1}^{m}(1 - j/m)\cos(j\theta)$$

denote the Fejer kernel of order m and recall that $F_m(\theta) \geq 0$ and $F_m(0) = m$. Let $\lambda = \cos\phi$ for $\phi \in [0, \pi]$ and define

$$K_\phi(\theta) := \frac{1}{2}\Big(F_m(r(\theta - \phi)) + F_m(r(\theta + \phi))\Big) - 1,$$

and notice that $K_\phi(\theta) \geq -1$,

$$K_\phi(\phi) \geq \frac{1}{2}(F_m(0) + 0) - 1 = m/2 - 1, \tag{3}$$

and

$$\begin{aligned} K_\phi(\theta) &= \sum_{j=1}^{m}(1 - j/m)\cos(jr(\theta - \phi)) + \cos(jr(\theta + \phi)) \\ &= \sum_{j=1}^{m} 2(1 - j/m)\cos(jr\phi)\cos(jr\theta) \\ &= 2\sum_{j=1}^{m}(1 - j/m)\cos(jr\phi)T_{jr}(\cos(\theta)). \end{aligned}$$

Let

$$f(x) := \sum_{j=1}^{m}(1 - j/m)(\cos(jr\phi) + 1)T_{jr}(x),$$

so that

$$f(\cos(\theta)) = \frac{K_\phi(\theta) + K_0(\theta)}{2}.$$

The first property is implied by (3) and $K_0(\theta) \geq -1$:

$$f(\lambda) = K_\phi(\phi)/2 + K_0(\phi)/2 \geq \frac{m}{4} - 1.$$

The second property holds for $x = \cos(\theta) \in [-1, 1]$ since $K_\phi(\theta), K_0(\theta) \geq -1$. For $x \notin [-1, 1]$ we observe that f is a nonnegative linear combination of Chebyshev polynomials of even degree, which are nonnegative outside $[-1, 1]$. For the third property, we observe that:

$$\begin{aligned}
\|f(A/2\sqrt{d})\|_{1\to\infty} &\leq \sum_{j=1}^{m} |(1 - j/m)(\cos(jr\phi) + 1)| \|T_{jr}(A/2\sqrt{d})\|_{1\to\infty} \\
&\leq 2(d-1)\left(\frac{1}{2d^{r/2}} + \frac{1}{2d^{2r/2}} + \dots\right) \\
&\leq \frac{2(d-1)}{d^{r/2}},
\end{aligned}$$

by Lemma 2.1, as desired.

If $\lambda \notin [-1, 1]$, then we simply use the polynomial f corresponding to $\lambda = 1$ (which by symmetry is the same as the one for $\lambda = -1$) Properties (2) and (3) continue to hold, and property (1) holds because f is a nonnegative linear combination of even degree Chebyshev polynomials, which are increasing on $[1, \infty)$ and decreasing on $(-\infty, 1]$. □

Proof of Theorem 1.2. Let λ be an eigenvalue of $A/2\sqrt{d}$ with normalized eigenvector v. Let f be the polynomial from Lemma 2.2 applied to λ, $m = \lceil 4/\varepsilon \rceil + 4$, and $r = \lceil g/2m \rceil - 1$ or $\lceil g/2m \rceil - 2$, whichever is even. Taking $K = f(A/2\sqrt{d})$, we then have:

$$\langle v1_S, Kv1_S \rangle \leq \|K\|_{1\to\infty} \|v1_S\|_1^2 \leq 2d \cdot d^{-\varepsilon g/4+1} \cdot |S| \|v1_S\|_2^2 = 2d^{2-\varepsilon g/4} |S| \varepsilon, \tag{4}$$

since $r \geq \varepsilon g/8 - 2$, by property (3) of Lemma 2.2.

On the other hand, decompose $v1_S$ as $av + bw$ for where w is a unit vector orthogonal to v. Observe that

$$a = \langle v1_S, v \rangle = \|v_S\|^2 = \varepsilon,$$

and

$$b^2 = \|v1_S\|^2 - a^2 = \varepsilon(1 - \varepsilon).$$

Since $\langle v, Kw \rangle = 0$, we have:

$$\begin{aligned}
\langle v1_S, Kv1_S \rangle &= a^2 \langle v, Kv \rangle + b^2 \langle w, Kw \rangle \\
&\geq a^2(1/\varepsilon) - b^2 \quad \text{by (1) and (2) of Lemma 2.2} \\
&= \varepsilon - \varepsilon(1 - \varepsilon) = \varepsilon^2.
\end{aligned}$$

Combining this with (4), we obtain:

$$|S| \geq \frac{d^{\varepsilon g/4} \varepsilon}{2d^2},$$

as desired. □

Remark 2.3 (Improvement in the Untempered Case)**.** *The proof of Lemma 2.2 is clearly wasteful with regards to eigenvalues of A outside* $[-2\sqrt{d}, 2\sqrt{d}]$*; in particular, noting that Chebyshev polynomials blow up exponentially outside* $[-1, 1]$*, one can considerably improve the approximation bound for untempered* λ *(see Remark 1.5), and obtain a significantly stronger delocalization result in this case.*

3 Lower Bound

In this section, we prove the Theorem 1.6, which shows that the logarithmic dependence on k and polynomial dependence on ε in Theorem 1.2 are sharp up to a $\log(1/\varepsilon)$ term.

The starting point is to observe that eigenvectors of finite trees already have good localization properties. For the remainder of the section, we will refer to a complete tree of some finite depth D (i.e., $D+1$ levels of vertices including the root) in which every non-leaf vertex has degree $d+1$ as a $d-$*ary tree*. Note that by symmetry, every eigenvalue of such a tree has an eigenvector which assigns the same value to every vertex in a level (i.e., set of vertices at a particular distance from the root) — we will refer to such vectors as *symmetric*.

We begin by recording some facts about eigenvalues and eigenvectors of $d-$ary trees. Recall that the eigenvalues of a $d-$ary tree are contained in the interval $(-2\sqrt{d}, 2\sqrt{d})$ [HLW06, Section 5].

Lemma 3.1 (Eigenvalues of $d-$ary Trees)**.** *The set of eigenvalues of any infinite sequence of distinct finite* $d-$*ary trees is dense in the interval* $(-2\sqrt{d}, 2\sqrt{d})$*.*

Proof. Let $T_1, T_2, \ldots, T_m, \ldots$ be an infinite sequence of $d-$ary trees. Let T be the infinite $(d+1)-$regular tree with root r and observe that there are sets $S_1 \subset S_2, \ldots$ such that T_m is the induced subgraph of T on S_m. Let A_m be the adjacency matrix of T_m and let A be the adjacency matrix of T. Assume for contradiction that there is a closed interval

$$I = [\lambda - \eta, \lambda + \eta] \subset (-2\sqrt{d}, 2\sqrt{d})$$

such that every A_m has no eigenvalues in I.

We derive a contradiction with the fact [HLW06, Theorem 5.2] that for $\lambda \in \mathrm{spec}(A) = [-2\sqrt{d}, 2\sqrt{d}]$, there is no vector $v \in \ell_2$ such that $(\lambda I - A_T)v = e_r$, where e_r is the indicator of r. Our assumption implies that $\|(\lambda I - A_m)^{-1}\| \leq \eta^{-1}$ for all m, so there must be a sequence of finite dimensional vectors $\{v_m\}$ with $v_m \in \mathbb{C}^{S_m}$ such that

$$(\lambda I - A_m)v_m = P_m(\lambda I - A)P_m^T v_m = P_m e_r,$$

where $P_m : \ell_2 \to \mathbb{C}^{S_m}$ is the restriction onto $\mathbb{C}^{S_m}$, and we have the uniform bound

$$\|P_m^T v_m\| = \|v_m\| \leq \eta^{-1}$$

for all m. Banach-Alaoglu implies that $\{P_m^T v_m\}$ must have a weakly convergent subsequence; let $v \in \ell_2$ be the weak limit of this subsequence, and note that v must satisfy:

$$e_j^T(\lambda I - A)v = \lim_{m\to\infty} e_j^T P_m^T P_m(\lambda I - A)P_m^T v_n = \lim_{m\to\infty} e_j^T P_m^T P_m e_r = e_j^T e_r,$$

for every vertex $j \in T$, so in fact we must have $(\lambda I - A)v = e_r$, which is impossible. ◻

Lemma 3.2 (Eigenvectors of $d-$ary Trees)**.** *Assume $d \geq 2$ and let T be a $d-$ary tree of depth D with root r. Let $S_0 = \{r\}, S_1, \dots, S_D \subset T$ be the vertices at levels $0, 1, \dots, D$ of the tree and let v be a symmetric eigenvector of its adjacency matrix with eigenvalue $\lambda = 2\sqrt{d}\cos\theta \in (-2\sqrt{d}, 2\sqrt{d})$. Then every pair of adjacent levels has approximately the same total ℓ_2^2 mass as the root:*

$$\Omega(\sin^2\theta) = \frac{\|v_{S_i}\|_2^2 + \|v_{S_{i+1}}\|_2^2}{\|v(r)\|_2^2} = O(1/\sin^2\theta).$$

Proof. Suppose v has value x_i for all vertices in S_i, and for convenience assume that the root has value $x_0 = 1$ (although this makes v un-normalized). The eigenvector equation at the non-leaf vertices yields the following quadratic recurrence:

$$\begin{aligned} \lambda x_0 &= (d+1)x_1, \\ \lambda x_i &= x_{i-1} + dx_{i+1} \quad 1 \leq i \leq D-1, \end{aligned}$$

which must be satisfied by any eigenvector (ignoring the boundary condition at the leaves). Since we are interested in the total ℓ_2^2 mass at each level, it will be more convenient to work with the quantities

$$m_0 = x_0 = 1 \qquad \text{and} \qquad m_i = \sqrt{|S_i|}x_i = \sqrt{(d+1)d^{i-1}}x_i \quad 1 \leq i \leq D,$$

which satisfy $m_i^2 = \|v_{S_i}\|_2^2$. Rewriting the recurrence in terms of the m_i, we obtain:

$$\begin{aligned} m_1 &= \frac{\lambda}{\sqrt{d+1}} m_0 \\ m_{i+1} &= \frac{\lambda m_i}{d} \cdot \sqrt{\frac{|S_{i+1}|}{|S_i|}} - \frac{m_{i-1}}{d} \cdot \sqrt{\frac{|S_{i+1}|}{|S_{i-1}|}} = \frac{\lambda}{\sqrt{d}} m_i - m_{i-1}, \quad 1 \leq i \leq D+1. \end{aligned}$$

Letting $\lambda = 2\sqrt{d}\cos\theta$ and writing the above in matrix form, we have

$$w_{i+1} := \begin{bmatrix} m_{i+1} \\ m_i \end{bmatrix} = \begin{bmatrix} 2\cos\theta & -1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} m_i \\ m_{i-1} \end{bmatrix} =: PDP^{-1}w_i,$$

since the matrix above is diagonalizable for $\theta \neq 0, \pi$, with

$$D := \begin{bmatrix} e^{i\theta} & 0 \\ 0 & e^{-i\theta} \end{bmatrix}, \quad P = \begin{bmatrix} 1 & 1 \\ e^{-i\theta} & e^{i\theta} \end{bmatrix}.$$

Since D is unitary we have $\|P^{-1}w_i\| = \|P^{-1}w_0\|$ for all i. Observe that

$$\|P\| \leq 2, \quad \|P^{-1}\| \leq \frac{1}{\sin\theta},$$

whence

$$\frac{\|w_i\|}{\|w_0\|} \in [\frac{\sin\theta}{2}, \frac{2}{\sin\theta}].$$

Noting that $m_1 \leq 2m_2$ and squaring yields the claim. □

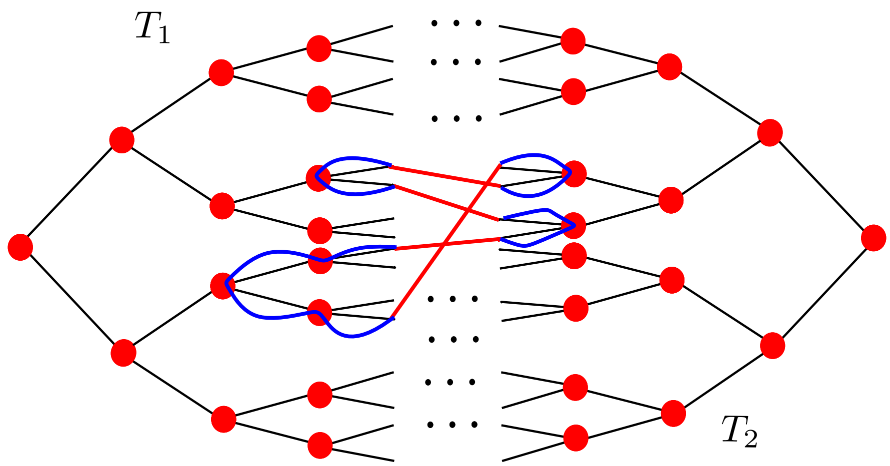

Figure 2: This describes the construction of gluing two trees by a random matching of the leaves. A cycle in the glued graph is illustrated where the blue paths denote the excursions into the trees and the red edges denote the jump from one tree to the other. Note that all the interior vertices in the trees do not have degree three and the roots only have degree two (since the third edge is not significant for the purposes of the illustration, it was omitted to avoid cluttering in the figure).

Let T be a $d-$ary tree of depth D. Choosing S to be the top $\lfloor \varepsilon D \cdot (d/d+1) \rfloor$ levels of T and applying the above lemma to any eigenvector with eigenvalue bounded away from $\pm 2\sqrt{d}$, we find that $\|v_S\|_2^2 = \Theta(\varepsilon)$ and $|S| = O((d+1)^{\varepsilon D}) = O(n^{\varepsilon})$, where n is the number of vertices in the tree. This is exactly the kind of localization we want for our lower bound. Unfortunately, finite $d-$ary trees are not regular because they have leaves. The rest of this section is devoted to showing that we can nonetheless embed these trees in $(d+1)-$regular graphs without disturbing their eigenvectors or creating any short cycles, thereby establishing Theorem 1.6. The main device in doing this is the following lemma which shows that it is possible to identify the leaves of two trees in a manner which does not introduce short cycles.

Lemma 3.3 (Pairing of Trees). *Suppose T_1 and T_2 are two $d-$ary trees of depth D, each with $n = (d+1)d^{D-1}$ leaf vertices V_1 and V_2. Then for sufficiently large n there is a bijection $\pi : V_1 \to V_2$ such that the graph obtained by identifying v with $\pi(v)$ has girth at least $\log_d(n)/4$.*

We defer the proof of this lemma. The second ingredient is:

Lemma 3.4 (Degree-Fixing Gadget). *For every degree $d \geq 3$ and sufficiently large n, there is a graph H on n vertices with the following properties:*

1. *H has one distinguished vertex of degree $d-2$ and the remaining vertices have degree d.*
2. *H has girth at least $\log_{d-1}(n)/3$.*

Proof. According to Corollary 2 of [MWW04], the number of $d-$regular graphs with girth at least g is asymptotic to

$$\frac{(dn)!}{(dn/2)!2^{dn/2}(d!)^n} \cdot \exp\left(-\sum_{r=1}^{g} \frac{(d-1)^r}{2r} + o(1)\right) \geq \exp\left(dn/2 - g(d-1)^r + o(1)\right),$$

whenever $(d-1)^{2g-1} = o(n)$. Taking $g = \log_{d-1}(n)/3 + 1$ we find that this condition is satisfied and the right hand side is positive for large enough n. Let G be a d−regular graph on n vertices with girth at least g. Let v be any vertex of G and let u_1, u_2 be two of its neighbors. Let H be the graph obtained by deleting the edges vu_1 and vu_2 and adding the edge u_1u_2. Observe that v has degree $d-2$ and every other vertex has degree d in H. Moreover, since we replaced an edge by a path of length 2, the length of every cycle decreases by at most 1, so H must have girth at least $\log_{d-1}(n)/3$, as desired. □

Proof of Theorem 1.6. Let $d \geq 2$, and let k be any integer larger than the n required for Lemmas 3.3 and 3.4 to apply. Suppose $\varepsilon \in (0, 1)$ is given. Let t be the largest integer such that $(d+1)d^{t-1} \leq k$. Choose $D - 1 = \lceil t/\varepsilon \rceil$ and let T_1 and T_2 be two disjoint d−ary trees with D levels. Let S_1 and S_2 be the sets of vertices consisting of the top t levels of T_1 and T_2 respectively.

Let λ be an eigenvalue of A_{T_1} and let f be the corresponding normalized symmetric eigenvector, (f to has the same value on every vertex within a level) as in Lemma 3.2. By Lemma 3.2, we know

$$\|f_{S_1}\|_2^2 = \|f_{S_2}\|_2^2 = \Omega_\lambda(\epsilon),$$

where the implicit constant is $\Omega(\sin^4\theta)$ for $\lambda = 2\cos\theta$.

Construct T_1' and T_2' by attaching d new *marked* vertices to each leaf of T_1 and T_2, respectively, so that they are d−ary trees of depth D, with $n = (d+1)d^{D-1}$ leaves each, corresponding to the marked vertices. Apply Lemma 3.3 to pair these marked leaves; call the resulting graph H. Notice that H is $d+1$-regular except for the marked vertices, which have degree two. Apply Lemma 3.4 with degree parameter $d+1$ and size n to obtain a graph W with a single distinguished vertex of degree $d-1$, all remaining vertices of degree $d+1$, and girth $\log_d(n)/3$. Finally, let G be the graph obtained by identifying each marked vertex with the distinguished vertex of W in its own copy of W (for a total of n copies).

Observe that G has girth at least $\log_d(n)/4 = \Omega(\log_d(k)/\varepsilon)$, since H has girth at least this much and attaching disjoint copies of W at single vertices does not create any new cycles.

Let v be the function equal to f on vertices of T_1, $-f$ on vertices of T_2, and zero elsewhere. We claim that v is an eigenvector of G with eigenvalue λ. To see this one has to verify the eigenvector equation at vertices of three kinds:

- At every vertex of T_1 and T_2 because all new neighbors of those vertices are assigned a value of 0 in v.
- It is also satisfied at the marked vertices, because every such vertex is adjacent to exactly one leaf in T_1 and one leaf in T_2. which have the same values with opposite signs,
- The remaining vertices in copies of W have value zero, so the eigenvector equation is trivially satisfied.

Observing that $\|v_{S_1 \cup S_2}\|_2^2 = \Omega_\lambda(\varepsilon)$ with $|S_1 \cup S_2| = 2k$ finishes the construction. Since this construction is valid for infinitely many n, Lemma 3.1 implies that the set of eigenvalues for which it works is dense in $(-2\sqrt{d}, 2\sqrt{d})$, as desired.

□

The rest of this section is devoted to proving Lemma 3.3.

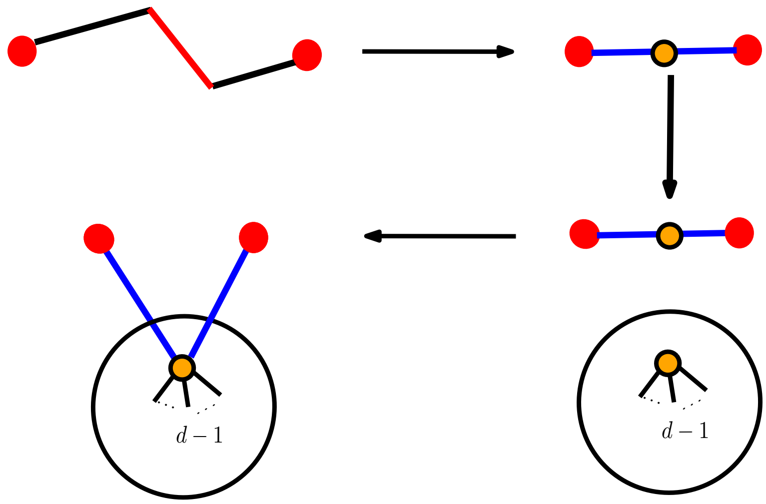

Figure 3: This illustrates the use of the degree correcting gadget. The two red vertices in the first figure denote two matched vertices in T_1 and T_2 respectively (see Figure 2). The matched leaf vertex (yellow) is then identified with a degree $d-1$ vertex in the gadget.

Probabilistic Model. Our construction is probabilistic, and inspired by the switching argument of [MWW04], but with much cruder estimates since we are not interested in precise asymptotic enumeration, but only in showing that a certain probability is not zero. Let T_1 and T_2 be two d−ary trees of depth D, each with exactly $n = (d+1)d^{D-1}$ leaf vertices, henceforth denoted V_1 and V_2. Consider the random graph G obtained by taking the union of T_1 and T_2 and a perfect matching between V_1 and V_2.

Graph-Theoretic Terminology. A *cycle* in a graph is an oriented closed walk with no repeated edges. We will consider cyclic shifts and reversals of a cycle to be the same cycle. A cycle in G can always be written as a sequence of alternating *matching edge traversals* e_i and *tree excursions* γ_i:

$$e_1, \gamma_1, e_2, \gamma_2, \ldots, e_k, \gamma_k$$

(or equivalently $\gamma_1, e_2, \gamma_2, \ldots, e_k, \gamma_k, e_1$), where the γ_i are simple paths in either T_1 or T_2 with endpoints at leaves, and γ_k ends where e_1 begins. We will follow the convention that $\gamma_1, \gamma_3, \ldots$ are excursions in T_1 and $\gamma_2, \gamma_4, \ldots$ are excursions in T_2. The total number of edges in the cycle will be called its *length*.

We begin by establishing some preliminary facts about short cycles in G.

Lemma 3.5 (Number and Overlaps of Short Cycles). *Let $c < 1/2$ be a constant. Then for sufficiently large n, with constant probability we have both of the following properties.*

1. *G does not contain two cycles which share an edge,*
2. *G contains at most $B = O(n^{c(1+o(1))})$ cycles of length at most L,*

where $L := 2c\log_d(n)$.

Proof. Let $v \in V_1$ be a leaf vertex. We will first show that

$$\mathbb{P}[v \text{ occurs in} \geq 1 \text{ cycle of length} \leq L] = O(n^{-1+c(1+o(1))}). \tag{5}$$

Call a cycle that occurs in G with nonzero probability a *potential cycle*. Every potential cycle consists of k matching traversals and k tree excursions for some even k. Observe that every excursion of length h has even length and consists of $h/2$ upward steps towards the root of the tree and $h/2$ downward steps back down to the leaves.

Given a starting vertex for the excursion, the upward steps are uniquely determined, and there are at most d choices for each of the downward steps (since backtracking is not allowed, and the root has degree $d+1$). Since there are at most n choices for each matching traversal given one of its endpoints, the total number of potential cycles containing v with exactly k matching traversals and excursions of lengths $h_1, \ldots, h_k$ is at most:

$$d^{h_1/2} \cdot n \cdot d^{h_2/2} \cdot n \ldots d^{h_k/2} = d^{(h_1 + \ldots h_k)/2} \cdot n^{k-1} \leq n^{k-1+c}, \tag{6}$$

since the last matching traversal is determined by the starting vertex v. Every such potential cycle fixes k matching edges, so the probability that it occurs in a random matching is at most $(n-k)!/n!$. Taking a union bound over all $k \leq L$, ordered partitions $h_1 + \ldots + h_k \leq L - k$, and potential cycles with those parameters, we have

$$\mathbb{P}[v \text{ occurs in a cycle of length} \leq L] \leq \sum_{k=1}^{L} e^{O(\sqrt{L})} k! \cdot \frac{n^{k-1+c}(n-k)!}{n!} = O(n^{-1+c(1+o(1))}).$$

Next, we use a similar argument to estimate the probability that any vertex v occurs in more than one cycle. Fix $v \in V_1$ and observe that a pair of potential cycles of length at most L both containing v can be specified by the following choices:

- Lengths $s, s' \leq L$, matching traversal counts $k, k' \leq L$, and tuples of excursion lengths $h_1, \ldots, h_k$ and $h'_1, \ldots, h'_{k'}$ for both cycles.
- The common matching edge e incident to v contained in both cycles.
- The excursions made in both cycles.
- The remaining $k-2$ and $k'-2$ matching edges in both cycles (noting that the final edge is not required once all excursions are specified).

Since any particular pair fixes $k + k' - 1$ matching edges, the probability that it occurs in G is at most $(n - (k + k' - 1))!/n!$. Bounding excursions as in (6) and taking a union bound, we have

$$\begin{aligned}
&\mathbb{P}[v \text{ occurs in} \geq 2 \text{ cycles of length} \leq L] \\
&\leq L^4 e^{O(\sqrt{L})} k!(k')! \cdot n \cdot n^{2c} \cdot n^{k+k'-4} \cdot \frac{(n-(k+k'-1))!}{n!} \\
&= O(n^{2c(1+o(1))-2}).
\end{aligned}$$

Taking a union bound over all vertices, we conclude that

$$\mathbb{P}[G \text{ contains two non-disjoint cycles of length} \leq L] = O(n^{2c(1+o(1))-1}) = o(1).$$

Since cycles in G are vertex disjoint if and only if they are edge disjoint, the first claim follows.

For the second claim we sum (5) over all v and apply Markov to obtain:

$$\mathbb{P}[|V_C| > O(n^{c(1+o(1))})] < 1/2,$$

where $|V_C|$ denotes the set of vertices contained in at least one cycle of length $\leq L$. Taking a union bound with our previous conclusion, we have that with probability 1/3 all cycles in G are disjoint and $|V_C| = O(n^{c(1+o(1))})$. Since every cycle contains at least one vertex, this gives the second claim. □

Proof of Lemma 3.3. Let $c = 1/4$ and $L := 2c\log_d n$, and choose n sufficiently large so that Lemma 3.5 applies with $B = O(n^{c(1+o(1))}) \leq n^{1/3}$. Let Γ be the set of graphs in the support of G such that both conditions of Lemma 3.5 are satisfied, and note that

$$|\Gamma| = \Omega(n!). \tag{7}$$

For integers $z_2, z_4, \ldots, z_L \leq B$ let

$$\Gamma(z_2, \ldots, z_L)$$

denote the subset of Γ containing graphs with exactly z_j cycles of length j, noting that in our model there are never any odd cycles.

Our goal is to show that $\Gamma(0, \ldots, 0)$ is not empty. Following [MWW04], our strategy will be to establish the following two claims

Claim 3.6. *There exists a $z^* \in [B]^L$ such that*

$$|\Gamma(z_2^*, \ldots, z_L^*)| \geq \exp(\Omega(n\log n)).$$

Claim 3.7. *For every $z \in [B]^L$ such that $z_k > 0$:*

$$\frac{|\Gamma(z_2, \ldots, z_k - 1, \ldots, z_L)|}{|\Gamma(z_2, \ldots, z_k, \ldots, z_L)|} = \Omega(n^{-3c}).$$

Iterating the above claims yields

$$|\Gamma(0, \ldots, 0)| \geq \exp(Cn\log n - O(\log(n)) \sum_{i=2,\ldots L} z_i) = \exp(Cn\log n - \Omega(n^{1/3}\log^2 n)) > 1,$$

so that with nonzero probability G has no cycles of length at most L. Contracting all matching edges shrinks the length of every cycle by at most a factor of 2, yielding the desired pairing.

To establish Claim 3.6, we observe that the tuple $z^* \in [B]^L$ which maxmizes $|\Gamma(z^*)|$ must have cardinality at least

$$\frac{\Omega(n!)}{B^L} \geq \Omega\left(\exp\left(n\log n - O(n) - O(\log^2 n)\right)\right).$$

For Claim 3.7 we use a switching argument. Given a graph $H \in \Gamma(z_2, \ldots, z_k, \ldots, z_L)$ with $z_k > 0$, a *forward switching* is defined as the following operation:

- Choose the lexicographically[2] first matching edge $e = st$ in the lexicographically first cycle C of length k in H.

[2] Fix an arbitrary ordering of edges and cycles.

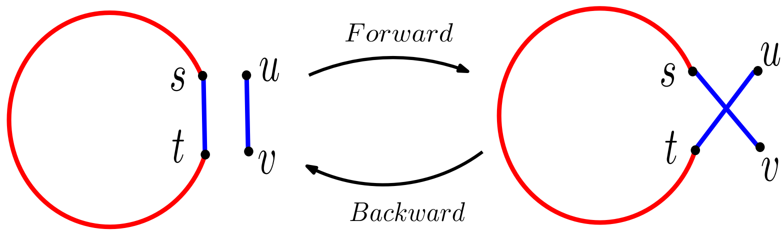

Figure 4: This illustrates the switching argument: forward switching takes a cycle and an edge and produces a path while the backward switching does the reverse.

- Choose any matching edge $f = uv$ in H at distance at least $2L$ from e which is not contained in any cycle of length at most L.
- Remove st and uv from the matching and add sv and ut.

Observe that since every edge is contained in exactly one cycle and f does not belong to any cycle, removing e and f destroys only the cycle C. Since the endpoints of e and f are at distance $2L$ in H, adding sv and ut does not create any cycles of length at most L. Thus, the outcome of a forward switching is a graph $H' \in \Gamma(z_2, \ldots, z_k - 1, \ldots, z_L)$, which has exactly the same cycle counts except with one less cycle of length k.

Let $\mathcal{F}(H)$ denote the set of foward switchings of a graph $H \in \Gamma(z_2, \ldots, z_L)$. Observe that the only choice in the switching is the choice of the second matching edge f. The number of matching edges contained in cycles of length at most L is bounded by $BL = o(n)$ and the number of edges within distance $2L$ of e is at most $(d+1)^{2L} = o(n)$. Therefore, for every $H \in \Gamma(z_2, \ldots, z_L)$, we have

$$|\mathcal{F}(H)| = \Omega(n).$$

We now investigate how many forward switchings can map to a given graph H' in $\Gamma(z_2, \ldots, z_k - 1, \ldots, z_L)$. Given such an H', a *backward switching* is defined as the following operation:

- Choose two vertices u and v at distance exactly $k+1$ in H', such that (a) the extreme edges ut and vs of the uv–path p are matching edges. (b) The distance between u and v in H' along any path other than p is at least L.
- Delete the edges ut and vs from the matching and add edges st and uv.

Observe that a backward switching always yields a graph $H \in \Gamma(z_2, \ldots, z_L)$, and that all graphs H with a forward switching equal to H' may be achieved in this manner. The number of backward switchings of any graph H' is upper bounded by

$$|\mathcal{B}(H')| \leq n \cdot (d+1)^{L+1} = O(n^{1+3c}),$$

where we have overcounted by ignoring the conditions (a) and (b) in the definition of a backward switching.

A double counting argument now yields:

$$\frac{|\Gamma(z_2, \ldots, z_k - 1, \ldots, z_L)|}{|\Gamma(z_2, \ldots, z_k, \ldots, z_L)|} \geq \frac{\min_H |\mathcal{F}(H)|}{\max_{H'} |\mathcal{B}(H')|} = \Omega(n^{-3c}),$$

as desired.

□

Acknowledgment. We would like to thank Mark Rudelson for a helpful conversation, and MSRI and the Simons Institute for the Theory of Computing, where this work was partially carried out.

References

[ABLS07] Noga Alon, Itai Benjamini, Eyal Lubetzky, and Sasha Sodin. Non-backtracking random walks mix faster. *Communications in Contemporary Mathematics*, 9(04):585–603, 2007.

[ALM15] Nalini Anantharaman and Etienne Le Masson. Quantum ergodicity on large regular graphs. *Duke Mathematical Journal*, 164(4):723–765, 2015.

[BHY] Roland Bauerschmidt, Jiaoyang Huang, and Horng-Tzer Yau. Local kesten–mckay law for random regular graphs, preprint (2016). *arXiv preprint arXiv:1609.09052*.

[BL13] Shimon Brooks and Elon Lindenstrauss. Non-localization of eigenfunctions on large regular graphs. *Israel Journal of Mathematics*, pages 1–14, 2013.

[Bro09] Simon Brooks. *Entropy bounds for quantum limits*. Princeton University, 2009.

[BS16] Ágnes Backhausz and Balázs Szegedy. On the almost eigenvectors of random regular graphs. *arXiv preprint arXiv:1607.04785*, 2016.

[DLL11] Yael Dekel, James R Lee, and Nathan Linial. Eigenvectors of random graphs: Nodal domains. *Random Structures & Algorithms*, 39(1):39–58, 2011.

[ERS17] Ronen Eldan, Miklós Z Rácz, and Tselil Schramm. Braess's paradox for the spectral gap in random graphs and delocalization of eigenvectors. *Random Structures & Algorithms*, 50(4):584–611, 2017.

[Gei13] Leander Geisinger. Convergence of the density of states and delocalization of eigenvectors on random regular graphs. *arXiv preprint arXiv:1305.1039*, 2013.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.

[Kah92] Nabil Kahale. On the second eigenvalue and linear expansion of regular graphs. In *Foundations of Computer Science, 1992. Proceedings., 33rd Annual Symposium on*, pages 296–303. IEEE, 1992.

[Lin06] Elon Lindenstrauss. Invariant measures and arithmetic quantum unique ergodicity. *Annals of Mathematics*, pages 165–219, 2006.

[MWW04] Brendan D McKay, Nicholas C Wormald, and Beata Wysocka. Short cycles in random regular graphs. *the electronic journal of combinatorics*, 11(1):R66, 2004.

[OVW16] Sean O'Rourke, Van Vu, and Ke Wang. Eigenvectors of random matrices: a survey. *Journal of Combinatorial Theory, Series A*, 144:361–442, 2016.

[RS94] Zeév Rudnick and Peter Sarnak. The behaviour of eigenstates of arithmetic hyperbolic manifolds. *Communications in Mathematical Physics*, 161(1):195–213, 1994.

[Rud17] Mark Rudelson. Delocalization of eigenvectors of random matrices, 2017.

[RV15] Mark Rudelson and Roman Vershynin. Delocalization of eigenvectors of random matrices with independent entries. *Duke Mathematical Journal*, 164(13):2507–2538, 2015.

[RV16] Mark Rudelson and Roman Vershynin. No-gaps delocalization for general random matrices. *Geometric and Functional Analysis*, 26(6):1716–1776, 2016.