Spectral Expanding Expanders

Gil Cohen ✉
Department of computer science, Tel Aviv University, Israel

Itay Cohen ✉
Department of computer science, Tel Aviv University, Israel

─── **Abstract** ───────────────────────────────────

Dinitz, Schapira, and Valadarsky [5] introduced the intriguing notion of *expanding expanders* – a family of expander graphs with the property that every two consecutive graphs in the family differ only on a small number of edges. Such a family allows one to add and remove vertices with only few edge updates, making them useful in dynamic settings such as for datacenter network topologies and for the design of distributed algorithms for self-healing expanders. [5] constructed explicit expanding-expanders based on the Bilu-Linial construction of spectral expanders [3]. The construction of expanding expanders, however, ends up being of *edge* expanders, thus, an open problem raised by [5] is to construct *spectral* expanding expanders (SEE).

In this work, we resolve this question by constructing SEE with spectral expansion which, like [3], is optimal up to a poly-logarithmic factor, and the number of edge updates is optimal up to a constant. We further give a simple proof for the existence of SEE that are close to Ramanujan up to a small additive term. As in [5], our construction is based on interpolating between a graph and its lift. However, to establish spectral expansion, we carefully weigh the interpolated graphs, dubbed partial lifts, in a way that enables us to conduct a delicate analysis of their spectrum. In particular, at a crucial point in the analysis, we consider the *eigenvectors* structure of the partial lifts.

1 Introduction

Expander graphs are among the most useful combinatorial objects in theoretical computer science, and in computer science in general. In theory, expanders proved to be pivotal in many groundbreaking results (e.g., [22, 1, 10, 6, 17, 19, 7, 4]). Informally, expanders are sparse undirected graphs that have many desirable pseudorandom properties.

There are several ways of defining the expansion of a graph. Taking the combinatorial perspective, one thinks of the edge- or vertex-expansion, whereas from the spectral point of view, the spectral expansion is considered. The latter coincides with the Markovian point of view as it captures the rate at which random walks converge. If one is willing to absorb some deterioration in parameters, it is possible to move from one definition to the next, and so in the non-extreme regime of parameters, and only there, the different definitions are, in a sense, equivalent. This work concerns with spectral expansion and so we recall the definition right away. For the formal definition of other notions of expansion, and for the relations between them, we refer the reader to the wonderful texts [9, 21, 20, 18].

Let G be an undirected graph with adjacency matrix $\mathbf{A}$. Since $\mathbf{A}$ is symmetric it has n real eigenvalues which we denote by $\lambda_1 \geq \cdots \geq \lambda_n$. The *spectral expansion*[1] of G is defined by $\lambda(G) \triangleq \max(\lambda_2, |\lambda_n|)$. As mentioned, the reason $\lambda(G)$ is of interest is mostly due to the fact that it captures the rate of converges of random walks on *regular* graphs. Indeed, for d-regular graphs the adjacency matrix is a simple normalization of the random walk matrix, $\mathbf{W} = \frac{1}{d}\mathbf{A}$. For many applications in theoretical computer science, and in particular in a typical work on expander graphs, restricting to regular graphs is a nonissue, and so the spectrum of the adjacency matrix is studied instead of that of the random walk matrix. For arbitrary undirected graphs, as those we will work with [2], the random walk matrix can be written as $\mathbf{W} = \mathbf{A}\mathbf{D}^{-1}$, where $\mathbf{D}$ is the diagonal matrix that encodes the degrees of the vertices of G. Since $\mathbf{W}$ is similar to the symmetric matrix $\mathbf{D}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}^{-\frac{1}{2}}$, its eigenvalues are real, denoted $\bar{\lambda}_1 \geq \cdots \geq \bar{\lambda}_n$. The parameter of interest, which determines the rate of converges, is then the *normalized spectral expansion*, defined by $\bar{\lambda}(G) = \max(\bar{\lambda}_2, |\bar{\lambda}_n|)$.

Typically, when working with expanders one cares not about one graph but rather about an infinite family of graphs $\mathcal{G} = (G_n)_{n \in \mathcal{I}}$ where, for each $n \in \mathcal{I}$, G_n is an undirected graph on n vertices. It is typically assumed that all graphs in the family have bounded degree d and, naturally, the denser $\mathcal{I}$ is in $\mathbb{N}$, the better. We denote $\lambda(\mathcal{G}) = \sup_n \lambda(G_n)$ and similarly for the normalized quantity $\bar{\lambda}(\mathcal{G}) = \sup_n \bar{\lambda}(G_n)$. For most applications, one requires the expander family $\mathcal{G}$ to be explicit which, in this work, means that given $n \in \mathcal{I}$, the graph G_n can be generated in poly(n)-time.

1.1 Expanding expanders

Dinitz, Schapira and Valadarsky [5] introduced a natural and intriguing aspect of families of expander graphs, which, informally captures the extent to which the family is "continuous". Formally, let $\mathcal{G} = (G_n = (V_n, E_n))_{n \in \mathcal{I}}$ be a family of expander graphs. Dinitz *et al.* considered the *expansion cost* $c(\mathcal{G})$ which is the number of edges one must add or remove, from any graph in the family so as to obtain the next. If we denote the least element in $\mathcal{I}$ that is larger than n by $\mathrm{next}_{\mathcal{I}}(n)$, then the expansion cost can be written as

$$c(\mathcal{G}) \triangleq \max_{n \in \mathcal{I}} \left| E_n \triangle E_{\mathrm{next}_{\mathcal{I}}(n)} \right|$$

if this maximum exists and ∞ otherwise.

Clearly, high $\delta(\mathcal{G}) \triangleq \max_{n \in \mathcal{I}} (\mathrm{next}_{\mathcal{I}}(n) - n)$ implies high expansion cost. Typically, we will think of $\mathcal{I}$ as a very dense set in $\mathbb{N}$, in particular, $\delta(\mathcal{G})$ will be bounded by some small universal constant (independent of the degree). Dinitz, Schapira and Valadarsky initiated the study of families of expander graphs with bounded expansion cost $c(\mathcal{G}) < \infty$. As noted by [5], for d-regular expanders, $c(\mathcal{G}) \geq \frac{3d}{2}$ and so the authors asked whether there is an infinite family of expanders whose expansion cost is bounded by some constant $c = c(d)$.

With such a family at hand, one can add and remove vertices with low cost in terms of edge updates while maintaining expansion. This stands in contrast to, say, a randomly sampled family in which the difference between expanders of consecutive size is linear in the number of vertices and, in particular, is unbounded. The motivation of [5] for studying such families, dubbed *expanding expanders*, originated from datacenter network topologies.

---

[1] There is some harmless inconsistency in the literature regarding the definition of spectral expansion. Some sources refer to $d - \lambda(G)$ as the spectral expansion. Others consider $1 - \frac{1}{d}\lambda(G)$. In some cases, it is only λ_2 that is considered.

[2] As we prove in the full version of the paper, in our setting, one must resort to the use of irregular graphs.

The different servers of a datacenters correspond to the vertices of a graph and the edges represent the wires connecting the servers. The degrees being bounded implies low cost in wiring, and the expansion of the graph is essential for avoiding traffic routed ineffectively. As datacenters grow regularly, low expansion cost translates to a low overhead in rewiring when adding a new server.

As another application, [5] showed how to obtain better distributed algorithms for self-healing expanders, improving upon a prior work by Pandurangan, Robinson, and Trehan [16]. Informally, self-healing expanders are expanders that can, distributively, fix themselves when vertices are added or removed. In this setting, the fact that the expander family is deterministically constructed makes the distributive task significantly simpler as there is no randomness that is needed to be communicated.

The main result of [5] is an explicit construction of a family of d-regular expanders where the expansion is with respect to edge-expansion. More concretely, it was shown how to maintain edge expansion of roughly $\frac{d}{3}$ while guaranteeing expansion cost of at most $\frac{5d}{2}$.

1.2   Our results

The construction of Dinitz, Schapira, and Valadarsky [5] for edge-expanding expanders is based on the work of Bilu and Linial [3] who analyzed the operation of "lifting" a spectral expander as a way of obtaining a graph on, say, twice the number of vertices, while maintaining spectral expansion. The [5] construction goes by way of interpolating between every two consecutive Bilu-Linial *spectral* expanders, assuring good *edge*-expansion.

Dinitz *et al.* observed that in their interpolation, some of the graphs in the family are only weak spectral expanders. That is, despite the fact that they are interpolating between spectral expanders, the construction ends up only having good edge expansion. Therefore, Dinitz *et al.* left open the question of whether spectral expanding expanders can be constructed (or exist for that matter). In this work we answer this question to the affirmative by constructing spectral expanding expanders via lifting which, spectrally, are essentially as good as the ones we are interpolating between.

▶ **Theorem 1** (Main result). *For every integer $d \geq 3$ there exists an explicit family of undirected graphs $\mathcal{G}$ such that all vertices of every graph in the family has degree bounded in $[d, 4d]$. The expansion cost $c(\mathcal{G}) = O(d)$, and the spectral- and normalized-spectral expansions are given by*

$$\lambda(\mathcal{G}) = O\left(\sqrt{d\log^3 d}\right), \quad \bar{\lambda}(\mathcal{G}) = O\left(\sqrt{\frac{\log^3 d}{d}}\right).$$

We wish to stress that, as our construction is of irregular graphs, hence, a bound on $\lambda(\mathcal{G})$ does not imply the bound on the normalized $\bar{\lambda}(\mathcal{G})$. As it turns out, our proof for the bound on $\bar{\lambda}(\mathcal{G})$ is significantly more involved than the proof for the unnormalized $\lambda(\mathcal{G})$, though the latter too requires careful analysis. In particular, the bound on $\lambda(\mathcal{G})$ is used for obtaining the bound on $\bar{\lambda}(\mathcal{G})$. The proofs of the two bounds are given in Section 4 and Section 5, and the straightforward derivation of Theorem 1 is then given in Section 6.

Theorem 1 makes the important tool of random walks on expanders available to applications that require the dynamical setting offered by expanding expanders.

Before moving on, we remark that the graphs in the family $\mathcal{G}$ that is constructed in Theorem 1 have multiple edges. As for the density of the family, $\delta(\mathcal{G}) = O(1)$ and $\min \mathcal{I} = O(d)$.

1.2.1 Expanding Ramanujan graphs

The bound obtained by Theorem 1 for $\lambda(\mathcal{G})$, and similarly for the normalized version, is off by a poly-logarithmic factor from the spectral expansion of a Ramanujan graph, namely, $2\sqrt{d-1}$ which is optimal up to a negligible additive term [15]. We leave it as an intriguing open problem to determine whether expanding Ramanujan graphs exist. It seems to us that the known algebraic constructions (e.g., [11]) are not adequate for the setting of expanding expanders. In their seminal work [12, 13, 14] Marcus, Spielman and Srivastava proved the existence of bipartite Ramanujan graphs. The question of whether expanding bipartite Ramanujan graphs exists is too an interesting one. It is unclear to us whether the proof technique of Marcus *et al.* is suitable in the expanding setting.

In the full version of the paper we give a simple proof for the existence of nearly-Ramanujan spectral expanding expanders. While our proof does not yield an explicit construction, the question of whether expanding Ramanujan graphs exist is of interest already on the combinatorial level.

▶ **Theorem 2.** *For every $\epsilon > 0$ and every even integer $d \geq 6$, there exists an infinite family $\mathcal{G}$ of d-regular spectral expanding expanders with expansion cost $c(\mathcal{G}) \leq 3d$ and spectral expansion $\lambda(\mathcal{G}) \leq 2\sqrt{d-1} + \varepsilon$.*

1.2.2 On spectral expanding expanders meeting the Ramanujan bound and the ubiquitous of Ramanujan graphs

We find the question on the existence of expanding Ramanujan graphs, as well as its bipartite analog, to be interesting also in the context of non-expanding expanders. Indeed, the existence of a family of expanding Ramanujan graphs would arguably be an indication for an affirmative answer to the fundamental open problem that asks whether a random d-regular graph, under a natural distribution, is Ramanujan with positive probability. The reasoning being that if Ramanujan graphs are sparse among graphs then they should have a certain underlying structure. That this structure happens to coincide with the structure required by expanding expanders seems far-fetched.

By the above discussion, proving the existence of spectral expanding expanders that attain the Ramanujan bound would be an indication for the ubiquitous of Ramanujan graphs. We remark, as further discussed in the full version of the paper, that for the construction of regular expanding expanders, as required in the context of Ramanujan graphs, one would need to employ a different paradigm than ours.

1.3 Organization

In Section 2 we give a proof overview of our results. The reader may freely skip this section and jump straight to Sections 4–6 in which the formal proof of Theorem 1 is given. The first section, Section 4, covers the bulk of the proof for the unnormalized case, Section 5 handles the normalized case, and in Section 6 we derive the proof of Theorem 1. Theorem 2 is proved in the full version of the paper. In the full version of the paper we further discuss the extent to which irregular graphs are necessary for constructing spectral expanding expanders. We prove that within a fairly general framework for constructing spectral expanding expanders, regularity must be abandoned.

2   Proof Overview

In this section, we give a brief account on some of the ideas that go into the proof of Theorem 1. As mentioned, bounding $\lambda(\mathcal{G})$ is simpler than (and used by) the proof for the bound on $\bar{\lambda}(\mathcal{G})$. Therefore, we start by discussing the unnormalized spectral expansion. In any case, both proofs are based on the well-known notion of a *lift* of a graph, as well as on our extension of this notion we dub *partial lifts*. These are covered in Section 2.1 and Section 2.2, respectively. A more extensive treatment of lifts can be found in [2] and references therein.

2.1   Lifts

For an integer $k \geq 2$, a *k-lift* of an undirected graph $G = (V, E)$ is an undirected graph $\widehat{G}$ on the vertex set $[k] \times V$ where each edge $\{u, v\} \in E$ induces k edges in $\widehat{G}$ that form a perfect matching between the vertices $[k] \times \{u\}$ and $[k] \times \{v\}$. The set $[k] \times \{v\}$ is called the *fiber* of v. Thus, a k-lift is determined by a choice of one perfect matching per edge $\{u, v\} \in E$ that is placed between the fibers of the two endpoints u, v. A slightly more formal treatment of the notion of a k-lift is given in Section 3.2.

A well-known fact is that the spectrum of G is inherited by that of $\widehat{G}$. Bilu and Linial [3] constructed explicit d-regular expanders by a repeated application of a 2-lift of some base graph (e.g., the clique on $d+1$ vertices). To this end, they proved that every d-regular graph has a 2-lift whose spectrum contains, on top of the eigenvalues of G, only eigenvalues that are bounded, in absolute value, by $O(\sqrt{d \log^3 d})$, thus forming a family of d-regular graphs, one for each size of the form $(d+1)2^k$, $k \in \mathbb{N}$, having spectral expansion $O(\sqrt{d \log^3 d})$. Bilu and Linial further gave an explicit construction based on a suitable derandomization of their existential proof.

We digress a bit and mention that Bilu and Linial conjectured that every d-regular graph has a 2-lift all of whose "new" eigenvalues are bounded, in absolute value, by $2\sqrt{d-1}$. This conjecture has been proved by Marcus, Spielman, and Srivastava [12] for the bipartite case.

2.2   Partial lifts

As mentioned, [5] obtained their result by interpolating between every two consecutive 2-lifts, guaranteeing that every graph between a pair of consecutive 2-lifts is a good edge-expander. Following [5], to prove Theorem 1, we also interpolate between consecutive lifts. However, proving *spectral* expansion require us to use a completely different proof technique. The underlying idea is to carefully weigh the interpolated graphs in a way that will allow us to argue about their *eigenvectors*. We dub these interpolated graphs *partial lifts*, and turn to define them.

▶ **Definition 3** (Partial lifts). *Let $G = (V, E)$ be an undirected simple graph with a k-lift $\widehat{G} = (\widehat{V} = [k] \times V, \widehat{E})$. Let (B, L) be a partition of V. The L-partial lift of G (with respect to $\widehat{G}$) is defined to be the undirected weighted graph $\widehat{G}_L = (\widehat{V}_L, \widehat{E}_L)$ whose vertex set consists of the vertices of B and the fibers of the vertices of L, namely, $\widehat{V}_L = B \cup ([k] \times L)$. The edge set $\widehat{E}_L$ is the union of the edges of three sets:*

1. *The edges of G connecting vertices in B.*
2. *The edges of $\widehat{G}$ connecting vertices in L.*
3. *For every edge $\{u, v\} \in E$ with $u \in B$ and $v \in L$, we add an edge of weight $\frac{1}{\sqrt{k}}$ between u and each of the vertices in the fiber of v. The edges from Items (1) and (2) have weight 1.*

The slightly more formal definition of a partial lift is given in Definition 5. Informally, we think of the vertices in B as vertices of the base graph G that are not yet lifted, and of those in L as the already lifted vertices. The edges from Items (1) and (2) form the corresponding induced graphs. The set of "hybrid" edges, appearing in Item (3), connect already-lifted and not-yet-lifted vertices, where the weight assigned to these edges is chosen in hindsight.

Note that $\widehat{G}_L$ is a simple weighted undirected graph. Moreover, $\widehat{G}_L$ interpolates between G and $\widehat{G}$ in the sense that $\widehat{G}_\emptyset = G$ and $\widehat{G}_V = \widehat{G}$. Already here we mention that for the proof of Theorem 1 we will be working with 4-lifts, or with any whole square number k for that matter, as then the $\frac{1}{\sqrt{k}}$ weight can be "simulated" without weights using parallel edges.

Before we proceed, we set some notation. We denote $b = |B|$, $\ell = |L|$, and further denote the number of vertices of $\widehat{G}_L$ by m, noting that $m = b + k\ell$. The number of vertices in G is denoted $n = b + \ell$. We denote the smallest eigenvalue of the adjacency matrix of an undirected graph H by $\lambda_{\min}(H)$. This will be convenient as the different graphs that we will be considering (G and $\widehat{G}_L$) will be on a different number of vertices.

2.3   Bounding the spectral expansion

Our main result with regards to the bound on the unnormalized spectral expansion is that for every L-partial lift of G (with respect to $\widehat{G}$) it holds that

$$\lambda_{\min}(\widehat{G}) \leq \lambda_{\min}(\widehat{G}_L) \leq \lambda_2(\widehat{G}_L) \leq \lambda_2(\widehat{G}). \tag{2.1}$$

This in particular implies that the spectral expansion of every L-partial lift is as good as the spectral expansion of the (fully) lifted graph, namely, $\lambda(\widehat{G}_L) \leq \lambda(\widehat{G})$. See Proposition 6 for the more complete statement.

To prove Equation (2.1), we consider the subspace $F^{\|}$ of $\mathbb{R}^m$ that consists of all vectors that are constants on the fibers of the lifted vertices. We denote the orthogonal complement of $F^{\|}$ by $F^{\perp}$, noting that it contains all vectors that sum up to zero on the fibers of the lifted vertices, and that vanish on the unlifted vertices. In the first step of the proof we characterize the eigenvectors of $\widehat{G}_L$ that lay inside $F^{\|}$. To do so, we order the vertices of $\widehat{G}_L$ such that the unlifted vertices, those in B, appear first. With this ordering, consider the $m \times n$ matrix

$$\mathbf{U} = \begin{pmatrix} \mathbf{I}_b & \mathbf{0} \\ \mathbf{0} & \frac{1}{\sqrt{k}}\mathbf{I}_\ell \\ \vdots & \vdots \\ \mathbf{0} & \frac{1}{\sqrt{k}}\mathbf{I}_\ell \end{pmatrix}.$$

We prove (see Lemma 8) that $\mathbf{U}\mathbf{x}$ is an eigenvector of $\widehat{G}_L$ if and only if $\mathbf{x}$ is an eigenvector of G. In fact, we weigh the hybrid edges as we did precisely for the purpose of making this statement true. At any rate, this accounts for n eigenvectors of $\widehat{G}_L$, all of which are contained in $F^{\|}$, whose eigenvalues are the same as those of G. In particular, every eigenvalue of G is an eigenvalue of $\widehat{G}_L$ with the same, or with higher multiplicity.

Since $\widehat{G}_L$ is symmetric, its eigenvectors are orthogonal to each other, and so the remaining eigenvectors of $\widehat{G}_L$ are contained in $F^{\perp}$. While we cannot argue that these correspond to eigenvectors of $\widehat{G}$, as one might have hoped, in the second step of the proof we show that these correspond to eigenvectors of some principal submatrix $\mathbf{M}$ of the adjacency matrix of $\widehat{G}$. This suffices for the purpose of bounding the eigenvalues as one can invoke the eigenvalue interlacing theorem (Theorem 4).

We stress that not every eigenvector of $\mathbf{M}$ induces an eigenvector of $\widehat{G}_L$, a fact that is crucial to the proof. Indeed, the crux of the proof is in showing that some "problematic" eigenvectors of $\mathbf{M}$ do not affect the spectrum of $\widehat{G}_L$. Although this is a key part of the proof, we cannot cover it without delving into more details, and so at this point we refer the reader to the formal treatment that is given in Section 4.

2.4 Bounding the normalized spectral expansion

Our main result regarding the normalized spectral expansion, which recall determines the rate of convergence of a random walk, is given by Proposition 11 and essentially states that assuming k is sufficiently small compared to $\lambda(G)$, for all $L \subseteq V$ it holds that

$$\bar{\lambda}(\widehat{G}_L) = O(k \cdot \bar{\lambda}(\widehat{G})). \tag{2.2}$$

We remark that in the normalized case, a stronger statement as in Equation (2.1) does not hold. Namely, $\lambda_2(\widehat{G}_L)$ depends on both $\lambda_2(\widehat{G})$ and $\lambda_{\mathsf{min}}(\widehat{G})$, and similarly for $\lambda_{\mathsf{min}}(\widehat{G}_L)$. Moreover, note that, unlike the unnormalized case, here k affects the bound, though the reader should keep in mind that for our construction of spectral expanding expanders, as given in Theorem 1, we will anyhow set k to 4. Another technical caveat worth mentioning is that we can only prove Equation (2.2) for a regular base graph G (which, again, suffices for the proof of Theorem 1). This is essentially because we need a good handle on an eigenvector of G that corresponds to its largest eigenvalue.

To discuss our proof strategy we introduce some notation. Let $\mathbf{M}_{\widehat{G}_L}$ be the adjacency matrix of $\widehat{G}_L$ and $\mathbf{W}_{\widehat{G}_L}$ be its random walk matrix. More precisely, if we denote by $\mathbf{D}_{\widehat{G}_L}$ the diagonal matrix that encodes the degrees of vertices in $\widehat{G}_L$ then $\mathbf{W}_{\widehat{G}_L} = \mathbf{M}_{\widehat{G}_L} \mathbf{D}_{\widehat{G}_L}^{-1}$. We first note that $\mathbf{z} = \mathbf{D}_{\widehat{G}_L} \mathbf{1}$ is an eigenvector of $\mathbf{W}_{\widehat{G}_L}$ with eigenvalue 1, and so to prove Equation (2.2) it suffices to bound the Rayleigh quotient, with respect to $\mathbf{W}_{\widehat{G}_L}$, of vectors orthogonal to $\mathbf{z}$. As $F^{\parallel}$ and $F^{\perp}$ are invariant subspaces of $\mathbf{W}_{\widehat{G}_L}$, it suffices to do so for each of these subspaces separately. In the first step of the proof, we use the result we already proved for the unnormalized case to deduce that

$$\forall \mathbf{x} \in F^{\perp} \quad \frac{\mathbf{x}^{\mathsf{T}} \mathbf{W}_{\widehat{G}_L} \mathbf{x}}{\mathbf{x}^{\mathsf{T}} \mathbf{x}} \le \sqrt{k} \cdot \bar{\lambda}(\widehat{G}),$$

which allows us to turn our focus to $F^{\parallel}$.

To bound the Rayleigh quotient of vectors in $F^{\parallel}$, we characterize the eigenvectors of $\mathbf{W}_{\widehat{G}_L}$ laying in $F^{\parallel}$ by the eigenvectors of another operator. Formally, if $\mathbf{M}_G$ is the adjacency matrix of the base graph G then, in Lemma 14, we prove that there exists a diagonal $n \times n$ matrix $\mathbf{D}$ such that a vector $\mathbf{x} \in \mathbb{R}^n$ is an eigenvector of $\mathbf{M}_G \mathbf{D}^{-1}$ if and only if $\mathbf{U}\mathbf{x}$ is an eigenvector of $\mathbf{W}_{\widehat{G}_L}$, and both correspond to the same eigenvalue. We stress that the matrix $\mathbf{D}$ is *not* the matrix encoding the degrees of G (as indeed it should somehow encode information about L) but rather it encodes the degree of vertices in the lifted graph, where from every fiber we take only one representative.

The above leaves us with the task of studying the eigenvalues of the matrix $\mathbf{M}_G \mathbf{D}^{-1}$ which, as eluded to above, "skews" the degrees of vertices in G according to the partial lift structure. The crux of the proof, which we will not be able to cover in this high level proof overview, boils down to bounding the sum of reciprocal of these skewed degrees

$$\sum_{v \in V} \frac{1}{\mathbf{D}_{v,v}}$$

(see Lemmas 19 and 20). We refer the reader to Section 5 for the formal treatment.

3 Preliminaries

We start by setting some fairly standard notation from spectral graph theory.

3.1 Spectral graph theory

The adjacency matrix of an undirected graph $G = (V, E)$ is denote by $\mathbf{M}_G$. Being real and symmetric, $\mathbf{M}_G$ has $n = |V|$ real eigenvalues which we denote by $\lambda_1(\mathbf{M}_G) \geq \cdots \geq \lambda_n(\mathbf{M}_G)$. For $i \in [n]$ we define $\lambda_i(G) = \lambda_i(\mathbf{M}_G)$, and write $\lambda_{\mathsf{min}}(G)$ for $\lambda_n(G)$. We refer to the eigenvectors of $\mathbf{M}_G$ as the eigenvectors of G. The *spectral expansion* of G is given by $\lambda(G) \triangleq \max(\lambda_2(G), |\lambda_n(G)|)$.

Let $\mathbf{D}_G$ be the *degrees matrix* of G, that is, the matrix that encodes the degrees of vertices in G (under the same order that they appear in $\mathbf{M}_G$). Assuming G has no isolated vertices, the random walk matrix of G, denoted $\mathbf{W}_G$, is defined by $\mathbf{W}_G = \mathbf{M}_G \mathbf{D}_G^{-1}$. Note that $\mathbf{W}_G$ has n real eigenvalues as it is similar to the symmetric matrix $\mathbf{D}_G^{-\frac{1}{2}} \mathbf{M}_G \mathbf{D}_G^{-\frac{1}{2}}$. We denote these by $\bar{\lambda}_1(G) \geq \cdots \geq \bar{\lambda}_n(G) \triangleq \bar{\lambda}_{\mathsf{min}}(G)$ and refer to them as the *normalized eigenvalues* of G. The *normalized spectral expansion* of G is given by $\bar{\lambda}(G) \triangleq \max(\bar{\lambda}_2(G), |\bar{\lambda}_n(G)|)$.

For a family $\mathcal{G} = (G_n)_{n \in \mathcal{I}}$ of expander graphs, we let $\lambda(\mathcal{G}) = \sup_{n \in \mathcal{I}} \lambda(G_n)$ if the maximum exists, and ∞ otherwise, and similarly define $\bar{\lambda}(\mathcal{G}) = \sup_{n \in \mathcal{I}} \bar{\lambda}(G_n)$.

We make use of the well-known fact that the eigenvalues of a real symmetric matrix interlace with the eigenvalues of any of its principal submatrices. For a proof see, e.g., [8], Theorem 9.1.1.

▶ **Theorem 4** (Eigenvalue Interlacing Theorem). *Let* $\mathbf{N}$ *be a real symmetric* $n \times n$ *matrix and let* $\mathbf{M}$ *be an* $m \times m$ *principal submatrix of* $\mathbf{N}$. *Then, for all* $i \in [m]$,

$$\lambda_i(\mathbf{N}) \geq \lambda_i(\mathbf{M}) \geq \lambda_{n-m+i}(\mathbf{N}).$$

3.2 Lifts

In contrast to the introductory part, from here on we define the notion of a lift in a somewhat more formal way, which is also easier to work with. To this end, we first recall the notion of graph orientation. Let $G = (V, E)$ be a simple undirected graph. An orientation of G is an assignment of a direction to each of its edges, resulting with a directed graph which we denote by $\vec{G} = (V, \vec{E})$. That is, for every undirected edge $\{u, v\}$ of G exactly one of (u, v), (v, u) is included in $\vec{E}$.

In what comes next, we consider maps $\pi : \vec{E} \to S_k$ where $\vec{E}$ is the edge (multi-)set of some oriented graph $\vec{G}$ and, as customary, S_k is the group of permutations on $[k]$. For ease of notation, we write $\pi_{u,v}$ instead of the more cumbersome expression $\pi((u, v))$.

Let $G = (V, E)$ be an undirected simple graph on n vertices, $\vec{G} = (V, \vec{E})$ an orientation of G, and let $\pi : \vec{E} \to S_k$ for some integer $k \geq 1$. The π-*lift* of $\vec{G}$ is the graph $\vec{G}_\pi = ([k] \times V, E_\pi)$ where for every $(u, v) \in \vec{E}$ we include the edges

$$\{(i, u), (\pi_{u,v}(i), v)\} \quad \text{for} \quad i = 1, 2, \ldots, k$$

in E_π. Note that regardless of the choice of orientation (and regardless of the choice of π), since G is simple so is $\vec{G}_\pi$. For $v \in V$, the set of vertices $[k] \times \{v\}$ of $\vec{G}_\pi$ is called the *fiber* of v. For ease of notation, from hereon we write G_π for $\vec{G}_\pi$ despite the fact that this graph depends on the chosen orientation.

We extend the map $\pi : \vec{E} \to S_k$ to the set $\{(v,u) \mid (u,v) \in \vec{E}\}$ as follows: For $(u,v) \in \vec{E}$ we set $\pi_{v,u} = \pi_{u,v}^{-1}$. With this, it is convenient to write down the adjacency matrix of G_π as follows. For $i,j \in [k]$ define the zero-one $n \times n$ matrix $\mathbf{M}_G^{i,j}$ by

$$(\mathbf{M}_G^{i,j})_{u,v} = 1 \quad \Longleftrightarrow \quad \{u,v\} \in E \text{ and } \pi_{u,v}(i) = j.$$

Then, the adjacency matrix $\mathbf{M}_{G_\pi}$ of G_π is the $k \times k$ block matrix, where block (i,j) is given by the $n \times n$ matrix $(\mathbf{M}_{G_\pi})_{i,j} = \mathbf{M}_G^{i,j}$. That is,

$$\mathbf{M}_{G_\pi} = \begin{pmatrix} \mathbf{M}_G^{1,1} & \mathbf{M}_G^{1,2} & \cdots & \mathbf{M}_G^{1,k} \\ \mathbf{M}_G^{2,1} & \mathbf{M}_G^{2,2} & \cdots & \mathbf{M}_G^{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{M}_G^{k,1} & \mathbf{M}_G^{k,2} & \cdots & \mathbf{M}_G^{k,k} \end{pmatrix}.$$

Note that $\mathbf{M}_G^{i,j} = (\mathbf{M}_G^{j,i})^\mathsf{T}$. A well-known fact about lifting is that the spectrum of the base graph G is contained in the spectrum of the lifted graph G_π. More precisely, if λ is an eigenvalue of G with multiplicity r then λ is an eigenvalue of G_π with multiplicity at least r. This is easily seen by noting that for every $i \in [k]$,

$$\sum_{j=1}^k \mathbf{M}_G^{i,j} = \mathbf{M}_G, \tag{3.1}$$

and so every eigenvector $\mathbf{x}$ of $\mathbf{M}_G$ induces the eigenvector $(\mathbf{x}, \ldots, \mathbf{x})$ of $\mathbf{M}_{G_\pi}$ with the same eigenvalue.

4 Bounding the Spectral Expansion

We start this section by introducing the notion of a *partial lift* and study its properties. Throughout, we make use of the notation from Section 3.2.

▶ **Definition 5.** Let $G = (V,E)$ be an undirected simple graph, and let $\vec{G} = (V, \vec{E})$ be an orientation of G. Let $\pi : \vec{E} \to S_k$ for some $k \geq 1$, and let (B,L) be a partition of the vertices of G. The *L-partial π-lift* of $\vec{G}$ is defined to be the undirected weighted graph $G_{\pi,L} = (V_{\pi,L}, E_{\pi,L})$ whose vertex set is $V_{\pi,L} = B \cup ([k] \times L)$. The edge set $E_{\pi,L}$ is the union of

$$E_B = \{\{u,v\} \in E \mid u,v \in B\},$$
$$E_L = \{\{(i,u),(j,v)\} \in E_\pi \mid u,v \in L\},$$

and

$$E_H = \{\{u,(i,v)\} \mid i \in [k] \text{ and } u \in B, v \in L \text{ s.t. } \{u,v\} \in E\},$$

with weight of $\frac{1}{\sqrt{k}}$ assigned to each edge in E_H. The edges in E_B, E_L have a unit weight assigned to them.

Note that E_B is the set of edges of the B-induced sub-graph of the base graph G, and E_L is the set of edges of the induced graph of G_π with respect to the fibers of the lifted vertices. The set $E_{\pi,L}$ contains the "hybrid" edges, connecting already-lifted and not-yet-lifted vertices where the weight assigned to these edges is chosen with a hindsight.

Note that $G_{\pi,L}$ is a weighted undirected simple graph. Moreover, $G_{\pi,L}$ interpolates between G and G_π in the sense that $G_{\pi,\emptyset} = G$ and $G_{\pi,V} = G_\pi$. Observe that, assuming G is d-regular, the weighted degree of every vertex in $G_{\pi,L}$ is in the range $[d/\sqrt{k}, \sqrt{k} \cdot d]$.

The main result of this section is the following proposition, which is the more complete and formal version of Equation (2.1) from the Proof Overview section.

▶ **Proposition 6.** Let $G = (V, E)$ be an undirected simple graph with orientation $\vec{G} = (V, \vec{E})$. Let $\pi : \vec{E} \to S_k$, and (B, L) a partition of V. Then,

$$\lambda_{\min}(G_\pi) \leq \lambda_{\min}(G_{\pi,L}) \leq \lambda_2(G_{\pi,L}) \leq \lambda_2(G_\pi) \leq \lambda_1(G_\pi) = \lambda_1(G_{\pi,L}).$$

Proof. Note that the non-trivial inequalities and equality, which we set to prove, are

$$\lambda_{\min}(G_\pi) \leq \lambda_{\min}(G_{\pi,L}), \tag{4.1}$$

$$\lambda_2(G_{\pi,L}) \leq \lambda_2(G_\pi), \tag{4.2}$$

$$\lambda_1(G_\pi) = \lambda_1(G_{\pi,L}). \tag{4.3}$$

Let $\mathbf{M}_G$ be the adjacency matrix of G where we order the rows and columns so that those corresponding to vertices in B appear first, namely,

$$\mathbf{M}_G = \begin{pmatrix} \mathbf{M}_B & \mathbf{M}_H \\ \mathbf{M}_H^{\mathsf{T}} & \mathbf{M}_L \end{pmatrix}.$$

Note that $\mathbf{M}_B, \mathbf{M}_L$ are the adjacency matrices of the B-induced and L-induced subgraphs of G, respectively. Observe that the adjacency matrix of $G_{\pi,L}$ is given by

$$\mathbf{M}_{G_{\pi,L}} = \begin{pmatrix} \mathbf{M}_B & \frac{1}{\sqrt{k}}\mathbf{M}_H & \cdots & \frac{1}{\sqrt{k}}\mathbf{M}_H \\ \frac{1}{\sqrt{k}}\mathbf{M}_H^{\mathsf{T}} & \mathbf{M}_L^{1,1} & \cdots & \mathbf{M}_L^{1,k} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{\sqrt{k}}\mathbf{M}_H^{\mathsf{T}} & \mathbf{M}_L^{k,1} & \cdots & \mathbf{M}_L^{k,k} \end{pmatrix},$$

where $\mathbf{M}_L^{i,j}$ is a slight abuse of notation (recall we defined $\mathbf{M}_H^{i,j}$ for a *graph* H) and is used as a shorthand for $\mathbf{M}_H^{i,j}$, H being the L-induced subgraph of G. Similar to Equation (3.1), we have

$$\forall i \in [k] \quad \sum_{j=1}^k \mathbf{M}_L^{i,j} = \mathbf{M}_L. \tag{4.4}$$

Let n, m be the number of vertices in G and in $G_{\pi,L}$, respectively. Denote $b = |B|, \ell = |L|$, and note that $n = b + \ell$ and $m = b + k\ell$. Define

$$F^{\|} = \left\{ \mathbf{x} \in \mathbb{R}^m \mid \mathbf{x}_{b+j} = \mathbf{x}_{b+\ell+j} = \cdots = \mathbf{x}_{b+(k-1)\ell+j} \quad \text{for } j = 1, 2, \ldots, \ell \right\}.$$

Informally, $F^{\|}$ is the space of vectors that are constant on the fibers of the lifted vertices, and are otherwise arbitrary. Let $F^{\perp}$ be the dual subspace of $F^{\|}$, namely,

$$F^{\perp} = \left\{ \mathbf{x} \in \mathbb{R}^m \mid \mathbf{x}_1 = \cdots = \mathbf{x}_b = 0 \text{ and } \sum_{i=0}^{k-1} \mathbf{x}_{b+i\ell+j} = 0 \quad \text{for } j = 1, 2, \ldots, \ell \right\}.$$

It is easy to verify, using Equation (4.4), that both $F^{\|}$ and $F^{\perp}$ are invariant subspaces of $\mathbf{M}_{G_{\pi,L}}$. Define the matrix $\mathbf{U} \in \mathbb{R}^{m \times n}$

$$\mathbf{U} = \begin{pmatrix} \mathbf{I}_b & \mathbf{0} \\ \mathbf{0} & \frac{1}{\sqrt{k}}\mathbf{I}_\ell \\ \vdots & \vdots \\ \mathbf{0} & \frac{1}{\sqrt{k}}\mathbf{I}_\ell \end{pmatrix}. \tag{4.5}$$

The following claim lists some useful, easy to prove, properties of $\mathbf{U}$.

$\triangleright$ **Claim 7.** $\mathbf{U}$ satisfies the following properties:

1. $\mathrm{Im}(\mathbf{U}) = F^{\|}$.
2. The right kernel of $\mathbf{U}$ is $\mathbf{0}$.
3. $\mathbf{U}^{\mathsf{T}}\mathbf{U} = \mathbf{I}_n$.
4. $\mathbf{U}\mathbf{U}^{\mathsf{T}}$ is the orthogonal projection to $F^{\|}$.
5. $\mathbf{U}^{\mathsf{T}}\mathbf{M}_{G_{\pi,L}}\mathbf{U} = \mathbf{M}_G$.

We can now easily characterize all the eigenvectors of $G_{\pi,L}$ laying in $F^{\|}$.

▶ **Lemma 8.** *For every* $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{x}$ *is an eigenvector of* G *corresponding to an eigenvalue* λ *if and only if* $\mathbf{U}\mathbf{x}$ *is an eigenvector of* $G_{\pi,L}$ *laying in* $F^{\|}$ *and corresponding to* λ.

Proof.

$$\mathbf{M}_G \mathbf{x} = \lambda \mathbf{x} \quad \Longleftrightarrow \quad \mathbf{U}^{\mathsf{T}}\mathbf{M}_{G_{\pi,L}}\mathbf{U}\mathbf{x} = \lambda \mathbf{x}$$
$$\Longleftrightarrow \quad \mathbf{U}\mathbf{U}^{\mathsf{T}}\mathbf{M}_{G_{\pi,L}}\mathbf{U}\mathbf{x} = \lambda \mathbf{U}\mathbf{x}$$
$$\Longleftrightarrow \quad \mathbf{M}_{G_{\pi,L}}\mathbf{U}\mathbf{x} = \lambda \mathbf{U}\mathbf{x},$$

where the first implication follows from Item (5) of Claim 7, the second from Item (2), and the last follows by Items (1) and (4) together with the invariance of $F^{\|}$ under $\mathbf{M}_{G_{\pi,L}}$. ◀

To summarize, we have found n eigenvectors of $\mathbf{M}_{G_{\pi,L}}$, all of which are contained in $F^{\|}$, whose eigenvalues correspond to those of $\mathbf{M}_G$. In particular, every eigenvalue of $\mathbf{M}_G$ is an eigenvalue of $\mathbf{M}_{G_{\pi,L}}$ with the same, or higher, multiplicity. Observe that $F^{\|}$ is defined by $(k-1)\ell$ linear constraints. Hence, its dimension is exactly $n = b + k\ell - (k-1)\ell$. We conclude that the characterized eigenvectors are exactly all the eigenvectors of $G_{\pi,L}$ in $F^{\|}$.

We proceed to explore the remaining eigenvectors of $\mathbf{M}_{G_{\pi,L}}$. Since $\mathbf{M}_{G_{\pi,L}}$ is symmetric, its eigenvectors are orthogonal to each other, and so the remaining eigenvectors of $\mathbf{M}_{G_{\pi,L}}$ are contained in $F^{\perp}$. We turn to prove that while we cannot argue that these correspond to eigenvectors of $\mathbf{M}_{G_\pi}$, as one might hope, they will correspond to eigenvectors of some principal submatrix of $\mathbf{M}_{G_\pi}$, at which point we can invoke the Eigenvalue Interlacing Theorem (see Theorem 4) so to bound the corresponding eigenvalues.

Take $\mathbf{x} \in F^{\perp}$ an eigenvector of $\mathbf{M}_{G_{\pi,L}}$ with eigenvalue λ. Then, $\mathbf{x} = (\mathbf{0}, \mathbf{z})$ for some nonzero $\mathbf{z} = (\mathbf{z}_1, \ldots, \mathbf{z}_k) \in \mathbb{R}^{k\ell}$ where $\sum_{i=1}^{k} \mathbf{z}_i = \mathbf{0} \in \mathbb{R}^{\ell}$. Therefore,

$$\lambda \mathbf{x} = \mathbf{M}_{G_{\pi,L}}\mathbf{x} = \begin{pmatrix} \mathbf{0} \\ \mathbf{M}_{\pi,L}\mathbf{z} \end{pmatrix},$$

where

$$\mathbf{M}_{\pi,L} = \begin{pmatrix} \mathbf{M}_L^{1,1} & \cdots & \mathbf{M}_L^{1,k} \\ \vdots & \ddots & \vdots \\ \mathbf{M}_L^{k,1} & \cdots & \mathbf{M}_L^{k,k} \end{pmatrix}.$$

Hence, $\mathbf{z}$ is an eigenvector of $\mathbf{M}_{\pi,L}$. That is, all eigenvectors of $\mathbf{M}_{G_{\pi,L}}$ that are contained in $F^{\perp}$ correspond to eigenvectors of $\mathbf{M}_{\pi,L}$.

We stress that not every eigenvector of $\mathbf{M}_{\pi,L}$ induces an eigenvector of $\mathbf{M}_{G_{\pi,L}}$, a fact that will be crucial in what follows. Indeed, the eigenvectors of $\mathbf{M}_{G_{\pi,L}}$ coming from $F^{\perp}$ have the special structure described above of being orthogonal to $\mathbf{1}$ on each fiber. This can also be seen by a dimension argument, noting that $\mathbf{M}_{\pi,L}$ has $k\ell$ eigenvectors, and together with the n eigenvectors that are induced from $\mathbf{M}_G$ these amount to $k\ell + n$ eigenvectors. However, $\mathbf{M}_{G_{\pi,L}}$ is a matrix of order $(k\ell + b) \times (k\ell + b)$, and so $n - b = \ell$ eigenvectors of $\mathbf{M}_{\pi,L}$ do not induce eigenvectors of $\mathbf{M}_{G_{\pi,L}}$. At any rate, for convenience, we summarize the analysis so far.

$\triangleright$ **Claim 9.** The spectrum of $\mathbf{M}_{G_{\pi,L}}$ consists of the spectrum of $\mathbf{M}_G$ with corresponding eigenvectors in $F^{\parallel}$ as well as of a subset of the spectrum of $\mathbf{M}_{\pi,L}$ with corresponding eigenvectors in $F^{\perp}$, where we remind the reader that we consider the spectrum as a multi-set so to track multiplicities correctly.

Note that $\mathbf{M}_{\pi,L}$ is a principal submatrix of $\mathbf{M}_{G_\pi}$ and so, by the Eigenvalue Interlacing Theorem (Theorem 4),

$$\lambda_{\mathsf{min}}(G_\pi) \leq \lambda_{\mathsf{min}}(\mathbf{M}_{\pi,L}) \leq \lambda_2(\mathbf{M}_{\pi,L}) \leq \lambda_2(G_\pi). \tag{4.6}$$

Further, recall that the spectrum of G_π contains that of G, in particular,

$$\lambda_{\mathsf{min}}(G_\pi) \leq \lambda_{\mathsf{min}}(G) \leq \lambda_2(G) \leq \lambda_2(G_\pi). \tag{4.7}$$

By putting together Claim 9, Equation (4.6) and Equation (4.7), we establish Equation (4.1). For proving Equation (4.2) we are left to prove that an eigenvector $\mathbf{x} \in F^{\perp}$ of $G_{\pi,L}$ cannot correspond to an eigenvalue $\lambda > \lambda_2(\mathbf{M}_{\pi,L})$. To this end, let $\mathbf{x} = (\mathbf{0}, \mathbf{z}) \in F^{\perp}$ be an eigenvector of $G_{\pi,L}$. Recall that for every $j \in [\ell]$,

$$\sum_{i=0}^{k-1} \mathbf{z}_{b+i\ell+j} = 0. \tag{4.8}$$

Assume for contradiction that $\mathbf{x}$ is an eigenvector of $G_{\pi,L}$ corresponding to an eigenvalue $\lambda > \lambda_2(\mathbf{M}_{\pi,L})$. Then, by the above discussion, $\mathbf{z}$ is an eigenvector of $\mathbf{M}_{\pi,L}$ with eigenvalue $\lambda > \lambda_2(\mathbf{M}_{\pi,L})$, meaning $\lambda = \lambda_1(\mathbf{M}_{\pi,L})$. As the vector corresponding the largest eigenvalue, $\mathbf{z}$ maximizes the Rayleigh quotient, we have that

$$\lambda_1(\mathbf{M}_{\pi,L}) = \frac{\mathbf{z}^{\mathsf{T}}\mathbf{M}_{\pi,L}\mathbf{z}}{\mathbf{z}^{\mathsf{T}}\mathbf{z}} = \max_{\mathbf{w} \neq \mathbf{0}} \frac{\mathbf{w}^{\mathsf{T}}\mathbf{M}_{\pi,L}\mathbf{w}}{\mathbf{w}^{\mathsf{T}}\mathbf{w}}.$$

Note, however, that the vector $|\mathbf{z}|$, which is obtained by taking the absolute value of every entry of $\mathbf{z}$, satisfies

$$\frac{|\mathbf{z}|^{\mathsf{T}}\mathbf{M}_{\pi,L}|\mathbf{z}|}{|\mathbf{z}|^{\mathsf{T}}|\mathbf{z}|} \geq \frac{\mathbf{z}^{\mathsf{T}}\mathbf{M}_{\pi,L}\mathbf{z}}{\mathbf{z}^{\mathsf{T}}\mathbf{z}}, \tag{4.9}$$

and so, as a maximizer of the Rayleigh quotient, $|\mathbf{z}|$ is also an eigenvector with eigenvalue $\lambda_1(\mathbf{M}_{\pi,L})$. However, by Equation (4.8), $\mathbf{z}$ and $|\mathbf{z}|$ are linearly independent. Indeed, there is a fiber on which $\mathbf{z}$ attains both a positive and negative entries. Hence, we have found two independent vectors corresponding to $\lambda_1(\mathbf{M}_{\pi,L})$, implying $\lambda_1(\mathbf{M}_{\pi,L}) = \lambda_2(\mathbf{M}_{\pi,L})$. This stands in contradiction to $\lambda_1(\mathbf{M}_{\pi,L}) = \lambda > \lambda_2(\mathbf{M}_{\pi,L})$. Putting this result together with Claim 9, Equation (4.6) and Equation (4.7) completes the proof of Equation (4.2). To prove Equation (4.3), we will use a general result on graph lifts.

▶ **Lemma 10.** *Let $G = (V, E)$ be an undirected simple graph with orientation $\vec{G} = (V, \vec{E})$. Let $\pi : \vec{E} \to S_k$. Then, $\lambda_1(G) = \lambda_1(G_\pi)$.*

Proof. By invoking Lemma 8 to $G_\pi = G_{\pi,V}$, we get that $\lambda_1(G)$ is an eigenvalue of G_π, which implies $\lambda_1(G) \leq \lambda_1(G_\pi)$. Proving $\lambda_1(G_\pi) \leq \lambda_1(G)$ will thus finish the proof. For any vector $\mathbf{x}$ on the vertices of G_π, take the vector $\mathbf{y}(\mathbf{x}) = \mathbf{y}$ on the vertices of G defined by $\mathbf{y}_v = \sqrt{\sum_{i=1}^{k} \mathbf{x}_{i,v}^2}$. Now, recall the definition of $|\mathbf{x}|$ and note that

$$\mathbf{y}^{\mathsf{T}}\mathbf{y} = \sum_{v \in V} \mathbf{y}_v^2 = \sum_{v \in V} \sum_{i=1}^{k} \mathbf{x}_{i,v}^2 = \mathbf{x}^{\mathsf{T}}\mathbf{x} = |\mathbf{x}|^{\mathsf{T}}|\mathbf{x}|.$$

Therefore,

$$
\begin{aligned}
\frac{\mathbf{y}^{\mathsf{T}}\mathbf{M}_G\mathbf{y}}{\mathbf{y}^{\mathsf{T}}\mathbf{y}} &= \frac{1}{\mathbf{y}^{\mathsf{T}}\mathbf{y}}\sum_{(u,v)\in\vec{E}}\mathbf{y}_u\mathbf{y}_v \\
&= \frac{1}{\mathbf{x}^{\mathsf{T}}\mathbf{x}}\sum_{(u,v)\in\vec{E}}\sqrt{\sum_{i=1}^{k}\mathbf{x}_{i,u}^2\cdot\sum_{i=1}^{k}\mathbf{x}_{i,v}^2} \\
&= \frac{1}{|\mathbf{x}|^{\mathsf{T}}|\mathbf{x}|}\sum_{(u,v)\in\vec{E}}\sqrt{\sum_{i=1}^{k}|\mathbf{x}_{i,u}|^2\cdot\sum_{i=1}^{k}|\mathbf{x}_{\pi_{u,v}(i),v}|^2} \\
&\geq \frac{1}{|\mathbf{x}|^{\mathsf{T}}|\mathbf{x}|}\sum_{(u,v)\in\vec{E}}\sum_{i=1}^{k}|\mathbf{x}_{i,u}||\mathbf{x}_{\pi_{u,v}(i),v}| \\
&= \frac{|\mathbf{x}|^{\mathsf{T}}\mathbf{M}_{G_\pi}|\mathbf{x}|}{|\mathbf{x}|^{\mathsf{T}}|\mathbf{x}|} \\
&\geq \frac{\mathbf{x}^{\mathsf{T}}\mathbf{M}_{G_\pi}\mathbf{x}}{\mathbf{x}^{\mathsf{T}}\mathbf{x}},
\end{aligned}
$$

where the third equality is just a reordering of the elements in the summation, as $\pi_{u,v}$ is a permutation on $[k]$. The first inequality follows by the Cauchy-Schwarz inequality, and the second inequality is as in Equation (4.9). Thus,

$$
\lambda_1(G_\pi) = \max_{\mathbf{x}\neq 0}\frac{\mathbf{x}\mathbf{M}_{G_\pi}\mathbf{x}}{\mathbf{x}^{\mathsf{T}}\mathbf{x}} \leq \max_{\mathbf{x}\neq 0}\frac{\mathbf{y}(\mathbf{x})\mathbf{M}_G\mathbf{y}(\mathbf{x})}{\mathbf{y}(\mathbf{x})^{\mathsf{T}}\mathbf{y}(\mathbf{x})} \leq \max_{\mathbf{z}\neq 0}\frac{\mathbf{z}\mathbf{M}_G\mathbf{z}}{\mathbf{z}^{\mathsf{T}}\mathbf{z}} = \lambda_1(G). \qquad \blacktriangleleft
$$

Combining the results we proved so far, we conclude that

$$
\lambda_1(G) \leq \lambda_1(G_{\pi,L}) \leq \lambda_1(G_\pi) = \lambda_1(G).
$$

Indeed, the first inequality follows by Lemma 8, the second follows by Claim 9 and by the fact that $\mathbf{M}_{\pi,L}$ is a principal submatrix of $\mathbf{M}_{G_{\pi,L}}$, together with the Eigenvalue Interlacing Theorem (Theorem 4). Lastly, the equality follows by Lemma 10, completing the proof of Proposition 6. $\qquad\blacktriangleleft$

5 Bounding the Normalized Spectral Expansion

In this section we bound the normalized eigenvalues of a partial lift. That is, we show that a random walk on a partial lift converges quickly given that the random walk on the (full) lift does so. Unlike the unnormalized case, we restrict ourselves to d regular base graphs.

▶ **Proposition 11.** *Let $G = (V, E)$ be an undirected simple d-regular graph having orientation $\vec{G} = (V, \vec{E})$. Let $\pi : \vec{E} \to S_k$, and (B, L) a partition of V. Assume that $\sqrt{k} \leq \frac{\lambda(G)}{3} + 1$. Then,*

$$
\bar{\lambda}(G_{\pi,L}) = O(k \cdot \bar{\lambda}(G_\pi)).
$$

Proof. Note that the degrees matrix of $G_{\pi,L}$, $\mathbf{D}_{G_{\pi,L}}$, is constant on every fiber, and so $F^{\|}$ and $F^{\perp}$ are invariant subspaces of $\mathbf{D}_{G_{\pi,L}}^{-1}$. As noted in the proof of Proposition 6, these are also invariant subspaces of $\mathbf{M}_{G_{\pi,L}}$, hence, also of $\mathbf{W}_{G_{\pi,L}} = \mathbf{M}_{G_{\pi,L}}\mathbf{D}_{G_{\pi,L}}^{-1}$.

Equation (4.2) and Equation (4.1) implies that $\lambda(G_{\pi,L}) \leq \lambda(G_\pi)$. This, together with the fact that the eigenvector corresponding to the largest eigenvalue of $G_{\pi,L}$ lays in $F^\parallel$, implies that

$$\mathbf{x}, \mathbf{y} \in F^\perp \quad \left|\mathbf{x}^\mathsf{T}\mathbf{M}_{G_{\pi,L}}\mathbf{y}\right| \leq \lambda(G_\pi)\|\mathbf{x}\|_2\|\mathbf{y}\|_2.$$

This, together with Item (4) of Claim 13, which we state below, yields the desired bound on the Rayleigh quotient for all vectors in $F^\perp$. Indeed, for every $\mathbf{x} \in F^\perp$, we have that

$$\frac{|\mathbf{x}^\mathsf{T}\mathbf{W}_{G_{\pi,L}}\mathbf{x}|}{\mathbf{x}^\mathsf{T}\mathbf{x}} = \frac{|\mathbf{x}^\mathsf{T}\mathbf{M}_{G_{\pi,L}}\mathbf{D}_{G_{\pi,L}}^{-1}\mathbf{x}|}{\mathbf{x}^\mathsf{T}\mathbf{x}} \leq \frac{\lambda(G_\pi)\cdot\|\mathbf{x}\|_2\cdot\|\mathbf{D}_{G_{\pi,L}}^{-1}\mathbf{x}\|_2}{\mathbf{x}^\mathsf{T}\mathbf{x}} \leq \frac{\sqrt{k}}{d}\cdot\lambda(G_\pi) = \sqrt{k}\cdot\bar{\lambda}(G_\pi).$$

We summarize this in the following corollary.

▶ **Corollary 12.** *The Rayleigh quotient of all eigenvectors laying in $F^\perp$, with respect to $\mathbf{W}_{G_{\pi,L}}$, are bounded by $\sqrt{k} \cdot \bar{\lambda}(G_\pi)$.*

Since $F^\parallel$ and $F^\perp$ are invariant subspaces of $\mathbf{W}_{G_{\pi,L}}$, we are left to analyze the vectors laying in $F^\parallel$. To this end, define the diagonal matrix $\mathbf{D} \in \mathbb{R}^{n\times n}$ by $\mathbf{D} = \mathbf{U}^\mathsf{T}\mathbf{D}_{G_{\pi,L}}\mathbf{U}$, where we recall the reader that the definition of $\mathbf{U}$ is in Equation (4.5). By this definition, since the degrees of vertices on the same fiber are equal, we get

$$\mathbf{D}_{u,u} = \begin{cases} \deg_{G_{\pi,L}}(u) & u \in B; \\ \deg_{G_{\pi,L}}((1,u)) & u \in L. \end{cases} \tag{5.1}$$

For a vertex v of G we define θ_v, the *cut degree* of v, to be

$$\theta_v = \begin{cases} |E_G(v,L)| & v \in B; \\ |E_G(v,B)| & v \in L, \end{cases} \tag{5.2}$$

where, for $S, T \subseteq V$, $|E_G(S,T)|$ is the sum of weights of edges between (S,T) in G. Denote $k_B = \sqrt{k} - 1$ and $k_L = \frac{1}{\sqrt{k}} - 1$. The following claim is easy to verify and is stated without a proof.

▷ **Claim 13.**   The matrices $\mathbf{D}_{G_{\pi,L}}$ and $\mathbf{D}$ have the following properties:
1. $\mathbf{D}_{G_{\pi,L}}\mathbf{U} = \mathbf{U}\mathbf{D}$ and $\mathbf{D}_{G_{\pi,L}}^{-1}\mathbf{U} = \mathbf{U}\mathbf{D}^{-1}$.
2. For a vertex $u \in B$ the value of $\mathbf{D}_{u,u}$ is given by

$$1 \cdot |E_G(u,B)| + \sqrt{k} \cdot |E_G(u,L)| = d + k_B\theta_u.$$

3. For a vertex $u \in L$ the value of $\mathbf{D}_{u,u}$ is given by

$$1 \cdot |E_G(u,L)| + \frac{1}{\sqrt{k}} \cdot |E_G(u,B)| = d + k_L\theta_u.$$

4. $(\mathbf{D}_{G_{\pi,L}})_{u,u} \in [\frac{1}{\sqrt{k}}d, \sqrt{k}d]$.
5. $\mathbf{D}_{u,u} \in [\frac{1}{\sqrt{k}}d, \sqrt{k}d]$.

▶ **Lemma 14.** *A vector $\mathbf{x} \in \mathbb{R}^n$ is an eigenvector of $\mathbf{M}_G\mathbf{D}^{-1}$ corresponding to eigenvalue λ if and only if $\mathbf{U}\mathbf{x}$ is an eigenvector of $\mathbf{W}_{G_{\pi,L}}$ corresponding to the same eigenvalue.*

Proof. By Items (5) and (2) of Claim 7,

$$\mathbf{M}_G \mathbf{D}^{-1} \mathbf{x} = \lambda \mathbf{x} \quad \Longleftrightarrow \quad \mathbf{U}^\mathsf{T} \mathbf{M}_{G_{\pi,L}} \mathbf{U} \mathbf{D}^{-1} \mathbf{x} = \lambda \mathbf{x}$$
$$\Longleftrightarrow \quad \mathbf{U} \mathbf{U}^\mathsf{T} \mathbf{M}_{G_{\pi,L}} \mathbf{U} \mathbf{D}^{-1} \mathbf{x} = \lambda \mathbf{U} \mathbf{x}.$$

Thus, by Item (1) of Claim 13,

$$\mathbf{M}_G \mathbf{D}^{-1} \mathbf{x} = \lambda \mathbf{x} \quad \Longleftrightarrow \quad \mathbf{U} \mathbf{U}^\mathsf{T} \mathbf{W}_{G_{\pi,L}} \mathbf{U} \mathbf{x} = \lambda \mathbf{U} \mathbf{x}.$$

By Item (1) of Claim 7, $\mathbf{U}\mathbf{x} \in F^\parallel$ and since $F^\parallel$ is an invariant subspace of $\mathbf{W}_{G_{\pi,L}}$, we get that $\mathbf{W}_{G_{\pi,L}} \mathbf{U}\mathbf{x} \in F^\parallel$. Item (4) of Claim 7 then implies that

$$\mathbf{U} \mathbf{U}^\mathsf{T} \mathbf{W}_{G_{\pi,L}} \mathbf{U} \mathbf{x} = \mathbf{W}_{G_{\pi,L}} \mathbf{U} \mathbf{x},$$

and so

$$\mathbf{M}_G \mathbf{D}^{-1} \mathbf{x} = \lambda \mathbf{x} \quad \Longleftrightarrow \quad \mathbf{W}_{G_{\pi,L}} \mathbf{U} \mathbf{x} = \lambda \mathbf{U} \mathbf{x},$$

as desired. ◀

Given Corollary 12 and Lemma 14, and since every vector in $F^\parallel$ is of the form $\mathbf{U}\mathbf{x}$ for some $\mathbf{x} \in \mathbb{R}^n$ (recall Item (1) of Claim 7), we can turn our focus to characterizing the eigenvalues of $\mathbf{M}_G \mathbf{D}^{-1}$. We do so by analyzing the eigenvalues of the symmetric matrix $\widetilde{\mathbf{M}}_G = \mathbf{D}^{-\frac{1}{2}} \mathbf{M}_G \mathbf{D}^{-\frac{1}{2}}$ as, note, it is similar to $\mathbf{M}_G \mathbf{D}^{-1}$. To start with, observe that $\mathbf{D}_{G_{\pi,L}} \mathbf{1}_m$ is an eigenvector of $\mathbf{W}_{G_{\pi,L}}$ corresponding to its largest eigenvalue, 1. Thus, Lemma 14 together with Item (4) of Claim 7 imply that $\mathbf{U}^\mathsf{T} \mathbf{D}_{G_{\pi,L}} \mathbf{1}$ is an eigenvector of $\mathbf{M}_G \mathbf{D}^{-1}$ corresponding to the eigenvalue 1. By Item (1) of Claim 13, $\mathbf{U}^\mathsf{T} \mathbf{D}_{G_{\pi,L}} = \mathbf{D} \mathbf{U}^\mathsf{T}$, and so the latter eigenvector can be written as $\mathbf{D} \mathbf{U}^\mathsf{T} \mathbf{1}$. As $\mathbf{M}_G \mathbf{D}^{-1}$ is similar to $\widetilde{\mathbf{M}}_G$, by the above, the latter has an eigenvector $\widetilde{\mathbf{x}}_1 = \mathbf{D}^{\frac{1}{2}} \mathbf{U}^\mathsf{T} \mathbf{1}$ corresponding to its largest eigenvalue, 1. Since $\widetilde{\mathbf{M}}_G$ is symmetric, an eigenvector corresponding to its second largest eigenvalue has to be orthogonal to $\widetilde{\mathbf{x}}_1$.

As we assume G is regular, $\mathbf{1}_n$ is an eigenvector corresponding to $\mathbf{M}_G$-s largest eigenvalue. For a vector $\mathbf{y}$, let $\mathbf{y}^\parallel = \langle \mathbf{1}_n, \mathbf{y} \rangle \cdot \mathbf{1}_n$ be the orthogonal projection of $\mathbf{y}$ onto $\mathbf{1}_n$. We write $\mathbf{y}^\perp = \mathbf{y} - \mathbf{y}^\parallel$ and turn to analyze the Rayleigh quotient of a vector $\mathbf{y} \perp \widetilde{\mathbf{x}}_1$. Let $\mathbf{z} = \mathbf{D}^{-\frac{1}{2}} \mathbf{y}$ and note that $\mathbf{y} \perp \widetilde{\mathbf{x}}_1$ if and only if $\mathbf{z} \perp \mathbf{D}^{\frac{1}{2}} \widetilde{\mathbf{x}}_1$. With this, we have that

$$\frac{\mathbf{y}^\mathsf{T} \widetilde{\mathbf{M}}_G \mathbf{y}}{\mathbf{y}^\mathsf{T} \mathbf{y}} = \frac{\mathbf{z}^\mathsf{T} \mathbf{M}_G \mathbf{z}}{\mathbf{z}^\mathsf{T} \mathbf{D} \mathbf{z}} = \frac{(\mathbf{z}^\parallel)^\mathsf{T} \mathbf{M}_G \mathbf{z}^\parallel}{\mathbf{z}^\mathsf{T} \mathbf{D} \mathbf{z}} + \frac{(\mathbf{z}^\perp)^\mathsf{T} \mathbf{M}_G \mathbf{z}^\perp}{\mathbf{z}^\mathsf{T} \mathbf{D} \mathbf{z}}. \tag{5.3}$$

Using Item (5) of Claim 13, it is clear that $\frac{d}{\sqrt{k}} \|\mathbf{z}\|^2 \leq \mathbf{z}^\mathsf{T} \mathbf{D} \mathbf{z}$. We can thus get a good bound on the size of the second summand in the right hand side of Equation (5.3), namely,

$$\left| \frac{(\mathbf{z}^\perp)^\mathsf{T} \mathbf{M}_G \mathbf{z}^\perp}{\mathbf{z}^\mathsf{T} \mathbf{D} \mathbf{z}} \right| \leq \frac{\lambda(G) \|\mathbf{z}\|^2}{\frac{d}{\sqrt{k}} \|\mathbf{z}\|^2} = \sqrt{k} \cdot \bar{\lambda}(G). \tag{5.4}$$

The first summand in the right hand side of Equation (5.3) equals to

$$\frac{(\mathbf{z}^\parallel)^\mathsf{T} \mathbf{M}_G \mathbf{z}^\parallel}{\mathbf{z}^\mathsf{T} \mathbf{D} \mathbf{z}} = d \cdot \frac{\|\mathbf{z}^\parallel\|^2}{\|\mathbf{z}\|_{\mathbf{D}}^2} \tag{5.5}$$

and is always non-negative. This already gives a bound on $\bar{\lambda}_{\mathsf{min}}(\widetilde{\mathbf{M}}_G)$, because for every $\mathbf{y} \perp \widetilde{\mathbf{x}}_1$, we have

$$\frac{\mathbf{y}^\mathsf{T} \widetilde{\mathbf{M}}_G \mathbf{y}}{\mathbf{y}^\mathsf{T} \mathbf{y}} = \frac{(\mathbf{z}^\parallel)^\mathsf{T} \mathbf{M}_G \mathbf{z}^\parallel}{\mathbf{z}^\mathsf{T} \mathbf{D} \mathbf{z}} + \frac{(\mathbf{z}^\perp)^\mathsf{T} \mathbf{M}_G \mathbf{z}^\perp}{\mathbf{z}^\mathsf{T} \mathbf{D} \mathbf{z}} \geq 0 - \sqrt{k} \cdot \bar{\lambda}(G).$$

We summarize this in the following corollary.

▶ **Corollary 15.** $\left|\lambda_{\min}(\widetilde{\mathbf{M}}_G)\right| \le \sqrt{k} \cdot \bar{\lambda}(G)$.

Going back to the first summand in the right hand side of Equation (5.3), per Equation (5.5), we are left to bound the quotient $\frac{\|\mathbf{z}^\|\|}{\|\mathbf{z}\|_\mathbf{D}}$. We start with the numerator. As $\mathbf{z} \perp \mathbf{D}^{\frac{1}{2}}\widetilde{\mathbf{x}}_1$, for every $\alpha \in \mathbb{R}$ we have that

$$\|\mathbf{z}^\|\| = \langle \mathbf{z}, \mathbf{1} \rangle = \left\langle \mathbf{z}, \mathbf{1} - \alpha \mathbf{D}^{\frac{1}{2}}\widetilde{\mathbf{x}}_1 \right\rangle = \left\langle \mathbf{D}^{\frac{1}{2}}\mathbf{z}, \mathbf{D}^{-\frac{1}{2}}\mathbf{1} - \alpha\widetilde{\mathbf{x}}_1 \right\rangle \le \|\mathbf{D}^{\frac{1}{2}}\mathbf{z}\|\|\mathbf{D}^{-\frac{1}{2}}\mathbf{1} - \alpha\widetilde{\mathbf{x}}_1\|. \quad (5.6)$$

Choosing the optimal α for the bound, $\alpha = \frac{\left\langle \mathbf{D}^{-\frac{1}{2}}\mathbf{1}, \widetilde{\mathbf{x}}_1 \right\rangle}{\|\widetilde{\mathbf{x}}_1\|^2}$, we get that

$$\|\mathbf{D}^{-\frac{1}{2}}\mathbf{1} - \alpha\widetilde{\mathbf{x}}_1\|^2 = \|\mathbf{D}^{-\frac{1}{2}}\mathbf{1}\|^2 - \frac{\left\langle \mathbf{D}^{-\frac{1}{2}}\mathbf{1}, \widetilde{\mathbf{x}}_1 \right\rangle^2}{\|\widetilde{\mathbf{x}}_1\|^2}. \quad (5.7)$$

We now turn to analyze each of the summands in the above expression. To this end, we introduce the following notations. Let $s = |(B, L)|$ be the size of the cut (B, L) in G, and let $e = s - \frac{b\ell}{n}d$. Note that, had G been a d-regular random graph, s would have equal to $\frac{b\ell}{n}d$ in expectation, and so we think of e as the "cut size error". It will also be convenient to consider e, b and ℓ-s normalized counterpart $\bar{e} = \frac{e}{nd}$, $\bar{b} = \frac{b}{n}$ and $\bar{\ell} = \frac{\ell}{n}$.

From this point, denote $\lambda = \lambda(G)$ and $\bar{\lambda} = \bar{\lambda}(G) = \frac{\lambda}{d}$. We further define $\mu = b + \sqrt{k}\ell$ and its normalized counterpart $\bar{\mu} = \frac{\mu}{\sqrt{mn}}$ (where this normalization is in hindsight). The analysis of Equation (5.7) is divided to the following three claims.

▷ **Claim 16.**

$$\|\widetilde{\mathbf{x}}_1\|^2 = \frac{2|E(G_{\pi,L})|}{m} = d\bar{\mu}^2 - (\sqrt{k} - 1)^2 \frac{e}{m},$$

where recall $|E(G_{\pi,L})|$ counts the number of edges, accounting for the weights.

▷ **Claim 17.** $\left\langle \mathbf{D}^{-\frac{1}{2}}\mathbf{1}, \widetilde{\mathbf{x}}_1 \right\rangle = \bar{\mu}$.

▷ **Claim 18.** $\|\mathbf{D}^{-\frac{1}{2}}\mathbf{1}\|^2 \le \frac{1+k\bar{\lambda}}{d}$.

Claim 16 and Claim 17 follow by a fairly straightforward calculation. Their proofs can be found in the full version of the paper. Claim 18, whose proof also appears in the full version, follows by the following two more substantial lemma.

▶ **Lemma 19.**

$$\sum_{v \in B} \frac{1}{\mathbf{D}_{v,v}} \le \frac{b}{d}\left(\frac{1}{1 + k_B\bar{\ell}} + \sqrt{k}\bar{\lambda}\right).$$

▶ **Lemma 20.**

$$\sum_{v \in L} \frac{1}{\mathbf{D}_{v,v}} \le \frac{\ell}{d}\left(\frac{1}{1 + k_L\bar{b}} + k\bar{\lambda}\right).$$

The proof of Lemma 19 and Lemma 20 can be found in the full version of the paper, and we proceed with the proof of Proposition 11. Using Claim 16 and Claim 17, we write

$$
\begin{aligned}
\frac{\left\langle \mathbf{D}^{-\frac{1}{2}}\mathbf{1}, \widetilde{\mathbf{x}}_1 \right\rangle^2}{\|\widetilde{\mathbf{x}}_1\|^2} &= \frac{1}{d} \cdot \frac{\bar{\mu}^2}{\bar{\mu}^2 - (\sqrt{k}-1)^2 \frac{e}{dm}} \\
&= \frac{1}{d} \cdot \frac{1}{1 - (\sqrt{k}-1)^2 \frac{e}{dm\bar{\mu}^2}} \\
&\geq \frac{1}{d}\left(1 + (\sqrt{k}-1)^2 \frac{e}{dm\bar{\mu}^2}\right) \\
&\geq \frac{1}{d}\left(1 - (\sqrt{k}-1)^2 \frac{|e|}{dm\bar{\mu}^2}\right).
\end{aligned}
\tag{5.8}
$$

Focusing on the error term,

$$
(\sqrt{k}-1)^2 \frac{|e|}{dm\bar{\mu}^2} = (\sqrt{k}-1)^2 \frac{|\bar{e}|}{(\mu/n)^2},
$$

observe that according to the expander mixing lemma, $|\bar{e}| \leq \bar{\lambda}\sqrt{\bar{b}\bar{\ell}} \leq \bar{\lambda}$. Additionally, $\frac{\mu}{n}$ is bounded below by 1. The error term is thus bounded above by $\bar{\lambda}(\sqrt{k}-1)^2$. Plugging the above back to Equation (5.8), we get

$$
\frac{\left\langle \mathbf{D}^{-\frac{1}{2}}\mathbf{1}, \tilde{\mathbf{x}}_1 \right\rangle^2}{\|\tilde{\mathbf{x}}_1\|^2} \geq \frac{1}{d}\cdot\left(1 - (\sqrt{k}-1)^2\bar{\lambda}\right).
\tag{5.9}
$$

By Equation (5.6), Equation (5.7), Claim 18, and Equation (5.9),

$$
\begin{aligned}
\|\mathbf{z}^{\|}\| &\leq \|\mathbf{D}^{\frac{1}{2}}\mathbf{z}\| \cdot \sqrt{\|\mathbf{D}^{-\frac{1}{2}}\mathbf{1}\|^2 - \frac{\left\langle \mathbf{D}^{-\frac{1}{2}}\mathbf{1}, \widetilde{\mathbf{x}}_1 \right\rangle^2}{\|\widetilde{\mathbf{x}}_1\|^2}} \\
&\leq \frac{\|\mathbf{D}^{\frac{1}{2}}\mathbf{z}\| \cdot \sqrt{\bar{\lambda}}}{\sqrt{d}} \cdot \sqrt{k + (\sqrt{k}-1)^2}.
\end{aligned}
\tag{5.10}
$$

By Equation (5.5) and Equation (5.10), using that $\|\mathbf{z}\|_{\mathbf{D}} = \|\mathbf{D}^{\frac{1}{2}}\mathbf{z}\|$,

$$
\frac{(\mathbf{z}^{\|})^{\mathsf{T}}\mathbf{M}_G\mathbf{z}^{\|}}{\mathbf{z}^{\mathsf{T}}\mathbf{D}\mathbf{z}} = d \cdot \frac{\|\mathbf{z}^{\|}\|^2}{\|\mathbf{z}\|_{\mathbf{D}}^2} \leq (k + (\sqrt{k}-1)^2)\bar{\lambda}.
$$

Thus, by Equation (5.3) and Equation (5.4),

$$
\lambda_2(\widetilde{\mathbf{M}}_G) \leq d \cdot \frac{\|\mathbf{z}^{\|}\|^2}{\|\mathbf{z}\|_{\mathbf{D}}^2} + \sqrt{k} \cdot \bar{\lambda}(G) \leq (k + \sqrt{k} + (\sqrt{k}-1)^2)\bar{\lambda} = O(k \cdot \bar{\lambda}).
$$

Combining this with Corollary 15, we obtain $\lambda(\widetilde{\mathbf{M}}_G) = O(k \cdot \bar{\lambda})$. Restating this result, for every $\mathbf{x} \in F^{\|}$ with $\mathbf{x} \perp \mathbf{D}_{G_{\pi,L}}\mathbf{1}_m$ we have $\mathbf{x}^{\mathsf{T}}\mathbf{W}_{G_{\pi,L}}\mathbf{x} = O(k \cdot \bar{\lambda})$. Recall $\bar{\lambda} = \bar{\lambda}(G) \leq \bar{\lambda}(G_\pi)$. Combining this and Corollary 12 we conclude $\bar{\lambda}(G_{\pi,L}) = O(k \cdot \bar{\lambda}(G_\pi))$, which completes the proof. ◀

6 Proof of Theorem 1

In this section we wrap it all up and prove Theorem 1.

Proof of Theorem 1. By repeatedly applying Corollary 3.1 of [3] to the base graph, denoted BL_0, which is the clique on $d+1$ vertices, we obtain an explicit family of d-regular expanders $\mathcal{BL} = (\mathrm{BL}_n)_{n \in \mathcal{I}}$, where $\mathcal{I} = \{(d+1) \cdot 2^i \mid i \in \mathbb{N}\}$ and $\lambda(\mathcal{BL}) = O(\sqrt{d \log^3 d})$. Moreover, as BL_0 is simple, all graphs in $\mathcal{BL}$ are simple. By regularity,

$$\bar{\lambda}(\mathcal{BL}) = \frac{1}{d}\lambda(\mathcal{BL}) = O\left(\sqrt{\frac{\log^3 d}{d}}\right).$$

Let $n : \mathbb{N} \to \mathcal{I}$ be the function defined by $n(i) = (d+1) \cdot 2^i$. That is, $n(i)$ is the number of vertices of the i-th graph in $\mathcal{BL}$. Observe that for every $i \in \mathcal{I}$, $\mathrm{BL}_{n(i+2)}$ is a π_i-lift of $\mathrm{BL}_{n(i)}$ with $k = 4$. Indeed, the 2-lift of a 2-lift is a 4-lift.

Fix $i \in \mathbb{N}$. Choose an arbitrary ordering on the vertices of $\mathrm{BL}_{n(i)}$ and denote the first j vertices, under this ordering, by $L_{n(i),j}$. Define $P_{n(i),j}$ to be the $L_{n(i),j}$-partial π_i-lift of $\mathrm{BL}_{n(i)}$. Note that the edge weights of $P_{n(i),j}$ are 1 and $\frac{1}{2}$. In order to avoid fractional edges, multiply every edge by 2 to obtain the unweighted graph with multiple edges $G_{n(i),j}$. Note that the latter is a graph on $n(i) + 3j$ vertices, and that $G_{n(i),n(i)} = G_{n(i+2),0}$. The family that we construct is given by $\mathcal{G} = (G_{n,j})_{(n,j) \in \mathcal{I}'}$, where the index set

$$\mathcal{I}' = \{(n(i), j) \mid i \in 2\mathbb{N} \text{ and } 0 \leq j \leq n(i)\}$$

is lexicographically ordered.

Per Definition 5, and as we duplicated all edges, the degrees of all vertices in this family are in the range $[2\frac{d}{\sqrt{4}}, 2\sqrt{4}d] = [d, 4d]$, as claimed. Proposition 6 and Proposition 11 readily imply the bounds on $\lambda(\mathcal{G})$ and $\bar{\lambda}(\mathcal{G})$, respectively. To conclude the proof, note that the expansion cost is $O(d)$.

◀

---- **References** ----

1 Miklós Ajtai, János Komlós, and Endre Szemerédi. Deterministic simulation in logspace. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 132–140, 1987.

2 Alon Amit and Nathan Linial. Random graph coverings. I. General theory and graph connectivity. *Combinatorica*, 22(1):1–18, 2002. `doi:10.1007/s004930200000`.

3 Yonatan Bilu and Nathan Linial. Lifts, discrepancy and nearly optimal spectral gap. *Combinatorica*, 26(5):495–519, 2006. `doi:10.1007/s00493-006-0029-7`.

4 Nikolas P. Breuckmann and Jens N. Eberhardt. Balanced product quantum codes. *IEEE Trans. Inform. Theory*, 67(10):6653–6674, 2021. `doi:10.1109/tit.2021.3097347`.

5 Michael Dinitz, Michael Schapira, and Asaf Valadarsky. Explicit expanding expanders. *Algorithmica*, 78(4):1225–1245, 2017. `doi:10.1007/s00453-016-0269-x`.

6 Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):Art. 12, 44, 2007. `doi:10.1145/1236457.1236459`.

7 Irit Dinur, Shai Evra, Ron Livne, Alexander Lubotzky, and Shahar Mozes. Locally testable codes with constant rate, distance, and locality. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pages 357–374, 2022.

8 Chris Godsil and Gordon F Royle. *Algebraic graph theory*, volume 207. Springer Science & Business Media, 2001.

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

10 Russell Impagliazzo, Noam Nisan, and Avi Wigderson. Pseudorandomness for network algorithms. In *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 356–364, 1994.

11 Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

12 Adam Marcus, Daniel A Spielman, and Nikhil Srivastava. Interlacing families I: Bipartite Ramanujan graphs of all degrees. In *2013 IEEE 54th Annual Symposium on Foundations of computer science*, pages 529–537. IEEE, 2013.

13 Adam W Marcus, Daniel A Spielman, and Nikhil Srivastava. Interlacing families IV: Bipartite Ramanujan graphs of all sizes. *SIAM Journal on Computing*, 47(6):2488–2509, 2018.

14 Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Finite free convolutions of polynomials. *Probab. Theory Related Fields*, 182(3-4):807–848, 2022. `doi:10.1007/s00440-021-01105-w`.

15 A. Nilli. On the second eigenvalue of a graph. *Discrete Math.*, 91(2):207–210, 1991. `doi:10.1016/0012-365X(91)90112-F`.

16 Gopal Pandurangan, Peter Robinson, and Amitabh Trehan. DEX: self-healing expanders. *Distrib. Comput.*, 29(3):163–185, 2016. `doi:10.1007/s00446-015-0258-3`.

17 Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4):Art. 17, 24, 2008. `doi:10.1145/1391289.1391291`.

18 Daniel A Spielman. Spectral and algebraic graph theory, 2019. URL: `http://cs-www.cs.yale.edu/homes/spielman/sagt/`.

19 Amnon Ta-Shma. Explicit, almost optimal, epsilon-balanced codes. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 238–251, 2017.

20 Luca Trevisan. Lecture notes on graph partitioning, expanders and spectral methods, 2017. University of California, Berkeley, `https://people.eecs.berkeley.edu/luca/books/expanders-2016.pdf`.

21 Salil P Vadhan. Pseudorandomness. *Foundations and Trends® in Theoretical Computer Science*, 7(1–3):1–336, 2012.

22 Leslie G. Valiant. Graph-theoretic properties in computational complexity. *J. Comput. System Sci.*, 13(3):278–285, 1976. `doi:10.1016/S0022-0000(76)80041-4`.