

Expansion in Matrix-Weighted Graphs[☆]

Jakob Hansen^{a,1}

^a*Department of Mathematics, Ohio State University, 231 West 18th Avenue, Columbus, OH 43210*

Abstract

A matrix-weighted graph is an undirected graph with a $k \times k$ positive semidefinite matrix assigned to each edge. There are natural generalizations of the Laplacian and adjacency matrices for such graphs, leading to a generalized notion of expansion. Extensions of some theorems about expansion hold for matrix-weighted graphs—in particular, an analogue of the expander mixing lemma and one half of a Cheeger-type inequality. These results lead to a definition of a matrix-weighted expander graph, and suggest the tantalizing possibility of families of matrix-weighted graphs with better-than-Ramanujan expansion.

Keywords: matrix-weighted graph, expander graph, cellular sheaf

2020 MSC: 05C48, 05C50

1. Introduction

A recent thread of investigation in spectral graph theory has been the extension of graph-theoretic concepts to higher dimensions. This extension may take place by raising the dimensionality of the underlying structure, as with the spectral theory for simplicial complexes and hypergraphs [1, 2, 3, 4]. However, it is also possible to raise the dimension of the algebraic components of interest: rather than consider \mathbb{R} -valued functions on the vertices of a graph, consider functions valued in higher-

[☆]This work was supported by funding from DARPA and the AFRL under agreement number FA8650-18-2-7840, the Office of the Assistant Secretary of Defense Research & Engineering through a Vannevar Bush Faculty Fellowship, ONR N00014-16-1-2010, and by the National Science Foundation under NSF-DMS #1547357.

Email address: hansen.612@osu.edu (Jakob Hansen)

¹Declarations of interest: none.

dimensional spaces. This extension allows us to define new classes of graph operators. The most famous of these is perhaps the *graph connection Laplacian*, which introduces a weighted orthogonal transformation corresponding to each edge. This has been used for dimensionality reduction and data analysis [5, 6], and various theoretical results including a Cheeger-type inequality [7], sparsification algorithms [8, 9], and results on the spectrum of random connection Laplacians [10].

A somewhat less well known higher-dimensional generalization is the *matrix-weighted graph*. Rather than assign an orthogonal matrix to each edge, a matrix-weighted graph assigns a positive semidefinite matrix to each edge. Matrix-weighted Laplacians in particular have seen development and use in the design and control of engineering systems [11, 12, 13].

Both connection graphs and matrix-weighted graphs are special cases of *cellular sheaves* [14]. These are algebraic structures attached to a graph (or a higher-dimensional combinatorial space) that describe consistency constraints for data parameterized by the graph. Graph connection Laplacians and matrix-weighted Laplacians are instances of *sheaf Laplacians* [15]. The cellular sheaf perspective can shed light on various phenomena arising in these more restricted domains.

This paper focuses on understanding the expansion properties of matrix-weighted graphs. Of the higher-dimensional extensions of graphs, these have the behavior most similar to that of the usual scalar-weighted graphs. (For instance, it is not even entirely clear what an appropriate definition of expansion is for connection graphs or other types of cellular sheaves.) Still, there are a number of subtle differences that add additional richness and interest to the theory in the matrix-weighted case.

We will first define matrix-weighted graphs and their paraphernalia—degrees, Laplacians, adjacency matrices, etc., as a generalization of standard objects from graph theory. We then introduce cellular sheaves and describe how matrix-weighted graphs are realized as sheaves. After a few examples, we explore the relationship between the spectra of matrix-weighted graphs and certain associated scalar-weighted graphs. We then prove a version of the expander mixing lemma for matrix-weighted graphs, as well as one half of a Cheeger inequality for regular matrix-weighted graphs, and show that the complementary inequality cannot hold. Finally, we propose a definition of a

matrix-weighted expander graph and discuss its implications.

2. Matrix-Weighted Graphs

2.1. Definitions

We will view a weighted graph as a structure built on top of an underlying unweighted, undirected graph. Let G be a graph with vertex set V and edge set E . We will write $v \triangleleft e$ for the vertex-edge incidence relation; that is, $v \triangleleft e$ if v is one of the endpoints of the edge e . A *weighting* on G is a function $w : E \rightarrow \mathbb{R}$, whose values we write w_e for $e \in E$, such that $w_e \geq 0$. For an edge $e = u \sim v$, we write $w_{uv} = w_e = w_{vu}$, and we can extend this by letting $w_{uv} = 0$ whenever there is no edge between u and v . One may represent a weighted graph by its adjacency matrix A , whose rows and columns are indexed by V , which has $A_{uv} = w_{uv}$. The weighted degree of a vertex v is $d_v = \sum_{v \triangleleft e} w_e = \sum_{u \in V} w_{uv}$. The adjacency matrix determines and is determined by the weighted Laplacian matrix $L = D - A$, where D is the diagonal matrix whose entries are the weighted degrees.

Matrix-weighted graphs are a generalization of this structure. Rather than assigning a nonnegative scalar w_e to each edge, we assign a $k \times k$ symmetric positive semidefinite matrix W_e . We can equivalently specify this as a symmetric function on pairs of vertices as before, letting $W_{uv} = W_e$ for $e = u \sim v$ and $W_{uv} = 0$ if there is no edge between u and v . A matrix-weighted graph may again be represented by its adjacency matrix. This is a block matrix with $k \times k$ blocks, whose block rows and columns are indexed by V , and where $A_{uv} = W_{uv}$. There is also a corresponding matrix-weighted Laplacian matrix $L = D - A$, defined blockwise analogously to the scalar-weighted version, with the degree matrix D having blocks on the diagonal equal to the block row (or column) sums of A . These matrices are interesting as generalizations of the constructions familiar from spectral graph theory.

We think of the matrix-weighted versions of the adjacency and Laplacian matrices as linear operators on the space of functions $V \rightarrow \mathbb{R}^k$. That is, these operators take as input an assignment of a vector in \mathbb{R}^k to each vertex of G and output an assignment of the same form. The action of a general matrix-weighted adjacency matrix or Laplacian

on $(\mathbb{R}^k)^V$ may be written vertexwise as

$$(Ax)_v = \sum_{u \in V} W_{uv} x_u \quad (1)$$

$$(Lx)_v = \sum_{u \in V} W_{uv} (x_v - x_u), \quad (2)$$

where we note that this is an expression relating vectors in \mathbb{R}^k . From this expression, it is easy to see that the kernel of L is at least k -dimensional, for it contains all constant functions $V \rightarrow \mathbb{R}^k$. If G is not connected, the kernel of L contains a direct summand of dimension k corresponding to each connected component of G . However, even if G is connected, the kernel of L may be more than k -dimensional. The matrix L is positive semidefinite, as will be easy to see by considerations in Section 2.2. Therefore, if we write its eigenvalues in increasing order, we have $0 = \lambda_1 = \dots = \lambda_k \leq \lambda_{k+1} \leq \dots$.

We will say that a matrix-weighted graph is *regular* if the vertexwise degree matrix $D_v = \sum_{u \in V} W_{uv}$ is the same for every vertex v . When necessary to avoid confusion, we will call D_v the *algebraic degree*, and the degree of the vertex in the underlying graph the *topological degree*. The “most regular” matrix-weighted graphs have algebraic degree equal to dI for some $d \in \mathbb{R}$; by an abuse of notation we will call these d -regular matrix-weighted graphs. The adjacency and Laplacian spectra of a d -regular matrix weighted graph have related eigenvalues: since the total degree matrix D is equal to dI , the eigenvalues of A are $\mu_i = d - \lambda_i$.

Just as with weighted graphs, it is often useful to normalize the Laplacian and adjacency matrices of matrix-weighted graphs. Since the degree matrices are positive semidefinite, they have square roots; we define the normalized Laplacian to be $\tilde{L} = D^{\dagger/2} L D^{\dagger/2}$, where $D^{\dagger/2}$ is the Moore-Penrose pseudoinverse of the square root of the degree matrix. We likewise define the normalized adjacency matrix to be $\tilde{A} = D^{\dagger/2} A D^{\dagger/2} = I - \tilde{L}$. If D is invertible, the block diagonal entries of \tilde{L} are copies of the $k \times k$ identity matrix. However, the off-diagonal block entries are not in general symmetric.

The scalar normalized Laplacian is useful in part because its spectrum is bounded above by a constant regardless of the size or degree distribution of the graph. The same holds for the matrix-weighted normalized Laplacian.

Proposition 2.1. *The eigenvalues of the normalized Laplacian of a matrix-weighted graph are bounded above by 2.*

PROOF. By the Courant-Fischer theorem, the largest eigenvalue of \tilde{L} is

$$\tilde{\lambda}_{\max} = \max_x \frac{\langle x, D^{\dagger/2} L D^{\dagger/2} x \rangle}{\langle x, x \rangle}.$$

Since any $x \in \ker D$ is also in $\ker L$ and hence is orthogonal to any eigenvector for $\tilde{\lambda}_{\max}$, we can restrict the domain of the maximization to get

$$\begin{aligned} \tilde{\lambda}_{\max} &= \max_{x \perp \ker D} \frac{\langle x, D^{\dagger/2} L D^{\dagger/2} x \rangle}{\langle x, x \rangle} = \max_{y \perp \ker D} \frac{\langle y, Ly \rangle}{\langle y, Dy \rangle} \\ &= \max_{y \perp \ker D} \frac{\sum_{u,v \in e} \langle y_u - y_v, W_e(y_u - y_v) \rangle}{\sum_v \sum_{v \in e} \langle y_v, W_e y_v \rangle} \\ &\leq \max_{y \perp \ker D} \frac{2 \sum_{u,v \in e} \langle y_u, W_e y_u \rangle + \langle y_v, W_e y_v \rangle}{\sum_v \sum_{v \in e} \langle y_v, W_e y_v \rangle} = 2. \end{aligned}$$

The bound is achieved when there exists a vector y such that $\langle y, Ly \rangle = 2\langle y, Dy \rangle$. As in the standard case, this occurs when the underlying graph is bipartite; in this case the choice of y that attains the bound is constant on each half of the partition, differing only by a sign across the bounds. However, this is not the only situation in which $\tilde{\lambda}_{\max} = 2$. The reader may find it instructive to construct other matrix-weighted graphs with $\tilde{\lambda}_{\max} = 2$.

Proposition 2.1 immediately implies that the adjacency spectrum of a d -regular matrix-weighted graph is contained in $[-d, d]$.

2.1.1. Notation

Throughout, G will be an underlying graph with vertex set V and edge set E . The graph will have n vertices and weight matrices will be $k \times k$. Regular matrix-weighted graphs will have (algebraic) degree d . Thus, the relevant matrices A , L , etc. will have size $kn \times kn$. Eigenvalues of the Laplacian will be denoted λ_i , in increasing order, while eigenvalues of the adjacency matrix will be denoted μ_i , in decreasing order.

2.2. Cellular Sheaves

Matrix-weighted graphs are instances of a more general structure on a graph: a *cellular sheaf*. We can understand their spectral theory in the context of a broader spectral theory of cellular sheaves.

Definition 2.1. Let G be a graph. A cellular sheaf \mathcal{F} on G consists of the following data:

1. A vector space $\mathcal{F}(v)$ for each vertex v of G , called the *stalk* over v
2. A vector space $\mathcal{F}(e)$ for each edge e of G , called the stalk over e , and
3. A linear map $\mathcal{F}_{v \triangleleft e} : \mathcal{F}(v) \rightarrow \mathcal{F}(e)$ for each incident vertex-edge pair $v \triangleleft e$ of G , called the *restriction map* from v to e .

Cellular sheaves describe systems of consistency relationships for data over graphs. Data may be assigned to vertices and edges, living in the stalks over these edges, and the restriction maps give conditions for consistency of this data.

Definition 2.2. Let \mathcal{F} be a cellular sheaf over a graph G . A *global section* x of \mathcal{F} is given by a choice of a vector $x_v \in \mathcal{F}(v)$ for each vertex v of G , such that for every edge $e = u \sim v$ of G , $\mathcal{F}_{v \triangleleft e} x_v = \mathcal{F}_{u \triangleleft e} x_u$.

Because these conditions are linear, the global sections of \mathcal{F} form a vector space, which we denote $H^0(G; \mathcal{F})$. The global sections of a sheaf are the collections of elements satisfying all the consistency conditions specified by the sheaf. We think of the space of section $H^0(G; \mathcal{F})$ as lying inside a larger space of assignments to vertices, which we denote

$$C^0(G; \mathcal{F}) = \bigoplus_v \mathcal{F}(v).$$

This is the space of *0-cochains* of \mathcal{F} ; it consists of all possible assignments to vertex stalks without reference to any consistency conditions. There is an analogous space of *1-cochains* consisting of assignments to edge stalks:

$$C^1(G; \mathcal{F}) = \bigoplus_e \mathcal{F}(e).$$

The space of global sections $H^0(G; \mathcal{F})$ is the kernel of a map $\delta : C^0(G; \mathcal{F}) \rightarrow C^1(G; \mathcal{F})$, called the *coboundary operator*. Given an orientation of the graph, the value of this operator on an oriented edge $e = u \rightarrow v$ is

$$(\delta x)_e = \mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u.$$

It is straightforward to see that $\delta x = 0$ if and only if $x \in H^0(G; \mathcal{F})$. The coboundary operator is a generalization of the signed incidence matrix of a graph.

The terminology associated with cellular sheaves is perhaps somewhat foreign. It originates in a more complex definition of sheaves used in geometry and topology (see, e.g., [16, 17]). The central idea of a sheaf as describing constraints for data parameterized by a space holds across these different instantiations. Cellular sheaves are a restriction of the concept to the discrete setting of regular cell complexes, which makes them particularly amenable to computation and applications [14]. We have further specialized to sheaves over graphs, which makes the constructions more accessible but also perhaps further obscures the reasoning for the terminology. A similar but algebraically dual definition of a sheaf on a graph was given by Friedman [18], who suggested the study of the Laplacians and adjacency matrices of such structures.

Thus far we have only required that the stalks of a cellular sheaf be abstract vector spaces. To develop the relationship between matrix-weighted graphs and cellular sheaves, each stalk must also have an inner product. Inner products on stalks extend to inner products on $C^0(G; \mathcal{F})$ and $C^1(G; \mathcal{F})$, and induce an adjoint δ^* to the coboundary operator. The *sheaf Laplacian* is then defined as $L_{\mathcal{F}} = \delta^* \delta$. This is a linear map $C^0(G; \mathcal{F}) \rightarrow C^0(G; \mathcal{F})$, computed vertexwise by

$$(L_{\mathcal{F}} x)_v = \sum_{u, v \triangleleft e} \mathcal{F}_{v \triangleleft e}^* (\mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u).$$

As a quadratic form, it is given by

$$\langle x, L_{\mathcal{F}} x \rangle = \langle \delta x, \delta x \rangle = \|\delta x\|^2 = \sum_{u, v \triangleleft e} \|\mathcal{F}_{v \triangleleft e} x_v - \mathcal{F}_{u \triangleleft e} x_u\|^2.$$

The Laplacian quadratic form measures how close a 0-cochain is to being a global section. Sheaf Laplacians are studied in greater generality in [15, 19].

How are matrix-weighted graphs related to cellular sheaves? We begin first by relating weighted graphs to weighted cellular sheaves. This relationship is mediated through the *constant sheaf* \mathbb{R} on a graph G . This sheaf has all vertex and edge stalks equal to \mathbb{R} , and all restriction maps the identity. The global sections of the constant sheaf are precisely the locally constant \mathbb{R} -valued functions on the vertices of G . A weighting on G corresponds to a choice of an inner product on each edge stalk:

$\langle x, y \rangle_e = w_e xy$ for $x, y \in \underline{\mathbb{R}}(e) = \mathbb{R}$. If we assign all vertex stalks the standard inner product $\langle x, y \rangle_v = xy$, the corresponding sheaf Laplacian is precisely the weighted graph Laplacian.

To extend this to matrix-weighted graphs, we need to reckon more carefully with the semidefiniteness of the weight matrices. If W_e is not positive definite, it does not define an inner product on \mathbb{R}^k , but only on $\text{im } W_e$. Given a matrix-weighted graph G with $k \times k$ weight matrices, we construct a sheaf \mathcal{F} with vertex stalks $\mathcal{F}(v) = \mathbb{R}^k$ and edge stalks $\mathcal{F}(e) = \text{im } W_e \subseteq \mathbb{R}^k$. The restriction map $\mathcal{F}_{v \triangleleft e}$ is the orthogonal projection $\mathbb{R}^k \rightarrow \text{im } W_e$. We give the vertex stalks the standard inner product on \mathbb{R}^k , and the edge stalks the inner product $\langle x, y \rangle_e = x^T W_e y$. It is easily checked that under the standard basis for \mathbb{R}^k the corresponding sheaf Laplacian is equal to the matrix-weighted Laplacian. Since the definition of the sheaf Laplacian is $L_{\mathcal{F}} = \delta^* \delta$, it is obvious that the matrix-weighted graph Laplacian is positive semidefinite.

The interpretation of matrix-weighted graphs in terms of weighted cellular sheaves gives them a coordinate-free description. We could define a matrix-weighted graph to be a weighted cellular sheaf \mathcal{F} with all vertex stalks equal to some vector space V , where for any edge $e = u \sim v$, the restriction maps $\mathcal{F}_{u \triangleleft e}$ and $\mathcal{F}_{v \triangleleft e}$ are equal to some map we will call ρ_e . If an orthonormal basis for V is chosen, the resulting sheaf Laplacian matrix will have the form of the Laplacian of a matrix-weighted graph. The edge weights W_e will be equal to $\rho_e^* \rho_e$. The adjacency matrix is then obtained from the Laplacian by $A = D - L$.

Under the direct definition, the matrix-weighted adjacency matrix is the primary object, and the Laplacian is generated therefrom. In the context of cellular sheaves, however, the Laplacian is the principal operator, and the adjacency matrix is extracted from it. For more general sheaves, the Laplacian matrix contains more information than the adjacency matrix.

For the remainder of this paper, we will adopt the elementary but less general terminology of matrix-weighted graphs. However, the sheaf-theoretic perspective has inspired and motivated this work, and can provide important insights into the deeper reasons for certain phenomena.

2.3. Examples

One frequently seen example of a matrix-weighted graph comes from the mechanical analysis of bar-and-joint structures. Given a collection of struts joined together at their ends, represented as a structure in \mathbb{R}^3 , consider the graph G with edges corresponding to struts and vertices corresponding to joints. We assign to each edge a scaled copy of the 3×3 matrix which computes the orthogonal projection onto the direction spanned by the corresponding strut in \mathbb{R}^3 . The scaling factor is a stiffness parameter representing the resistance of the strut to compression or tension. The Laplacian of this matrix-weighted graph is the *stiffness matrix* of the truss. As a quadratic form, it represents the amount of work done under an infinitesimal deformation of the structure.

This physical interpretation allows us to quickly conclude that the kernel of the Laplacian contains more than simply the constant functions $V \rightarrow \mathbb{R}^3$. These constant functions correspond to infinitesimal translations; the fact that they are in the kernel of L is the physical fact that translations of a truss do not cause it to deform and hence require no expenditure of energy. But rigid rotations of the truss also cause no deformation, and so the infinitesimal generators of these rotations must also correspond to vectors in the kernel of L . The kernel of L is therefore at least 6-dimensional. These bar-and-joint structures give a class of nontrivial examples of connected matrix-weighted graphs with a Laplacian kernel of dimension greater than k .

An essentially identical example has been studied for specific graphs representing molecular structures, under the name “vibrational spectrum” [20], so called because the eigenfunctions of the matrix-weighted Laplacian correspond (up to first order) to vibrational modes of the molecule. The vibrational spectrum of a symmetric graph with a symmetric embedding in \mathbb{R}^3 is strongly constrained by representation theoretic considerations.

Other instances of matrix-weighted graphs arise in the engineering control literature. Examples include certain systems of coupled oscillators [11], differential observations of networked systems [12], and distributed coordination for autonomous agents [13]. Many of these motivating examples are quite concrete, but very little theoretical work has been done exploring the algebraic and spectral properties of matrix-weighted graphs. One exception to this pattern is [21], which constructed effective

resistance matrices for matrix-weighted graphs.

2.4. Relationships between scalar- and matrix-weighted graphs

There is a straightforward way to turn any weighted graph into a matrix-weighted graph for any block size k : simply let the matrix-valued weights be $W_{uv} = w_{uv}I_{k \times k}$. The corresponding matrix-weighted adjacency and Laplacian matrices are then given by $A \otimes I_{k \times k}$ and $L \otimes I_{k \times k}$, where the tensor product of operators is realized by the Kronecker product on matrices.

Conversely, given a matrix-weighted graph (G, W) , we can construct a scalar-weighted graph $(G, \text{tr } W)$ by letting $w_e = \text{tr}(W_e)$ for all edges e of G . This construction is invariant to an orthogonal change of basis of the vertex stalks in the cellular sheaf definition. The Laplacian and adjacency spectral radii of (G, W) are controlled by the spectral radii of $(G, \text{tr } W)$.

Proposition 2.2. *Let (G, W) be a matrix-weighted graph with n vertices and $k \times k$ weights, with Laplacian L_W , and let $L_{\text{tr } W}$ be the Laplacian of $(G, \text{tr } W)$. If $\lambda_1(L) \leq \lambda_2(L) \leq \dots$ are the eigenvalues of the matrix L , then*

$$\sum_{i=1}^k \lambda_{k+i}(L_W) \leq \lambda_2(L_{\text{tr } W}) \leq \lambda_n(L_{\text{tr } W}) \leq \sum_{i=1}^k \lambda_{(n-1)k+i}(L_W).$$

PROOF. Let x be a unit eigenvector of $L_{\text{tr } W}$ corresponding to the eigenvalue $\lambda_2(L_{\text{tr } W})$. Let $\{e_1, \dots, e_k\}$ be an orthonormal basis for \mathbb{R}^k , and consider the orthogonal vectors $x \otimes e_i$, which are naturally in the domain of L_W . Note that $\|x \otimes e_i\| = 1$. Further, for any constant \mathbb{R}^k -valued function $y = a\mathbb{1} \otimes s$ on the vertices of G , $\langle x \otimes e_i, y \rangle = a \langle x, \mathbb{1} \rangle \langle e_i, s \rangle = 0$, so $x \otimes e_i$ is orthogonal to the eigenspace of L_W corresponding to the first k eigenvalues. Thus by a generalized form of the Courant-Fischer theorem, we

have

$$\begin{aligned}
\sum_{i=1}^k \lambda_{k+i}(L_W) &\leq \sum_{i=1}^k \langle x \otimes e_i, L_W x \otimes e_i \rangle \\
&= \sum_{i=1}^k \sum_{u,v \in e} \langle (x \otimes e_i)_v - (x \otimes e_i)_u, W_e((x \otimes e_i)_v - (x \otimes e_i)_u) \rangle \\
&= \sum_{u,v \in e} (x_v - x_u)^2 \sum_{i=1}^k \langle e_i, W_e e_i \rangle \\
&= \sum_{u,v \in e} \text{tr}(W_e)(x_v - x_u)^2 = \langle x, L_{\text{tr} W} x \rangle = \lambda_2(L_{\text{tr} W}).
\end{aligned}$$

The same calculation applied to an eigenvector for $\lambda_n(L_{\text{tr} W})$ gives the upper bound.

An immediate corollary is that $\lambda_{k+1}(L_W) \leq \frac{1}{k} \lambda_2(L_{\text{tr} W})$ and $\lambda_{nk}(L_W) \geq \frac{1}{k} \lambda_n(L_{\text{tr} W})$.

The analogous bound for the adjacency eigenvalues is proved by exactly the same method. For dI -regular matrix-weighted graphs the bound implied by Proposition 2.2 and the fact that $A = dI - L$ is stronger, since it constrains μ_{k+1} rather than μ_1 .

Proposition 2.3. *Let (G, W) be a matrix-weighted graph on n vertices with $k \times k$ weights, with adjacency matrix A_W , and let $A_{\text{tr} W}$ be the adjacency matrix of $(G, \text{tr} W)$. If $\mu_1(A) \geq \mu_2(A) \geq \dots \geq$ are the eigenvalues of the matrix A , then*

$$\sum_{i=1}^k \mu_i(A_W) \geq \mu_1(A_{\text{tr} W}) \geq \mu_n(A_{\text{tr} W}) \geq \sum_{i=1}^k \mu_{(n-1)k+i}(A_W).$$

3. An Expander Mixing Lemma

The expander mixing lemma is a well-known result, perhaps first explicitly proven in [22] (although it has arisen in many different guises), connecting the number of edges between a pair of subsets of a graph and its adjacency spectrum. For a d -regular graph with n vertices, it states that for any two subsets of vertices S, T , the number of edges between S and T , $e(S, T)$, satisfies

$$\left| e(S, T) - \frac{d|S||T|}{n} \right| \leq |\mu_2| \sqrt{|S||T| \left(1 - \frac{|S|}{n}\right) \left(1 - \frac{|T|}{n}\right)},$$

where μ_2 is the nontrivial eigenvalue of A_G of largest modulus.

When applied to weighted graphs, the edge count $e(S, T)$ is the sum of weights of edges between S and T . Similarly, for matrix weighted graphs, we define $E(S, T) = \sum_{s \in S, t \in T} W_{st}$, so that the edge count becomes a positive semidefinite matrix. If we let I_S be the $kn \times k$ block matrix with blocks

$$(I_S)_v = \begin{cases} I_{k \times k} & v \in S \\ 0 & v \notin S \end{cases}$$

and similarly for I_T , it is easy to see that for a matrix-weighted graph (G, W) , $E(S, T) = I_S^T A I_T$. This fact allows us to generalize the standard proof of the expander mixing lemma to d -regular matrix-weighted graphs.

Lemma 3.1. *Let (G, W) be a d -regular matrix-weighted graph on n vertices, with $k \times k$ weight matrices. Denote the adjacency eigenvalues of G by $d = \mu_1 = \dots = \mu_k \geq \mu_{k+1} \geq \dots$, and let $|\mu| = \max\left(\sum_{i=1}^k \mu_{k+i}, \sum_{i=1}^k |\mu_{(n-1)k+i}|\right)$. If S and T are subsets of the vertices of G , the matrix-weighted edge count $E(S, T)$ satisfies*

$$\left| \text{tr}(E(S, T)) - \frac{kd|S||T|}{n} \right| \leq |\mu| \sqrt{|S||T| \left(1 - \frac{|S|}{n}\right) \left(1 - \frac{|T|}{n}\right)} \quad (3)$$

and the eigenvalues of $E(S, T) - \frac{k|S||T|}{n}$ have magnitude at most

$$\max(|\mu_{k+1}|, |\mu_{kn}|) \sqrt{|S||T| \left(1 - \frac{|S|}{n}\right) \left(1 - \frac{|T|}{n}\right)}. \quad (4)$$

PROOF. The first inequality follows directly from Proposition 2.2 and the standard expander mixing lemma. Note that $\text{tr}(E(S, T))$ for the matrix weighting W on G is equal to $e(S, T)$ for the weighting $\text{tr} W$. Thus, if $|\mu(A_{\text{tr} W})|$ is the magnitude of the largest nontrivial adjacency eigenvalue of $(G, \text{tr} W)$,

$$\left| \text{tr}(E(S, T)) - \frac{kd|S||T|}{n} \right| \leq |\mu(A_{\text{tr} W})| \sqrt{|S||T| \left(1 - \frac{|S|}{n}\right) \left(1 - \frac{|T|}{n}\right)}.$$

We use the fact that $|\mu(A_{\text{tr} W})| = \max(|d - \lambda_2(L_{\text{tr} W})|, |d - \lambda_n(L_{\text{tr} W})|)$ to apply the trace bound, finding that

$$|\mu(A_{\text{tr} W})| \leq \max\left(\sum_{i=1}^k \mu_{k+i}, \sum_{i=1}^k |\mu_{kn-i+1}|\right).$$

For the second inequality we must mimic the proof of the standard expander mixing lemma. We use the fact that $E(S, T) = I_S^T A_G I_T$, and decompose these indicator matrices appropriately. Let $I_S^\perp = I_S - \frac{|S|}{n} I_G$ and $I_T^\perp = I_T - \frac{|T|}{n} I_G$. This gives an orthogonal decomposition of I_S and I_T in the following strong sense: every column of I_G is orthogonal to every column of I_S^\perp and every column of I_T^\perp . Further, any two columns selected from one of I_G , I_S^\perp , and I_T^\perp have disjoint supports and hence are orthogonal as well. We therefore have

$$\begin{aligned} E(S, T) &= I_S^T A_G I_T \\ &= \left(\frac{|S|}{n} I_G + I_S^\perp \right)^T A_G \left(\frac{|T|}{n} I_G + I_T^\perp \right) \\ &= \frac{|S||T|}{n^2} I_G^T A_G I_G + \frac{|S|}{n} I_G^T A_G I_T^\perp + (I_S^\perp)^T A_G \frac{|T|}{n} I_G + (I_S^\perp)^T A_G I_T^\perp. \end{aligned}$$

Every column of I_G is an eigenvector of A_G with eigenvalue d , so that, for instance $(I_S^\perp)^T A_G I_G = d(I_S^\perp)^T I_G = 0$, due to the orthogonality relations between these matrices. Thus, the two middle terms vanish, and the first term is equal to $\frac{d|S||T|}{n} I_{k \times k}$. Combining these simplifications gives

$$E(S, T) - \frac{d|S||T|}{n} I_{k \times k} = (I_S^\perp)^T A_G I_T^\perp. \quad (5)$$

We therefore need to bound the eigenvalues of $(I_S^\perp)^T A_G I_T^\perp$. Since this matrix is symmetric, its eigenvalues are bounded in magnitude by the operator norm $\|(I_S^\perp)^T A_G I_T^\perp\|$, which is bounded above by $|\mu_{k+1}| \|I_S^\perp\| \|I_T^\perp\|$. The matrices I_S^\perp and I_T^\perp have orthogonal columns, so their operator norm is equal to the norm of any column. Since $\|(I_S^\perp)_i\|^2 + \|\frac{|S|}{n}(I_G)_i\|^2 = \|(I_S)_i\|^2$, we have

$$\|(I_S^\perp)_i\| = \sqrt{|S| - \frac{|S|^2}{n^2} n} = \sqrt{|S| \left(1 - \frac{|S|}{n} \right)},$$

and similarly for $\|(I_T^\perp)_i\|$. Substituting these values for the operator norms gives the bound in (4).

The two bounds given in Lemma 3.1 are incomparable. The spectral bound (4) implies a weaker inequality on $\text{tr}(E(S, T))$ than (3) gives. On the other hand, the trace bound implies weaker constraints on the eigenvalues of $E(S, T) - \frac{d|S||T|}{n} I_{k \times k}$ than the spectral

bound does. The second bound is perhaps the most interesting, as it is not directly implied by a reduction of (G, W) to a scalar-weighted graph.

One interpretation of the standard expander mixing lemma is that for a d -regular graph with small $|\mu_2|$, the number of edges between two subsets is not far from the expected number of edges between two such subsets in a random d -regular graph. Similarly, the matrix-weighted expander mixing lemma says that d -regular matrix-weighted graphs with small $|\mu_2|$ have properties similar to those of a random d -regular graph with matrix weights $I_{k \times k}$.

The name “expander mixing lemma” arises from the use of this result to prove theorems about mixing times of random walks on regular graphs. While it is possible to construct stochastic processes that might justly be termed “random walks” associated with matrix-weighted graphs (and cellular sheaves in general), Lemma 3.1 does not seem to have much relevance to their behavior. It may be that this lemma does control the behavior of other sorts of dynamical processes on a matrix-weighted graph—perhaps the spread of information under a diffusion-like process.

3.1. Irregular matrix-weighted graphs

The standard expander mixing lemma has an extension to non-regular graphs. Like isoperimetric inequalities for irregular graphs, it replaces the simple count of vertices in a subset with the *volume* of the subset: the sum of degrees of those vertices. That is, $\text{vol}(S) = \sum_{s \in S} d_s$. The irregular expander mixing lemma for a scalar-weighted graph G is then captured in the formula

$$\left| E(S, T) - \frac{\text{vol}(S) \text{vol}(T)}{\text{vol}(G)} \right| \leq |\tilde{\mu}_2| \sqrt{\text{vol}(S) \text{vol}(T) \left(1 - \frac{\text{vol}(S)}{\text{vol}(G)}\right) \left(1 - \frac{\text{vol}(T)}{\text{vol}(G)}\right)},$$

where $|\tilde{\mu}_2|$ is the magnitude of the largest nontrivial eigenvalue of the normalized adjacency matrix $\tilde{A} = D^{-1/2}AD^{-1/2}$ of G .

For a matrix-weighted graph, we define the volume of a set S of vertices similarly:

$$\text{vol}(S) = \sum_{s \in S} D_s = \sum_{s \in S} \sum_{s \leq e} W_e.$$

Lemma 3.2. *Let (G, W) be a matrix-weighted graph with n vertices and $k \times k$ weight matrices. If S and T are subsets of the vertices of G , then*

$$|\operatorname{tr}(E(S, T) - V(S, T))| \leq |\mu_{k+1}| \sqrt{\operatorname{tr}(\operatorname{vol}(S) - V(S, S)) \operatorname{tr}(\operatorname{vol}(T) - V(T, T))}, \quad (6)$$

where $V(A, B) = \operatorname{vol}(A) \operatorname{vol}(G)^{-1} \operatorname{vol}(B)$ and $1 = \tilde{\mu}_1 = \dots = \tilde{\mu}_k \geq |\tilde{\mu}_{k+1}| \geq \dots$ are the eigenvalues of the normalized adjacency matrix \tilde{A}_W of (G, W) ordered by decreasing absolute value.

PROOF. Define the $kn \times k$ matrix ψ whose $k \times k$ blocks consist of the diagonal blocks of $D^{1/2}$. The columns of ψ are all eigenvectors of \tilde{A} with eigenvalue 1. We further define the matrices ψ_S and ψ_T , where the blocks of ψ corresponding to vertices not in S or T have been set to zero. Then we have

$$E(S, T) = I_S^T A I_T = \psi_S^T D^{-1/2} A D^{-1/2} \psi_T = \psi_S^T \tilde{A} \psi_T.$$

We can also calculate $\operatorname{vol}(S)$ and $\operatorname{vol}(T)$ from ψ_S and ψ_T :

$$\operatorname{vol}(S) = I_S^T D I_S = \psi_S^T \psi_S = \psi_S^T \psi.$$

Following the pattern from the proof of the regular expander mixing lemma, we decompose $\psi_S = \psi \operatorname{vol}(G)^{-1} \operatorname{vol}(S) + \psi_S^\perp$. These two terms satisfy a sort of orthogonality:

$$\begin{aligned} (\psi_S^\perp)^T \psi \operatorname{vol}(G)^{-1} \operatorname{vol}(S) &= (\psi_S - \psi \operatorname{vol}(G)^{-1} \operatorname{vol}(S))^T \psi \operatorname{vol}(G)^{-1} \operatorname{vol}(S) \\ &= \operatorname{vol}(S) \operatorname{vol}(G)^{-1} \operatorname{vol}(S) - \operatorname{vol}(S) \operatorname{vol}(G)^{-1} \operatorname{vol}(G) \operatorname{vol}(G)^{-1} \operatorname{vol}(S) = 0. \end{aligned}$$

The individual columns of these two matrices do not satisfy a nice orthogonality relation, however, which means we will only be able to obtain a bound on the trace of $E(S, T)$, not its eigenvalues. We have

$$\begin{aligned} E(S, T) &= (\psi_S)^T \tilde{A} \psi_T \\ &= (\psi \operatorname{vol}(G)^{-1} \operatorname{vol}(S) + \psi_S^\perp)^T \tilde{A} (\psi \operatorname{vol}(G)^{-1} \operatorname{vol}(T) + \psi_T^\perp) \\ &= \operatorname{vol}(S) \operatorname{vol}(G)^{-1} \psi \tilde{A} \psi \operatorname{vol}(G)^{-1} \operatorname{vol}(T) + (\psi_S^\perp)^T \tilde{A} \psi_T^\perp \\ &= \operatorname{vol}(S) \operatorname{vol}(G)^{-1} \operatorname{vol}(T) + (\psi_S^\perp)^T \tilde{A} \psi_T^\perp, \end{aligned}$$

and hence

$$E(S, T) - \text{vol}(S) \text{vol}(G)^{-1} \text{vol}(T) = (\psi_S^\perp)^T \tilde{A} \psi_T^\perp. \quad (7)$$

Taking the trace and absolute value gives

$$\begin{aligned} |\text{tr}(E(S, T) - \text{vol}(S) \text{vol}(G)^{-1} \text{vol}(T))| &\leq |\text{tr}((\psi_S^\perp)^T \tilde{A} \psi_T^\perp)| \\ &\leq \|\psi_S^\perp\|_F \|\tilde{A} \psi_T^\perp\|_F \\ &\leq |\tilde{\mu}_{k+1}| \|\psi_S^\perp\|_F \|\psi_T^\perp\|_F. \end{aligned}$$

The norms in this formula are, e.g.,

$$\begin{aligned} \|\psi_S^\perp\|_F &= \text{tr} \left[(\psi_S - \psi \text{vol}(G)^{-1} \text{vol}(S))^T (\psi_S - \psi \text{vol}(G)^{-1} \text{vol}(S)) \right] \\ &= \text{tr} \left[\text{vol}(S) - \text{vol}(S) \text{vol}(G)^{-1} \text{vol}(S) \right]. \end{aligned}$$

Combining these calculations gives the inequality (6).

In the case that G is actually regular, this inequality is looser than (3). It amounts to replacing, e.g., $\sum_{i=1}^k \mu_{k+i}$ with $k |\mu|_{k+1}$ in that formula.

4. Isoperimetric Inequalities

The expander mixing lemma is one canonical inequality comparing combinatorial measures of expansion (the density of edges between two subsets of vertices) with spectral measures of expansion (the largest nontrivial eigenvalue of the adjacency matrix). Another important inequality is the Cheeger inequality, which connects the Cheeger constant of a graph with the second eigenvalue of the (normalized) Laplacian. Letting $h(S) = \frac{E(S, V \setminus S)}{\min(\text{vol}(S), \text{vol}(V \setminus S))}$ and $h_G = \min_S h(S)$, the Cheeger inequality states that

$$\frac{\tilde{\lambda}_2}{2} \leq h_G \leq \sqrt{2\tilde{\lambda}_2}, \quad (8)$$

where $\tilde{\lambda}_2$ is the second-smallest eigenvalue of the normalized Laplacian of G [23, ch. 2]. This is known as an isoperimetric inequality, due to the analogy with the classical problem of controlling the perimeter of a subset of \mathbb{R}^2 in terms of its area. Here, the perimeter is represented by the (weighted) number of edges leaving a subset of vertices,

while the area of that subset is given by the sum of vertex degrees. In the case of a d -regular graph, this is simply proportional to the number of vertices.

A generalization of the Cheeger constant to matrix-weighted graphs is most straightforward for dI -regular weightings, as this simplifies the interpretation of the denominator. The correct generalization of this ratio is unclear for irregular graphs. For a subset S of vertices of a dI -regular matrix-weighted graph, we define two Cheeger ratios:

$$h^{\text{tr}}(S) = \frac{\text{tr } E(S, V \setminus S)}{d \min(|S|, |V \setminus S|)} \quad (9)$$

$$h^{\leq}(S) = \frac{E(S, V \setminus S)}{d \min(|S|, |V \setminus S|)}. \quad (10)$$

These lead to two Cheeger constants

$$h_G^{\text{tr}} = \min_{S \subseteq V} h^{\text{tr}}(S) \quad (11)$$

$$h_G^{\leq} = \inf_{S \subseteq V} h^{\leq}(S). \quad (12)$$

This second Cheeger constant is defined as an infimum in the set of symmetric positive semidefinite matrices under the Loewner order, where $A \leq B$ if $B - A$ is positive semidefinite. Since this is only a partial order, there may not exist a set S of vertices such that $h_G^{\leq} = h^{\leq}(S)$.

Proposition 4.1. *Let (G, W) be a dI -regular matrix-weighted graph with $k \times k$ weight matrices. Then*

$$h_G^{\text{tr}} \geq \frac{1}{2d} \sum_{i=1}^k \lambda_{k+i} \quad (13)$$

$$h_G^{\leq} \geq \frac{\lambda_{k+1}}{2d} I, \quad (14)$$

where $0 = \lambda_1 = \dots = \lambda_k \leq \lambda_{k+1} \leq \dots$ are the eigenvalues of the Laplacian of (G, W) .

PROOF. The first inequality is a direct consequence of the relationship between (G, W) and $(G, \text{tr } W)$ given in Proposition 2.2. Since $\text{tr}(E(S, V \setminus S))$ is equal to the total weight of edges between S and $V \setminus S$ in $(G, \text{tr } W)$, we apply the standard Cheeger bound to obtain, for every S , $h^{\text{tr}}(S) \geq \frac{1}{2d} \lambda_2(\text{tr } W)$. We then apply the relation $\lambda_2(\text{tr } W) \geq \sum_{i=1}^k \lambda_{k+i}$ to obtain the bound.

The second bound is only slightly more involved. For a vertex subset S of G , we let $x^S \in \mathbb{R}^V$ be the vector with

$$x_v^S = \begin{cases} |V \setminus S| & v \in S \\ |S| & v \notin S \end{cases}.$$

Then x is orthogonal to the constant vector $\mathbf{1}$ and if $|S| < |V \setminus S|$,

$$\frac{(x \otimes I)^T L(x \otimes I)}{x^T x} = \frac{E(S, V \setminus S)}{|S||V \setminus S|} \leq 2dh^\leq(S).$$

Meanwhile, the Courant-Fischer theorem implies that for any $y \in \mathbb{R}^V$ orthogonal to $\mathbf{1}$,

$$\lambda_{k+1} I \leq \frac{(y \otimes I)^T L(y \otimes I)}{\|y\|^2}.$$

Taking the infimum over the relevant sets, we then have

$$\frac{\lambda_{k+1}}{2d} I \leq \frac{1}{2d} \inf_{y \perp \mathbf{1}} \frac{(y \otimes I)^T L(y \otimes I)}{\|y\|^2} \leq \inf_{S \subset V} \frac{1}{2d} \frac{(x^S \otimes I)^T L(x^S \otimes I)}{\|x^S\|^2} \leq h_G^\leq.$$

These bounds correspond to the easy-to-prove side of the standard Cheeger inequality. Unfortunately, analogous upper bounds on h_G in terms of the spectrum of L do not exist. Specifically, there are no upper bounds of the form $h_G^{\text{tr}} \leq f(\lambda_{2k})$, where $f(0) = 0$, nor of the form $h_G^\leq \leq F(\lambda_{2k})$, where F is the zero matrix when $\lambda_{2k} = 0$. To see this, consider the matrix-weighted graph G in Figure 1. The weight matrices correspond to the edge labels as follows:

$$a : \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad b : \begin{bmatrix} \frac{1}{4} & \frac{\sqrt{3}}{4} \\ \frac{\sqrt{3}}{4} & \frac{3}{4} \end{bmatrix} \quad c : \begin{bmatrix} \frac{1}{4} & -\frac{\sqrt{3}}{4} \\ -\frac{\sqrt{3}}{4} & \frac{3}{4} \end{bmatrix}. \quad (15)$$

This graph is regular and has algebraic degree $\frac{3}{2}$. Any two of these weight matrices sum to a full-rank matrix, and removing any set of edges with the same weights leaves a connected graph. Therefore, for any set S of vertices of G , $E(S, V \setminus S)$ is full rank. Hence, $h_G^\leq \geq \alpha I$ for some $\alpha > 0$ and $h_G^{\text{tr}} > 0$. However, we can calculate that the zero eigenvalue of the Laplacian of G has multiplicity four, so $\lambda_{2k} = 0$, meaning that our putative spectral upper bound on h_G must be zero. The conclusion to be drawn is that unlike the case for scalar-weighted graphs, combinatorial measures of expansion in matrix-weighted graphs are in general weaker than spectral measures of

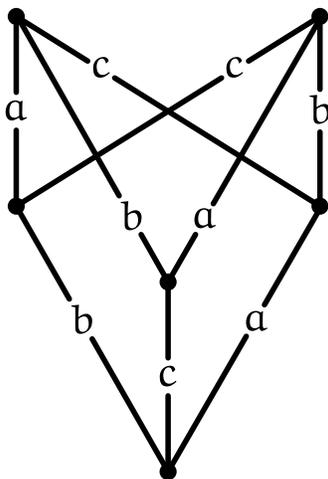


Figure 1: A counterexample to a spectral upper bound on the matrix-weighted Cheeger constants

expansion. One cannot ensure that eigenvalues of the matrix-weighted Laplacian are bounded away from zero by controlling a Cheeger constant (at least one of the form we have considered here).

5. Expander Sheaves

These expansion-related bounds for matrix-weighted graphs suggest that we attempt to generalize expander graphs to the matrix-weighted setting. Expander graphs are typically defined as unweighted graphs, so a generalization allowing matrix weights may seem slightly contradictory. However, many constructions of expander graphs end up producing graphs which may have multiple edges between a pair of vertices, which amounts to allowing positive integer weights. One may think of this as allowing a sequence of combinatorial decisions about where to place edges in the graph. We extend this to the matrix-weighted setting by adding an extra choice: that of a subspace of \mathbb{R}^k for each edge. Such a subspace might be generated by iteratively choosing atomic elements of the lattice of subspaces of \mathbb{R}^k .

A precise definition is as follows:

Definition 5.1. Let (G, W) be a d -regular matrix-weighted graph. We say that it is a

matrix-weighted η -expander if all its weight matrices are orthogonal projections $\mathbb{R}^k \rightarrow \mathbb{R}^k$ and all nontrivial eigenvalues of its adjacency matrix are at most $d - \eta$ in magnitude.

There is the immediate question of how to construct a regular matrix-weighted graph with projection-valued weights, regardless of its spectral properties. The trivial construction is obvious: take a regular unweighted graph, and assign each edge the identity matrix. A more interesting approach is to note that the condition that $D_v = \sum_{v \preceq e} W_e = dI$ is the same as the condition for the relevant matrices W_e to form a *tight fusion frame* with frame constant d . Fusion frames are a generalization of the notion of frame from harmonic analysis [24]. They are typically defined as collections of subspaces of $V_i \leq \mathbb{R}^k$ such that any vector $x \in \mathbb{R}^k$ is uniquely determined by its projections onto V_i for all i . Equivalently, a fusion frame may be defined as a collection of orthogonal projections on \mathbb{R}^k that sum to an invertible operator. Tight fusion frames are those for which these orthogonal projections sum to a scalar multiple of the identity.

It is a nontrivial result that tight fusion frames exist [25]. In particular, for $r \geq \lceil k/\ell \rceil + 2$, there exists a tight fusion frame in \mathbb{R}^k consisting of r subspaces of dimension ℓ , while for $r \leq \lceil k/\ell \rceil$, no tight fusion frames of this form exist.

We can use a nontrivial fusion frame to construct nontrivial matrix-weighted graphs with projection-valued weights. Let G be an r -regular graph with an r -edge coloring, and take a tight fusion frame in \mathbb{R}^k with r subspaces of dimension ℓ . Assign one element of the fusion frame to each edge color of G ; these will become the matrix weights. The resulting matrix-weighted graph has degree $r\ell/k$. Note that this degree may not be an integer.

A matrix-weighted graph constructed in this way need not have any particular expansion properties. Indeed, its Laplacian may have a large kernel. However, nontrivial individual examples of these matrix-weighted expanders do exist. Consider the graph shown in Figure 2. The underlying graph is 4-regular, and is 4-edge colored. The weights are given by the matrices in (15), with d corresponding to the identity matrix. Thus, the four-element fusion frame used is given by three one-dimensional subspaces in \mathbb{R}^2 together with \mathbb{R}^2 itself. The resulting matrix-weighted graph is regular, with algebraic degree $5/2$. Numerical calculations show that the nontrivial adjacency eigen-

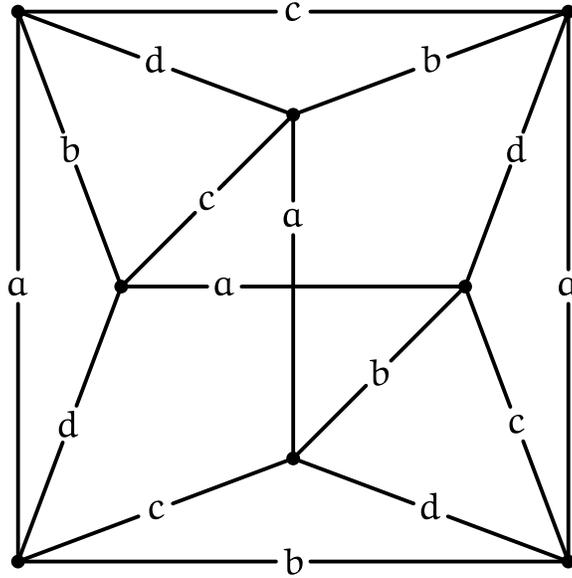


Figure 2: A nontrivial matrix-weighted expander

values of this graph lie between -2.406 and 1.803 , giving it a two-sided expansion constant of $\eta = 0.094$. While this particular expansion constant is nothing to write home about, significantly better expansion may be possible in general.

The Alon-Boppana bound [26] gives a constraint on the spectral expansion of an infinite family of graphs. The second adjacency eigenvalue μ_2 of a d -regular graph is bounded below by $2\sqrt{d-1} - o(1)$. Is there a similar bound for matrix-weighted graphs? Take a r -regular graph with $k \times k$ matrix weights which are orthogonal projections of rank ℓ , and hence has matrix-degree $r\ell/k$. If we take the trace of weights, we get a scalar-weighted graph whose edge weights are all ℓ . Its adjacency matrix is ℓ times the adjacency matrix of the underlying graph. The Laplacian trace bound (2.2) implies that $k\mu_2(A_W) \geq \ell\mu_2(A_G)$, so

$$\mu_2(A_W) \geq 2\frac{\ell}{k}\sqrt{r-1} - o(1).$$

The algebraic degree of this matrix-weighted graph is $d = r\ell/k$, so the bound is $\mu_2(A_W) \geq 2d\sqrt{r-1}/r$. For $2 < d < r$, $\sqrt{r-1}/r \leq \sqrt{d-1}/d$, and so $2d\sqrt{r-1}/r \leq 2\sqrt{d-1}$. Since this bound is less restrictive on μ_2 , it may be possible for a family of matrix-weighted expander graphs to exhibit better-than-Ramanujan expansion for a

given algebraic degree. To be clear, we have not here shown that this is the case; we have only failed to rule it out using one line of argument. However, other approaches to extending the Alon-Boppana bound to matrix-weighted graphs give the same results.

Such a property may be useful for the design of communications networks. Expander graphs were initially introduced in part to study the design of fault-tolerant networks. They have since found use in the design of distributed consensus algorithms. The convergence rate of the algorithm depends on the spectral properties of the network, and hence Ramanujan graphs give optimal networks for a given amount of local connectivity. The algebraic degree of a matrix-weighted expander represents the total amount of communication a node must carry on with its neighbors in order to advance another step in the algorithm. Better expansion constants for a given algebraic degree mean faster convergence for the same amount of communication.

6. Conclusion

Matrix-weighted graphs are an expressive generalization of undirected graphs, and expand the concern of spectral graph theory to operators acting on higher-dimensional spaces of functions. Expansion in matrix-weighted graphs has more subtle behavior than in standard graphs. We have shown that spectral measures of expansion control combinatorial measures of expansion, as in the expander mixing lemma and one side of the Cheeger inequality. However, we do not have a converse combinatorial condition for a matrix-weighted graph to have good spectral expansion.

There is a converse to the expander mixing lemma for scalar-weighted graphs [27]. Its proof was a byproduct of a construction of families of expander graphs with nearly optimal spectral expansion. It would be interesting to know whether a converse similarly holds for matrix-weighted graphs. This would offer some level of control over the spectral properties of matrix-weighted graphs in terms of a combinatorial measure of expansion. The failure to exist of a spectral upper bound on the Cheeger constant suggests that a converse to the expander mixing lemma may be similarly false.

The problem of constructing infinite families of matrix-weighted expanders offers many interesting challenges. Standard methods for constructing expander graphs do

not readily generalize to the matrix-weighted case. Even the problem of choosing kernels of weights so that the Laplacian kernel has dimension k —what in the sheaf theoretic language might be termed an “approximation to the constant sheaf”—is a subtle problem. Solving these combinatorial problems will require insights about graphs, lattices of subspaces, and fusion frames.

Other aspects of the theory of matrix-weighted graphs are also interesting. Some results may be obtained as specializations of results in the spectral theory of cellular sheaves [15], but this frontier has only just begun to be explored. Combining the broad perspective of cellular sheaves with the particular properties of matrix-weighted graphs gives a fruitful field for further investigation.

- [1] O. Parzanchevski, High Dimensional Expanders, Ph.D. thesis, Hebrew University of Jerusalem (2013).
- [2] J. Steenbergen, Towards a Spectral Theory for Simplicial Complexes, Ph.D. thesis, Duke University (2013).
- [3] J. Cooper, A. Dutle, Spectra of uniform hypergraphs, *Linear Algebra and its Applications* 436 (9) (2012) 3268–3292.
- [4] A. Louis, Hypergraph Markov Operators, Eigenvalues and Approximation Algorithms, in: *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, STOC '15*, Association for Computing Machinery, Portland, Oregon, USA, 2015, pp. 713–722.
- [5] A. Singer, H.-T. Wu, Vector Diffusion Maps and the Connection Laplacian, *Communications in Pure and Applied Mathematics* 65 (8).
- [6] H.-T. Wu, Embedding Riemannian manifolds by the heat kernel of the connection Laplacian, *Advances in Mathematics* 304 (2017) 1055–1079.
- [7] A. S. Bandeira, A. Singer, D. A. Spielman, A Cheeger inequality for the graph connection Laplacian, *SIAM Journal on Matrix Analysis and Applications* 34 (4) (2013) 1611–1630.

- [8] W. Zhao, M. Kempton, F. Chung, Ranking and sparsifying a connection graph, *Internet Mathematics* 10 (1) (2014) 1553.
- [9] R. Kyng, Y. T. Lee, R. Peng, S. Sachdeva, D. A. Spielman, Sparsified Cholesky and multigrid solvers for connection Laplacians, in: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing, STOC '16*, Association for Computing Machinery, Cambridge, MA, USA, 2016, pp. 842–850.
- [10] N. El Karoui, H.-t. Wu, Graph connection Laplacian and random matrices with random blocks, *Information and Inference: A Journal of the IMA* 4 (1) (2015) 1–44.
- [11] S. E. Tuna, Synchronization under matrix-weighted Laplacian, *Automatica* 73 (2016) 76–81.
- [12] S. E. Tuna, Observability Through a Matrix-Weighted Graph, *IEEE Transactions on Automatic Control* 63 (7) (2018) 2061–2074.
- [13] M. H. Trinh, C. Van Nguyen, Y.-H. Lim, H.-S. Ahn, Matrix-weighted consensus and its applications, *Automatica* 89 (2018) 415–419.
- [14] J. Curry, Sheaves, cosheaves, and applications, Ph.D. thesis, University of Pennsylvania (2014).
- [15] J. Hansen, R. Ghrist, Toward a spectral theory of cellular sheaves, *Journal of Applied and Computational Topology* 3 (4) (2019) 315–358. arXiv:1808.01513.
- [16] M. Kashiwara, P. Schapira, Sheaves on Manifolds, no. 292 in *Grundlehren Der Mathematischen Wissenschaften*, Springer-Verlag Berlin Heidelberg, 1990.
- [17] R. Hartshorne, *Algebraic Geometry*, Vol. 52 of *Graduate Texts in Mathematics*, Springer New York, New York, NY, 1977.
- [18] J. Friedman, Sheaves on graphs, their homological invariants, and a proof of the Hanna Neumann conjecture, *Memoirs of the American Mathematical Society* 233 (1100).

- [19] J. Hansen, Laplacians of Cellular Sheaves: Theory and Applications, Ph.D. Thesis, University of Pennsylvania (2020).
- [20] F. R. K. Chung, S. Sternberg, Laplacian and vibrational spectra for homogeneous graphs, *Journal of Graph Theory* 16 (6) (1992) 605–627.
- [21] F. Atik, R. B. Bapat, M. Rajesh Kannan, Resistance matrices of graphs with matrix weights, *Linear Algebra and its Applications* 571 (2019) 41–57.
- [22] N. Alon, F. R. K. Chung, Explicit construction of linear sized tolerant networks, *Discrete Mathematics* 72 (1) (1988) 15–19.
- [23] F. Chung, *Spectral Graph Theory*, AMS, 1992.
- [24] P. G. Casazza, G. Kutyniok, F. Philipp, Introduction to Finite Frame Theory, in: P. G. Casazza, G. Kutyniok (Eds.), *Finite Frames*, Birkhäuser Boston, Boston, 2013, pp. 1–53.
- [25] P. G. Casazza, M. Fickus, D. G. Mixon, Y. Wang, Z. Zhou, Constructing tight fusion frames, *Applied and Computational Harmonic Analysis* 30 (2) (2011) 175–187.
- [26] A. Nilli, On the second eigenvalue of a graph, *Discrete Mathematics* 91 (2) (1991) 207–210.
- [27] Y. Bilu, N. Linial, Lifts, discrepancy and nearly optimal spectral gap, *Combinatorica* 26 (5) (2006) 495–519.