

# Toward a Spectral Theory of Cellular Sheaves

Jakob Hansen · Robert Ghrist

Received: date / Accepted: date

**Abstract** This paper outlines a program in what one might call *spectral sheaf theory* — an extension of spectral graph theory to cellular sheaves. By lifting the combinatorial graph Laplacian to the Hodge Laplacian on a cellular sheaf of vector spaces over a regular cell complex, one can relate spectral data to the sheaf cohomology and cell structure in a manner reminiscent of spectral graph theory. This work gives an exploratory introduction, and includes results on eigenvalue interlacing, complex sparsification, effective resistance, synchronization, and sheaf approximation. These results and subsequent applications are prefaced by an introduction to cellular sheaves and Laplacians.

**Keywords** Cohomology · Cellular sheaf theory · Spectral graph theory · Effective resistance · Eigenvalue interlacing

**Mathematics Subject Classification (2000)** MSC 55N30 · MSC 05C50

## 1 Introduction

In spectral graph theory, one associates to a combinatorial graph additional algebraic structures in the form of square matrices whose spectral data is

---

J. Hansen  
Department of Mathematics  
University of Pennsylvania  
David Rittenhouse Lab.  
209 South 33rd Street  
Philadelphia, PA 19104-6395  
E-mail: jhansen@math.upenn.edu

R. Ghrist  
Departments of Mathematics and Electrical & Systems Engineering  
University of Pennsylvania  
David Rittenhouse Lab.  
209 South 33rd Street  
Philadelphia, PA 19104-6395  
E-mail: ghris@math.upenn.edu

then investigated and related to the graph. These matrices come in several variants, most particularly degree and adjacency matrices, Laplacian matrices, and weighted versions thereof. In most cases, the size of the implicated matrix is based on the vertex set, while the structure of the matrix encodes data carried by the edges.

To say that spectral graph theory is useful is an understatement. Spectral methods are key in such disparate fields as data analysis [BN03, CL06], theoretical computer science [HLW06, CS11], probability theory [LP16], control theory [Bull8], numerical linear algebra [ST14], coding theory [Spi96], and graph theory itself [Chu92, BH12].

Much of spectral graph theory focuses on the Laplacian, leveraging its unique combination of analytic, geometric, and probabilistic interpretations in the discrete setting. This is not the complete story. Many of the most well-known and well-used results on the spectrum of the graph Laplacian return features that are neither exclusively geometric nor even combinatorial in nature, but rather more qualitative. For example, it is among the first facts of spectral graph theory that the multiplicity of the zero eigenvalue of the graph Laplacian enumerates connected components of the graph, and the relative size of the smallest nonzero eigenvalue in a connected graph is a means of approximate dis-connectivity. Such features are topological, and it is here that the present inquiry begins.

There is another branch of mathematics in which Laplacians hold sway: Hodge theory. This is the slice of algebraic geometry that uses Laplacians on (complex) Riemannian manifolds to characterize global features. The classical initial result is that one recovers the cohomology of the manifold as the kernel of the Laplacian on differential forms [AMR88]. For example, the kernel of the Laplacian on 0-forms ( $\mathbb{R}$ -valued functions) returns the rank of  $H^0$ , the 0-th cohomology group (with coefficients in  $\mathbb{R}$ ), whose dimension the number of connected components. In spirit, then, Hodge theory categorifies elements of spectral graph theory.

Hodge theory, like much of algebraic topology, survives the discretization from Riemannian manifolds to (weighted) cell complexes [Eck45, Fri98]. The classical boundary operator for a cell complex and its formal adjoint combine to yield a generalization of the graph Laplacian which, like the Laplacian of Hodge theory, acts on higher dimensional objects (cellular cochains, as opposed to differential forms). The kernel of this discrete Laplacian is isomorphic to the cellular cohomology of the complex with coefficients in the reals, generalizing the connectivity detection of the graph Laplacian in grading zero. As such, the spectral theory of the discrete Laplacian offers a geometric perspective on algebraic-topological features of higher-dimensional complexes. This is an active area of research [Par13, Ste13, HJ13].

This is not the end. Our aim is a generalization of both spectral graph theory and discrete Hodge theory which ties in to recent developments in topological data analysis. The past two decades have witnessed a burst of activity in computing the homology of cell complexes (and sequences thereof) to extract robust global features, leading to the development of specialized

tools, such as persistent homology, barcodes, and more, as descriptors for cell complexes [Car13, EH10, KMM04, OPT+17].

Topological data analysis is evolving rapidly. One particular direction of evolution concerns a change in perspective from working with cell complexes as topological spaces in and of themselves to focusing instead on data over a cell complex — viewing the cell complex as a base on which data to be investigated resides. For example, one can consider scalar-valued data over cell complexes, as occurs in weighted networks and complexes; or sensor data, as occurs in target enumeration problems [CGR12]. Richer data involves vector spaces and linear transformations, as with recent work in cryo-EM [HS11] and synchronization problems [Ban15]. Recent work in TDA points to the appropriate generalization of these and related data structures over topological spaces. This is the theory of *sheaves*.

We will work exclusively with cellular sheaves [Cur14]. Fix a (regular, locally finite) cell complex — a triangulated surface will suffice for purposes of imagination. A cellular sheaf of vector spaces is, in essence, a data structure on this domain, associating local data (in the form of vector spaces) to cells and compatibility relations (linear transformations) to cells of incident ascending dimension. Data over vertices is sent to data over incident edges; data over edges is sent to data over incident 2-cells, etc. As a trivial example, the constant sheaf assigns a rank-one vector space to each cell and identity isomorphisms according to boundary faces. More interesting is the cellular analogue of a vector bundle: a cellular sheaf which assigns a fixed vector space of dimension  $n$  to each cell and isomorphisms as linear transformations (with specializations to  $O(n)$  or  $SO(n)$  as desired).

The data assigned to a cellular sheaf naturally arranges into a cochain complex. As such, cellular sheaves possess a Laplacian that specializes to the graph Laplacian and the Hodge Laplacian for the constant sheaf. For cellular sheaves of real vector spaces, a spectral theory — an examination of the eigenvalues and eigenvectors of the sheaf Laplacian — is natural, motivated, and, to date, unexamined apart from a few special cases (see §3.5).

This paper sketches an emerging spectral theory for cellular sheaves. Given the motivation as a generalization of spectral graph theory, we will often specialize to cellular sheaves over a 1-dimensional cell complex (that is, a graph, occasionally with loops or multiple edges between vertex pairs). This is mostly for the sake of simplicity and initial applications, as zero- and one-dimensional homological invariants are the most readily applicable. However, as the theory is general, we occasionally point to higher-dimensional side-quests.

The plan of this paper is as follows. In §2, we cover the necessary topological and algebraic preliminaries, including definitions of cellular sheaves. Next, §3 gives definitions of the various matrices involved in the extension of spectral theory to cellular sheaves. Section 4 uses these to explore issues related to harmonic functions and cochains on sheaves. In §5, we extend various elementary results from spectral graph theory to cellular sheaves. The subsequent two sections treat more sophisticated topics, effective resistance (§6) and the Cheeger inequality (§7), for which we have some preliminary results. We conclude with

outlines of potential applications for the theory in §8 and directions for future inquiry in §9.

The results and applications we sketch are at the beginnings of the subject, and a great deal more in way of fundamental and applied work remains.

This paper has been written in order to be readable without particular expertise in algebraic topology beyond the basic ideas of cellular homology and cohomology. Category-theoretic terminology is used sparingly and for concision. Given the well-earned reputation of sheaf theory as difficult for the non-specialist, we have provided an introductory section with terminology and core concepts, noting that much more is available in the literature [Bre97, KS90]. Our recourse to the cellular theory greatly increases simplicity, readability, and applicability, while resonating with the spirit of spectral graph theory. There are abundant references available for the reader who requires more information on algebraic topology [Hat01], applications thereof [EH10, Ghr14], and cellular sheaf theory [Cur14, Ghr14].

## 2 Preliminaries

### 2.1 Cell complexes

**Definition 1.** A *regular cell complex* is a topological space  $X$  with a decomposition into subspaces  $\{X_\alpha\}_{\alpha \in P_X}$  satisfying the following conditions:

1. For each  $x \in X$ , every sufficiently small neighborhood of  $x$  intersects finitely many  $X_\alpha$ .
2. For all  $\alpha, \beta$ ,  $\overline{X_\alpha} \cap X_\beta \neq \emptyset$  only if  $X_\beta \subseteq \overline{X_\alpha}$ .
3. Every  $X_\alpha$  is homeomorphic to  $\mathbb{R}^{n_\alpha}$  for some  $n_\alpha$ .
4. For every  $\alpha$ , there is a homeomorphism of a closed ball in  $\mathbb{R}^{n_\alpha}$  to  $\overline{X_\alpha}$  that maps the interior of the ball homeomorphically onto  $X_\alpha$ .

Condition (2) implies that the set  $P_X$  has a poset structure, given by  $\beta \leq \alpha$  iff  $X_\beta \subseteq \overline{X_\alpha}$ . This is known as the *face poset* of  $X$ . The regularity condition (4) implies that all topological information about  $X$  is encoded in the poset structure of  $P_X$ . For our purposes, we will identify a regular cell complex with its face poset, writing the incidence relation  $\beta \trianglelefteq \alpha$ . The class of posets that arise in this way can be characterized combinatorially [Bjö84]. For our purposes, a morphism of cell complexes is a morphism of posets between their face incidence posets that arises from a continuous map between their associated topological spaces. In particular, morphisms of simplicial and cubical complexes will qualify as morphisms of regular cell complexes.

The class of regular cell complexes includes simplicial complexes, cubical complexes, and so-called multigraphs (as 1-dimensional cell complexes). As nearly every space that can be characterized combinatorially can be represented as a regular cell complex, these will serve well as a default class of spaces over which to develop a combinatorial spectral theory of sheaves. We note that the spectral theory of complexes has heretofore been largely restricted to the

study of simplicial complexes [SB+18]. A number of our results will specialize to results about the spectra of regular cell complexes by consideration of the constant sheaf.

## 2.2 Cellular sheaves

Let  $X$  be a regular cell complex. A cellular sheaf attaches data spaces to the cells of  $X$  together with relations that specify when assignments to these data spaces are consistent.

**Definition 2.** A *cellular sheaf* of vector spaces on a regular cell complex  $X$  is an assignment of a vector space  $\mathcal{F}(\sigma)$  to each cell  $\sigma$  of  $X$  together with a linear transformation  $\mathcal{F}_{\sigma \triangleleft \tau}: \mathcal{F}(\sigma) \rightarrow \mathcal{F}(\tau)$  for each incident cell pair  $\sigma \triangleleft \tau$ . These must satisfy both an identity relation  $\mathcal{F}_{\sigma \triangleleft \sigma} = \text{id}$  and the composition condition:

$$\rho \triangleleft \sigma \triangleleft \tau \Rightarrow \mathcal{F}_{\rho \triangleleft \tau} = \mathcal{F}_{\sigma \triangleleft \tau} \circ \mathcal{F}_{\rho \triangleleft \sigma}.$$

The vector space  $\mathcal{F}(\sigma)$  is called the *stalk* of  $\mathcal{F}$  at  $\sigma$ . The maps  $\mathcal{F}_{\sigma \triangleleft \tau}$  are called the *restriction maps*.

For experts, this definition at first seems only reminiscent of the notion of sheaves familiar to topologists. The depth of the relationship is explained in detail in [Cur14], but the essence is this: the data of a cellular sheaf on  $X$  specifies spaces of local sections on a cover of  $X$  given by open stars of cells. This translates in two different ways into a genuine sheaf on a topological space. One may either take the Alexandrov topology on the face incidence poset of the complex, or one may view the open stars of cells and their natural refinements a basis for the topology of  $X$ . There then exists a natural completion of the data specified by the cellular sheaf to a constructible sheaf on  $X$ .

One may compress the definition of a cellular sheaf to the following: If  $X$  is a regular cell complex with face incidence poset  $P_X$ , viewed as a category, a cellular sheaf is a functor  $\mathcal{F}: P_X \rightarrow \mathbf{Vect}_{\mathbb{k}}$  to the category of vector fields over a field  $\mathbb{k}$ .

**Definition 3.** Let  $\mathcal{F}$  be a cellular sheaf on  $X$ . A *global section*  $x$  of  $\mathcal{F}$  is a choice  $x_\sigma \in \mathcal{F}(\sigma)$  for each cell  $\sigma$  of  $X$  such that  $x_\tau = \mathcal{F}_{\sigma \triangleleft \tau} x_\sigma$  for all  $\sigma \triangleleft \tau$ . The space of global sections of  $\mathcal{F}$  is denoted  $\Gamma(X; \mathcal{F})$ .

Perhaps the simplest sheaf on any complex is the constant sheaf with stalk  $\mathbb{V}$ , which we will denote  $\underline{\mathbb{V}}$ . This is the sheaf with all stalks equal to  $\mathbb{V}$  and all restriction maps equal to the identity.

### 2.2.1 Cosheaves

In many situations it is more natural to consider a dual construction to a cellular sheaf. A *cellular cosheaf* preserves stalk data but reverses the direction of the face poset, and with it, the restriction maps.

**Definition 4.** A cellular cosheaf of vector spaces on a regular cell complex  $X$  is an assignment of a vector space  $\mathcal{F}(\sigma)$  to each cell  $\sigma$  of  $X$  together with linear maps  $\mathcal{F}_{\sigma \triangleleft \tau} : \mathcal{F}(\tau) \rightarrow \mathcal{F}(\sigma)$  for each incident cell pair  $\sigma \triangleleft \tau$  which satisfies the identity ( $\mathcal{F}_{\sigma \triangleleft \sigma} = \text{id}$ ) and composition condition:

$$\rho \triangleleft \sigma \triangleleft \tau \Rightarrow \mathcal{F}_{\rho \triangleleft \tau} = \mathcal{F}_{\rho \triangleleft \sigma} \circ \mathcal{F}_{\sigma \triangleleft \tau}.$$

More concisely, a cellular cosheaf is a functor  $P_X^{\text{op}} \rightarrow \mathbf{Vect}_{\mathbb{k}}$ . The contravariant functor  $\text{Hom}(\bullet, \mathbb{k}) : \mathbf{Vect}_{\mathbb{k}}^{\text{op}} \rightarrow \mathbf{Vect}_{\mathbb{k}}$  means that every cellular sheaf  $\mathcal{F}$  has a dual cosheaf  $\hat{\mathcal{F}}$  whose stalks are  $\text{Hom}(\mathcal{F}(\sigma), \mathbb{k})$ .

### 2.2.2 Homology and Cohomology

The cells of a regular cell complex have a natural grading by dimension. By regularity of the cell complex, this grading can be extracted from the face incidence poset as the height of a cell in the poset. This means that a cellular sheaf has a graded vector space of cochains

$$C^k(X; \mathcal{F}) = \bigoplus_{\dim(\sigma)=k} \mathcal{F}(\sigma).$$

To develop this into a chain complex, we need a boundary operator and a notion of orientation — a signed incidence relation on  $P_X$ . This is a map  $[\bullet : \bullet] : P_X \times P_X \rightarrow \{0, \pm 1\}$  satisfying the following conditions:

1. If  $[\sigma : \tau] \neq 0$ , then  $\sigma \triangleleft \tau$  and there are no cells between  $\sigma$  and  $\tau$  in the incidence poset.
2. For any  $\sigma \triangleleft \tau$ ,  $\sum_{\gamma \in P_X} [\sigma : \gamma][\gamma : \tau] = 0$ .

Given a signed incidence relation on  $P_X$ , there exist coboundary maps  $\delta^k : C^k(X; \mathcal{F}) \rightarrow C^{k+1}(X; \mathcal{F})$ . These are given by the formula

$$\delta^k|_{\mathcal{F}(\sigma)} = \sum_{\dim(\tau)=k+1} [\sigma : \tau] \mathcal{F}_{\sigma \triangleleft \tau},$$

or equivalently,

$$(\delta^k x)_\tau = \sum_{\dim(\sigma)=k} [\sigma : \tau] \mathcal{F}_{\sigma \triangleleft \tau}(x_\sigma).$$

Here we use subscripts to denote the value of a cochain in a particular stalk; that is,  $x_\sigma$  is the value of the cochain  $x$  in the stalk  $\mathcal{F}(\sigma)$ .

It is simple to show that  $\delta^k \circ \delta^{k-1} = 0$ , so that these coboundary maps define a cochain complex and hence a cohomology of cellular sheaves. In particular,  $H^0(X; \mathcal{F})$  is naturally isomorphic to  $\Gamma(X; \mathcal{F})$ , the space of global sections. An analogous construction defines a homology theory for cosheaves. Cosheaf homology may be thought of as dual to sheaf cohomology in a Poincaré-like sense. That is, frequently the natural analogue of degree zero sheaf cohomology is degree  $n$  cosheaf homology.

There is a relative version of cellular sheaf cohomology. Let  $A$  be a subcomplex of  $X$ . There is a natural subspace of  $C^k(X; \mathcal{F})$  consisting of cochains which vanish on stalks over cells in  $A$ . The coboundary of a cochain which vanishes on  $A$  also vanishes on  $A$ , since any cell in  $A^{(k+1)}$  has only cells in  $A^{(k)}$  on its boundary. Therefore we get a subcomplex  $C^\bullet(X, A; \mathcal{F})$  of  $C^\bullet(X; \mathcal{F})$ . The cohomology of this subcomplex is the relative sheaf cohomology  $H^\bullet(X, A; \mathcal{F})$ .

### 2.2.3 Sheaf morphisms

**Definition 5.** If  $\mathcal{F}$  and  $\mathcal{G}$  are sheaves on a cell complex  $X$ , a sheaf *morphism*  $\varphi : \mathcal{F} \rightarrow \mathcal{G}$  is a collection of maps  $\varphi_\sigma : \mathcal{F}(\sigma) \rightarrow \mathcal{G}(\sigma)$  for each cell  $\sigma$  of  $X$ , such that for any  $\sigma \triangleleft \tau$ ,  $\varphi_\tau \circ \mathcal{F}_{\sigma \triangleleft \tau} = \mathcal{G}_{\sigma \triangleleft \tau} \circ \varphi_\sigma$ . Equivalently, all diagrams of the following form commute:

$$\begin{array}{ccc} \mathcal{F}(\sigma) & \xrightarrow{\varphi_\sigma} & \mathcal{G}(\sigma) \\ \mathcal{F}_{\sigma \triangleleft \tau} \downarrow & & \downarrow \mathcal{G}_{\sigma \triangleleft \tau} \\ \mathcal{F}(\tau) & \xrightarrow{\varphi_\tau} & \mathcal{G}(\tau) \end{array}$$

Functoriality assures that a sheaf morphism  $\varphi : \mathcal{F} \rightarrow \mathcal{G}$  induces maps  $\varphi^k : C^k(X; \mathcal{F}) \rightarrow C^k(X; \mathcal{G})$  which commute with the coboundary maps, resulting in the induced maps on cohomology  $H^k \varphi : H^k(X; \mathcal{F}) \rightarrow H^k(X; \mathcal{G})$ .

### 2.2.4 Sheaf operations

There are several standard operations that act on sheaves to produce new sheaves.

**Definition 6** (Direct sum). If  $\mathcal{F}$  and  $\mathcal{G}$  are sheaves on  $X$ , their *direct sum*  $\mathcal{F} \oplus \mathcal{G}$  is a sheaf on  $X$  with  $(\mathcal{F} \oplus \mathcal{G})(\sigma) = \mathcal{F}(\sigma) \oplus \mathcal{G}(\sigma)$ . The restriction maps are  $(\mathcal{F} \oplus \mathcal{G})_{\sigma \triangleleft \tau} = \mathcal{F}_{\sigma \triangleleft \tau} \oplus \mathcal{G}_{\sigma \triangleleft \tau}$ .

**Definition 7** (Tensor product). If  $\mathcal{F}$  and  $\mathcal{G}$  are sheaves on  $X$ , their *tensor product*  $\mathcal{F} \otimes \mathcal{G}$  is a sheaf on  $X$  with  $(\mathcal{F} \otimes \mathcal{G})(\sigma) = \mathcal{F}(\sigma) \otimes \mathcal{G}(\sigma)$ . The restriction maps are  $(\mathcal{F} \otimes \mathcal{G})_{\sigma \triangleleft \tau} = \mathcal{F}_{\sigma \triangleleft \tau} \otimes \mathcal{G}_{\sigma \triangleleft \tau}$ .

**Definition 8** (Pullback). If  $f : X \rightarrow Y$  is a morphism of cell complexes and  $\mathcal{F}$  is a sheaf on  $Y$ , the *pullback*  $f^* \mathcal{F}$  is a sheaf on  $X$  with  $f^* \mathcal{F}(\sigma) = \mathcal{F}(f(\sigma))$  and  $(f^* \mathcal{F})_{\sigma \triangleleft \tau} = \mathcal{F}_{f(\sigma) \triangleleft f(\tau)}$ .

**Definition 9** (Pushforward). The full definition of the pushforward of a cellular sheaf is somewhat more categorically involved than the previous constructions. If  $f : X \rightarrow Y$  is a morphism of cell complexes and  $\mathcal{F}$  is a sheaf on  $X$ , the *pushforward*  $f_* \mathcal{F}$  is a sheaf on  $Y$  with stalks  $f_* \mathcal{F}(\sigma)$  given as the limit  $\lim_{\sigma \triangleleft f(\tau)} \mathcal{F}(\tau)$ . The restriction maps are induced by the restriction maps of  $\mathcal{F}$ , since whenever  $\sigma \triangleleft \sigma'$ , the cone for the limit defining  $f_* \mathcal{F}(\sigma)$  contains the cone for the limit defining  $f_* \mathcal{F}(\sigma')$ , inducing a unique map  $f_* \mathcal{F}(\sigma) \rightarrow f_* \mathcal{F}(\sigma')$ .

In this paper, we will only work with pushforwards over locally injective cell maps, that is, where the open stars of each cell in  $f^{-1}(\sigma)$  are disjoint. In this case,  $f_* \mathcal{F}(\sigma) \simeq \bigoplus_{\sigma' \in f^{-1}(\sigma)} \mathcal{F}(\sigma')$ , and  $(f_* \mathcal{F})_{\sigma \triangleleft \tau} = \bigoplus_{(\sigma' \triangleleft \tau') \in f^{-1}(\sigma \triangleleft \tau)} \mathcal{F}_{\sigma' \triangleleft \tau'}$ .

Those familiar with the definitions of pushforward and pullback for sheaves over topological spaces will note a reversal of fates when we define sheaves over cell complexes. Here the pullback is simple to define, while the pushforward is more involved. This complication arises because cellular sheaves are in a sense defined pointwise rather than over open sets.

### 3 Definitions

#### 3.1 Weighted cellular sheaves

Let  $\mathbb{k} = \mathbb{R}$  or  $\mathbb{C}$ . A weighted cellular sheaf is a cellular sheaf with values in  $\mathbb{k}$ -vector spaces where the stalks have additionally been given an inner product structure. One may view such as a functor  $P_X \rightarrow \mathbf{Hilb}_{\mathbb{k}}$ .

Because vector space direct sums and tensor products of inner product spaces have natural inner products, the direct sum and tensor product of weighted cellular sheaves have natural weighted structures. Similarly, the pullbacks and pushforwards of a weighted sheaf have natural weighted structures.

There are some categorical subtleties to be aware of when we pass from cellular sheaves to weighted cellular sheaves. The space of global sections of a cellular sheaf is defined in categorical terms as the limit of the functor  $X \rightarrow \mathbf{Vect}$  defining the sheaf. This defines the space of global sections up to unique isomorphism. We might want a weighted space global sections to be a sort of limit in  $\mathbf{Hilb}$  which is defined up to unique *unitary* isomorphism. The categorical structure allowing us to define unitarity is the fact that  $\mathbf{Hilb}$  is a *dagger category*, and the corresponding concept of limit is the *dagger limit*. Unfortunately,  $\mathbf{Hilb}$  does not have all dagger limits; in particular, pullbacks over spans of noninjective maps do not exist [HK18]. As a result, there is no single canonical way to define an inner product on the space of global sections of a cellular sheaf. There are two approaches that seem most natural, however. One is to view the space of global sections as  $\ker \delta^0$  with its natural inner product given by inclusion into  $C^0$ . The other is to view this space as a subspace of  $\bigoplus_{\sigma} \mathcal{F}(\sigma)$ . We will generally take the view that global sections are a subspace of  $C^0$ ; that is, we will weight  $\Gamma(X; \mathcal{F})$  by its natural isomorphism with  $\mathcal{H}^0(X; \mathcal{F})$ .

An inner product on a vector space  $V$  induces an isomorphism with its dual  $V^* = \text{Hom}(V, \mathbb{k})$ . In particular, this means that a weighted sheaf gives an isomorphism between the stalks of a sheaf and its dual cosheaf. This is reminiscent of the bisheaves recently introduced by MacPherson and Patel [MP18]. However, the structure maps  $\mathcal{F}(\sigma) \rightarrow \hat{\mathcal{F}}(\sigma)$  will rarely commute with the restriction and extension maps as required by the definition of the bisheaf—this only holds in general if all restriction maps are unitary. The bisheaf construction is meant to give a generalization of local systems, and as such fits better with our discussion of discrete vector bundles in §3.4.



The inner products on stalks of  $\mathcal{F}$  extend by the orthogonal direct sum to inner products on  $C^k(X; \mathcal{F})$ , permitting a straightforward definition of the sheaf Laplacians.

### 3.2 The sheaf Laplacian

Given a chain complex of inner product spaces  $C^0 \rightarrow C^1 \rightarrow \dots$  we can construct the Hodge Laplacian  $\Delta = (\delta + \delta^*)^2 = \delta^* \delta + \delta \delta^*$ . This operator is naturally graded into components  $\Delta^k : C^k \rightarrow C^k$ , with  $\Delta^k = (\delta^k)^* \delta^k + \delta^{k-1} (\delta^{k-1})^*$ . This operator can be further separated into *up*- (coboundary) and *down*- (boundary) Laplacians  $\Delta_+^k = (\delta^k)^* \delta^k$  and  $\Delta_-^k = \delta^{k-1} (\delta^{k-1})^*$  respectively.

The basic and well-known theorem of discrete Hodge theory is the following:

**Theorem 1.** *Let  $C^0 \rightarrow C^1 \rightarrow \dots$  be a chain complex of inner product spaces, with Hodge Laplacians  $\Delta^k$ . Then  $\ker \Delta^k \cong H^k(C^\bullet)$ .*

*Proof.* By definition,  $H^k(C^\bullet) = \ker \delta^k / \text{im } \delta^{k-1}$ . In an inner product space,  $\ker \delta^k / \text{im } \delta^{k-1}$  is isomorphic to the orthogonal complement of  $\text{im } \delta^{k-1}$  in  $\ker \delta^k$ , which we may write  $(\ker \delta^k) \cap (\text{im } \delta^{k-1})^\perp = (\ker \delta^k) \cap (\ker (\delta^{k-1})^*)$ . So it suffices to show that  $\ker \Delta^k = (\ker \delta^k) \cap (\ker (\delta^{k-1})^*)$ . Note that  $\ker \delta^k = \ker (\delta^k)^* \delta^k = \ker \Delta_+^k$  and similarly for  $\Delta_-^k$ . So we need to show that  $\ker(\Delta_+^k + \Delta_-^k) = \ker \Delta_+^k \cap \ker \Delta_-^k$ , which will be true if  $\text{im } \Delta_+^k \cap \text{im } \Delta_-^k = 0$ . But this is true because  $\text{im } \Delta_+^k = \text{im } (\delta^k)^* = (\ker \delta^k)^\perp$  and  $\text{im } \Delta_-^k = \text{im } \delta^{k-1} \subseteq \ker \delta^k$ .  $\square$

The upshot of this theorem is that the kernel of  $\Delta^k$  gives a set of canonical representatives for  $H^k(C^\bullet)$ . This is commonly known as the space of *harmonic cochains*, denoted  $\mathcal{H}^k(C^\bullet)$ . In particular, the proof above implies that there is an orthogonal decomposition  $C^k = \mathcal{H}^k \oplus \text{im } \delta^{k-1} \oplus \text{im } (\delta^k)^*$ .

When the chain complex in question is the complex of  $\mathcal{F}$ -cochains, the Hodge construction produces the sheaf Laplacians. The Laplacian which is easiest to study and most immediately interesting is the degree-0 Laplacian, which is a generalization of the graph Laplacian. We can represent it as a symmetric block matrix with blocks indexed by the vertices of the complex. The entries on the diagonal are  $\Delta_{v,v}^0 = \sum_{v \leq e} \mathcal{F}_{v \leq e}^* \mathcal{F}_{v \leq e}$  and the entries on the off-diagonal are  $\Delta_{u,v}^0 = -\mathcal{F}_{u \leq e}^* \mathcal{F}_{v \leq e}$ , where  $e$  is the edge between  $v$  and  $u$ . Laplacians of other degrees have similar block structures.

The majority of results in combinatorial spectral theory have to do with up-Laplacians. We will frequently denote these  $L^k$  by analogy with spectral graph theory, where  $L$  typically denotes the (non-normalized) graph Laplacian. In particular, we will further elide the index  $k$  when  $k = 0$ , denoting the graph sheaf Laplacian by simply  $L$ . A subscript will be added when necessary to identify the sheaf.

The Laplacian of a labeled weighted graph determines the graph. The analogous result is not true of sheaves on a graph. For instance, the two sheaves




Fig. 1 Two nonisomorphic sheaves with the same Laplacian.

in Figure 3.2 both have Laplacian

$$\begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}$$

but are not unitarily isomorphic. More pithily, one cannot hear the shape of a sheaf. There are several sources for the lossiness of the sheaf Laplacian representation. One is that restriction maps may be the zero morphism, effectively allowing for edges that are only attached to one vertex. More generally, restriction maps may fail to be full rank, which means that it is impossible to identify the dimensions of edge stalks from the Laplacian.

### 3.2.1 Harmonic cochains

The elements of  $\ker \Delta^k = \mathcal{H}^k$  are known as harmonic  $k$ -cochains. More generally, a  $k$ -cochain may be harmonic on a subcomplex:

**Definition 10.** A  $k$ -cochain  $x$  of a sheaf  $\mathcal{F}$  on a cell complex  $X$  is *harmonic* on a set  $S$  of  $k$ -cells if  $(\Delta_{\mathcal{F}}^k x)|_S = 0$ .

When  $k = 0$  and  $\mathcal{F}$  is the constant sheaf, this can be expressed as a local averaging property: For each  $v \in S$ ,  $x_v = \frac{1}{d(v)} \sum_{u \sim v} x_u$ , where  $\sim$  indicates edge-incidence.

### 3.2.2 Identifying sheaf Laplacians

Degree-zero sheaf Laplacians for sheaves with stalks of dimension at most 1 are known as matrices of *factor width* equal to two, and are precisely the class of symmetric generalized diagonally dominant matrices [BCPT05]. These are matrices  $L$  for which there exists a positive diagonal matrix  $D$  such that  $DLD$  is diagonally dominant. Indeed, the fact that sheaves are not in general determined by their Laplacians is in part a consequence of the nonuniqueness of width-two factorizations.

### 3.3 The normalized Laplacian and choice of weights

Many results in spectral graph theory rely on a normalized version of the standard graph Laplacian, which is typically defined via the diagonal degree matrix  $D$ , letting  $\mathcal{L} = D^{-1/2}LD^{-1/2}$ . This definition preserves the Laplacian as a symmetric matrix, but it obscures the true meaning of the normalization. The normalized Laplacian is the standard Laplacian with a different choice of weights for the vertices. The matrix  $D^{-1/2}LD^{-1/2}$  is similar to  $D^{-1}L$ , which is self adjoint with respect to the inner product  $\langle x, y \rangle = x^T Dy$ . In this interpretation, each vertex is weighted proportional to its degree. Viewing the normalization process as a reweighting of cells leads to the natural definition of normalized Laplacians for simplicial complexes given by Horak and Jost [HJ13].

Indeed, following Horak and Jost's definition for simplicial complexes, we propose the following extension to sheaves.

**Definition 11.** Let  $\mathcal{F}$  be a weighted cellular sheaf defined on a regular cell complex  $X$ . We say  $\mathcal{F}$  is *normalized* if for every cell  $\sigma$  of  $X$  and every  $x, y \in \mathcal{F}(\sigma) \cap (\ker \delta)^\perp$ ,  $\langle \delta x, \delta y \rangle = \langle x, y \rangle$ .

**Lemma 1** *Given a weighted sheaf  $\mathcal{F}$  on  $X$ , it is always possible to reweight  $\mathcal{F}$  to a normalized version.*

*Proof.* Note that if  $X$  has dimension  $k$ , the normalization condition is automatically satisfied for all cells  $\sigma$  of dimension  $k$ . Thus, starting at cells of dimension  $k - 1$ , we recursively redefine the inner products on stalks. If  $\sigma$  is a cell of dimension  $k - 1$ , let  $\Pi_\sigma$  be the orthogonal projection  $\mathcal{F}(\sigma) \rightarrow \mathcal{F}(\sigma) \cap \ker \delta$ . Then define the normalized inner product  $\langle \bullet, \bullet \rangle_\sigma^N$  on  $\mathcal{F}(\sigma)$  to be given by  $\langle x, y \rangle_\sigma^N = \langle \delta(\text{id} - \Pi_\sigma)x, \delta(\text{id} - \Pi_\sigma)y \rangle + \langle \Pi_\sigma x, \Pi_\sigma y \rangle$ . It is clear that this reweighted sheaf satisfies the condition of Definition 11 for cells of dimension  $k$  and  $k - 1$ . We may then perform this operation on cells of progressively lower dimension to obtain a fully normalized sheaf.  $\square$

Note that there is an important change of perspective here: we do not normalize the Laplacian of a sheaf, but rather normalize the sheaf itself.

If we apply this process to a sheaf  $\mathcal{F}$  on a graph, this has an immediate interpretation in terms of the original sheaf Laplacian. Let  $D$  be the block diagonal of the standard sheaf Laplacian, and note that for  $x \perp \ker L$ ,  $\langle x, Dx \rangle$  is the reweighted inner product on  $C^0(G; \mathcal{F})$ . In particular, the adjoint of  $\delta$  with respect to this inner product has the form  $D^\dagger \delta^T$ , where  $D^\dagger$  is the Moore-Penrose pseudoinverse of  $D$ , so that the matrix form of the reweighted Laplacian with respect to this inner product is  $D^\dagger L$ . Changing to the standard basis then gives  $\mathcal{L} = D^{\dagger/2}LD^{\dagger/2}$ .

### 3.4 Discrete vector bundles

A subclass of sheaves of particular interest are those where all restriction maps are invertible. These sheaves have been the subject of significantly more study

than the general case, since they extend to locally constant sheaves on the geometric realization of the cell complex. The Riemann-Hilbert correspondence describes an equivalence between locally constant sheaves (or cosheaves) on  $X$ , local systems on  $X$ , vector bundles on  $X$  with a flat connection, and representations of the fundamental group of  $X$ . When we represent a local system by a cellular sheaf or cosheaf, we will call it a discrete vector bundle.

One way to understand the space of 0-cochains of a discrete vector bundle is as representing a subspace of the sections of a geometric realization of the associated flat vector bundle, defined by linear interpolation over higher-dimensional cells. The coboundary map can be seen as a sort of discretization of the connection, whose flatness is manifest in the fact that  $\delta^2 = 0$ .

Discrete vector bundles have some subtleties when we study their Laplacians. The sheaf-cosheaf duality corresponding to a local system, given by taking inverses of restriction maps, is not in general the same as the duality induced by an inner product on stalks. Indeed, these duals are only the same when the restriction maps are orthogonal — their adjoints must be their inverses.

The inner product on stalks of a cellular sheaf has two roles: it gives a relative weight to vectors in each stalk, but via the restriction maps also gives a relative weight to cells in the complex. This second role complicates our interpretation of certain sorts of vector bundles. For instance, one might wish to define an  $O(n)$  discrete vector bundle on a graph to be one where all restriction maps are orthogonal. However, from the perspective of the degree-0 Laplacian, a uniform scaling of the inner product on an edge does not change the orthogonality of the bundle, but instead in some sense changes the length of the edge, or perhaps the amount of emphasis we give to discrepancies over that edge. So a discrete  $O(n)$ -bundle should be one where the restriction maps on each cell are *scalar multiples* of orthonormal maps.

That is, for each cell  $\sigma$ , we have a positive scalar  $\alpha_\sigma$ , such that for every  $\sigma \triangleleft \tau$ , the restriction map  $\mathcal{F}_{\sigma \triangleleft \tau}$  is an orthonormal map times  $\alpha_\tau / \alpha_\sigma$ . One way to think of this is as a scaling of the inner product on each stalk of  $\mathcal{F}$ . Frequently, especially when dealing with graphs, we set  $\alpha_\sigma = 1$  when  $\dim(\sigma) = 0$ , but this is not necessary.

The rationale for this particular definition is that in the absence of a basis, inner products are not absolutely defined, but only in relation to maps in or out of a space. Scaling the inner product on a vector space is meaningless except in relation to a given collection of maps, which it transforms in a uniform way.

As a special case of this definition, it will be useful to think about weighted versions of the constant sheaf. These are isomorphic to the ‘true’ constant sheaf, but not unitarily so. Weighted constant sheaves on a graph are analogous to weighted graphs. The distinction between the true constant sheaf and weighted versions arises because it is often convenient to think of the sections of a cellular sheaf as a subspace of  $C^0(X; \mathcal{F})$ . As a result, we often only want our sections to be constant on 0-cells, allowing for variation up to a scalar multiple on higher-dimensional cells. This notion will be necessary in §8.2 when we discuss approximations of a cellular sheaf.

### 3.5 Comparison with previous constructions

Friedman, in [Fri15], gave a definition of a sheaf<sup>1</sup> on a graph, developed a homology theory, and suggested constructing sheaf Laplacians and adjacency matrices. The suggestion that one might develop a spectral theory of sheaves on graphs has remained until now merely a suggestion.

The graph connection Laplacian, introduced by Singer and Wu in [SW12], is simply the sheaf Laplacian of an  $O(n)$ -vector bundle over a graph. This construction has attracted significant interest from a spectral graph theory perspective, including the development of a Cheeger-type inequality [BSS13] and a study of random walks and sparsification [CZ12]. Connection Laplacian methods have proven enlightening in the study of synchronization problems. Others have approached the study of vector bundles, and in particular line bundles, over graphs without reference to the connection Laplacian, studying analogues of spanning trees and the Kirchhoff theorems [Ken11, CCK13]. Other work on discrete approximations to connection Laplacians of manifolds has analyzed similar matrices [Man07].

Gao, Brodski, and Mukherjee developed a formulation in which the graph connection Laplacian is explicitly associated to a flat vector bundle on the graph and arises from a twisted coboundary operator [GBM16]. This coboundary operator is not a sheaf coboundary map and has some difficulties in its definition. These arise from a lack of freedom to choose the basis for the space of sections over an edge of the graph. Further work by Gao uses a sheaf Laplacian-like construction to study noninvertible correspondences between probability distributions on surfaces [Gao16].

Wu et al. [WRWX18] have recently proposed a construction they call a weighted simplicial complex and studied their Laplacians. These are cellular cosheaves where all stalks are equal to a given vector space and restriction maps are scalar multiples of the identity. Their work discusses the cohomology and Hodge theory of weighted simplicial complexes, but did not touch on issues related to the Laplacian spectrum.

## 4 Harmonicity

As a prelude to results about the spectra of sheaf Laplacians, we will discuss issues related to harmonic cochains on sheaves. While these do not immediately touch on the spectral properties of the Laplacian, they are closely bound with its algebraic properties.

### 4.1 Harmonic Extension

**Proposition 1.** *Let  $B \subseteq X$  be a subcomplex and let  $x|_B \in C^k(B; \mathcal{F})$  be an  $\mathcal{F}$ -valued  $k$ -cochain specified on  $B$ . If  $H^k(X, B; \mathcal{F}) = 0$ , then there exists a*

<sup>1</sup> In our terminology, Friedman's sheaves are cellular cosheaves.

unique cochain  $x \in C^k(X; \mathcal{F})$  which restricts to  $x|_B$  on  $B$  and is harmonic on  $S = X \setminus B$ .

*Proof.* A matrix algebraic formulation suffices. Representing  $\Delta_{\mathcal{F}}^k$  in block form as partitioned by  $B$  and  $S$ , the relevant equation is

$$\begin{bmatrix} \Delta_{\mathcal{F}}^k(S, S) & \Delta_{\mathcal{F}}^k(S, B) \\ \Delta_{\mathcal{F}}^k(B, S) & \Delta_{\mathcal{F}}^k(B, B) \end{bmatrix} \begin{bmatrix} x|_S \\ x|_B \end{bmatrix} = \begin{bmatrix} 0 \\ y \end{bmatrix}.$$

Since  $y$  is indeterminate, we can ignore the second row of the matrix, giving the equation  $\Delta_{\mathcal{F}}^k(S, S)x|_S + \Delta_{\mathcal{F}}^k(S, B)x|_B = 0$ . We can write  $\Delta_{\mathcal{F}}^k(S, S) = (\delta^k|_S)^* \delta^k|_S + ((\delta^{k-1})^*|_S)^* (\delta^{k-1})^*|_S$ , which is very close to the  $k$ -th Hodge Laplacian of the relative cochain complex

$$\dots \rightarrow C^{k-1}(X, B; \mathcal{F}) \rightarrow C^k(X, B; \mathcal{F}) \rightarrow C^{k+1}(X, B; \mathcal{F}) \rightarrow \dots$$

Letting  $\pi_S$  be the orthogonal projection  $C^i(X; \mathcal{F}) \rightarrow C^i(X, B; \mathcal{F})$ , we may write the Hodge Laplacian of the relative complex as

$$\Delta^k(X, B; \mathcal{F}) = \pi_S (\delta^k)^* \delta^k \pi_S^* + \pi_S \delta^{k-1} \pi_S^* \pi_S (\delta^{k-1})^* \pi_S^*.$$

Meanwhile, we can write the submatrix

$$\Delta^k(S, S) = \pi_S (\delta^k)^* \delta^k \pi_S^* + \pi_S \delta^{k-1} (\delta^{k-1})^* \pi_S^*.$$

It is then immediate that  $\ker(\Delta^k(S, S)) \subseteq \ker(\Delta^k(X, B; \mathcal{F}))$ , so that  $\Delta_{\mathcal{F}}^k(S, S)$  is invertible if  $H^k(X, B; \mathcal{F}) = 0$ .  $\square$

If we restrict to up- or down-Laplacians, a harmonic extension always exists, even if it is not unique. This is because, for instance,  $\text{im}(\delta^k|_S)^* \delta^k|_B \subseteq \text{im}(\delta^k|_S)^* \delta^k|_S$ . In particular, this implies that harmonic extension is always possible for 0-cochains, with uniqueness if and only if  $H^0(X, B; \mathcal{F}) = 0$ .

## 4.2 Kron Reduction

*Kron reduction* is one of many names given to a process of simplifying graphs with respect to the properties of their Laplacian on a boundary. If  $G$  is a connected graph with a distinguished set of vertices  $B$ , which we consider as a sort of boundary of  $G$ , Proposition 1 shows that there is a harmonic extension map  $E : \mathbb{R}^B \rightarrow \mathbb{R}^{V(G)}$ . It is then possible to construct a graph  $G'$  on  $B$  such that for every function  $x$  on the vertices of  $G'$ , we have  $L_{G'}x = \pi_B L_G E(x)$ , where  $\pi_B$  is the orthogonal projection map  $\mathbb{R}^{V(G)} \rightarrow \mathbb{R}^B$ . Indeed,  $E(x)|_S = -L_G(S, S)^{-1} L_G(S, B)x$ , so

$$L_{G'}x = \pi_B L_G E(x) = L_G(B, B)x - L_G(B, S)L_G(S, S)^{-1} L_G(S, B)x,$$

and hence

$$L_{G'} = L_G(B, B) - L_G(B, S)L_G(S, S)^{-1} L_G(S, B),$$

that is,  $L_{G'}$  is the Schur complement of the  $(B, B)$  block of  $L_G$ . It is also the Laplacian of a graph on  $B$ :

**Fig. 2** A sheaf illustrating the general impossibility of Kron reduction.

**Theorem 2** (see [DB13]). *If  $L_G$  is the Laplacian of a connected graph  $G$ , and  $B$  a subset of vertices of  $G$ , then  $L_{G'} = L_G(B, B) - L_G(B, S)L_G(S, S)^{-1}L_G(S, B)$  is the Laplacian of a graph with vertex set  $B$ .*

A physically-inspired way to understand this result (and a major use of Kron reduction in practice) is to view it as reducing a network of resistors given by  $G$  to a smaller network with node set  $B$  that has the same electrical behavior on  $B$  as the original network. In this guise, Kron reduction is a high-powered version of the  $Y$ - $\Delta$  and star-mesh transforms familiar from circuit analysis. Further discussion of Kron reduction and its various implications and applications may be found in [DB13].

Can we perform Kron reduction on sheaves? That is, given a sheaf  $\mathcal{F}$  on a graph  $G$  with a prescribed boundary  $B$ , can we find a sheaf  $\mathcal{F}_B$  on a graph with vertex set  $B$  only such that for every  $x \in C^0(B; \mathcal{F}_B)$  we have  $L_{\mathcal{F}_B}x = \pi_{C^0(B; \mathcal{F}_B)}L_{\mathcal{F}}E(x)$ , where  $E(x)$  is the harmonic extension of  $x$  to  $G$ ?

The answer is, in general, no. Suppose we want to remove the vertex  $v$  from our graph, *i.e.*,  $B = G \setminus \{v\}$ . Let  $D_v = \sum_{v \leq e} \mathcal{F}_{v \leq e}^* \mathcal{F}_{v \leq e} = L_{v, v}$ . To eliminate the vertex  $a$  we apply the condition  $(L_{\mathcal{F}}(x, E(x)))(v) = 0$ , and take a Schur complement, replacing  $L(B, B)$  with  $L(B, B) - L(B, v)D_v^{-1}L(v, B)$ . This means that we add to the entry  $L(w, w')$  the map  $\mathcal{F}_{w \leq e}^* \mathcal{F}_{v \leq e} D_v^{-1} \mathcal{F}_{v \leq e'}^* \mathcal{F}_{w' \leq e'}$ , where  $e$  is the edge between  $v$  and  $w$ , and  $e'$  the edge between  $v$  and  $w'$ . This does not in general translate to a change in the restriction maps for the edge between  $w$  and  $w'$ . In general, Kron reduction is not possible for sheaves.

In particular, if  $s$  is a section of  $\mathcal{F}$ , its restriction to  $B$  must be a section of  $\mathcal{F}_B$ , and if  $s$  is not a section, its restriction to  $B$  cannot be a section of  $\mathcal{F}_B$ . But we can construct sheaves with a space of sections on the boundary that cannot be replicated with a sheaf on the boundary vertices only. For instance, take the star graph with three boundary vertices, with stalks  $\mathbb{R}$  over boundary vertices and edges, and  $\mathbb{R}^2$  over the internal vertex. Take as the restriction maps from the central vertex restriction onto the first and second components, and addition of the two components. See Figure 4.2 for an illustration.

Note that a global section of this sheaf is determined by its value on the central vertex. If we label the boundary vertices counterclockwise starting at the top, the space of global sections for  $\mathcal{F}_B$  must have as a basis

$$x_1 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \quad x_2 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}.$$

But there is no sheaf on  $K_3$  which has the same global sections as this sheaf. To see this, note that if  $x_1$  is a section, the map from  $\mathcal{F}(v_1)$  to  $\mathcal{F}(v_3)$  must be the zero map, and similarly for the map from  $\mathcal{F}(v_2)$  to  $\mathcal{F}(v_3)$ . Similarly, if  $x_2$  is a section, the maps  $\mathcal{F}(v_1) \rightarrow \mathcal{F}(v_2)$  and  $\mathcal{F}(v_3) \rightarrow \mathcal{F}(v_2)$  must be zero. But this already shows that the vector  $[1 \ 0 \ 0]^T$  must be a section, giving  $\mathcal{F}_B$

a three-dimensional space of sections. The problem is that the internal node allows for constraints between boundary nodes that cannot be expressed by purely pairwise interactions. This fact is a fundamental obstruction to Kron reduction for general sheaves.

However, there is a sheaf Kron reduction for sheaves with vertex stalks of dimension at most 1. This follows from the identification of the Laplacians of such sheaves as the matrices of factor width two in §3.2.2.

**Theorem 3.** *The class of matrices of factor width at most two is closed under taking Schur complements.*

*Proof.* By Theorems 8 and 9 of [BCPT05], a matrix  $L$  has factor width at most two if and only if it is symmetric and generalized weakly diagonally dominant with nonnegative diagonal, that is, there exists a positive diagonal matrix  $D$  such that  $DL D$  is weakly diagonally dominant. Equivalently, these are the symmetric positive semidefinite generalized weakly diagonally dominant matrices. The class of generalized weakly diagonally dominant matrices coincides with the class of  $H$ -matrices, which are shown to be closed under Schur complements in [JS05]. Similarly, the class of symmetric positive definite matrices is closed under Schur complements, so the intersection of the two classes is also closed.  $\square$

### 4.3 Maximum Modulus Theorem

Harmonic 0-cochains of an  $O(n)$ -bundle satisfy a local averaging property which leads directly to a maximum modulus principle.

**Lemma 1.** *Let  $G$  be a graph with an  $O(n)$ -bundle  $\mathcal{F}$ , where the vertex weights  $\alpha_v$  as discussed in §3.4 are constant. If  $x \in C^0(G; \mathcal{F})$  is harmonic at a vertex  $v$ , then*

$$x_v = \frac{1}{d_v} \sum_{\substack{v, w \triangleleft e \\ v \neq w}} \mathcal{F}_{v \triangleleft e}^* \mathcal{F}_{w \triangleleft e} x_w,$$

where  $d_v = \sum_{v \triangleleft e} \|\mathcal{F}_{v \triangleleft e}\|^2 = \sum_{v \triangleleft e} \alpha_e^2$ .

*Proof.* The block row of  $L_{\mathcal{F}}$  corresponding to  $v$  has entries  $-\mathcal{F}_{v \triangleleft e}^* \mathcal{F}_{w \triangleleft e}$  off the diagonal and  $\sum_{v \triangleleft e} \mathcal{F}_{v \triangleleft e}^* \mathcal{F}_{v \triangleleft e} = \sum_{v \triangleleft e} \|\mathcal{F}_{v \triangleleft e}\|^2 \text{id}_{\mathcal{F}(v)}$  on the diagonal. The harmonicity condition is then

$$d_v x(v) - \sum_{\substack{v, w \triangleleft e \\ v \neq w}} \mathcal{F}_{v \triangleleft e}^* \mathcal{F}_{w \triangleleft e} x_w = 0.$$

$\square$



**Theorem 4** (Maximum modulus principle). *Let  $G$  be a graph, and  $B$  be a thin subset of vertices of  $G$ ; that is,  $G \setminus B$  is connected, and every vertex in  $B$  is connected to a vertex not in  $B$ . Let  $\mathcal{F}$  is an  $O(n)$ -bundle on  $G$ , with  $x \in C^0(G; \mathcal{F})$  harmonic on  $G \setminus B$ . Then if  $x$  attains its maximum modulus on  $G \setminus B$ , it has constant modulus.*

*Proof.* Let  $v \in G \setminus B$  and suppose  $\|x_v\| \geq \|x_w\|$  for all  $w \in G$ . Then this holds in particular for neighbors of  $v$ , so that we have

$$\begin{aligned} \|x_v\| &= \frac{1}{d_v} \left\| \sum_{\substack{v, w \triangleleft e \\ v \neq w}} \mathcal{F}_{v \triangleleft e}^* \mathcal{F}_{w \triangleleft e} x_w \right\| \leq \frac{1}{d_v} \sum_{\substack{v, w \triangleleft e \\ v \neq w}} \|\mathcal{F}_{v \triangleleft e}^* \mathcal{F}_{w \triangleleft e} x_w\| \\ &= \frac{1}{d_v} \sum_{\substack{v, w \triangleleft e \\ v \neq w}} \alpha_e^2 \|x_w\| \leq \frac{1}{d_v} \sum_{v \triangleleft e} \alpha_e^2 \|x_v\| = \|x_v\|, \end{aligned}$$

Equality holds throughout, forcing the modulus of  $x$  to be constant on its immediate neighborhood. The same holds for any vertex not in  $B$ , due to the fact that there exists a path between them avoiding  $B$ , so  $\|x\|$  is constant on  $G \setminus B$ . It is then also constant on  $B$  by the condition that every vertex in  $B$  be adjacent to a vertex in  $G \setminus B$ . So any harmonic function that attains its maximum modulus on  $G \setminus B$  has constant modulus.  $\square$

**Corollary.** *Let  $B$  be a thin subset of vertices of  $G$ . If  $x \in C^0(G; \mathcal{F})$  is harmonic on  $G \setminus B$ , then it attains its maximum modulus on  $B$ .*

The constant sheaf on a graph is an  $O(n)$ -bundle, so this result gives a maximum modulus principle for harmonic functions on the vertices of a graph. A slightly stronger result in this vein, involving maxima and minima of  $x$ , is discussed in [Sun08]. The thinness condition for  $B$  is not strictly necessary for the corollary to hold — there are a number of potential weakenings of the condition. For instance, we might simply require that there exists some  $w \in B$  such that for every vertex  $v \in G \setminus B$  there exists a path from  $v$  to  $w$  not passing through  $B$ .

## 5 Spectra of Sheaf Laplacians

The results in this section are straightforward generalizations and extensions of familiar results from spectral graph theory. Most are not particularly difficult, but they illustrate the potential for lifting nontrivial notions from graphs and complexes to sheaves.

It is useful to note a few basic facts about the spectra of Laplacians arising from Hodge theory.

**Proposition 2.** *The nonzero spectrum of  $\Delta^k$  is the disjoint union of the nonzero spectra of  $\Delta_+^k$  and  $\Delta_-^k$ .*

*Proof.* If  $x$  is an eigenvector of  $\Delta_+^k$  with nonzero eigenvalue  $\lambda$ , it is orthogonal to  $\ker \Delta_+^k$ . In particular, then, because  $\ker \Delta_+^k$  contains  $\text{im } \Delta_-^k$ , we see that  $x \perp \text{im } \Delta_-^k$  and hence  $x \in \ker \Delta_-^k$ . Thus  $x$  is an eigenvector of  $\Delta_-^k$  with eigenvalue  $\lambda$ . A symmetric argument shows that all eigenvectors of  $\Delta_-^k$  with nonzero eigenvalues are also eigenvectors of  $\Delta_+^k$ .  $\square$

**Proposition 3.** *The nonzero spectra of  $\Delta_+^k$  and  $\Delta_-^{k+1}$  are the same.*

*Proof.* We have  $\Delta_+^k = (\delta^k)^* \delta^k$  and  $\Delta_-^{k+1} = \delta^k (\delta^k)^*$ . The eigendecompositions of these matrices are determined by the singular value decomposition of  $\delta^k$ , and the nonzero eigenvalues are precisely the squares of the nonzero singular values of  $\delta^k$ .  $\square$

One reason for the study of the normalized graph Laplacian is that its spectrum is bounded above by 2. A similar result holds for up-Laplacians of normalized simplicial complexes [HJ13]: the eigenvalues of the degree- $k$  up-Laplacian of a normalized simplicial complex are bounded above by  $k + 2$ . This fact extends to normalized sheaves on simplicial complexes.

**Proposition 4.** *Suppose  $\mathcal{F}$  is a normalized sheaf on a simplicial complex  $X$ . The eigenvalues of the degree  $k$  up-Laplacian  $L_{\mathcal{F}}^k$  are bounded above by  $k + 2$ .*

*Proof.* By the Courant-Fischer theorem, the largest eigenvalue of  $L_{\mathcal{F}}^k$  is equal to

$$\begin{aligned} \sup_{x \in C^k(X; \mathcal{F})} \frac{\langle x, L_{\mathcal{F}}^k x \rangle}{\langle x, x \rangle} &= \sup_{x \perp \ker \delta^k} \frac{\langle \delta^k x, \delta^k x \rangle}{\sum_{\dim \sigma = k} \langle \delta^k x_{\sigma}, \delta^k x_{\sigma} \rangle} \\ &= \sup_{x \perp \ker \delta^k} \frac{\sum_{\dim \tau = k+1} \sum_{\sigma, \sigma' \triangleleft \tau} [\sigma : \tau][\sigma' : \tau] \langle \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma}, \mathcal{F}_{\sigma' \triangleleft \tau} x_{\sigma'} \rangle}{\sum_{\dim \sigma = k} \sum_{\sigma \triangleleft \tau} \langle \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma}, \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma} \rangle}. \end{aligned}$$

Note that for  $\sigma \neq \sigma'$ ,

$$\begin{aligned} [\sigma : \tau][\sigma' : \tau] \langle \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma}, \mathcal{F}_{\sigma' \triangleleft \tau} x_{\sigma'} \rangle &\leq \| \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma} \| \| \mathcal{F}_{\sigma' \triangleleft \tau} x_{\sigma'} \| \\ &\leq \frac{1}{2} (\| \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma} \|^2 + \| \mathcal{F}_{\sigma' \triangleleft \tau} x_{\sigma'} \|^2) \end{aligned}$$

by the Cauchy-Schwarz inequality. In particular, then, each term of the numerator is bounded above by

$$\sum_{\sigma \triangleleft \tau} \| \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma} \|^2 + \frac{1}{2} \sum_{\sigma \neq \sigma' \triangleleft \tau} (\| \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma} \|^2 + \| \mathcal{F}_{\sigma' \triangleleft \tau} x_{\sigma'} \|^2) = (k+2) \sum_{\sigma \triangleleft \tau} \| \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma} \|^2.$$

Meanwhile, the denominator is equal to  $\sum_{\dim \tau = k+1} \sum_{\sigma \triangleleft \tau} \| \mathcal{F}_{\sigma \triangleleft \tau} x_{\sigma} \|^2$ , so the Rayleigh quotient is bounded above by  $k + 2$ .  $\square$

### 5.1 Eigenvalue Interlacing

**Definition 12.** Let  $A, B$  be  $n \times n$  matrices with real spectra. Let  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$  be the eigenvalues of  $A$  and  $\mu_1 \leq \mu_2 \leq \dots \leq \mu_n$  be the eigenvalues of  $B$ . We say the eigenvalues of  $A$  are  $(\mathbf{p}, \mathbf{q})$ -interlaced with the eigenvalues of  $B$  if for all  $k$ ,  $\lambda_{k-p} \leq \mu_k \leq \lambda_{k+q}$ . (We let  $\lambda_k = \lambda_1$  for  $k < 1$  and  $\lambda_k = \lambda_n$  for  $k > n$ .)

The eigenvalues of low-rank perturbations of symmetric positive semidefinite matrices are related by interlacing. The following is a standard result:

**Theorem 5.** *Let  $A$  and  $B$  be positive semidefinite matrices, with  $\text{rank } B = t$ . Then the eigenvalues of  $A$  are  $(t, 0)$ -interlaced with the eigenvalues of  $A - B$ .*

*Proof.* Let  $\mu_k$  be the  $k$ -th largest eigenvalue of  $A - B$  and  $\lambda_k$  the  $k$ -th largest eigenvalue of  $A$ . By the Courant-Fischer theorem, we have

$$\begin{aligned} \mu_k &= \min_{\dim Y=k} \left( \max_{y \in Y, y \neq 0} \frac{\langle y, Ay \rangle - \langle y, By \rangle}{\langle y, y \rangle} \right) \\ &\geq \min_{\dim Y=k} \left( \max_{y \in Y \cap \ker B, y \neq 0} \frac{\langle y, Ay \rangle}{\langle y, y \rangle} \right) \\ &\geq \min_{\dim Y=k-t} \left( \max_{y \in Y, y \neq 0} \frac{\langle y, Ay \rangle}{\langle y, y \rangle} \right) = \lambda_{k-t} \end{aligned}$$

and

$$\begin{aligned} \lambda_k &= \min_{\dim Y=k} \left( \max_{y \in Y, y \neq 0} \frac{\langle y, Ay \rangle}{\langle y, y \rangle} \right) \\ &\geq \min_{\dim Y=k} \left( \max_{y \in Y, y \neq 0} \frac{\langle y, Ay \rangle - \langle y, By \rangle}{\langle y, y \rangle} \right) = \mu_k. \end{aligned}$$

□

This result is immediately applicable to the spectra of sheaf Laplacians under the deletion of cells from their underlying complexes. The key part is the interpretation of the difference of the two Laplacians as the Laplacian of a third sheaf.<sup>2</sup> Let  $\mathcal{F}$  be a sheaf on  $X$ , and let  $C$  be an upward-closed set of cells of  $X$ , with  $Y = X \setminus C$ . The inclusion map  $i : Y \rightarrow X$  induces a restriction of  $\mathcal{F}$  onto  $Y$ , the pullback sheaf  $i^* \mathcal{F}$ . Consider the Hodge Laplacians  $\Delta_{\mathcal{F}}^k$  and  $\Delta_{i^* \mathcal{F}}^k$ . If  $C$  contains cells of dimension  $k$ , these matrices are different sizes, but we can derive a relationship by padding  $\Delta_{i^* \mathcal{F}}^k$  with zeroes. Equivalently, this is the degree- $k$  Laplacian of  $\mathcal{F}$  with the restriction maps incident to cells in  $C$  set to zero.

**Proposition 5.** *Let  $\mathcal{G}$  be the sheaf on  $X$  with the same stalks as  $\mathcal{F}$  but with all restriction maps between cells not in  $C$  set to zero. The eigenvalues of  $\Delta_{i^* \mathcal{F}}^k$  are  $(t, 0)$ -interlaced with the eigenvalues of  $\Delta_{\mathcal{F}}^k$ , where  $t = \text{codim } H^k(X; \mathcal{G}) = \dim C^k(X; \mathcal{F}) - \dim H^k(X; \mathcal{G})$ .*

<sup>2</sup> Such subtle moves are part and parcel of a sheaf-theoretic perspective.

Similar results can be derived for the up- and down-Laplacians. Specializing to graphs, interlacing is also possible for the normalized graph sheaf Laplacian. The Rayleigh quotient for the normalized Laplacian  $\mathcal{L}_{i^*\mathcal{F}}$  is

$$\frac{\langle x, D_{i^*\mathcal{F}}^{-1/2} L_{i^*\mathcal{F}} D_{i^*\mathcal{F}}^{-1/2} x \rangle}{\langle x, x \rangle} = \frac{\langle y, L_{i^*\mathcal{F}} y \rangle}{\langle y, D_{i^*\mathcal{F}} y \rangle} = \frac{\langle y, L_{\mathcal{F}} y \rangle - \langle y, L_{\mathcal{G}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle - \langle y, D_{\mathcal{G}} y \rangle}.$$

We then have

$$\begin{aligned} \mu_k &= \min_{\dim Y=k} \left( \max_{y \in Y, y \neq 0} \frac{\langle y, L_{\mathcal{F}} y \rangle - \langle y, L_{\mathcal{G}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle - \langle y, D_{\mathcal{G}} y \rangle} \right) \\ &\geq \min_{\dim Y=k} \left( \max_{y \in Y \cap H^0(X; \mathcal{G}), y \neq 0} \frac{\langle y, L_{\mathcal{F}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle - \langle y, D_{\mathcal{G}} y \rangle} \right) \\ &\geq \min_{\dim Y=k} \left( \max_{y \in Y \cap H^0(X; \mathcal{G}), y \neq 0} \frac{\langle y, L_{\mathcal{F}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle} \right) \\ &\geq \min_{\dim Y=k-t} \left( \max_{y \in Y, y \neq 0} \frac{\langle y, L_{\mathcal{F}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle} \right) = \lambda_{k-t} \end{aligned}$$

$$\begin{aligned} \mu_k &= \max_{\dim Y=n-k+1} \left( \min_{y \in Y, y \neq 0} \frac{\langle y, L_{\mathcal{F}} y \rangle - \langle y, L_{\mathcal{G}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle - \langle y, D_{\mathcal{G}} y \rangle} \right) \\ &\leq \max_{\dim Y=n-k+1} \left( \min_{y \in Y \cap H^0(X; \mathcal{G}), y \neq 0} \frac{\langle y, L_{\mathcal{F}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle - \langle y, D_{\mathcal{G}} y \rangle} \right) \\ &\leq \max_{\dim Y=n-k+1} \left( \min_{y \in Y \cap H^0(X; \mathcal{G}), y \neq 0} \frac{\langle y, L_{\mathcal{F}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle} \right) \\ &\leq \max_{\dim Y=n-k-t+1} \left( \min_{y \in Y, y \neq 0} \frac{\langle y, L_{\mathcal{F}} y \rangle}{\langle y, D_{\mathcal{F}} y \rangle} \right) = \lambda_{k+t} \end{aligned}$$

This shows that the eigenvalues of the normalized Laplacians are  $(t, t)$ -interlaced.

## 5.2 Sheaf morphisms

**Proposition 6.** *Suppose  $\varphi : \mathcal{F} \rightarrow \mathcal{G}$  is a morphism of weighted sheaves on a regular cell complex  $X$ . If  $\varphi^{k+1}$  is a unitary map, then  $L_{\mathcal{F}}^k = (\varphi^k)^* L_{\mathcal{G}}^k \varphi^k$ .*

*Proof.* The commutativity condition  $\varphi^{k+1} \delta^{\mathcal{F}} = \delta^{\mathcal{G}} \varphi^k$  implies that  $(\delta^{\mathcal{F}})^* (\varphi^{k+1})^* \varphi^{k+1} \delta^{\mathcal{F}} = (\varphi^k)^* (\delta^{\mathcal{G}})^* \delta^{\mathcal{G}} \varphi^k = (\varphi^k)^* L_{\mathcal{G}}^k \varphi^k$ . Thus if  $(\varphi^{k+1})^* \varphi^{k+1} = \text{id}_{C^1(\mathcal{F})}$ , we have  $L_{\mathcal{F}}^k = (\varphi^k)^* L_{\mathcal{G}}^k \varphi^k$ . This condition holds if  $\varphi^{k+1}$  is unitary.  $\square$

An analogous result holds for the down-Laplacians of  $\mathcal{F}$ , and these combine to a result for the full Hodge Laplacians.

### 5.3 Cell complex morphisms

The following constructions are restricted to *locally injective* cellular morphisms, those whose geometric realizations are locally injective. Under these morphisms, cells map to cells of the same dimension and the preimage of the star of a cell is a disjoint union of stars. The sheaf Laplacian is then invariant under pushforwards:

**Proposition 7.** *Let  $X$  and  $Y$  be cell complexes, and let  $f : X \rightarrow Y$  be a locally injective cellular morphism. If  $\mathcal{F}$  is a sheaf on  $X$ , the  $k$ th coboundary Laplacian corresponding to  $f_*\mathcal{F}$  on  $Y$  is the same (up to a unitary change of basis) as the  $k$ th coboundary Laplacian of  $\mathcal{F}$  on  $X$ .*

**Corollary.** *The sheaves  $\mathcal{F}$  and  $f_*\mathcal{F}$  are isospectral for the coboundary Laplacian.*

*Proof.* There is a canonical isometry  $f_k : C^k(X, \mathcal{F}) \rightarrow C^k(Y, f_*\mathcal{F})$ , which is given on stalks by the obvious inclusion  $f_\sigma : \mathcal{F}(\sigma) \rightarrow f_*\mathcal{F}(f(\sigma)) = \bigoplus_{f(\tau)=f(\sigma)} \mathcal{F}(\tau)$ . For  $\sigma \triangleleft \sigma'$ ,  $f_\sigma$  commutes with the restriction map  $\mathcal{F}_{\sigma \triangleleft \sigma'}$  and hence  $f_i$  commutes with the coboundary map. But this implies that:

$$L_{f_*\mathcal{F}}^k = (\delta_{f_*\mathcal{F}}^k)^* \delta_{f_*\mathcal{F}}^k = (\delta_{f_*\mathcal{F}}^k)^* f_{(k+1)}^* f_{(k+1)} \delta_{f_*\mathcal{F}}^k = f_k^* (\delta_{\mathcal{F}}^k)^* \delta_{\mathcal{F}}^k f_k = f_k^* L_{\mathcal{F}}^k f_k.$$

□

General locally injective maps behave nicely with sheaf pushforwards, and covering maps behave well with sheaf pullbacks.

**Proposition 8.** *Let  $f : C \rightarrow X$  be a covering map of cell complexes, with  $\mathcal{F}$  a sheaf on  $X$ . Then for any  $k$ , the spectrum of  $L_{\mathcal{F}}^k$  is contained in the spectrum of  $L_{f_*\mathcal{F}}^k$ .*

*Proof.* Consider the lifting map  $\varphi : C^k(X; \mathcal{F}) \rightarrow C^k(C; f_*\mathcal{F})$  given by  $x \mapsto x \circ f_k$ . This map commutes with  $\delta$  and  $\delta^*$ . The commutativity with  $\delta$  follows immediately from the proof of the contravariant functoriality of cochains. The commutativity with  $\delta^*$  is more subtle, and relies on the fact that  $f$  is a covering map.

For  $y \in C^k(C; f_*\mathcal{F})$  and  $x \in C^{k+1}(X; \mathcal{F})$ , we have

$$\begin{aligned} \langle y, \delta^* \varphi x \rangle &= \langle \delta y, \varphi x \rangle = \sum_{\substack{\sigma', \tau' \in P_C \\ \sigma' \triangleleft \tau'}} [\sigma' : \tau'] \langle f^* \mathcal{F}_{\sigma' \triangleleft \tau'}(y_{\sigma'}), (\varphi x)_{\tau'} \rangle \\ &= \sum_{\substack{\sigma, \tau \in P_X \\ \sigma \triangleleft \tau}} [\sigma : \tau] \sum_{\sigma' \in f^{-1}(\sigma)} \langle \mathcal{F}_{\sigma \triangleleft \tau}(y_{\sigma'}), x_\tau \rangle \\ &= \sum_{\substack{\sigma, \tau \in P_X \\ \sigma \triangleleft \tau}} [\sigma : \tau] \langle \mathcal{F}_{\sigma \triangleleft \tau}(i^* y)_\sigma, x_\tau \rangle = \langle \delta \varphi^* y, x \rangle = \langle y, \varphi \delta^* x \rangle. \end{aligned}$$

Now, if  $L_{\mathcal{F}}^k x = \lambda x$ , we have  $L_{f_*\mathcal{F}}^k \varphi x = (\delta_{f_*\mathcal{F}}^k)^* \delta_{f_*\mathcal{F}}^k \varphi x = \varphi (\delta_{\mathcal{F}}^k)^* \delta_{\mathcal{F}}^k x = \varphi L_{\mathcal{F}}^k x = \lambda \varphi x$ , so  $\lambda$  is an eigenvalue of  $L_{f_*\mathcal{F}}^k$ . □

Even if  $f : Y \rightarrow X$  is not quite a covering map, it is still possible to get some information about the spectrum of  $f^*\mathcal{F}$ . For instance, for dimension-preserving cell maps with uniform fiber size we have a bound on the smallest nontrivial eigenvalue of the pullback:

**Proposition 9.** *Suppose  $f : Y \rightarrow X$  is a dimension-preserving map of regular cell complexes such that for  $\dim(\sigma) = d$ ,  $|f^{-1}(\sigma)| = \ell_d$  is constant, and let  $\mathcal{F}$  be a sheaf on  $X$ . If  $\lambda_k(\mathcal{F})$  is the smallest nontrivial eigenvalue of  $L_{\mathcal{F}}^d$ , then  $\lambda_k(\mathcal{F}) \geq \frac{\ell_d}{\ell_{d+1}} \lambda_k(f^*\mathcal{F})$ .*

*Proof.* Let  $x$  be an eigenvector corresponding to  $\lambda_k(\mathcal{F})$  and note that since every fiber is the same size, the pullback  $\xi = f^*x$  preserves the inner product up to a scaling. That is, if  $\xi$  is a  $d$ -cochain,  $\langle \xi x, \xi y \rangle = \ell_d \langle x, y \rangle$ . This means that the pullback of  $x$  is orthogonal to the pullback of any cochain in the kernel of  $L_{\mathcal{F}}$ . Therefore, we have

$$\lambda_k(\mathcal{F}) = \frac{\langle \delta x, \delta x \rangle}{\langle x, x \rangle} = \frac{\ell_d \langle \xi \delta x, \xi \delta x \rangle}{\ell_{d+1} \langle \xi x, \xi x \rangle} = \frac{\ell_d \langle \delta \xi x, \delta \xi x \rangle}{\ell_{d+1} \langle \xi x, \xi x \rangle} \geq \frac{\ell_d}{\ell_{d+1}} \lambda_k(f^*\mathcal{F}).$$

□

#### 5.4 Product Complexes

If  $X$  and  $Y$  are cell complexes, their product  $X \times Y$  is a cell complex with cells  $\sigma \times \tau$  for  $\sigma \in X$ ,  $\tau \in Y$ , and incidence relations  $(\sigma \times \tau) \triangleleft (\sigma' \times \tau')$  whenever  $\sigma \triangleleft \sigma'$  and  $\tau \triangleleft \tau'$ . The dimension of  $\sigma \times \tau$  is  $\dim(\sigma) + \dim(\tau)$ . The complex  $X \times Y$  possesses projection maps  $\pi_X$  and  $\pi_Y$  onto  $X$  and  $Y$ .

**Definition 13.** If  $\mathcal{F}$  and  $\mathcal{G}$  are sheaves on  $X$  and  $Y$ , respectively, their *product* is the sheaf  $\mathcal{F} \boxtimes \mathcal{G} = \pi_X^* \mathcal{F} \otimes \pi_Y^* \mathcal{G}$ . Equivalently, we have  $(\mathcal{F} \boxtimes \mathcal{G})(\sigma \times \tau) = \mathcal{F}(\sigma) \otimes \mathcal{G}(\tau)$  and  $(\mathcal{F} \boxtimes \mathcal{G})_{\sigma \times \tau \triangleleft \sigma' \times \tau'} = \mathcal{F}_{\sigma \triangleleft \sigma'} \otimes \mathcal{G}_{\tau \triangleleft \tau'}$ .

**Proposition 10.** *If  $L_{\mathcal{F}}$  and  $L_{\mathcal{G}}$  are the degree-0 Laplacians of  $\mathcal{F}$  and  $\mathcal{G}$ , the degree-0 Laplacian of  $\mathcal{F} \boxtimes \mathcal{G}$  is  $L_{\mathcal{F} \boxtimes \mathcal{G}} = \text{id}_{C^0(X; \mathcal{F})} \otimes L_{\mathcal{G}} + L_{\mathcal{F}} \otimes \text{id}_{C^0(Y; \mathcal{G})}$ .*

*Proof.* The vector space  $C^1(X \times Y; \mathcal{F} \boxtimes \mathcal{G})$  has a natural decomposition into two subspaces: one generated by stalks of the form  $\mathcal{F}(v) \otimes \mathcal{G}(e)$  for  $v$  a vertex of  $X$  and  $e$  an edge of  $Y$ , and another generated by stalks of the opposite form  $\mathcal{F}(e) \otimes \mathcal{G}(v)$ . This induces an isomorphism

$$C^1(X \times Y; \mathcal{F} \boxtimes \mathcal{G}) \cong (C^0(X; \mathcal{F}) \otimes C^1(Y; \mathcal{G})) \oplus (C^1(X; \mathcal{F}) \otimes C^0(Y; \mathcal{G})).$$

Then the coboundary map of  $\mathcal{F} \boxtimes \mathcal{G}$  can be written as the block matrix

$$\delta_{\mathcal{F} \boxtimes \mathcal{G}} = \begin{bmatrix} \text{id}_{C^0(X; \mathcal{F})} \otimes \delta_{\mathcal{G}} \\ \delta_{\mathcal{F}} \otimes \text{id}_{C^0(Y; \mathcal{G})} \end{bmatrix}.$$

A quick computation then gives  $L_{\mathcal{F} \boxtimes \mathcal{G}} = \delta_{\mathcal{F} \boxtimes \mathcal{G}}^* \delta_{\mathcal{F} \boxtimes \mathcal{G}} = \text{id}_{C^0(X; \mathcal{F})} \otimes \delta_{\mathcal{G}}^* \delta_{\mathcal{G}} + \delta_{\mathcal{F}}^* \delta_{\mathcal{F}} \otimes \text{id}_{C^0(Y; \mathcal{G})} = \text{id}_{C^0(X; \mathcal{F})} \otimes L_{\mathcal{G}} + L_{\mathcal{F}} \otimes \text{id}_{C^0(Y; \mathcal{G})}$ . □

**Corollary.** *If the spectrum of  $L_{\mathcal{F}}$  is  $\{\mu_i\}_i$  and the spectrum of  $L_{\mathcal{G}}$  is  $\{\lambda_j\}_j$ , then the spectrum of  $L_{\mathcal{F} \boxtimes \mathcal{G}}$  is  $\{\mu_i + \lambda_j\}_{ij}$ .*

For higher degree Laplacians, this relationship becomes more complicated. For instance, the degree-1 up-Laplacian is computed as follows:

$$\delta_{\mathcal{F} \boxtimes \mathcal{G}}^1 = \begin{bmatrix} \text{id}_{C^0(X; \mathcal{F})} \otimes \delta_{\mathcal{G}}^1 & 0 \\ \delta_{\mathcal{F}}^0 \otimes \text{id}_{C^1(Y; \mathcal{G})} & \text{id}_{C^1(X; \mathcal{F})} \otimes \delta_{\mathcal{G}}^0 \\ 0 & \delta_{\mathcal{F}}^1 \otimes \text{id}_{C^0(Y; \mathcal{G})} \end{bmatrix}.$$

$$L_{\mathcal{F} \boxtimes \mathcal{G}}^1 = \begin{bmatrix} \text{id}_{C^0(X; \mathcal{F})} \otimes L_{\mathcal{G}}^1 + L_{\mathcal{F}}^0 \otimes \text{id}_{C^1(Y; \mathcal{G})} & (\delta_{\mathcal{F}}^0)^* \otimes \delta_{\mathcal{G}}^0 \\ \delta_{\mathcal{F}}^0 \otimes (\delta_{\mathcal{G}}^0)^* & \text{id}_{C^1(X; \mathcal{F})} \otimes L_{\mathcal{G}}^0 + L_{\mathcal{F}}^1 \otimes \text{id}_{C^0(Y; \mathcal{G})} \end{bmatrix}.$$

Because  $L_{\mathcal{F} \boxtimes \mathcal{G}}^1$  is given by a block matrix in terms of various Laplacians and coboundary maps of  $\mathcal{F}$  and  $\mathcal{G}$ , computing its spectrum is more involved. When  $X$  and  $Y$  are graphs, this simplifies significantly and it is possible to compute the spectrum in terms of the spectra of  $\mathcal{F}$  and  $\mathcal{G}$ .

**Proposition 11.** *Suppose  $X$  and  $Y$  are graphs, with  $\mathcal{F}$  and  $\mathcal{G}$  sheaves on  $X$  and  $Y$ . If  $v_{\mathcal{F}}$  is an eigenvector of  $L_{\mathcal{F}}^0$  with eigenvalue  $\lambda$  and  $v_{\mathcal{G}}$  an eigenvector of  $L_{\mathcal{G}}^0$  with eigenvalue  $\mu$ , then the vector  $v_{\mathcal{F} \boxtimes \mathcal{G}} = \begin{bmatrix} \sqrt{\frac{\lambda}{\mu}} v_{\mathcal{F}} \otimes \delta_{\mathcal{G}}^0 v_{\mathcal{G}} \\ \sqrt{\frac{\mu}{\lambda}} \delta_{\mathcal{F}}^0 v_{\mathcal{F}} \otimes v_{\mathcal{G}} \end{bmatrix}$  is an eigenvector of  $L_{\mathcal{F} \boxtimes \mathcal{G}}^1$  with eigenvalue  $\lambda + \mu$ .*

*Proof.* A computation.

$$\begin{aligned} L_{\mathcal{F} \boxtimes \mathcal{G}}^1 v_{\mathcal{F} \boxtimes \mathcal{G}} &= \begin{bmatrix} L_{\mathcal{F}}^0 \otimes \text{id}_{C^1(Y; \mathcal{G})} & (\delta_{\mathcal{F}}^0)^* \otimes \delta_{\mathcal{G}}^0 \\ \delta_{\mathcal{F}}^0 \otimes (\delta_{\mathcal{G}}^0)^* & \text{id}_{C^1(X; \mathcal{F})} \otimes L_{\mathcal{G}}^0 \end{bmatrix} \begin{bmatrix} \sqrt{\frac{\lambda}{\mu}} v_{\mathcal{F}} \otimes \delta_{\mathcal{G}}^0 v_{\mathcal{G}} \\ \sqrt{\frac{\mu}{\lambda}} \delta_{\mathcal{F}}^0 v_{\mathcal{F}} \otimes v_{\mathcal{G}} \end{bmatrix} \\ &= \begin{bmatrix} \left( \sqrt{\frac{\lambda}{\mu}} + \sqrt{\frac{\mu}{\lambda}} \right) \lambda v_{\mathcal{F}} \otimes \delta_{\mathcal{G}}^0 v_{\mathcal{G}} \\ \left( \sqrt{\frac{\lambda}{\mu}} + \sqrt{\frac{\mu}{\lambda}} \right) \mu \delta_{\mathcal{F}}^0 v_{\mathcal{F}} \otimes v_{\mathcal{G}} \end{bmatrix} = (\lambda + \mu) \begin{bmatrix} \sqrt{\frac{\lambda}{\mu}} v_{\mathcal{F}} \otimes \delta_{\mathcal{G}}^0 v_{\mathcal{G}} \\ \sqrt{\frac{\mu}{\lambda}} \delta_{\mathcal{F}}^0 v_{\mathcal{F}} \otimes v_{\mathcal{G}} \end{bmatrix}. \end{aligned}$$

□

A simpler way to obtain nearly the same result is to recall that  $(\delta_{\mathcal{F} \boxtimes \mathcal{G}}^1)^* \delta_{\mathcal{F} \boxtimes \mathcal{G}}^1$  and  $\delta_{\mathcal{F} \boxtimes \mathcal{G}}^1 (\delta_{\mathcal{F} \boxtimes \mathcal{G}}^1)^*$  have the same spectrum up to the multiplicity of zero. But when  $X$  and  $Y$  are graphs,

$$\delta_{\mathcal{F} \boxtimes \mathcal{G}}^1 (\delta_{\mathcal{F} \boxtimes \mathcal{G}}^1)^* = L_{\mathcal{F}}^0 \otimes \text{id}_{C^1(Y; \mathcal{G})} + \text{id}_{C^1(X; \mathcal{F})} \otimes L_{\mathcal{G}}^0.$$

However, for higher-dimensional complexes and higher-degree Laplacians, there appears to be no general simple formula giving the spectrum of  $\mathcal{F} \boxtimes \mathcal{G}$  in terms of the spectra of  $\mathcal{F}$  and  $\mathcal{G}$ . Indeed, we strongly suspect that no such formula can exist, *i.e.*, that the spectrum of  $\mathcal{F} \boxtimes \mathcal{G}$  is not in general determined by the spectra of  $\mathcal{F}$  and  $\mathcal{G}$ .

## 6 Effective Resistance

Effective resistance is most naturally defined for weighted cosheaves, but since every weighted sheaf has a canonical dual weighted cosheaf, the following definition applies to weighted sheaves as well.

**Definition 14.** Let  $\mathcal{F}$  be a cosheaf on a cell complex  $X$ , and let  $a, b \in \ker \partial_{\mathcal{F}}^k$  be homologous  $k$ -cycles. The *effective resistance*  $R_{\text{eff}}(a, b)$  is given by the solution to the optimization problem

$$\min_{c \in C_{k+1}(X; \mathcal{F})} \|c\|^2 \quad \text{s.t.} \quad \partial_{k+1} c = b - a. \quad (1)$$

**Proposition 12.** *Cosheaf effective resistance may be computed by the formula  $R_{\text{eff}}(a, b) = \langle b - a, (L_{\mathcal{F}}^k)^\dagger(b - a) \rangle$ , where  $L_{\mathcal{F}}^k$  is the coboundary Laplacian  $\partial_{k+1} \partial_{k+1}^*$ .*

*Proof.* By a standard result about matrix pseudoinverses,  $(L_{\mathcal{F}}^k)^\dagger = (\partial_{k+1}^\dagger)^* \partial_{k+1}^\dagger$ , so  $\langle b - a, (L_{\mathcal{F}}^k)^\dagger(b - a) \rangle = \langle \partial_{k+1}^\dagger(b - a), \partial_{k+1}^\dagger(b - a) \rangle = \|\partial_{k+1}^\dagger(b - a)\|^2$ . But if  $b - a \in \text{im } \partial^{k+1}$ , then  $\partial_{k+1}^\dagger(b - a)$  is the minimizer of the optimization problem (1).  $\square$

The condition that  $a$  and  $b$  be homologous becomes trivial if  $H_k(X; \mathcal{F}) = 0$ . Note that if  $\mathcal{F}$  is the constant cosheaf on a graph, a 0-cycle supported on a vertex  $v$  is homologous to a 0-cycle supported on a vertex  $w$  if their values are the same. This means that the definition of cosheaf effective resistance recovers the definition of graph effective resistance.

For any cosheaf on a cell complex, we can get a measure of effective resistance over a  $(k + 1)$ -cell  $\sigma$  by using the boundary map restricted to  $\sigma$ . Any choice of  $c \in \mathcal{F}(\sigma)$  gives an equivalence class of pairs of homologous  $k$ -cycles supported inside the boundary of  $\sigma$ . That is, we can decompose  $\partial c$  into the sum of two  $k$ -cycles in a number of equivalent ways. For instance, if  $\sigma$  is a 1-simplex with two distinct incident vertices, there is a natural decomposition of  $\partial c$  into a sum of two 0-cycles, one supported on each vertex. This gives a quadratic form on  $\mathcal{F}(\sigma)$ :  $R_{\text{eff}}(\sigma)(x) = \langle \partial x, L_{\mathcal{F}}^\dagger \partial x \rangle$ . The choice of decomposition does not affect the quadratic form. Of course, by the inner product pairing, this can be represented as a matrix  $(\partial|_{\mathcal{F}(\sigma)})^* L_{\mathcal{F}}^\dagger \partial|_{\mathcal{F}(\sigma)}$ . In particular, this gives a notion of effective resistance over an edge for a cosheaf on a graph.

### 6.1 Sparsification

Graph effective resistance has also attracted attention due to its use in graph sparsification. It is possible to use effective resistance to construct spectral sparsifiers of sheaves on cell complexes as well.

**Theorem 6.** *Let  $X$  be a regular cell complex of dimension  $d$  and  $\mathcal{F}$  a cosheaf on  $X$  with  $\dim C_{d-1}(X; \mathcal{F}) = n$ . Given  $\epsilon > 0$  there exists a subcomplex  $X' \subseteq X$  with the same  $(d-1)$ -skeleton and  $O(\epsilon^{-2} n \log n)$   $d$ -cells, together with a cosheaf  $\mathcal{F}'$  on  $X'$  such that  $(1 - \epsilon)L_{\mathcal{F}}^{d-1} \preceq L_{\mathcal{F}'}^{d-1} \preceq (1 + \epsilon)L_{\mathcal{F}}^{d-1}$ .*



*Proof.* If  $\ker L_{\mathcal{F}}^{d-1} = 0$ , then an equivalent condition to the conclusion is

$$\lambda_{\max}((L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} L_{\mathcal{F}'}^{d-1} (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}}) \leq 1 + \epsilon$$

and

$$\lambda_{\min}((L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} L_{\mathcal{F}'}^{d-1} (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}}) \geq 1 - \epsilon.$$

If  $\ker L_{\mathcal{F}}^{d-1}$  is nontrivial, we use the pseudoinverse of  $L_{\mathcal{F}}^{d-1}$  and restrict to the orthogonal complement of the kernel. This only offers notational difficulties, so in the following we will calculate as if the kernel were trivial.

Consider the restrictions of  $\partial_d$  to each  $d$ -cell  $\sigma$ , and note that

$$\sum_{\dim(\sigma)=d} (\partial_d|_{\sigma})(\partial_d|_{\sigma})^* = L_{\mathcal{F}}^{d-1}.$$

For each  $d$ -cell  $\sigma$ , we will choose  $\sigma$  to be in  $X'$  with probability

$$p_{\sigma} = \min(1, 4\epsilon^{-2} \log(n) \operatorname{tr}(R_{\text{eff}}(\sigma))).$$

If  $\sigma$  is chosen to be in  $X'$ , we choose its extension maps to be  $\mathcal{F}'_{\sigma \triangleleft \bullet} = \frac{1}{\sqrt{p_{\sigma}}} \mathcal{F}_{\sigma \triangleleft \bullet}$ .

Let  $A_{\sigma}$  be independent Bernoulli random variables with  $\mathbb{E}[A_{\sigma}] = p_{\sigma}$  and let  $X_{\sigma} = \frac{1}{p_{\sigma}} A_{\sigma} (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} (\partial_d|_{\sigma})(\partial_d|_{\sigma})^* (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}}$ . Note that by construction  $\sum_{\sigma} X_{\sigma} = (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} L_{\mathcal{F}'}^{d-1} (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}}$  and

$$\mathbb{E} \left[ \sum_{\sigma} X_{\sigma} \right] = (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} L_{\mathcal{F}'}^{d-1} (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} = \operatorname{id}_{C_{d-1}(X; \mathcal{F})}.$$

We wish to show that the eigenvalues of  $\sum_{\sigma} X_{\sigma}$  are close to those of its expectation, for which we use a matrix Chernoff bound proven in [Tro12]. This bound requires a bound on the norms of  $X_{\sigma}$ :

$$\begin{aligned} \|X_{\sigma}\| &\leq \frac{1}{p_{\sigma}} \|(L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} (\partial_d|_{\sigma})(\partial_d|_{\sigma})^* (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}}\| \\ &\leq \frac{1}{p_{\sigma}} \operatorname{tr}((L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} (\partial_d|_{\sigma})(\partial_d|_{\sigma})^* (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}}) \\ &\leq \frac{1}{p_{\sigma}} \operatorname{tr}((\partial_d|_{\sigma})^* (L_{\mathcal{F}}^{d-1})^{\dagger} \partial_d) = \frac{\operatorname{tr}(R_{\text{eff}}(\sigma))}{p_{\sigma}}. \end{aligned}$$

We can *a priori* subdivide any  $d$ -cells necessary to have small enough effective resistance so that  $p_{\sigma} < 1$  for all  $\sigma$ . This does not affect the asymptotics, so we can assume that  $p_{\sigma} < 1$ . This means that  $\|X_{\sigma}\| \leq \frac{\epsilon^2}{4 \log n}$ . Our matrix Chernoff bound then gives

$$\begin{aligned} \mathbb{P} \left[ \lambda_{\min}((L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} L_{\mathcal{F}'}^{d-1} (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}}) \leq 1 - \epsilon \right] &\leq n \exp \left( \frac{-4\epsilon^2}{2} \epsilon^{-2} \log n \right) = n^{-1} \\ \mathbb{P} \left[ \lambda_{\max}((L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}} L_{\mathcal{F}'}^{d-1} (L_{\mathcal{F}}^{d-1})^{-\frac{1}{2}}) \geq 1 + \epsilon \right] &\leq n \exp \left( \frac{-4\epsilon^2}{3} \epsilon^{-2} \log n \right) = n^{-1/3}. \end{aligned}$$

When  $n$  is not trivially small there is a high probability of this happening.

We now check the expected number of  $d$ -cells in  $X'$ . This is

$$\sum_{\sigma} p_{\sigma} \leq 4\epsilon^{-2} \log n \sum_{\sigma} \operatorname{tr}(R_{\text{eff}}(\sigma)),$$

and

$$\begin{aligned} \sum_{\sigma} \operatorname{tr}(R_{\text{eff}}(\sigma)) &= \sum_{\sigma} \operatorname{tr}((\partial|_{\sigma})^*(L_{\mathcal{F}}^{d-1})^{-1}\partial|_{\sigma}) = \sum_{\sigma} \operatorname{tr}((L_{\mathcal{F}}^{d-1})^{-1}\partial_{\sigma}(\partial|_{\sigma})^*) \\ &= \operatorname{tr}\left((L_{\mathcal{F}}^{d-1})^{-1} \sum_{\sigma} \partial_{\sigma}(\partial|_{\sigma})^*\right) = \operatorname{tr}((L_{\mathcal{F}}^{d-1})^{-1}L_{\mathcal{F}}^{d-1}) \leq n. \end{aligned}$$

A standard Chernoff bound argument with the Bernoulli random variables determining whether each cell is included then shows that the number of  $d$ -cells is concentrated around its expectation and thus can be chosen to be  $O(\epsilon^{-2}n \log n)$ .  $\square$

The proof given here follows the outline of the proof given by Spielman in [Spi15], simplifying the initial version of the proof which used a sampling of edges with replacement. Theorem 6 generalizes a number of theorems on sparsification of graphs and simplicial complexes [ST11, CZ12, OPW17]; however, it is not the most general sparsification theorem. Indeed, the core argument does not rely on the cell complex structure, but only on the decomposition of the Laplacian into a sum of matrices, one corresponding to each cell. More general, and stronger, theorems about sparsifying sums of symmetric positive semidefinite matrices have been proven, such as the following from [SHS16]:

**Theorem 7** (Silva et al. 2016). *Let  $B_1, \dots, B_m$  be symmetric, positive semidefinite matrices of size  $n \times n$  and arbitrary rank. Set  $B := \sum_i B_i$ . For any  $\epsilon \in (0, 1)$ , there is a deterministic algorithm to construct a vector  $y \in \mathbb{R}^m$  with  $O(n/\epsilon^2)$  nonzero entries such that  $y$  is nonnegative and*

$$B \preceq \sum_{i=1}^m y_i B_i \preceq (1 + \epsilon)B.$$

*The algorithm runs in  $O(mn^3/\epsilon^2)$  time. Moreover, the result continues to hold if the input matrices  $B_1, \dots, B_m$  are Hermitian and positive semidefinite.*

The sheaf theoretic perspective, though not the most general or powerful possible, nevertheless maintains both a great deal of generality along with a geometric interpretation of sparsification in terms of an effective resistance. This geometric interpretation may then be pursued to develop efficient methods for approximating the effective resistance and hence fast algorithms for sparsification of cell complexes and cellular sheaves atop them.

## 7 The Cheeger inequality

### 7.1 Cheeger Inequality for $O(n)$ -bundles

Bandeira, Singer, and Spielman proved an analogue to the graph Cheeger inequality for  $O(n)$ -bundles [BSS13]. Their goal was to give guarantees on the performance of a spectral method for finding an approximate section to a principal  $O(n)$ -bundle over a graph. They defined the *frustration* of a 0-cochain  $x$  to be

$$\eta(x) = \frac{\langle x, Lx \rangle}{\langle x, Dx \rangle} = \frac{\langle x, \mathcal{L}x \rangle}{\langle x, x \rangle},$$

and showed that any 0-cochain can be rounded to one with controlled frustration. Given a cochain  $x$  and a threshold  $\kappa \geq 0$ , let  $x^\kappa$  be the cochain whose value at a vertex  $v$  is  $x_v/\|x_v\|$  if  $\|x_v\|^2 \geq \kappa$  and is zero otherwise. For any such  $x$  there exists a  $\kappa$  such that  $\eta(x^\kappa) \leq \sqrt{10\eta(x)}$ . This immediately implies that

$$\lambda_1(\mathcal{L}) \leq \min_{\|x_v\|=1 \text{ or } 0} \eta(x) \leq \sqrt{10\lambda_1(\mathcal{L})}, \quad (2)$$

where  $\mathcal{L}$  is the normalized degree-0 Laplacian of the  $O(n)$ -bundle.

A natural question is whether this theorem extends to a more general class of sheaves on graphs. One reasonable candidate for extension is the class of sheaves on graphs where all restriction maps are partial isometries. One might view these as  $O(n)$ -bundles where the edge stalks have been reduced in dimension by an orthogonal projection. However, the cochain rounding approach does not work for these sheaves, as the following simple counterexample shows. Let  $G$  be a graph with two vertices and one edge, and let  $\mathcal{F}(v_1) = \mathcal{F}(v_2) = \mathbb{R}^2$  and  $\mathcal{F}(e) = \mathbb{R}$ . Then let  $\mathcal{F}_{v_1 \leq e} = [1 \ 0]$  and  $\mathcal{F}_{v_2 \leq e} = [\frac{1}{2} \ \frac{\sqrt{3}}{2}]$ . Then let  $x_{v_1} = \begin{bmatrix} \frac{1}{2} \\ 0 \end{bmatrix}$  and  $x_{v_2} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ . Then  $\eta(x) = 0$ , but  $\eta(x^u) > 0$  for any choice of  $u < 1$ . This means that there cannot exist any function  $f : \mathbb{R} \rightarrow \mathbb{R}$  with  $f(0) = 0$  such that  $\eta(x^u) \leq f(\eta(x))$ .

This example does not immediately show that the Cheeger inequality (2) is false for this class of sheaves, since this sheaf does have a section of stalkwise norm 1, but it does offer a counterexample to the key lemma in the proof. Indeed, a more complicated family of counterexamples exists with sheaves that have no global sections. These counterexamples show that an approach based on variational principles and rounding is unlikely to prove an analogue of the results of Bandeira et al. for more general classes of sheaves.

### 7.2 Structural Cheeger Inequality

Many extensions of the graph Cheeger inequality view it from the perspective of a constrained optimization problem over cochains. This is the origin of the Cheeger inequality for  $O(n)$ -bundles, and of the higher-dimensional

Cheeger constants proposed by Gromov, Linial and Meshulam, and others [LM06, Gro10, PRT16]. However, a sheaf gives us more structure to work with than simply cochains.

The traditional Cheeger inequality for graphs is frequently stated as a graph cutting problem: what is the optimal cut balancing the weight of edges removed with the sizes of the resulting partition? If we take the constant sheaf on a graph  $G$ , we can represent a cut of  $G$  by setting some restriction maps to 0, or, more violently, setting the edge stalks on the cut to be zero-dimensional. Thus, a potential analogue to the Cheeger constant for sheaves might be an optimal perturbation to the structure of the sheaf balancing the size of the perturbation with the size of the support of a new global section that the perturbation induces.

For instance, we might measure the size of a perturbation of the sheaf's restriction maps in terms of the square of the Frobenius norm  $\|\cdot\|_F$  of the coboundary matrix. If we minimize  $\|\delta_{\mathcal{F}} - \delta_{\mathcal{F}'}\|_F^2$ , a natural relaxation to the space of all matrices shows that this value is greater than  $\lambda_1(L_{\mathcal{F}})$ .

## 8 Toward Applications

The increase in abstraction and technical overhead implicit in lifting spectral graph theory to cellular sheaves is nontrivial. However, given the utility of spectral graph theory in so many areas, the generalization to sheaves would appear to be a good investment over time. As this work is an initial survey of the landscape, we will not dwell on any one potential application too long here. Rather, we sketch some possible avenues and provide plausibility arguments.

### 8.1 Distributed Consensus

Graph Laplacians and adjacency matrices play an important role in the study and design of distributed dynamical systems. This begins with the observation that the dynamical system

$$\dot{x} = -Lx$$

is local with respect to the graph structure: the only terms that influence  $\dot{x}_v$  are  $x_w$  for  $w$  adjacent to  $v$ . Further, if the graph is connected, diagonalization of  $L$  shows that the flows of this dynamical system converge to the average of the initial condition. A similar observation holds for sheaves on graphs:

**Proposition 13.** *Let  $\mathcal{F}$  be a sheaf on a cell complex  $X$ . The dynamical system  $\dot{x} = -\Delta_{\mathcal{F}}^k x$  has as its space of equilibria  $\mathcal{H}^k(X; \mathcal{F})$ , the space of harmonic  $k$ -cochains of  $\mathcal{F}$ . The trajectory of this dynamical system initialized at  $x_0$  converges exponentially quickly to the orthogonal projection of  $x_0$  onto  $\mathcal{H}^k(X; \mathcal{F})$ .*

In particular, this result implies that a distributed system can reach consensus on the nearest global section to an initial condition. One application

of this idea is to reduce the total amount of communication needed for a system to reach consensus. Suppose we have a distributed system on  $G$  where each node has state in  $\mathbb{R}^d$  for some large  $d$  and the nodes are required to reach consensus. The traditional approach is to use a discretized version of the Laplacian dynamics of the constant  $\mathbb{R}^d$  sheaf on  $G$ . But this requires the communication of the entire  $d$ -dimensional state vector over each edge at each time step. The possibility of using a different sheaf on  $G$  that has the same space of global sections to reduce the communication cost is appealing. Such a sheaf would comprise a certain *approximation* to the constant sheaf.

## 8.2 Approximations to Cellular Sheaves

A sheaf of vector spaces can be thought of as a distributed system of linear transformations, and its cohomology  $H^\bullet$  consists of equivalence classes of solutions to systems based on these constraints. From this perspective, questions of approximation — of sheaves and of sheaf cohomology — take on especial relevance. The question of approximating global sections to a given sheaf has appeared in, *e.g.*, Robinson [Rob17, Rob18].

Questions of approximating sheaves are equally interesting. Given the relative lack of investigation, the following definition is perhaps premature; nevertheless, it is well-motivated by the problem of distributed consensus as well as notions of cellular approximation in algebraic topology.

**Definition 15.** Let  $X$  be a regular cell complex, and let  $\mathcal{G}$  be a sheaf on  $X$ . We say that a sheaf  $\mathcal{F}$  on  $X$  is a *k-approximation* to  $\mathcal{G}$  if there exists a sheaf morphism  $a : \mathcal{G} \rightarrow \mathcal{F}$  which is an isomorphism on stalks over cells of degree at most  $k$ , and which induces an isomorphism  $H^i(X; \mathcal{G}) \rightarrow H^i(X; \mathcal{F})$  for all  $i \leq k$ .

If  $\mathbb{V}$  is a vector space, we denote the constant sheaf with stalk  $\mathbb{V}$  by  $\underline{\mathbb{V}}$ , and say that  $\mathcal{F}$  is an *approximation to the constant sheaf* if  $\mathcal{F}$  is an approximation to  $\underline{\mathbb{V}}$ .

This definition is reminiscent of cellular approximation methods from algebraic topology. A space  $X$  may be  $k$ -approximated by a cell complex  $Y$ , via a morphism  $Y \rightarrow X$  inducing an isomorphism on homotopy groups up to degree  $k$ . Here we approximate a sheaf on a cell complex by one with the same cohomology in degrees up to  $k$ .

In particular, a 0-approximation to  $\mathcal{F}$  has the same vertex stalks and the same space of global sections as  $\mathcal{F}$ . In what follows, we will work over a graph  $G$ , and all approximations will be 0-approximations.

**Proposition 14.** *If  $\mathcal{F}$  is an approximation to  $\underline{\mathbb{V}}$ , then it is isomorphic to a sheaf with vertex stalks  $\mathbb{V}$  where the restriction maps  $\mathcal{F}_{v \leq e} : \mathbb{V} \rightarrow \mathcal{F}(e)$  and  $\mathcal{F}_{w \leq e} : \mathbb{V} \rightarrow \mathcal{F}(e)$  are equal for all  $v \sim w$ .*

*Proof.* Note that because  $a : \underline{\mathbb{V}} \rightarrow \mathcal{F}$  is an isomorphism on vertex stalks,  $\mathcal{F}$  is clearly isomorphic to a sheaf with vertex stalks  $\mathbb{V}$ . For every edge  $e = (v, w)$

we have the diagram

$$\begin{array}{ccc}
 \mathbb{V} & \xrightarrow{\text{id}} & \mathbb{V} \\
 \text{id} \downarrow & & \downarrow \mathcal{F}_{v \triangleleft e} \\
 \mathbb{V} & \xrightarrow{a_e} & \mathcal{F}(e) \\
 \text{id} \uparrow & & \uparrow \mathcal{F}_{w \triangleleft e} \\
 \mathbb{V} & \xrightarrow{\text{id}} & \mathbb{V}
 \end{array},$$

and the only way it can commute is if  $\mathcal{F}_{v \triangleleft e} = \mathcal{F}_{w \triangleleft e} = a_e$ .  $\square$

The proof of this proposition shows that specifying an approximation to  $\underline{\mathbb{V}}$  is the same as specifying a morphism  $a_e : \mathbb{V} \rightarrow \mathcal{F}(e)$  for each edge  $e$  of  $G$ . Further, in order to produce an approximation to  $\underline{\mathbb{V}}$ , the  $a_e$  must assemble to a map  $a : C^1(G; \underline{\mathbb{V}}) \rightarrow C^1(G; \mathcal{F}) = \bigoplus_{e \in E} \mathcal{F}(e)$  such that  $\ker(a \circ \delta_{\underline{\mathbb{V}}}) = \ker \delta_{\underline{\mathbb{V}}}$ . This holds if  $\ker a$  is contained in a complement to  $\text{im } \delta$ ; equivalently, the projection map  $\pi : C^1(G; \underline{\mathbb{V}}) \rightarrow H^1(G; \underline{\mathbb{V}})$  must be an isomorphism when restricted to  $\ker a$ .

This suggests a way to construct an approximation to the constant sheaf. Choose a subspace  $K_e$  of  $\mathbb{V}(e)$  for each edge  $e$  of  $G$  and define  $a_e$  to be the projection map  $\mathbb{V} \rightarrow \mathbb{V}/K_e$ . If  $\bigoplus_{e \in E} K_e$  has the same dimension in  $H^1(G; \underline{\mathbb{V}})$  as in  $C^1(G; \underline{\mathbb{V}})$ , then  $a = \bigoplus_{e \in E} a_e$  defines the edge maps giving an approximation to  $\underline{\mathbb{V}}$ . (The vertex maps may be taken to be the identity.)

One necessary condition for the maps  $a_e$  to form an approximation to the constant sheaf is as follows:

**Proposition 15.** *Let  $K_e \leq \mathbb{V}$  be a subspace for each edge  $e$  of  $G$ . Suppose the quotient maps  $a_e : \mathbb{V} \rightarrow \mathbb{V}/K_e$  assemble to form an approximation to the constant sheaf as in Proposition 14. Then for every cutset  $C$  of  $G$ ,  $\bigcap_{e \in C} K_e = 0$ .*

*Proof.* Let  $C$  be a cutset of  $G$  for which  $\bigcap_{e \in C} K_e \neq 0$ , so that there is some nonzero  $x_C \in \bigcap_{e \in C} K_e$ .  $C$  partitions  $G$  into at least two subgraphs. Pick one such subgraph  $H$ , and construct a cochain  $x \in C^0(G; \mathcal{F})$  by letting  $x_v = x_C$  for  $v \in H$ , and  $x_v = 0$  for  $v \notin H$ . It is easy to see that  $\delta_{\mathcal{F}} x = 0$  but  $x$  is not constant, so that  $\mathcal{F}$  is not an approximation to the constant sheaf.  $\square$

The converse is not true; the question of when a collection  $\{K_e\}$  produces an approximation to the constant sheaf appears surprisingly subtle.

Once we put weights on an approximation to the constant sheaf, there are bounds on its Laplacian spectrum in terms of the spectrum of the corresponding weighted constant sheaf. Recall from §3.4 that a weighted constant sheaf is isomorphic to one where the restriction map  $\underline{\mathbb{V}}_{v \triangleleft e}$  is equal to  $(\alpha_e/\alpha_v) \text{id}$ . We will for the current purposes assume that  $\alpha_v = 1$  for all vertices. In this situation, the component maps  $a_e$  of the approximation  $a : \underline{\mathbb{V}} \rightarrow \mathcal{F}$  are not equal to the restriction maps  $\mathcal{F}_{v \triangleleft e}$ , but are scalar multiples.

**Theorem 8.** *Let  $a : \mathbb{V} \rightarrow \mathcal{F}$  be an approximation to  $\mathbb{V}$ , where the component maps of  $a$  are orthogonal projections, and the edge stalks of  $\mathcal{F}$  have constant dimension  $k$ . Then if  $\lambda_{\mathbb{V}}$  is the smallest nontrivial eigenvalue of  $L_{\mathbb{V}}$  and  $\lambda_{\mathcal{F}}$  is the smallest nontrivial eigenvalue of  $\mathcal{F}$ , we have  $\lambda_{\mathcal{F}} \leq \frac{k}{\dim \mathbb{V}} \lambda_{\mathbb{V}}$ .*

*Proof.* Consider  $L_G$ , the weighted Laplacian of the weighted graph corresponding to the weighting of  $\mathbb{V}$ . Let  $w$  be a unit eigenvector of  $L_G$  corresponding to the smallest nontrivial eigenvalue  $\lambda_2^G$ . Consider the edge projection matrix  $a = \text{diag } a_e$ , with  $D_e = a_e^* a_e$ . Then  $L_{\mathcal{F}} = \delta_{\mathbb{V}}^* \text{diag}(D_e) \delta_{\mathbb{V}}$ . Choose  $x$  uniformly at random from the unit sphere in  $\mathbb{V}$ , and note that  $\|w \otimes x\| = 1$  and  $w \otimes x$  is orthogonal to the space of constant cochains on  $\mathbb{V}$ . Consider the expectation  $\mathbb{E}[\langle w \otimes x, L_{\mathcal{F}}(w \otimes x) \rangle]$ . This is

$$\begin{aligned} & \mathbb{E} \left[ \sum_{v, v' \triangleleft e} \langle w_v x - w_{v'} x, D_e(w_v x - w_{v'} x) \rangle \right] \\ &= \sum_{v, v' \triangleleft e} (w_v - w_{v'})^2 \mathbb{E}[\langle x, D_{ij} x \rangle] = \sum_{v, v' \triangleleft e} (w_v - w_{v'})^2 \frac{\text{tr}(D_e)}{\dim \mathbb{V}}. \end{aligned}$$

Since  $a_e$  is an orthogonal projection of rank  $k$ ,  $\text{tr}(D_e) = k$ . Therefore,

$$\mathbb{E}[\langle w \otimes x, L_{\mathcal{F}}(w \otimes x) \rangle] = \frac{k}{\dim \mathbb{V}} \lambda_2^G.$$

Since this is an expectation over a set of unit vectors we see that there must exist some unit norm  $x$  such that  $(w \otimes x)^T L_{\mathcal{F}}(w \otimes x) \leq \frac{k}{\dim \mathbb{V}} \lambda_2^G$ , and hence  $\lambda \leq \frac{k}{\dim \mathbb{V}} \lambda_{\mathbb{V}}$  since  $w \otimes x$  is orthogonal to the space of constant cochains.  $\square$

The preceding result in fact holds if the edge maps  $a_e$  are of norm at most 1. If  $\lambda_{\mathcal{F}}^{\max}$  is the largest eigenvalue of  $L_{\mathcal{F}}$ , we have by a nearly identical argument  $\lambda_{\mathcal{F}}^{\max} \geq \frac{k}{\dim \mathbb{V}} \lambda_{\mathbb{V}}^{\max}$ .

Using an approximation to the constant sheaf reduces the amount of communication necessary for a single step of a discrete-time consensus scheme, so if bandwidth in each time step is strictly limited, this can enable consensus where it was not previously possible. However, if bandwidth is not bounded, but merely incurs some cost of communication, we need to balance the cost of communication with the speed of convergence, and study the spectral properties of approximations to the constant sheaf. Numerical simulations show that it is in fact possible to reduce the total amount of communication necessary for consensus up to a given threshold, but further work is necessary to develop ways to quickly produce an approximation to the constant sheaf with good spectral properties.

### 8.3 Synchronization

The term *synchronization* in the context of problems with data on graphs is exemplified in work by Singer on the angular alignment problem [Sin11].

The concept was developed further by Bandeira in his dissertation [Ban15]. The general idea is to recover information about some set of parameters from knowledge of their pairwise relationships. The general formulation, due to Bandeira, is as follows: Given a group  $G$ , a graph  $X$ , and a function  $f_{ij} : G \rightarrow \mathbb{R}$  for each edge  $i \sim j$  of  $X$ , find a function  $g : V(X) \rightarrow G$  minimizing  $\sum_{i \sim j} f_{ij}(g(i)g(j)^{-1})$ .

Often, the functions  $f_{ij}$  are chosen such that they have a unique minimum at a given element  $g_{ij} \in G$ . One may view this as originating from a  $G$ -principal bundle on a graph, with the group elements  $g_{ij}$  defining transition maps. The desired solution is a section of the bundle, determined up to a self-action of  $G$ , but may not exist if there is error in the measured  $g_{ij}$ . As a result, we seek an error-minimizing solution, where the error is measured by the functions  $f_{ij}$ .

Sheaves on graphs offer a broader formulation: synchronization is the problem of finding a global section, or approximate global section, of an observed sheaf on a graph. By choosing sheaves valued in a suitable category, we can recover the group-theoretic formulation. There is something of a gap between the natural formulation of many synchronization problems and a sheaf valued in vector spaces; bridging that gap for synchronization over  $O(d)$  is the goal of [BSS13].

The explicit sheaf-theoretic formulation of the data of a synchronization problem also suggests a different approach to solving such a problem. If a synchronization sheaf has a global section, as is the case when the data are internally consistent and uncorrupted by noise, finding that section is trivial. The traditional approach to synchronization takes these transition functions as they are, and seeks an approximate section to the sheaf. On the other hand, we might try to denoise the measured relationships themselves using the condition of cycle-consistency. That is, given an observed sheaf, find the nearest sheaf supporting a global section. A structural Cheeger inequality as discussed in §7.2 would give spectral insights into this problem. Deeper understanding would come from study of the moduli spaces of cellular sheaves.

#### 8.4 Consistent Clustering

As suggested by Gao et al. [GBM16], the data of a sheaf on a graph is useful for more than recovering a global section. The problem of clustering objects where the similarity measure comes from an explicit matching or transformation gives extra information. If we stipulate that objects within a cluster should have consistent transformations along cycles, the problem of clustering becomes the problem of partitioning a graph into subgraphs, each of which supports the appropriate space of global sections.

Similar ideas arise in [Gao16], which considers correspondences between surfaces produced using the soft Procrustes distance. These correspondences are maps between probability distributions on the vertex sets of discretized surfaces. When these surfaces are meshes with varying numbers of vertices, these maps are not invertible, but by construction they are represented by



doubly stochastic matrices, and the analogue of the inverse for such a map is simply the transpose of its corresponding matrix. These sorts of geometric correspondences are natural to consider in the context of geometric morphometrics, the field devoted to studying and classifying species based on their geometric properties.

Gao constructs a matrix he calls the graph horizontal Laplacian, together with normalized versions he uses to formulate a fiber bundle version of the diffusion maps algorithm for dimensionality reduction. The graph horizontal Laplacian is related to a map of graphs  $X \rightarrow G$ , where the fibers over vertices of  $G$  are discrete. A weighting on the edges of  $X$  induces a matrix-valued weighting on the edges of  $G$ . This produces a weighted adjacency matrix  $W$  of  $G$ , from which the graph horizontal Laplacian is generated by  $L^H = D - W$ , where  $D$  is the diagonal matrix necessary to make  $L^H$  have row sums equal to zero. This is in fact equivalent to the sheaf Laplacian of the pushforward of the weighted constant sheaf on  $X$ , and as a consequence of Proposition 7, is simply a block subdivision of the Laplacian of  $X$ . This sheaf on  $G$  can then be normalized to construct a diffusion map embedding of the vertices of  $G$ , as well as an embedding of the vertices of  $X$ . When applied to the surface correspondence problem, the eigenvectors of the resulting sheaf Laplacian serve to partition the surfaces into automatically determined landmarks or regions of interest.

Approaching these notions of partitioning, partial sections, and noninvertible matchings from a sheaf-theoretic perspective offers new tools and clarifies the problems in question. There are clearly many questions and applications to investigate, and the tools of spectral sheaf theory should be of use here.

## 9 Closing Questions

There are numerous interesting open questions in a spectral sheaf theory. We highlight a few below, with comments.

### 9.1 Metrics on the space of cellular sheaves

Interleaving-type constructions have been used to define metrics on the space of constructible sheaves. However, these rely on explicit geometric information about the sheaves and their underlying spaces. Working with weighted cellular sheaves may make it possible to define useful distances that rely only on the combinatorial and algebraic structure of the sheaves. What are the most useful metrics on the space of sheaves? How do they interact with the sheaf Laplacians and their spectra? How does this shed light on the moduli space of cellular sheaves?

## 9.2 Developing a Cheeger inequality

Following the discussion in §7.2, a structural Cheeger inequality for sheaves is connected to questions of the moduli space of sheaves. Such an inequality would describe how the spectral properties of a sheaf interact with its distance to the nearest sheaf with a nontrivial global section. Can a Cheeger inequality emerge from approximations to the constant sheaf, seeking 0-cochains of small coboundary which are constant on large sets of vertices?

## 9.3 Interactions with the derived category

The standard way of understanding sheaf cohomology is through the derived category of complexes of sheaves [GM03]. We may replace a sheaf by an injective resolution and take the cohomology of the complex of sheaves. What is the relationship between a weighted sheaf and its injective resolutions, and how do the resulting Laplacians connect with the Hodge Laplacian defined on the cochain complex? What results can be proven about their spectra? How do they interact with the standard sheaf operations? Is there a consistent way to add weights to the derived category of sheaves? We should not expect the answers to these questions to be unique due to the dagger categorical issues discussed in §3.1, but there may be constructions which are nevertheless appealing.

## 9.4 Random walks

Chung and Zhao [CZ12] considered random walks on discrete  $O(n)$ -bundles, including a definition of a sort of *PageRank* algorithm. Is it possible to define random walks on sheaves, and to what extent are such related to the Laplacian? Is there an analogous PageRank algorithm for sheaves?

## 9.5 Cones and directedness

How does one model directedness and asymmetric relations on sheaves? Sheaves of cones and sheaf cohomology taking values in categories of cones have proven useful in recent applications of sheaf theory to problems incorporating directedness [GK17,KS18]. Such methods, though promising, may be noncommutative, using semigroups and semimodules to encode the directedness, which, in turn, pushes the boundaries of existing methods in sheaf theory and non-abelian sheaf cohomology. It is fitting that such extensions of the theory would arise from novel applications in an emerging spectral theory for sheaves over directed graphs.

## References

- [AMR88] Ralph Abraham, Jerrold E. Marsden, and Tudor Ratiu. *Manifolds, Tensor Analysis, and Applications*. Springer-Verlag, New York Inc. 1988.
- [Ban15] Afonso Bandeira. *Convex Relaxations for Certain Inverse Problems on Graphs*. PhD thesis, Princeton University, 2015.
- [BCPT05] Erik G. Boman, Doron Chen, Ojas Parekh, and Sivan Toledo. On Factor Width and Symmetric H-matrices. *Linear Algebra and its Applications*, 405(1):239–248, 2005.
- [BH12] Andries E. Brouwer and Willem H. Haemers. *Spectra of Graphs*. Universitext. Springer, 2012.
- [Bjö84] Anders Björner. Posets, Regular CW Complexes and Bruhat Order. *European Journal of Combinatorics*, 5(1):7–16, 1984.
- [BN03] Mikhail Belkin and Partha Niyogi. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural Computation*, 15:1373–1396, 2003.
- [Bre97] Glen E. Bredon. *Sheaf Theory*. Number 170 in Graduate Texts in Mathematics. Springer, 2 edition, 1997.
- [BSS13] Afonso S. Bandeira, Amit Singer, and Daniel A. Spielman. A Cheeger Inequality for the Graph Connection Laplacian. *SIAM Journal on Matrix Analysis and Applications*, 34(4):1611–1630, 2013.
- [Bul18] Francesco Bullo. *Lectures on Network Systems*. CreateSpace, 2018.
- [Car13] G. Carlsson. The shape of data. In *Foundations of Computational Mathematics, Budapest 2011*, volume 403 of *London Math. Soc. Lecture Note Ser.*, pages 16–44. Cambridge Univ. Press, Cambridge, 2013.
- [Cat16] Michael Catanzaro. *A Topological Study Of Stochastic Dynamics On CW Complexes*. PhD thesis, Wayne State University, 2016.
- [CCK13] Michael J. Catanzaro, Vladimir Y. Chernyak, and John R. Klein. On Kirchhoff’s theorems with coefficients in a line bundle. *Homology, Homotopy and Applications*, 15(2):267–280, 2013.
- [CGR12] Justin Curry, Robert Ghrist, and Michael Robinson. Euler Calculus with Applications to Signals and Sensing. In *Proceedings of the Symposium in Applied Mathematics*. AMS, 2012.
- [Chu92] Fan Chung. *Spectral Graph Theory*. AMS, 1992.
- [CL06] Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006.
- [CS11] Dragoš Cvetcović and Slobodan Simić. Graph spectra in Computer Science. *Linear Algebra and its Applications*, 434(6):1545–1562, 2011.
- [Cur14] Justin Curry. *Sheaves, Cosheaves, and Applications*. PhD thesis, University of Pennsylvania, 2014.
- [CZ12] Fan Chung and Wenbo Zhao. Ranking and sparsifying a connection graph. In *International Workshop on Algorithms and Models for the Web-Graph*, Lecture Notes in Computer Science. Springer, 2012.
- [DB13] Florian Dörfler and Francesco Bullo. Kron Reduction of Graphs with Applications to Electrical Networks. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(1), January 2013.
- [Eck45] Beno Eckmann. Harmonische Funktionen und Randwertaufgaben in einem Komplex. *Commentarii Mathematici Helvetici*, 17(1):240–255, 1945.
- [EH10] Herbert Edelsbrunner and John Harer. *Computational Topology: an Introduction*. American Mathematical Society, Providence, RI, 2010.
- [Fri98] Joel Friedman. Computing Betti Numbers via Combinatorial Laplacians. *Algorithmica*, 21(4):331–346, 1998.
- [Fri15] Joel Friedman. Sheaves on Graphs, Their Homological Invariants, and a Proof of the Hanna Neumann Conjecture. *Memoirs of the American Mathematical Society*, 233(1100), January 2015.
- [Gao16] Tingran Gao. The Diffusion Geometry of Fibre Bundles. *arXiv:1602.02330*, 2016.
- [GBM16] Tingran Gao, Jacek Brodzki, and Sayan Mukherjee. The Geometry of Synchronization Problems and Learning Group Actions. *arXiv:1610.09051*, 2016.

- [Ghr14] Robert Ghrist. *Elementary Applied Topology*. CreateSpace, 2014.
- [GK17] Robert Ghrist and Sanjeevi Krishnan. Positive Alexander duality for pursuit and evasion. *SIAM Journal of Applied Algebra and Geometry*, 1(1):308–327, 2017.
- [GM03] Sergei I. Gelfand and Yuri I. Manin. *Methods of Homological Algebra*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, second edition, 2003.
- [Gro10] Mikhail Gromov. Singularities, Expanders and Topology of Maps. Part 2: From Combinatorics to Topology via Algebraic Isoperimetry. *Geometric and Functional Analysis*, 20:416–526, 2010.
- [Hat01] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.
- [HJ13] Danijela Horak and Jürgen Jost. Spectra of combinatorial Laplace operators on simplicial complexes. *Advances in Mathematics*, 244:303–336, 2013.
- [HK18] Chris Heunen and Matti Karvonen. Limits in dagger categories. *arXiv:1803.06651*, 2018.
- [HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- [HS11] Ronny Hadani and Amit Singer. Representation Theoretic Patterns in Three-Dimensional Cryo-Electron Microscopy II - The Class Averaging Problem. *Foundations of Computational Mathematics*, 11(5):589–616, 2011.
- [JS05] Charles R. Johnson and Ronald L. Smith. Closure Properties. In *The Schur Complement and Its Applications*, number 4 in Numerical Methods and Algorithms, pages 111–136. Springer, 2005.
- [Ken11] Richard Kenyon. Spanning Forests and the Vector Bundle Laplacian. *The Annals of Probability*, 39(5):1983–2017, 2011.
- [KMM04] Tomasz Kaczynski, Konstantin Mischaikow, and Marian Mrozek. *Computational Homology*, volume 157 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2004.
- [KS90] Masaki Kashiwara and Pierre Schapira. *Sheaves on Manifolds*. 1990.
- [KS18] Masaki Kashiwara and Pierre Schapira. Persistent homology and microlocal sheaf theory. *arXiv:1705.00955*, 2018.
- [LM06] Nathan Linial and Roy Meshulam. Homological Connectivity of Random 2-Complexes. *Combinatorica*, 26(4):475–487, 2006.
- [LP16] Russell Lyons and Yuval Peres. *Probability on Trees and Networks*. Number 42 in Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2016.
- [Man07] Tatiana Mantuano. Discretization of vector bundles and rough Laplacian. *Asian Journal of Mathematics*, 11(4):671–698, 2007.
- [MP18] Robert MacPherson and Amit Patel. Persistent Local Systems. 2018.
- [OPW17] Braxton Osting, Sourabh Palande, and Bei Wang. Towards Spectral Sparsification of Simplicial Complexes Based on Generalized Effective Resistance. 2017.
- [OPT+17] Nina Otter, Mason Porter, Ulrica Tillmann, Peter Grindrod, and Heather Harrington. A roadmap for the computation of persistent homology. *EPJ Data Science* 6(17), 2017.
- [Par13] Ori Parzanchevski. *High Dimensional Expanders*. PhD thesis, Hebrew University of Jerusalem, 2013.
- [PRT16] Ori Parzanchevski, Ron Rosenthal, and Ran J. Tessler. Isoperimetric Inequalities in Simplicial Complexes. *Combinatorica*, 36(2):195–227, 2016.
- [Rob17] Michael Robinson. Sheaves are the canonical data structure for sensor integration. *Information Fusion*, 36:208–224, 2017.
- [Rob18] Michael Robinson. Assignments to sheaves of pseudometric spaces. *arXiv:1805.08927*, 2018.
- [SB+18] Michael T. Schaub, Austin R. Benson, Paul Horn, Gabor Lippner, and Ali Jadbabaie. Random Walks on Simplicial Complexes and the normalized Hodge Laplacian. *arXiv:1807.05044*, 2018.
- [SHS16] Marcel K. De Carli Silva, Nicholas J. A. Harvey, and Christiane M. Sato. Sparse Sums of Positive Semidefinite Matrices. *ACM Transactions on Algorithms*, 12(1):9:1–9:17, 2016.

- 
- [Sin11] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30:20–36, 2011.
- [Spi96] Daniel A. Spielman. Linear-time encodable and decodable error-correcting codes. *IEEE Transactions on Information Theory*, 42(6):1723–1731, 1996.
- [Spi15] Daniel Spielman. Course Notes: Spectral Graph Theory. <http://www.cs.yale.edu/homes/spielman/561/>, 2015.
- [ST11] Daniel Spielman and Shang-Hua Teng. Spectral Sparsification of Graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.
- [ST14] Daniel A. Spielman and Shang-Hua Teng. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM Journal on Matrix Analysis and Applications*, 35(3):835–885, 2014.
- [Ste13] John Steenbergen. *Towards a Spectral Theory for Simplicial Complexes*. PhD thesis, Duke University, 2013.
- [Sun08] Toshikazu Sunada. Discrete Geometric Analysis. *Proceedings of Symposia in Pure Mathematics*, January 2008.
- [SW12] Amit Singer and Hau-Tieng Wu. Vector Diffusion Maps and the Connection Laplacian. *Communications in Pure and Applied Mathematics*, 65(8), 2012.
- [Tro12] Joel A. Tropp. User-Friendly Tail Bounds for Sums of Random Matrices. *Foundations of Computational Mathematics*, 12:389–434, 2012.
- [WRWX18] Chengyuan Wu, Shiquan Ren, Jie Wu, and Keli Xia. Weighted (Co)homology and Weighted Laplacian. *arXiv:1804.06990*, 2018.