CELLULAR SHEAVES OF HILBERT SPACES

Julian Joseph Gould

A DISSERTATION

in

Mathematics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the
Degree of Doctor of Philosophy

2025

Supervisor of Dissertation

Robert Ghrist, Andrea Mitchell University Professor of Mathematics

Graduate Group Chairperson

Angela Gibney, Presidential Professor of Mathematics

Dissertation Committee

Robert Ghrist, Andrea Mitchell University Professor of Mathematics
Mona Merling, Associate Professor of Mathematics
Paige Randall North, Assistant Professor of Mathematics and Computer Science, Utrecht University

ACKNOWLEDGMENTS

Growing up, my father regularly told me to "take math until you just can't take it anymore." This guidance has carried me further than either of us could have anticipated. As I complete my doctoral studies, I am deeply grateful to those who have supported me throughout this journey.

To my advisor, Robert Ghrist, I would not be finishing graduate school if it weren't for you. I know for a fact that I would have dropped out of this program without your encouragement and insight. After one of my darkest nights, you told me that whenever I don't believe in myself, to take it out on credit from the belief you have in me. I used that credit card a lot! Thank you for not charging interest. I have appreciated getting your perspective on mathematics and teaching over the past several years. You have been more generous to me than I have deserved. Thank you.

To my friend and collaborator, Miguel Lopez, you rekindled my passion for mathematics. After I had fallen out of love with homotopy type theory, it was so refreshing to begin a series of new projects with you—at the time just for fun. Your impact on this thesis cannot be understated. While our time as academic siblings has ended, I look forward to friendship and collaboration for years to come. Thank you.

To my closest math friends, Andres Mejia, Maxine Calle, Hunter Stufflebeam, and Eben Blaisdell, my most cherished memories of graduate school are with you all. As we scatter to different cities, I miss spending time with you all inside of DRL. I will never forget your jokes, your conversations, your dinners, your group chats, and most of all, your friendship. Thank you.

To my professors-turned-friends, Scott Weinstein and Robin Pemantle, thank you for the experiences I've had inside and outside of your classrooms. Your mentorship has profoundly shaped my approach to both teaching and research. I look forward to spending time with you both for many years to come. Thank you.

To my mentor, Paige North, you taught me more math than any other single person. I am in awe of your depth of knowledge. I appreciate your mentorship, your insights, and conversation. While my mathematical interests shifted away from homotopy type theory, I am enormously grateful for my time under your tutelage. Thank you.

To the people who keep the Penn math graduate program running, Reshma Tanna, Monica Pallanti, Nichole Battle-Walker, Paula Scarborough, and Robin Toney, thank you for your ceaseless work and endless kindness. I will miss your smiling faces and wonderful conversations in the math office. Thank you.

Finally, I want to thank the people closest to me. To my family, Chris, Brian and Natalie, I am enormously grateful for your love, support, and kindness. There are so many people in this world who can’t call their parents or sibling at any time of day, in any state of mind, and get a pep talk. To my partner, Maia, thank you for supporting me, for keeping me sane, for taking breaks with me, and for making the thesis writing process possible and tolerable. You saw the light at the end of the tunnel when I couldn’t. You all have helped me through my lowest moments, and gotten me back on my feet. I look forward to celebrating with you at my highest. Thank you.

ABSTRACT

CELLULAR SHEAVES OF HILBERT SPACES

Julian Joseph Gould

Robert Ghrist

This dissertation extends the theory of cellular sheaves from finite-dimensional to infinite-dimensional Hilbert spaces, thereby broadening the scope of cellular sheaf theory through the incorporation of functional and non-smooth analytic techniques. While classical cellular sheaves, particularly weighted cellular sheaves valued in finite-dimensional Hilbert spaces, have found applications in network analysis, opinion dynamics, and neural networks, some applications naturally require sheaves valued in infinite-dimensional spaces.

The passage from finite to infinite dimensions introduces fundamental complications that necessitate careful theoretical development. When restriction maps are unbounded operators with partial domains, the composition of morphisms requires precise domain considerations, cochain complexes may fail to satisfy the standard hypotheses for cohomology theory, and even elementary sheaf operations become problematic. This work systematically addresses these challenges through a trio of technical tools: the restriction categories of Cockett and Lack [26], the formalism of Hilbert complexes as developed by Brüning and Lesch [17], and the analysis of block operators between direct sums of Hilbert spaces.

The central construction of this thesis is the Hilbert sheaf. Pre-Hilbert sheaves are introduced as functors from combinatorially well-behaved acyclic categories to the category of Hilbert spaces and unbounded operators. While these objects generalize weighted cellular sheaves directly, they may exhibit pathological behavior. The dissertation therefore identifies necessary conditions for well-behaved objects, leading to the definition of Hilbert sheaves proper. A Hilbert sheaf is a pre-Hilbert sheaf whose associated coboundary operators are closable, ensuring the formation of genuine Hilbert complexes.

The theoretical framework encompasses several key developments. First, it establishes conditions under which Hilbert sheaves admit meaningful cohomology groups and spectral theory. Second, it identifies distinguished classes including bounded Hilbert sheaves (where all restriction maps are bounded) and closed Hilbert sheaves (where coboundary operators have closed range), each possessing favorable computational and theoretical properties. Third, it develops

dynamical systems on these sheaves, including heat flows, wave propagation, and nonlinear diffusion processes, which serve as tools for the study of consensus problems.

This work establishes cellular sheaves of Hilbert spaces as a rigorous mathematical framework in the intersection of algebraic topology, functional analysis, and applied mathematics, opening new avenues for the analysis of complex systems with infinite-dimensional local structure.

CONTENTS

LIST OF FIGURES

Part I

PRELIMINARIES

1

CELLULAR SHEAVES

This introductory chapter serves two distinct purposes. First, we review the established theory of cellular sheaves and their cohomology, following the work of Shepard [112] and Curry [33]. Second, we simultaneously extend this technology from the traditional setting of partially ordered sets to the more general setting of graded acyclic categories (GACs). This generalization enables a systematic treatment of network sheaves on graphs with self-loops and, more broadly, allows us to track not just that cells are glued together, but how they are glued.

Section 1.1 introduces acyclic categories—posets that permit multiple morphisms between objects while maintaining acyclicity. When equipped with a grading, these GACs provide the appropriate domain for cellular sheaves that admit cohomology and spectral theory. Section 1.2 shows how GACs arise from cell complexes, introducing weakly-regular cell structures and establishing an isomorphism between face categories and discrete exit path categories. A suitable generalization of signed incidence relations from posets to GACs is given by assigning parities to morphisms rather than object pairs. This structure enables the alternating signs necessary for cohomological cancellations while accommodating multiple parallel morphisms. With these foundations, Section 1.4 defines cellular sheaves as functors $\mathcal{F} : \mathcal{P} \to \mathcal{D}$ from a graded acyclic category $\mathcal{P}$ with a signed incidence structure to a data category $\mathcal{D}$. The grading and signed incidence combine to yield coboundary operators whose composition vanishes, producing well-defined cohomology groups that are invariant under different choices of signed incidence structure.

1.1 ACYCLIC CATEGORIES

An acyclic category $\mathcal{P}$ can be viewed as a "poset with extra arrows." Like a partially ordered set, it has no non-trivial end-to-end loops: every endomorphism $f : x \to x$ is the identity. Unlike a poset, however, the homset $\mathcal{P}(x, y)$ may contain multiple distinct morphisms, and composable chains may merge and branch. This added flexibility makes acyclic categories a natural language for encoding nonbinary incidence data. When gluing objects with an ambient notion of dimension, such as cells of a CW complex, the resulting acyclic category can be graded, assigning a rank to every object capturing how far up the chain of the category the object lives. For a comprehensive introduction to acyclic categories, see [75, Chapter 10].

Definition 1.1.1 (Acyclic category [75])**.** A category $\mathcal{A}$ is **acyclic** if it satisfies the following conditions:

(i) $\mathcal{A}(X, X) = \{\mathrm{id}_X\}$ for every object X.

(ii) If $f : x \to y$ is an isomorphism, then $y = x$ and $f = \mathrm{id}_x$.

A morphism $f : x \to y$ in an acyclic category $\mathcal{A}$ is **indecomposable** if f cannot be written as a composition of two non-identity morphisms.

Acyclic categories may be thought of as an ordered structure that mildly generalize **partially ordered sets** (posets). Every poset $(P, \leqslant)$ can be viewed as a (small) acyclic $\mathcal{P}$ with objects $\mathrm{Ob}(\mathcal{P}) = P$ and morphisms:

$$\mathcal{P}(x, y) \text{ contains a unique morphism } \iff x \leqslant y.$$

A general acyclic category $\mathcal{A}$ allows for distinct parallel morphisms between a pair of objects. One may think of $\mathcal{A}$ as a poset that allows for multiple distinct witnesses of $x \leqslant y$. Or equivalently, a poset in which one object can be greater than another in multiple ways. Each acyclic category $\mathcal{A}$ has an **underlying poset** structure defined by

$$x \leqslant y \iff \mathcal{A}(x, y) \text{ is inhabited.}$$

Equivalently, one may make a choice of a thin, wide subcategory of $\mathcal{A}$ to find a categorical representation of the underlying poset. Hence to say that $x \leqslant y$ in an acyclic category asserts the existence of a morphism $x \to y$, but does not give any information about how many distinct morphisms live in $\mathcal{A}(x, y)$, nor how they compose with other morphisms. In the interest of preserving the intuition of an acyclic category as a generalization of a poset, we will regularly use $\mathcal{P}$ to refer to an acyclic category.

To further borrow from the language of posets, we say that y **covers** x in an acyclic category $\mathcal{P}$, written $x \lhd_1 y$, if y covers x in the underlying poset. That is, $x \lhd_1 y$ if and only if $x \leqslant y$ and there is no point z such that $x \lneq z \lneq y$. When y covers x, we call each morphism in $\mathcal{P}(x, y)$ a **covering morphism**. As a shorthand, when $f : x \to y$ is a covering morphism, we will write $f \in (x \lhd_1 y)$. It is clear that every covering morphism is indecomposable, but the converse need not be true.

We may form a category **AcycCat** of small acyclic categories whose objects are acyclic categories and whose morphisms are functors.

Definition 1.1.2. Let **AcycCat** be a category consisting of the following data.

- **Objects.** An object of **AcycCat** is a small acyclic category $\mathcal{P}$.
- **Morphisms.** A morphism $f : \mathcal{P} \to \mathcal{Q}$ is a functor.

Equivalently, **AcycCat** is the full subcategory of **Cat** whose objects are acyclic.

1.1.1 *Graded acyclic categories*

Arbitrary acyclic categories are insufficiently structured for the theory of cellular sheaves. The underlying poset of an acyclic category may contain densely ordered segments such as those found in $\mathbb{Q}$ and $\mathbb{R}$, or have limit points like the ordinal $\omega + 1$. Such partial orders lack the combinatorial structure required for the algebraic topology we wish to capture. We will stipulate that our acyclic categories come equipped with a certain combinatorial map, called a grading. Our definition directly generalizes that of a graded poset.

Let $(\mathbb{N}, \leqslant)$ denote the natural numbers equipped with their usual ordering. This is a poset, and hence an acyclic category.

Definition 1.1.3 (Graded acyclic category)**.** Let $\mathcal{P}$ be an acyclic category. A **grading** on $\mathcal{P}$ is a functor $r : \mathcal{P} \to (\mathbb{N}, \leqslant)$ that satisfies the following conditions.

(i) If $\bigcup_{y \in \mathcal{P}} \mathcal{P}(y, x) = \{\mathrm{id}_x\}$, then $r(x) = 0$.

(ii) If y covers x in $\mathcal{P}$, then $r(y) = r(x) + 1$.

The value $r(x) \in \mathbb{N}$ is the **rank**, **grade**, or **dimension** of x, and the pair $(\mathcal{P}, r)$ a **graded acyclic category** (GAC).

An acyclic category admits at most one grading; when a grading exists on $\mathcal{P}$, it can be constructed inductively by first assigning rank 0 to each object x with $\bigcup_{y \in \mathcal{P}} \mathcal{P}(y, x) = \{\mathrm{id}_x\}$, and assigning $r(y) = n + 1$ if and only if y covers an object of rank n. When the grading is clear from context, we will conflate a GAC $(\mathcal{P}, r)$ with its underlying acyclic category $\mathcal{P}$.

Notation 1.1.4. When $x \leqslant y$ in a GAC and $r(y) = r(x) + k$ for a given $k \geqslant 0$, we write $x \lhd_k y$. Moreover, for $f : x \to y$ with $x \lhd_k y$, we adopt the shorthand $\mathrm{Cov}(f)$ for the collection of sequences of composable covering morphisms $(f_1, \dots, f_k)$ such that $f_k \circ \cdots \circ f_1 = f$.

Example 1.1.5. The following are examples of graded acyclic categories:

1. Any graded poset is a graded acyclic category.
2. Specifically, any finite distributive lattice is graded by its **height function** $h : \mathcal{P} \to \mathbb{N}$ that sends an element $x \in \mathcal{P}$ to the length of the longest increasing chain $x_0 \lneq x_1 \lneq \cdots \lneq x$. However, not all finite posets are graded. Consider the pentagonal lattice N_5, consisting of five elements $\{\bot, a, b, c, \top\}$ ordered such that $\bot < a < b < \top$ and $\bot < c < \top$, but c is incomparable to a and b. This lattice cannot be graded, as the (a, b)-half of the pentagon yields $r(\top) = r(\bot) + 3$, while the c-half yields $r(\top) = r(\bot) + 2$.

3. The natural numbers with their usual ordering are graded by the identity map. Indeed, this is the only infinite graded total order up to isomorphism.

4. A **quiver** is a multi-graph with directed edges (self loops allowed). To each quiver $\mathcal{Q} = (\mathcal{V}, \mathcal{E})$, we may construct a **free category**, whose objects are exactly the vertices V, with a morphism $v \to w$ for each path from v to w along edges in the quiver. This construction induces a functor

$$\mathrm{Free} : \mathbf{Quiv} \to \mathbf{Cat}.$$

   Call a quiver acyclic if there is no non-trivial path $v \rightsquigarrow v$ for any vertex v. The free category $\mathrm{Free}(Q)$ is an acyclic category if and only if the quiver Q is acyclic. Moreover, $\mathrm{Free}(Q)$ is a non-empty GAC if and only if the following conditions hold.

   (i) There is at least one **minimal** vertex with no incoming edges.

   (ii) For every vertex v, every path $m \rightsquigarrow v$ where m is a minimal vertex has the same length.

   The rank $r(v)$ is exactly the unique path length from a minimal vertex.

5. Given two graded acyclic categories $(\mathcal{P}, r)$ and $(\mathcal{Q}, s)$, the product category $\mathcal{P} \times \mathcal{Q}$ inherits a grading by $(p, q) \mapsto r(p) + s(q)$. In the event that $\mathcal{P}$ and $\mathcal{Q}$ are graded posets, the resulting order and grading correspond to that of the **Pareto order** $(p, q) \leqslant_{\mathcal{P}\times\mathcal{Q}} (p', q')$ if and only if $p \leqslant_{\mathcal{P}} p'$ and $q \leqslant_{\mathcal{Q}} q'$ on $\mathcal{P} \times \mathcal{Q}$.

6. The face poset $\mathrm{Fc}(G)$ of a regular cell complex G is graded by the dimensions of the faces (Section 1.2).

Definition 1.1.6. Let $(\mathcal{P}, r)$ be a graded acyclic category. $\mathcal{P}$ is **levelwise-finite** if, for every $n \in \mathbb{N}$, there are finitely many objects $x \in \mathcal{P}$ such that $r(x) = n$, and for each pair of objects $x, y \in \mathcal{P}$, the homset $\mathcal{P}(x, y)$ is finite.

1.1.2 *Morphisms of GACs*

Let $(\mathcal{P}, r)$ and $(\mathcal{Q}, s)$ be graded acyclic categories. While every functor $\phi : \mathcal{P} \to \mathcal{Q}$ defines a morphism of acyclic categories, such a natural transformation need not respect the gradings of $\mathcal{P}$ and $\mathcal{Q}$. There are a few important classes of grade-respecting maps for the theory of cellular sheaves.

Notation 1.1.7. Let $\mathcal{P}$ be an acyclic category. For an object $x \in \mathcal{P}$, we let $(\downarrow x)$ denote the **down-set** of x, consisting of the full subcategory of $\mathcal{P}$ containing all $y \in \mathcal{P}$ such that $y \leqslant x$ in the underlying poset. Similarly, we let $\mathrm{st}(x)$ denote the **star** of x, consisting of the full subcategory of $\mathcal{P}$ containing all objects $y \geqslant x$ above x in the underlying poset structure.

Definition 1.1.8. Let $(\mathcal{P}, r)$ and $(\mathcal{Q}, s)$ be GACs. A functor $\phi : \mathcal{P} \to \mathcal{Q}$ is a **cellular map** if the following conditions hold for all objects $\sigma \in \mathcal{P}$.

(i) $s(\phi(\sigma)) \leqslant r(\sigma)$.

(ii) $\{\phi(y) : y \in (\downarrow \sigma)\} = (\downarrow \phi(\sigma))$.

Definition 1.1.9. Let $\phi : (\mathcal{P}, r) \to (\mathcal{Q}, s)$ be a cellular map.

- ϕ is a **cellular homeomorphism** if ϕ is an isomorphism of categories.
- ϕ is an **inclusion** if ϕ is injective (on both objects and hom-sets).
- ϕ is a **covering map** if for each object $y \in \mathcal{Q}$, the pre-image of the star $\phi^{-1}(\mathrm{st}(y))$ is a disjoint union of isomorphic copies of $\mathrm{st}(y)$, each of which is mapped isomorphically onto $\mathrm{st}(y)$ by ϕ.

Remark 1.1.10. The nomenclature of "cellular," "covering," and "homeomorphism" are meant to evoke the similar concepts for cellular structures (Definition 1.2.17).

1.2 CELL COMPLEXES

Acyclic categories are readily found in the theory of cell complexes and stratified spaces, through the identification of cells with objects and gluing data with morphisms. Such acyclic categories, subject to additional regularity constraints, serve as a fruitful setting for the theory of cellular sheaves. In the literature on cellular sheaves, it is common to use the face poset of a regular cell complex or a simplicial complex as the domain for a functor [33, 45, 54, 102]. We briefly review these concepts, as well as introduce new a class of cellular decompositions, akin to regular cell structures, that generate a suitable graded acyclic category for the theory of cellular sheaves.

1.2.1 *Regular cell complexes*

Let D^n and $\overline{D^n}$ denote the open and closed balls of dimension n respectively for $n \geqslant 1$. We adopt the convention that $D^0 = \overline{D^0}$ is the one-point space.

Definition 1.2.1 (Regular cell structure)**.** Let X be a Hausdorff space. A **regular cell structure** on X consists of the following data.

- A partition of X into disjoint open **cells** $\{X_\alpha : \alpha \in P_X\}$, each endowed with a **dimension** $d_\alpha \in \mathbb{N}$.
- For every cell X_α, a continuous **attaching map** $\phi_\alpha \colon \overline{D^{d_\alpha}} \to \overline{X_\alpha}$, where $\overline{X_\alpha}$ is the closure of X_α in X.

These data must satisfy the following axioms.

(i) **Locally finite.** Every point of X has a neighborhood that meets only finitely many cells.

(ii) **Frontier.** If $\overline{X_\alpha} \cap X_\beta \neq \varnothing$ then $X_\beta \subseteq \overline{X_\alpha}$.

(iii) **Homeomorphic attaching maps.** Each ϕ_α is a homeomorphism of pairs $(\overline{D^{d_\alpha}}, D^{d_\alpha}) \xrightarrow{\cong} (\overline{X_\alpha}, X_\alpha)$, i.e. it is a homeomorphism $\overline{D^{d_\alpha}} \cong \overline{X_\alpha}$ whose restriction to the open ball is the homeomorphism $D^{d_\alpha} \cong X_\alpha$ specified above.

Notation 1.2.2. A regular cell structure for a topological space X refers to the full the collection $(X_\alpha, \phi_\alpha)_{\alpha \in P_X}$ of cells and attaching maps. We will often adopt the shorthand $X = (X_\alpha, \phi_\alpha)_{\alpha \in P_X}$, and utilize the notation $|X|$ for the underlying topological space. A cell X_α will often be identified with its index $\alpha \in P_X$.

While this definition of a regular cell complex, due to MacPherson [85] and transmitted through Curry [33], is somewhat non-standard, the attaching maps and cells are equivalent to those of a locally-finite **regular CW complex**—a locally-finite CW complex whose attaching maps are homeomorphisms.

Proposition 1.2.3. *Let X be a Hausdorff space. The data of a regular cell structure $(X_\alpha, \phi_\alpha)_{\alpha \in P_X}$ for X is equivalent to the data of a locally-finite regular CW-structure for X.*

Proof. It is straightforward to check that a locally-finite regular CW structure defines the data of a regular cell complex. Conversely, suppose that $X = \bigcup_{\alpha \in P_X} X_\alpha$ is a regular cell complex structure with attaching maps $\phi_\alpha : \overline{D^{d_\alpha}} \to \overline{X_\alpha}$. By [57, Proposition A.2], to confirm that these maps define a CW-structure for X, we must verify the following conditions.

1. Each ϕ_α restricts to a homeomorphism from D^{d_α} onto its image, these images are all distinct as we vary $\alpha \in P_X$, and their union is X.
2. For each α, the image $\phi_\alpha(\partial D^{d_\alpha})$ is contained in the union of a finite number of cells of dimension less than d_α.
3. X is topologized in the weak topology with respect to its cells. That is, a subset $A \subseteq X$ is closed if and only if $A \cap \overline{X_\alpha}$ is closed for each $\alpha \in P_X$.

Condition 1 is immediate from the definition of a cell complex. Condition 2 is also follows straightforwardly from the frontier condition and local finiteness.

Condition 3, which says that X is topologized with the weak topology, is only slightly more complicated to check. If $A \subseteq X$ is closed in X, then $A \cap \overline{X_\alpha}$ is closed by definition. Conversely, suppose that $A \cap \overline{X_\alpha}$ is closed in $\overline{X_\alpha}$ for each $\alpha \in P_X$. For a point $x \in X \backslash A$, we may take a neighborhood U_x that intersects finitely many cells $X_{\alpha_1}, \ldots, X_{\alpha_n}$ of X. In each cell X_{α_j}, since $A \cap \overline{X_{\alpha_j}}$ is closed in $\overline{X_{\alpha_j}}$, we may take an open set $V_j \subseteq X$ such that $V_j \cap A \cap \overline{X_{\alpha_j}}$ is empty. Taking the intersection

$$W_x = U_x \cap \bigcap_{j=1}^{n} V_j$$

gives a new open set in X, which contains x and is contained in $X \backslash A$. Hence we may write $X \backslash A$ as the union $\bigcup_x W_x$, proving that A is closed, and that X is a regular CW complex. □

Given a cell structure X with cells $\{X_\alpha : \alpha \in P_X\}$, the structure of the decomposition of X into cells induces a poset structure (and hence an acyclic category structure) on the index set P_X. For indices $\alpha, \beta \in P_X$, the ordering is given by:

$$\alpha \leqslant \beta \iff X_\alpha \subseteq \overline{X_\beta}.$$

The frontier condition ensures that this is a partial order structure. We call this order structure $(P_X, \leqslant)$ the **face incidence poset** (or simply the **face poset**) of X. In this poset structure, when $\alpha \lneq \beta$ and there is no cell $\alpha \lneq \gamma \lneq \beta$, we say that β **covers** α.

Remark 1.2.4. Regular cell complexes are less prevalent in algebraic topology than the usual CW complexes. The main practical difference is that regular cell complexes are more rigid and structured, giving the cellular decomposition an extremely well-behaved combinatorial flavor. If we build out attaching maps, the cells attach to each other in a clean way, with no "pinching" or "folding" allowed. This makes them easier to work with computationally and more suitable for certain applications. We catalog a few of these useful structural properties here. Given a regular cell complex X with face poset P_X, the following hold.

1. P_X is naturally graded by the dimensions of the underlying cells. One may check that for any cell X_α of dimension $n+1$, the frontier condition enforces that the image of the boundary $\partial \overline{D^{d_\alpha}}$ under the attaching map ϕ_α is a union of closures of cells of dimension n. Moreover, there can only be finitely many cells in this image by local finiteness.

2. P_X has the **diamond property**[1]; every closed interval of length 2 in P_X has exactly four elements. That is, if $\alpha \lhd_2 \gamma$ in P_X, then there are exactly two distinct elements $\beta_1, \beta_2 \in P_X$ such that $\alpha \lneq \beta_j \lneq \gamma$.

3. The work of Björner [12], as well as Danaraj and Klee [35], allows face posets of regular cell structures to be directly identified directly from their poset structure. In particular, given a poset $\mathcal{P}$, if after attaching a bottom element $\perp$, $\mathcal{P} \cup \{\perp\}$ is both thin and **shellable**—a technical condition that captures some of the global gluing properties of regular cell complexes, like how k-cells can always be thought of as being glued to $(k-1)$-cells.

4. The topological space $|X|$ can be reconstructed (up to homeomorphism) from its face poset. Hence, we may safely conflate a regular cell complex with it face poset. See [52, Proposition 1.1.1] for details.

Regular cell maps also come equipped with a class of combinatorially well-behaved maps between them.

Definition 1.2.5 (Regular cell map)**.** Let $\{X_\alpha\}_{\alpha \in P_X}$ and $\{Y_\beta\}_{\beta \in P_Y}$ be regular cell complexes. A **regular cell map** is a continuous map $f : |X| \to |Y|$ such that the following conditions hold.

(i) Each cell X_α of X is mapped by f surjectively onto a cell Y_β of Y, with $\dim(X_\alpha) \geqslant \dim(Y_\beta)$.

(ii) The restriction $f\big|_{X_\alpha} : X_\alpha \to Y_\beta$ factors as

$$X_\alpha \xrightarrow{\cong} \mathbb{R}^{\dim(X_\alpha)} \xrightarrow{P} \mathbb{R}^{\dim(Y_\beta)} \xrightarrow{\cong} Y_\beta \quad ,$$

where P is an orthogonal projection map.

If a regular cell map f is a homeomorphism, it restricts to a cell-wise homeomorphism. We call such a map a **regular cell homeomorphism**.

Each regular cell map $f : X \to Y$ induces a poset map $P_X \to P_Y$ via $\alpha \mapsto \beta$ if and only if f maps X_α onto Y_β. However, not every poset map arises from a regular cell map. For more details on regular cell complexes, their properties, and maps between them, see [29, 43].

1.2.2 *Weakly-regular cell structures*

There are many topological spaces with natural cell structures that are not regular. For example, consider the usual decomposition of the torus and the Klein bottle, as shown in Figure 1, as cell

1 Such a poset is commonly referred to as **thin** in the literature [59] [60] (dating back to Björner [12]), but we use the term "diamond property" to disambiguate from the notion of a **thin category** [126].

structures with a single vertex, two edges, and one face. These are not regular cell structures as the 2-cells are not glued homeomorphically along their boundary. While these spaces do admit regular cell structures, they require more cells. Nonetheless, these cell structures fit into a more general schema of weakly-regular cell structures, obtained by weakening the requirement of attaching maps being homeomorphisms on their boundaries.

Figure 1: Non-regular cell structures for the torus (left) and Klein bottle (right)

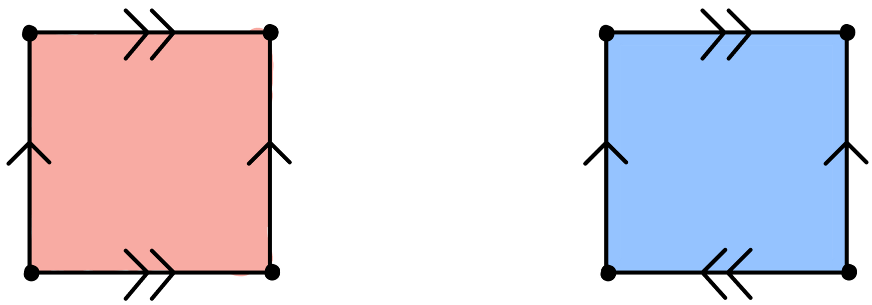

Definition 1.2.6 (Weakly-regular cell structure)**.** Let X be a Hausdorff space. A **weakly-regular cell structure** on X consists of the following data.

- A partition of X into disjoint cells $\{ X_\alpha : \alpha \in P_X \}$, each endowed with a dimension $d_\alpha \in \mathbb{N}$.
- For each cell X_α, a regular cell decomposition $\mathcal{K}_\alpha$ of the closed ball $\overline{D^{d_\alpha}}$ containing a single cell of top dimension d_α.
- For each cell X_α, a continuous attaching map

$$\phi_\alpha \colon \overline{D^{d_\alpha}} \to \overline{X_\alpha}$$

  that maps the interior D^{d_α} (the unique cell of dimension d_α in $\mathcal{K}_\alpha$) homeomorphically onto X_α.

These data must satisfy the following axioms.

(i) **Locally finite.** Every point of X has a neighborhood that meets only finitely many cells.

(ii) **Frontier.** If $\overline{X_\alpha} \cap X_\beta \neq \varnothing$ then $X_\beta \subseteq \overline{X_\alpha}$.

(iii) **Cell-wise homeomorphism.** For every cell $\sigma \in \mathcal{K}_\alpha$ the restriction $\phi_\alpha|_\sigma \colon \sigma \xrightarrow{\cong} X_\beta$ is a homeomorphism onto a unique cell X_β of the same dimension as σ.

Remark 1.2.7. The cell-wise homeomorphism condition (plus the axiom of the frontier) ensures that the attaching map ϕ_α attaches X_α along the closures of cells of dimension $d_\alpha - 1$, albeit not necessarily homeomorphically. This will preserve the graded structure of face poset while allowing for more general attaching maps.

Remark 1.2.8. This notion of a weakly-regular cell complex is similar to, but distinct from that of a **semi-regular** CW-complex [67]. There is a different generalization of regular cell structures given in the work of Shepard [112] and Curry [33]. These authors instead define a **cell complex** by replacing Part iii in the definition of a regular cell complex with the requirement that the cells $\{X_\alpha : \alpha \in P_X\} \cup \{\infty\}$ are the cells of a regular cell complex structure for the one-point compactification of X. There is no containment relationship between weakly-regular cell complexes and the cell complexes of Curry and Shepard.

Notation 1.2.9. Let $\mathcal{K}_\alpha$ be the regular cell structure for a closed d_α-ball in a weakly-regular cell structure. We let $\partial\mathcal{K}_\alpha$ denote the collection of cells in $\mathcal{K}_\alpha$ that make up the boundary $\partial\overline{D^{d_\alpha}}$. $\partial\mathcal{K}_\alpha$ contains every cell of $\mathcal{K}_\alpha$ except for the unique top-dimensional cell. For a cell $\sigma \in \partial\mathcal{K}_\alpha$, we further denote the collection of cells $\tau \neq \sigma$ such that $\tau \subseteq \overline{\sigma}$ by $\partial\sigma$.

Example 1.2.10. Every regular cell structure is a weakly-regular cell structure. In particular, for each cell X_α, the regular cell structure $\partial\mathcal{K}_\alpha$ for the boundary $\partial\overline{D^{d_\alpha}}$ may be found by identifying $\mathcal{K}_\alpha$ with the cells of X in the image $\phi_\alpha(\partial\overline{D^{d_\alpha}})$. The converse is not true, as shown in the next example.

Example 1.2.11. Consider the usual CW-structure for the circle S^1 with one 0-cell and one 1-cell.

This is not a regular cell structure, as the boundary of the unique 1-cell is not glued homeomorphically. It also fails to be a cell complex in the sense of Shepard, as the one-point compactification adds a disjoint point and fails to fix the irregularity. However, when we decompose the boundary of the 1-cell as a pair of 0-cells, we see that this is a weakly-regular cell structure. On the other hand, the usual CW-structure for the n-sphere with one 0-cell and one n-cell for $n \geqslant 2$ fails to be a weakly regular cell complex.

While weakly-regular cell structures are more flexible than regular cell structures, they are no more general topologically.

Proposition 1.2.12. *Let X be a Hausdorff space. If X admits a weakly-regular cell structure, it admits a cell regular structure.*

Proof. Let $(X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_{\alpha \in P_X}$ be a weakly-regular cell structure on X. We may refine this weakly-regular cell structure to a regular cell structure as follows. For each index $\alpha \in P_X$, take O_α to be a new 0-cell at the center of D^{d_α}, and for each $\sigma \in \partial\mathcal{K}_\alpha$, a $(\dim(\sigma) + 1)$-cell:

$$Y_{\alpha,\sigma} = \{tx + (1-t)O_\alpha \,:\, x \in \sigma,\, t \in (0,1)\}.$$

The collection $\partial\mathcal{K}_\alpha \cup \{O_\alpha\} \cup \{Y_{\alpha,\sigma} \,:\, \sigma \in \partial\mathcal{K}_\alpha\}$ is a regular cell structure for the close ball $\overline{D^{d_\alpha}}$. Removing the boundary cells $\partial\mathcal{K}_\alpha$ similarly gives a regular cell structure for the open ball. Let $\hat{O}_\alpha := \phi_\alpha(O_\alpha)$, and $\hat{Y}_{\alpha,\sigma} := \phi_\alpha(Y_{\alpha,\sigma})$. We claim that the cells

$$\bigcup_{\alpha \in P_X} \left(\{\hat{O}_\alpha\} \cup \{\hat{Y}_{\alpha,\sigma} \,:\, \sigma \in \partial\mathcal{K}_\alpha\}\right) ,$$

with attaching maps given by restricting the attaching maps for X_α, form a regular cell structure for X. We check the axioms.

(i) **Locally finite.** Let $x \in X$. There is a neighborhood $U \ni x$ that intersects only finitely many cells X_α. Since the boundary of the d_α-ball is compact, $\mathcal{K}_\alpha$ has only finitely many cells, and U can only intersect finitely many $\hat{Y}_{\alpha,\sigma}$. Hence the new structure is locally finite.

(ii) **Frontier.** This follows from the fact that the axiom of the frontier was satisfied in the original decomposition X_α, as well as in the regular cell-decompositions of the closed ball $\overline{D^{d_\alpha}}$.

(iii) **Homeomorphic attaching maps.** For each cell $\hat{Y}_{\alpha,\sigma}$, the attaching map $\psi_{\alpha,\sigma}$ is given by the restriction $\phi_\alpha|_{\overline{Y_{\alpha,\sigma}}}$, where we have identified the open ball $D^{\dim(\hat{Y}_{\alpha,\sigma})}$ with $Y_{\alpha,\sigma}$ itself. $\phi_{\alpha,\sigma}$ is hence a homeomorphism of $Y_{\alpha,\sigma}$ onto $\hat{Y}_{\alpha,\sigma}$ by definition. Moreover, the boundary of $Y_{\alpha,\sigma}$ inside of $\overline{D^{d_\alpha}}$ consists of the cell $\sigma \in \partial\mathcal{K}_\alpha$, O_α, and cells of the form $Y_{\alpha,\tau}$ where $\tau \leqslant \sigma$ in $\mathcal{K}_\alpha$. Every cell other than σ is mapped homeomorphically into $X_\alpha \subseteq X$, and σ is mapped homeomorphically onto its image by the definition of a weakly-regular cell complex. Therefore $\psi_{\alpha,\sigma}$ is a homeomorphism of pairs, as required.

□

The cell-wise homeomorphism condition imposes a strict combinatorial structure on the closed-ball decompositions $\{\mathcal{K}_\alpha\}_{\alpha \in P_X}$. The crux of the combinatorial structure is the following lemma.

Lemma 1.2.13. *Let X_α be a cell in a weakly regular cell structure with attaching map $\phi_\alpha : \overline{D^{d_\alpha}} \to \overline{X_\alpha}$. Let $\mathring{\phi}_\alpha := \phi_\alpha|_{D^{d_\alpha}}$ denote the restriction of ϕ_α to the interior of the ball. The cell structure $\mathcal{K}_\alpha$ can be determined from $\mathring{\phi}_\alpha$.*

Proof. For each point $x \in \partial D^{d_\alpha}$, fix a path $\gamma_x : [0,1) \to D^{d_\alpha}$ such that $\lim_{t \to 1} \gamma_x(t) = x$. The induced path $\mathring{\phi}_\alpha \circ \gamma_x$ in X_α has a limit in $\overline{X_\alpha}$. For a boundary-cell $X_\tau \subseteq \overline{X_\sigma}$, set:

$$S_\tau := \{x \in \partial D^{d_\alpha} \ : \ \lim_{t \to 1} \mathring{\phi}_\alpha(\gamma_x(t)) \in X_\tau\}.$$

Since $\mathring{\phi}_\alpha$ is a surjection, S_τ is non-empty. Moreover, since $\mathring{\phi}_\alpha$ is the restriction of a cell-wise homeomorphism, $S_\tau \subseteq \partial D^{d_\alpha}$ can be decomposed into path-connected components, each of which is homeomorphic to X_τ. Each connected component $Y \subseteq S_\tau$ is a cell in $\mathcal{K}_\alpha$. By a direct inductive argument on cell-dimension, the attaching map $\psi : \overline{D^{d_\tau}} \to \overline{Y}$ can be determined as well. □

This lemma essentially enforces a weak uniqueness condition on weakly-regular cell structures. Given a cell X_α in a weakly-regular cell structure, we cannot change the boundary-decomposition $\mathcal{K}_\alpha$ without also changing the interior of the attaching map ϕ_α. This forces the following combinatorial corollary.

Corollary 1.2.14. *Let* $(X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_{\alpha \in P_X}$ *be a weakly-regular cell structure. Let* $\partial\mathcal{K}_\alpha$ *denote the cells of* $\mathcal{K}_\alpha$ *making up the boundary of the sphere* $\partial\overline{D^{d_\alpha}}$. *Let* $\sigma \in \partial\mathcal{K}_\alpha$ *be a boundary cell, and suppose* σ *is mapped homeomorphically onto the cell* X_τ *by* ϕ_α. *The closure* $\overline{\sigma} \subseteq \partial\overline{D^{d_\alpha}}$ *inherits a regular cell structure from* $\partial\mathcal{K}_\alpha$, *denoted* $\mathcal{K}_\sigma$. *Moreover, the restriction* $\phi_\alpha\big|_{\overline{\sigma}} : \overline{\sigma} \to \overline{X_\tau}$ *factors as*

$$\begin{array}{ccc} \overline{\sigma} & \xrightarrow{\psi_\sigma} & \overline{D^{d_\beta}} \\ & {\scriptstyle \phi_\alpha|_{\overline{\sigma}}} \searrow & \downarrow {\scriptstyle \phi_\beta} \\ & & \overline{X_\beta} \end{array}$$

where ψ_σ *is a regular cell homeomorphism from the cell structure on* $\overline{\sigma}$ *inherited from* $\mathcal{K}_\alpha$, *and* $\overline{D^{d_\beta}}$ *has the cell structure of* $\mathcal{K}_\beta$. *Moreover, any such regular cell homeomorphism induces the same correspondence at the level of cells.*

Suppose we have a trio of cells $X_\alpha, X_\beta, X_\gamma$ such that $X_\alpha \subseteq \partial X_\beta$ and $X_\beta \subseteq \partial X_\gamma$. Let $\sigma \in \partial\mathcal{K}_\beta$ be a cell mapped homeomorphically by ϕ_β onto X_α, and $\tau \in \partial\mathcal{K}_\gamma$ be a cell mapped homeomorphically by ϕ_γ onto X_β. By Corollary 1.2.14, there a regular cell-homeomorphism $\psi : \overline{\tau} \to (\overline{D^{d_\beta}}, \mathcal{K}_\beta)$. Let $\rho := \psi^{-1}(\sigma)$ denote the unique cell in $\overline{\tau}$ that is mapped homeomorphically onto σ by ψ. We call this cell ρ, which is homeomorphic to X_α, the **cell in** $\overline{\tau}$ **over** σ.

1.2.2.1 *Maps of weakly-regular cell complexes*

We may generalize regular cell maps to act between weakly regular cell complexes.

Definition 1.2.15. Let $(X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_{\alpha \in P_X}$ and $(Y_\beta, \psi_\beta, \mathcal{J}_\beta)_{\beta \in P_Y}$ be weakly-regular cell structures. A **weakly-regular cell map** is a continuous function $f : |X| \to |Y|$ such that the following conditions hold.

(i) Each cell X_α of X is mapped by f surjectively onto a cell Y_β of Y, with $\dim(X_\alpha) \geqslant \dim(Y_\beta)$.

(ii) There is a regular cell map $\hat{f}_{\alpha\beta} : \overline{D^{d_\alpha}} \to \overline{D^{d_\beta}}$ such that the following diagram commutes:

$$\begin{array}{ccc} \overline{D^{d_\alpha}} & \xrightarrow{\hat{f}_{\alpha\beta}} & \overline{D^{d_\beta}} \\ {\scriptstyle \phi_\alpha}\downarrow & & \downarrow{\scriptstyle \psi_\beta} \\ \overline{X_\alpha} & \xrightarrow[f|_{\overline{X_\alpha}}]{} & \overline{Y_\beta} \end{array}$$

Weakly-regular cell maps have the following combinatorial property.

Proposition 1.2.16. *Suppose* $f : |X| \to |Y|$ *is a weakly regular cell map. If* f *maps* X_α *surjectively onto* Y_β*, then* $f(\overline{X}_\alpha) = \overline{Y}_\beta$.

Proof. When f is a regular cell map of regular cell complexes, f has the desired property [52, Proposition 1.1.3]. Therefore $\hat{f}_{\alpha\beta}$ maps $\overline{D^{d_\alpha}}$ surjectively onto $\overline{D^{d_\beta}}$, and $f(\overline{X}_\alpha) = \overline{Y}_\beta$. □

There are several classes of weakly-regular cell maps that will prove important for sheaf operations on Hilbert spaces. We outline them here.

Definition 1.2.17. Let $f : (X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_{\alpha \in P_X} \to (Y_\beta, \psi_\beta, \mathcal{J}_\beta)_{\beta \in P_Y}$ be a weakly regular cell map.

- f is a **homeomorphism** if the underlying map of topological spaces $f : |X| \to |Y|$ is a homeomorphism. Such a map is necessarily a cell-wise homeomorphism.
- f is an **inclusion** if $f : |X| \to |Y|$ is injective. Such a map is necessarily injective on the cells of X.
- For a cell $\beta \in P_Y$, let the **star** of β, denoted $\mathrm{st}(\beta)$, denote the collection of all cells Y_γ of Y for which Y_β is a face. The map f is a **covering map** if for each cell β, the preimage $f^{-1}(\mathrm{st}(\beta))$ is a disjoint union of homeomorphic copies of $\mathrm{st}(\beta)$, each of which is mapped homeomorphically onto $\mathrm{st}(\beta)$ by f.

Remark 1.2.18. Each weakly-regular cellular map (resp. homeomorphism / injection / covering map) $f : |X| \to |Y|$ functorially induces a cellular map (resp. homeomorphism / injection / covering map) $\mathrm{Fc}(f) : \mathrm{Fc}(X) \to \mathrm{Fc}(Y)$.

1.2.2.2 *Face categories*

There are multiple ways to associate an order structure to the cells of a weakly-regular cell structure X. First, one can form the usual face poset, where $\alpha \leqslant \beta$ if and only if $X_\alpha \subseteq \overline{X_\beta}$. However, this structure no longer uniquely specifies the topological space $|X|$ up to homeomorphism. As a simple example, consider the following poset, with height corresponding to cell-dimension.

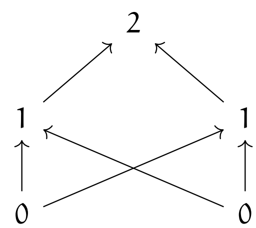

While the poset specifies that the 0-cells and 1-cells are glued into a circle with two arcs, the gluing of the 2-cell is ambiguous. The 2-cell could be glued so as to form a disk, or it could be glued so as to wrap around the boundary circle n times for any $n \geqslant 2$. This generates an infinite family of non-homeomorphic topological spaces, each with the same face poset.

The essential problem captured by the previous example is that while the face poset can capture when one cell of X is part of the boundary of another, it fails to capture the multiplicity. The face poset cannot see *how* a cell X_β is included as part of the image of the attaching map ϕ_α. To capture non-boolean gluing data, we may instead associate the structure of an acyclic category to the cell structure, with parallel morphisms capturing the multiplicity of the inclusion of a cell in the boundary of another, and the compositional structure capturing the orientations. The key to the construction is the boundary compatibility imposed by Corollary 1.2.14.

Definition 1.2.19 (Face category)**.** Let $(X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_{\alpha \in P_X}$ be a weakly-regular cell structure for a Hausdorff space X. The associated **face category** $\mathrm{Fc}(X)$ consists of the following data.

- **Objects.** Each index $\alpha \in P_X$ is an object of $\mathrm{Fc}(X)$.
- **Morphisms.** For each α, the identity id_α is the unique morphism in $\mathrm{Fc}(X)(\alpha, \alpha)$. For $\alpha \neq \beta$, there is a unique morphism in $\mathrm{Fc}(X)(\alpha, \beta)$ for every cell σ in the boundary decomposition $\partial\mathcal{K}_\beta$ that maps homeomorphically onto X_α. Equivalently,

$$\mathrm{Fc}(X)(\alpha, \beta) = \left\{\text{connected components of } \phi_\beta^{-1}(X_\alpha) \text{ in } \partial\overline{D^{d_\beta}}\right\}.$$

- **Composition.** Identity maps are formal and compose in the necessary way. Given a pair of composable morphisms $\alpha \xrightarrow{\sigma} \beta \xrightarrow{\tau} \gamma$ for distinct α, β, γ, the composition $\tau \circ \sigma$ is defined to be the cell $\rho \in \partial\overline{\tau}$ over $\sigma \in \partial\mathcal{K}_\beta$, as guaranteed by Corollary 1.2.14.

- **Associativity.** Consider a sequence of composable morphisms

$$X_\alpha \xrightarrow{\sigma} X_\beta \xrightarrow{\tau} X_\gamma \xrightarrow{\rho} X_\delta\,,$$

and pick regular cellular homeomorphisms $\psi_\tau : \overline{\tau} \xrightarrow{\cong} \mathcal{K}_\beta$, and $\psi_\rho : \overline{\rho} \xrightarrow{\cong} \mathcal{K}_\gamma$, and $\psi_{\rho\circ\tau} : \overline{\rho\circ\tau} \xrightarrow{\cong} \mathcal{K}_\beta$ from Corollary 1.2.14. It is straightforwardly verified that the following diagram commutes

$$\begin{array}{ccc} \overline{\rho\circ\tau} & \xrightarrow{\psi_{\rho\circ\tau}} & \mathcal{K}_\beta \\ & {\scriptstyle \psi_\rho|_{\overline{\rho\circ\tau}}} \searrow \qquad \nearrow {\scriptstyle \psi_\tau} & \\ & \overline{\tau} & \end{array},$$

where we have identified $\overline{D^{d_\beta}}$ with its regular cell structure $\mathcal{K}_\beta$. It follows that $(\rho\circ\tau)\circ\sigma = \rho\circ(\tau\circ\sigma)$, verifying associativity.

Example 1.2.20. We return to the torus and the Klein bottle, each parameterized with one vertex, two edges, and a single face, as shown in Figure 2.

Figure 2: Non-regular cell structures for the torus (left) and Klein bottle (right)

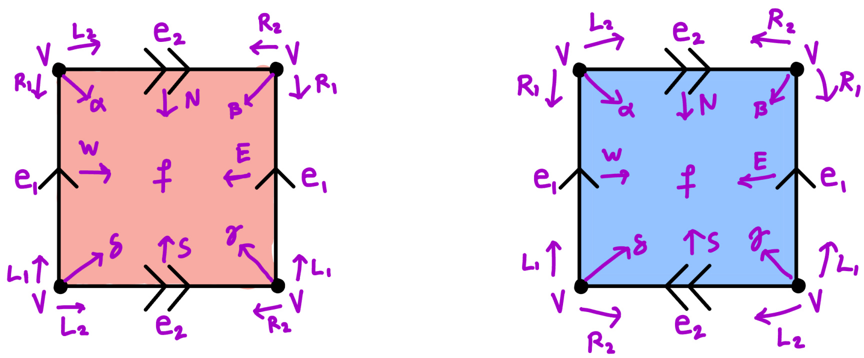

Both of these cell structures are glued with the same multiplicities, and give rise to the same labeled quiver of objects and morphisms:

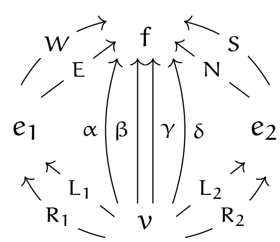

However, the morphisms compose differently in the two face categories. For example, in the torus (left), we see that $\delta = W\circ L_1 = S\circ L_2$, while in the Klein bottle, we have that $\delta = W\circ L_1 = S\circ R_2$.

Hence the compositional structure of the face category encodes the orientation information of how the top and bottom edges of the squares are glued together.

The face category $\mathrm{Fc}(X)$ of a weakly-regular cell structure X is a natural generalization of the usual face-poset of a regular cell structure. Indeed, when X is a regular cell complex, the face category $\mathrm{Fc}(X)$ is exactly the usual face poset. This follows from the fact that for a regular cell complex, a cell X_α in the boundary of X_β is included into the boundary in exactly one way. The face category shares some of the same nice order properties as the face poset, such as having a grading.

Proposition 1.2.21. *The face category* $\mathrm{Fc}(X)$ *of a weakly-regular cell structure is a GAC, with grading given by cell dimension.*

Proof. The morphism structure of $\mathrm{Fc}(X)$ makes the category acyclic. The grading follows from the fact that each cell of dimension $n+1$ is attached along the closures of cells of dimension n. Every morphism $\alpha \to \beta$ with $d_\beta = d_\alpha + n$ factors through a cell of dimension $d_\alpha + j$ for each $1 \leqslant j \leqslant n-1$. Hence every indecomposable morphism is a covering morphism, and $\mathcal{P}$ inherits a grading. $\square$

The face category of a weakly-regular cell structure also generalizes some of the combinatorial features of the face poset. For example, consider this variant of the diamond property.

Proposition 1.2.22. *Let* $\mathrm{Fc}(X)$ *be the face category of a weakly-regular cell structure* $(X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_\alpha$*. Let* $\alpha \lhd_2 \gamma$*, and* $\rho : \alpha \to \gamma$*. There are exactly two distinct composable pairs of covering morphisms in* $\mathrm{Cov}(\rho)$*.*

Proof. The map $\rho : \alpha \to \gamma$ corresponds to a cell $\rho \in \partial\mathcal{K}_\gamma$ of dimension d_α that is mapped by ϕ_γ onto X_α. In the regular cell structure $\mathcal{K}_\gamma$, there are exactly two cells τ_1, τ_2 such that $\alpha \lhd_1 \tau_j \lhd_1 \gamma$, where we have identified γ with the open cell of dimension d_γ in $\mathcal{K}_\gamma$. We get a pair of maps $\tau_j : \beta_j \to \gamma$, where X_{β_1} and X_{β_2} denote the (not necessarily distinct) cells of X that τ_1 and τ_2 are mapped to by ϕ_γ. Every pair of composable arrows through a cell of dimension $d_\alpha + 1$ that compose to ρ must include τ_1 or τ_2. For each τ_j, there is a unique cell $\sigma_j \in \mathcal{K}_{\beta_j}$ such that $\tau_j \sigma_j = \rho$, proving the result. $\square$

Finally, we get a reconstruction result analogous to that of regular cell structures.

Theorem 1.2.23. *Let* $(X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_{\alpha \in P_X}$ *be a weakly-regular cell structure for* $|X|$*. From the face category* $\mathrm{Fc}(X)$*, we can reconstruct the topological space* X *up to homeomorphism.*

Proof. We construct a topological space $|\mathrm{Fc}(X)|$ inductively by dimension. Let X^n denote the n-skeleton of X:

$$X^n := \bigcup_{\alpha \,:\, d_\alpha \leqslant n} X_\alpha \,.$$

At the n^{th} step of the inductive construction, we will build a space $|\operatorname{Fc}(X)|^n$ and a homeomorphism $\Phi^n : |\operatorname{Fc}(X)^n| \xrightarrow{\cong} X^n$.

As the base case, take $|\operatorname{Fc}(X)|^0$ to be the discrete topological space on the set of points $\{\alpha \in \operatorname{Fc}(X) \,:\, d_\alpha = 0\}$, where d_α is the rank of the object $\alpha \in \operatorname{Fc}(X)$. This space is trivially homeomorphic to X^0 under the homeomorphism $\Phi^0 : \alpha \mapsto X_\alpha$.

For the inductive step, suppose we have homeomorphism $\Phi^n : |\operatorname{Fc}(X)|^n \to X^n$. Fix an object $\sigma \in \operatorname{Fc}(X)$ of rank $d_\sigma = n+1$. First, we determine the regular cell structure $\mathcal{K}_\sigma$ on $\overline{D^{n+1}}$. For each distinct morphism $p : \alpha \to \sigma$ in $\operatorname{Fc}(X)$, there is a distinct cell of dimension d_α in $\mathcal{K}_\sigma$. Moreover, given a pair of cells $\alpha \xrightarrow{p} \sigma$ and $\beta \xrightarrow{q} \sigma$, we have an incidence $p \leqslant q$ in $\mathcal{K}_\sigma$ if and only if there is an $s : \alpha \to \beta$ such that $q \circ s = p$. Since $\mathcal{K}_\sigma$ is a regular cell structure, the decomposition $\mathcal{K}_\sigma$ for $\overline{D^{n+1}}$ can be reconstructed up to homeomorphism from the data of this poset [52, 86]. Without loss of generality, suppose we have reconstructed $\mathcal{K}_\sigma$ exactly.

Pick a regular cell map $f_\sigma : \partial D^{n+1} \to |\operatorname{Fc}(X)|^n$ that maps each cell $p : \alpha \to \sigma$ in $\partial\mathcal{K}_\sigma$ homeomorphically onto the cell $|\operatorname{Fc}(X)|_\alpha \subseteq |\operatorname{Fc}(X)|$. Such a map must exist by the definition of a weakly-regular cell complex. Consider the following commutative diagram in the category **Top** :

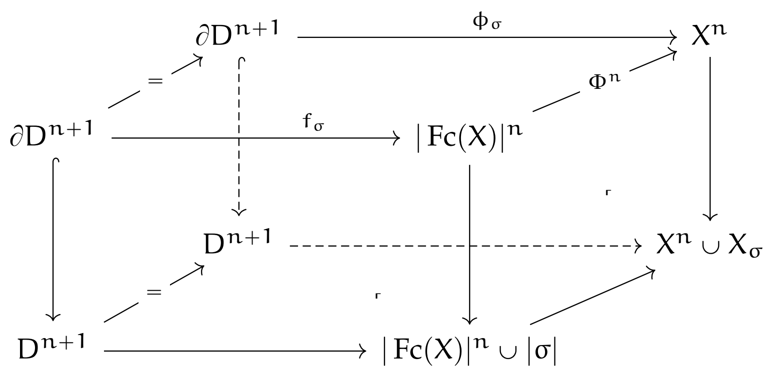

The front and back faces of the cube are pushout squares. Since three labeled arrows from the front face to the back face are homeomorphisms, there is an induced homeomorphism $|\operatorname{Fc}(X)|^n \cup |\sigma| \to X^n \cup X_\sigma$. Repeating this process for all cells of dimension $n+1$ allows us to extend Φ^n to a homeomorphism $\Phi^{n+1} : |\operatorname{Fc}(X)|^{n+1} \xrightarrow{\cong} X^{n+1}$. □

Remark 1.2.24. The crux of this argument is that the compositional structure of the face category encodes all the gluing information of the attaching maps, such as the "orientation" of the gluing. For example, when considering the torus and the Klein bottle (Example 1.2.20), the constructions diverge exactly when the two-cell is glued in with different orientations along the southern edge.

1.2.2.3 *An exit path perspective*

We have constructed the face category of a weakly regular cell complex in a combinatorial manner, with morphisms capturing inclusions of cells into the boundaries of other cells. We may instead link the face category to the theory of exit path categories of stratified spaces. In short, an exit path in a stratified space is a path γ that only moves from cells of lower dimension to cells of higher dimension. The space of exit paths, up to homotopies which preserve the exit-path property, yield a category of exit paths. The exit path categories (and their opposite entrance path categories) are prevalent in discrete Morse theory [75, 93], and are invaluable for representations of constructible sheaves and stacks on stratified spaces [33, 81, 83] following the work of Treumann [118].

Definition 1.2.25 (Tagged weakly-regular cell structure)**.** Let $X = (X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_\alpha$ be a weakly-regular cell structure. A **tagging** for X is the choice of a point $x_\alpha \in X_\alpha$ for each index α. We call the collection $(X_\alpha, x_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_\alpha$ a **tagged weakly-regular cell complex**.

Definition 1.2.26 (Discrete exit path)**.** Let $X = (X_\alpha, x_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_\alpha$ be a tagged weakly-regular cell structure. A **discrete exit path** in X is a Moore path $\gamma : [0, T] \to |X|$ satisfying the following conditions.

(i) $\gamma(0) = x_\alpha$ for some α.

(ii) $\gamma(T) = x_\beta$ for some β.

(iii) The function $t \mapsto \dim(\gamma(t))$ is weakly increasing, where $\dim(x)$ is the dimension of the unique cell of X containing x.

Remark 1.2.27. The "discrete" qualifier is not meant to imply that anything about the path γ is discrete. Instead, we are limiting the endpoints of our paths to the discrete set of points $\{x_\alpha\}_\alpha$. This is in contrast to the usual definition of an exit path in a stratified space, which can begin and end at any point.

Definition 1.2.28 (exit path homotopy)**.** Let $\gamma : [0, T_0] \to |X|$ and $\eta : [0, T_1] \to |X|$ be exit paths from x_α to x_β in a tagged weakly-regular cell structure. An **exit path homotopy** $H : \gamma \Rightarrow \eta$ from γ to η is a continuous map $H : [0, 1] \times [0, \infty) \to |X|$ that satisfies the following conditions.

(i) $H(0, t) = \gamma(t \wedge T_0)$ for all $t \geqslant 0$.

(ii) $H(1, t) = \eta(t \wedge T_1)$ for all $t \geqslant 0$.

(iii) $H(s, -) : [0, \infty) \to |X|$ is an eventually-constant exit path for all $s \in [0, 1]$.

That is, H is a fixed-endpoint Moore-homotopy from γ to η such that every slice is an exit path.

Given a pair of discrete exit paths γ, η with the same endpoints, if there is an exit path homotopy $H : \gamma \Rightarrow \eta$, we say that γ and η are **exit path homotopic**, and write $\gamma \sim \eta$. This is an equivalence relation, and we denote the equivalence class of γ by $[\gamma]$.

Definition 1.2.29 (Discrete exit path category)**.** Let $X = (X_\alpha, x_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_\alpha$ be a tagged weakly-regular cell structure. The **discrete exit path category** of X, denoted $\mathbf{DExit}_X$, is the category consisting of the following data.

- **Objects.** The objects of $\mathbf{DExit}_X$ are the tags $\{x_\alpha\}_\alpha$.
- **Morphisms.** The set of morphisms from x_α to x_β is the set of exit path homotopy classes of discrete exit paths from x_α to x_β. That is,

$$\mathbf{DExit}_X(x_\alpha, x_\beta) = \big\{[\gamma] \,:\, \gamma \text{ is a discrete exit path from } x_\alpha \text{ to } x_\beta\big\}.$$

- **Composition.** Given $x_\alpha \xrightarrow{[\gamma]} x_\beta \xrightarrow{[\eta]} x_\delta$, the composition $[\eta] \circ [\gamma]$ is given by $[\eta * \gamma]$, where $\eta * \gamma$ denotes the usual concatenation of Moore paths.

This is easily seen to be a valid category. Moreover, since $\dim(\gamma(t))$ must be weakly increasing in t and each cell X_α is contractible, $\mathbf{DExit}_X$ is an acyclic category. In fact, it's a GAC.

Proposition 1.2.30. $\mathbf{DExit}_X$ *is a graded acyclic category, graded by* $r(x_\alpha) = d_\alpha$.

Proof. To confirm that $\mathbf{DExit}_X$ is graded by dimension, we must show that every morphism $[\gamma] : x_\alpha \to x_\delta$ with $d_\delta = d_\alpha + n$ factors as

$$x_\alpha \xrightarrow{[\gamma_1]} x_{\beta_1} \xrightarrow{[\gamma_2]} \cdots \xrightarrow{[\gamma_{n-1}]} x_{\beta_{n-1}} \xrightarrow{[\gamma_n]} x_\delta \quad (\text{with } [\gamma] : x_\alpha \to x_\delta), $$

where $d_{\beta_j} = d_\alpha + j$. If $n = 1$, this vacuous, so suppose $n \geqslant 2$.

The representative discrete exit path $\gamma : [0, T] \to |X|$ from x_α to x_δ lifts to a unique path $\hat{\gamma}$ in $\overline{D^{d_\delta}}$ such that $\phi_\delta(\hat{\gamma}(t)) = \gamma(t)$. For each cell $\sigma \in \mathcal{K}_\delta$ Let $X_{\zeta(\sigma)}$ denote the cell of X onto which σ is mapped homeomorphically. Tag each $\sigma \in \mathcal{K}_\delta$ with $y_\sigma := \phi_\delta^{-1}(x_{\zeta(\sigma)})$. The path $\hat{\gamma}$ is a discrete exit path with respect to this tagged weakly-regular cell structure. Moreover, inside the ball $\overline{D^{d_\delta}}$, the path $\hat{\gamma}$ is exit path homotopic to a discrete exit path $\hat{\eta}$ which passes through the tags $y_{\sigma_1}, \ldots, y_{\sigma_{n-1}}$ in order. Pushing forward by ϕ_δ yields a discrete exit path $\eta \sim \gamma$ in X that passes through $x_{\beta_j} := x_{\zeta(\sigma_j)}$. Taking η_j to be the sub-path from $x_{\beta_{j-1}}$ to x_{β_j}, we find that $[\gamma] = [\eta_n] \circ [\eta_{n-1}] \circ \cdots \circ [\eta_1]$. □

Theorem 1.2.31. *Let $X = (X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_\alpha$ be a weakly-regular cell structure equipped with a tagging $\{x_\alpha\}_\alpha$. There is an isomorphism of categories:*

$$\mathbf{DExit}_X \cong \mathrm{Fc}(X)\,.$$

Proof. The isomorphism $\Psi : \mathrm{Fc}(X) \to \mathbf{DExit}_X$ acts on objects and morphisms as follows.

- **Objects.** $\Psi(\alpha) := x_\alpha$.
- **Morphisms.** For an identity morphism id_α, take $\Psi(\mathrm{id}_\alpha)$ to be the exit path homotopy class of a constant path on x_α. For a non-identity morphism $\sigma : \alpha \to \beta$, take the straight-line path $\hat{\gamma}_\sigma : [0,1] \to \overline{D^{d_\beta}}$ from $\phi_\beta^{-1}(x_\alpha) \in \sigma \subseteq \partial D^{d_\beta}$ to $\phi_\beta^{-1}(x_\beta) \in D^{d_\beta}$. The map $\gamma_\sigma(t) := \phi_\beta(\hat{\gamma}_\sigma(t))$ is a discrete exit path from x_α to x_β. We take $\Psi(\sigma) := [\gamma_\sigma]$.

To verify that Ψ is a functor, we may essentially repeat the argument of Proposition 1.2.30. Since the domain and codomain are graded, it suffices to check that $\Psi(\tau \circ \sigma) = \Psi(\tau) \circ \Psi(\sigma)$ when $\alpha \xrightarrow{\sigma} \beta \xrightarrow{\tau} \delta$ are a composable pair of covering morphisms. The path $\gamma_\sigma * \gamma_\tau$ lifts a unique path $\hat{\eta}$ in $\mathcal{K}_\delta$ that follows straight-line segments $\phi_\delta^{-1}(x_\alpha) \rightsquigarrow \hat{x}_\beta \rightsquigarrow \phi_\delta^{-1}(x_\delta)$, where $\hat{x}_\delta := \phi_\delta^{-1}(x_\delta)$, and $\hat{x}_\alpha$ and $\hat{x}$ denote the unique points in $\phi_\delta^{-1}(x_\alpha) \cap (\tau \circ \sigma)$ and $\phi_\delta^{-1}(x_\beta) \cap \tau$ respectively. The discrete exit path $\hat{\eta}$ with respect to the corresponding tagging on $\mathcal{K}_\delta$ is exit path homotopic to the straight-line path $\hat{\gamma}_{\tau\circ\sigma} : \hat{x}_\alpha \rightsquigarrow \hat{x}_\delta$. The pushforward of this path by ϕ_δ is exit path homotopic to $\gamma_\sigma * \gamma_\tau$, and is exactly the path $\gamma_{\tau\circ\sigma}$. This proves associativity.

Ψ is clearly a bijection on objects, but we still must verify that Ψ is full and faithful. Since $\mathrm{Fc}(X)$ and $\mathbf{DExit}_X$ are graded, it suffices to verify $\Psi : \mathrm{Fc}(X)(\alpha, \beta) \to \mathbf{DExit}_X(x_\alpha, x_\beta)$ is a bijection when β covers α. This follows from how $\phi_\beta^{-1}(x_\alpha)$ picks out a distinct point in every cell of $\mathcal{K}_\beta$ that is mapped to X_α by ϕ_β. Two discrete exit paths with respect to the induced tagging on $\mathcal{K}_\delta$ are exit path homotopic if and only if they begin and end at the same point. This easily yields a one-to-one correspondence $\mathrm{Fc}(X)(\alpha, \beta) \cong \mathbf{DExit}_X(x_\alpha, x_\beta)$. ☐

Remark 1.2.32. This theorem also shows that the choice of the tagging $\{x_\alpha \in X_\alpha\}_\alpha$ does not impact the categorical structure of $\mathbf{DExit}_X$. Hence, we may safely discuss the exit path category on an *un-tagged* weakly-regular cell structure without ambiguity.

It will also be useful to define **discrete entrance paths** on a tagged weakly-regular cell structure, and the **discrete entrance path category**. A Moore path $\gamma : [0,T] \to |X|$ is a discrete entrance path if and only if the reversed path $\gamma^{\mathrm{rev}}(t) := \gamma(T - t)$ is a discrete exit path. Hence, the quantity $\dim(\gamma(t))$ is weakly *decreasing* in t instead of increasing. Two entrance paths γ and η are **entrance path homotopic** if there is a fixed-endpoint Moore homotopy $H : \gamma \Rightarrow \eta$ that is an entrance path on each slice. Analogously to the discrete exit path category, the discrete entrance path category $\mathbf{DEnt}_X$ again has entrance path homotopy classes of discrete entrance paths for objects, with composition given by $[\gamma] \circ [\eta] = [\eta * \gamma]$. By identifying $\gamma^{op} = \gamma^{rev}$, there is a canonical isomorphism $\mathbf{DEnt}_X \cong \mathbf{DExit}_X^{op}$.

While the entrance path category $\mathbf{DEnt}_X$ is still an acyclic category, it is not necessarily graded. In particular, if X has cells of every dimension, then $\mathbf{DEnt}_X$ will have no minimal objects with respect to the underlying poset structure, and will fail to have a grading.

1.3 SIGNED INCIDENCE STRUCTURES

In order to add and subtract maps along morphisms in a GAC in a manner yielding the correct cancellations for cohomology, we need to assign parities to the covering morphisms. This assignment is a mild generalization of the signed incidence relation on a poset [33, 45, 52, 54], which we now briefly review.

Definition 1.3.1 (signed incidence relation on a poset)**.** A **signed incidence relation** on a graded poset $(\mathcal{P}, r)$ is a map $[- : -] : \mathcal{P} \times \mathcal{P} \to \{-1, 0, 1\}$ that satisfies the following conditions.

(i) $[x : y] \neq 0$ if and only if y covers x.

(ii) For any $x \leqslant z$, $\sum_{y \in \mathcal{P}} [x : y][y : z] = 0$.

(iii) For any $e \in \mathcal{P}$ of rank 1, there are exactly two cells v_0, v_1 covered by e, and $[v_0 : e] = -[v_1 : e]$.

Remark 1.3.2. We make a few remarks on this definition.

1. Condition (i) ensures that the data of a signed incidence relation is determined by how it acts on covering pairs $x \lhd_1 y$.

2. Condition (ii) is capturing information about intervals of length 2 in $\mathcal{P}$. The specified sum $\sum_{z \in \mathcal{P}} [x : z][z : y]$ is trivially equal to 0 whenever $r(z) \neq r(x) + 2$. Moreover, when $r(z) = r(x) + 2$, the sum reduces to:

$$\sum_{y \in \mathcal{P}} [x : y][y : z] = \sum_{y \,:\, x \lhd_1 y \lhd_1 z} [x : y][y : z] \,.$$

3. Condition (iii) is not always included in the definition of a signed incidence relation. Following [54], we adopt this convention of opposite parities on one-cells to eventually enforce a correspondence between global sections of cellular sheaves and the kernel of a coboundary operator.

Our first task is to generalize this definition to a graded acyclic category. Incidence algebras for acyclic categories have been studied [99], but without a clear analogue to the signed incidence relation above. However, the observation that a morphisms in an acyclic category plays the role of an interval in a poset leads to a suitable generalization by changing the domain of the incidence structure from pairs of objects to arrows in the acyclic category.

Definition 1.3.3 (signed incidence structure on a GAC)**.** A **signed incidence structure** on a graded acyclic category $(\mathcal{P}, r)$ is a map $\epsilon : \mathrm{Mor}(\mathcal{P}) \to \{-1, 0, 1\}$ that satisfies the following conditions.

(i) **Supported on covering morphisms.** $\epsilon(x \xrightarrow{f} y) \neq 0$ if and only if $x \lhd_1 y$.

(ii) **Coboundary condition.** For any $x \leqslant z$ and $h : x \to z$,

$$\sum_{g \circ f = h} \epsilon(f)\epsilon(g) = 0$$

where the sum ranges over pairs of composable arrows $x \xrightarrow{f} \bullet \xrightarrow{g} z$ such that $g \circ f = h$.

(iii) **1-cell condition.** For any $e \in \mathcal{P}$ of rank 1, there are exactly two covering morphisms f_0, f_1 with codomain e, and $\epsilon(f_0) = -\epsilon(f_1)$.

Remark 1.3.4. As in the case of signed incidence structures on posets, the coboundary condition reduces to a statement about "intervals" of length 2. $\epsilon : \mathrm{Mor}(\mathcal{P}) \to \{-1, 0, 1\}$ satisfies the coboundary condition if and only if for every pair $x \lhd_2 y$ and map $h : x \to y$, the sum

$$\sum_{(f,g) \in \mathrm{Cov}(h)} \epsilon(f)\epsilon(g) = 0\,.$$

In order for the sum in the coboundary condition to exist, $\mathrm{Cov}(h)$ must be finite for every such h.

In the event that the graded acyclic category $\mathcal{P}$ is itself a poset, a signed incidence structure on $\mathcal{P}$ induces a signed incidence relation under the correspondence:

$$[x : y] = \begin{cases} \epsilon(x \to y) & \text{if } \mathcal{P}(x, y) \text{ has a unique inhabitant} \\ 0 & \text{else.} \end{cases}$$

Not every GAC admits a signed incidence structure. To witness this failure, one may look to the many posets that fail to admit signed incidence relations, such as the following poset graded by height.

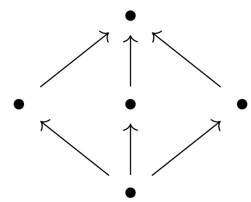

This poset does not admit a signed incidence relation, as there are three distinct composable pairs of morphisms from the bottom object to the top object.

The face posets of regular cell structures form a broad class of posets that admit signed incidence relations [33, Section 6.1.1]. Similarly, the face categories of weakly-regular cell structures admit signed incidence structures. We now give an explicit algorithm for finding a signed incidence structure on $Fc(X)$ for a weakly-regular cell structure X.

Algorithm 1.3.5. Let $X = (X_\alpha, \phi_\alpha, \mathcal{K}_\alpha)_\alpha$ be a weakly-regular cell structure with face category $Fc(X)$. Construct $\epsilon : Mor(Fc(x)) \to \{-1, 0, 1\}$ as follows.

1. For each cell X_α of dimension $d_\alpha \geqslant 1$, fix a homeomorphism $D^{d_\alpha} \cong \mathbb{R}^{d_\alpha}$, and an orientation for $\mathbb{R}^{d_\alpha}$ in the form of an ordered basis $\{b_1, \dots, b_{d_\alpha}\}$. This data induces an orientation on D^{d_α}, and consequently on $X_\alpha \cong D^{d_\alpha}$ when pushed forward by $\phi_\alpha|_{D^{d_\alpha}}$.

2. If X_α is a 1-cell, the orientation on D^1 corresponds to a direction for the edge $\bullet \longrightarrow \bullet$. Suppose the source vertex is attached to v_0 and the target vertex to v_1 by ϕ_α. Then the covering morphisms $v_0 \to \alpha$ and $v_1 \to \alpha$ are assigned -1 and $+1$ respectively.

3. If X_α is a d_α-cell for $d_\alpha \geqslant 2$, each cell $\sigma \in \partial\mathcal{K}_\alpha$ of dimension $d_\sigma = d_\alpha - 1$ inherits an orientation from the orientation $\{b_1, \dots, b_{d_\alpha}\}$. In particular, we adopt the *outward normal* convention [80]. The attaching map ϕ_α maps σ homeomorphically onto a cell X_β. If the orientation on σ induced from the orientation of D^{d_α}, when pushed forward by ϕ_α, agrees with the orientation on X_β from D^{d_β}, assign to the covering morphism σ the value $\epsilon(\sigma) := 1$. If these orientation disagree, instead assign $\epsilon(\sigma) = -1$.

4. All other morphisms τ are assigned $\epsilon(\tau) = 0$.

At the end of the process, we have a signed incidence structure.

Proposition 1.3.6. *The assignment $\epsilon : Mor(Fc(X)) \to \{-1, 0, 1\}$ constructed in Algorithm 1.3.5 is a signed incidence structure.*

Proof. The resulting assignment ϵ is clearly supported on covering morphisms, and satisfies the 1-cell condition. We need to check the coboundary condition.

For a given cell X_α, let $O(\alpha)$ denote the chosen orientation on X_α. For a boundary cell $\sigma \in \partial\mathcal{K}_\alpha$ which is mapped homeomorphically onto X_β by ϕ_α, let $O_\sigma(\beta)$ denote the orientation induced on X_β by X_α through the cell σ. Finally, given a pair of orientations O_1 and O_2 on the same cell, write:

$$E(O_1, O_2) = \begin{cases} 1 & \text{if the orientations agree} \\ -1 & \text{else.} \end{cases}$$

Consider a pair of cells X_α and X_γ with $d_\gamma = d_\alpha + 2$, and fix a map $\rho : \alpha \to \gamma$ in $\mathrm{Fc}(X)$. By Proposition 1.2.22, there are exactly two distinct composable pairs of arrows $\alpha \xrightarrow{\sigma_j} \beta_j \xrightarrow{\tau_j} \gamma$ with $\alpha \lhd_1 \beta_j \lhd_1 \gamma$ and $\tau_j \circ \sigma_j = \rho$, $j = 1, 2$. To prove that ϵ is a signed incidence structure, it suffices to show that $\epsilon(\tau_1)\epsilon(\sigma_1) = -\epsilon(\tau_2)\epsilon(\sigma_2)$. The desired equation may be re-written as

$$E\big(O_{\tau_1}(\beta_1), O(\beta_1)\big)E\big(O_{\sigma_1}(\alpha), O(\alpha)\big) = -E\big(O_{\tau_2}(\beta_2), O(\beta_2)\big)E\big(O_{\sigma_2}(\alpha), O(\alpha)\big)\,.$$

Notice that if we flip the orientation $O(\beta_1)$ to an opposite orientation, then we also flip the orientation $O_{\sigma_1}(\alpha)$. Consequently, both terms in the left hand side would change sign, leaving the left hand side unchanged. The left hand side is therefore independent of $O(\beta_1)$, and we may assume without loss of generality that $E\big(O_{\tau_1}(\beta_1), O(\beta_1)\big) = 1$. Similarly, we may show the right hand side is independent of $O(\beta_2)$, and assume without loss of generality that $E\big(O_{\tau_2}(\beta_2), O(\beta_2)\big) = 1$. It now suffices to show that

$$E\big(O_{\sigma_1}(\alpha), O(\alpha)\big) = -E\big(O_{\sigma_2}(\alpha), O(\alpha)\big)\,.$$

But this follows easily from the fact β_1 and β_2 lie on opposite sides of α inside of $\mathcal{K}_\gamma$. Since both X_{β_1} and X_{β_2} have orientations that agree with that of X_γ, Hence σ_1 and σ_2 must induce opposite orientations on X_α, proving the desired equality. □

1.4 CELLULAR SHEAVES

With graded acyclic categories and signed incidence structures in hand, we are finally able to define cellular sheaves.

Definition 1.4.1. Let $\mathcal{P}$ be a levelwise-finite graded acyclic category that admits a signed incidence structure. A **cellular sheaf** on $\mathcal{P}$ is a functor $\mathcal{F} : \mathcal{P} \to \mathcal{D}$ for some data category $\mathcal{D}$. We call the $\mathcal{D}$-object $\mathcal{F}(x) \in \mathcal{D}$ the **stalk** over x for each object $x \in \mathcal{P}$. Meanwhile the $\mathcal{D}$-morphism $\mathcal{F}_f := \mathcal{F}(x \xrightarrow{f} y) : \mathcal{F}(x) \to \mathcal{F}(y)$ is called the **restriction map** over $f : x \to y$.

When $\mathcal{P}$ is a poset, the at-most-uniquely inhabited homsets enforce a friendly composition property for the restriction maps. Namely, if $x \leqslant y \leqslant z$ in the poset $\mathcal{P}$, then $\mathcal{F}_{y \leqslant z} \circ \mathcal{F}_{x \leqslant y} = \mathcal{F}_{x \leqslant z}$. This follows from the uniqueness of morphisms in homsets.

Equivalently, cellular sheaves can be defined in terms of cells and covering-morphisms exclusively. That is, a cellular sheaf on $\mathcal{P}$ is a choice of an object $\mathcal{F}(x) \in \mathcal{D}$ for each $x \in \mathcal{P}$, and a $\mathcal{D}$-morphism $\mathcal{F}_f : \mathcal{F}(x) \to \mathcal{F}(y)$ whenever $f : x \to y$ is a covering morphism in $\mathcal{P}$. From this information, all other restriction maps can be determined by composing chains of covering morphisms.

Definition 1.4.2. Let $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathcal{D}$ be cellular sheaves defined on the same acyclic category $\mathcal{P}$. A **sheaf morphism** $\phi : \mathcal{F} \to \mathcal{G}$ is a natural transformation.

With sheaf morphisms, we may form a category of cellular sheaves $\mathbf{Shv}(\mathcal{P}; \mathcal{D}) := [\mathcal{P}, \mathcal{A}]$.

1.4.1 *On the definition of cellular sheaves*

There are a variety of different definitions of "cellular sheaf" in the literature. The most common definitions either take the domain of a cellular sheaf $\mathcal{F}$ to always be the face poset of a regular cell structure X [30, 31, 45, 52, 54, 55, 102], or allow the domain to be a more-or-less arbitrary poset [20, 33, 40, 66]. The first definition has a generally more topological flavor. By restricting to the face posets of regular cell structures, this definition keeps cellular sheaves closer to their origins in algebraic topology. Sheaves have been used to study topological spaces with great success since Leray's work as a prisoner of war in Nazi Germany [89]. Straying too far afield, especially when unnecessary, might disconnect cellular sheaves from their natural place in the history of mathematical thought. Moreover, the poset inherits a natural grading by dimension and a signed incidence relation via orientations, as discussed in Section 1.3. This additional structure allows for a rich cohomology theory, as well as the introduction of dynamics and spectral theory. Finally, face posets of regular cell complexes are better behaved than general graded partially ordered sets that admit signed incidence structures, and cellular sheaves defined on them will be better behaved as well.

The second definition, on the other hand, is more combinatorial in nature. While still topological, a cellular sheaf on an arbitrary poset allows one to extend certain aspects of the theory of cellular sheaves, most notably those that do not require an appeal to the graded structure. This comes at the expense of the clear connection to topological spaces and cohomology. Adjectives can be added to get access to those tools as needed and desired.

Our definition of a cellular sheaf in this thesis deviates substantially from both of the preceding definitions. On the topological v. combinatorial divide, we aim to strike a balance. When our graded acyclic category is a poset, we are more-or-less exactly restricting to the class of posets in which we can discuss cohomology and Laplacian dynamics. However, by divorcing cellular sheaves from those posets that specifically come from the faces of regular cell complexes, we are treating cellular sheaves as a fundamentally combinatorial object.

This choice is further justified by our focus on network sheaves—cellular sheaves defined over graphs. While graphs can be viewed as cell structures, it is more natural to think of of them as graded posets directly; they are combinatorial—not topological—structures. For applications in applied mathematics, the cell structure perspective requires additional conceptual machinery that may obscure the underlying combinatorial nature.

More strikingly, our definition of a cellular sheaf generalizes from posets to acyclic categories. This generalization serves a concrete purpose. When looking at network sheaves, by allowing for multiple morphisms between objects, we are able to accommodate graphs with self-loops. Such graphs sometimes admit meaningful interpretations in applications of cellular sheaves. While graphs with self loops can be studied in the existing framework of cellular sheaf theory easily through subdivision or ad-hoc methods, defining cellular sheaves over acyclic categories provides a natural systematic framework for doing so. The fact that many results about cellular sheaves on posets can be lifted to a more general acyclic categorical setting is an additional benefit.

1.4.2 *Cellular sheaves are sheaves*

Definition 1.4.1 may initially appear disconnected from classical sheaf theory. There are at least three apparent issues.

1. No topology is explicitly present in the definition—classical nor Grothendieck.
2. A sheaf on a category $\mathcal{C}$ is a contravariant functor $\mathcal{F} : \mathcal{C}^{op} \to \mathcal{D}$ that satisfies the sheaf condition. However, a cellular sheaf is a functor with domain $\mathcal{P}$, not $\mathcal{P}^{op}$ as one would expect.
3. The sheaf condition is notably absent from this definition.

In [33, Section 4.2], Curry cleanly resolved all three of these issues for cellular sheaves on posets through the use of the **(upper) Alexandrov topology** [2] [4]. Given a partially ordered set (or more generally a preordered set) $\mathcal{P}$, the order structure induces a topology on $\mathcal{P}$ whose open sets are exactly upward-closed sets with respect to the order. That is, $U \subseteq \mathcal{P}$ is open if and only if whenever $x \leqslant y$ and $x \in U$, then $y \in U$ as well. This topology has a basis given by the collection of upsets $\uparrow x := \{y \in \mathcal{P} : y \geqslant x\}$ for all $x \in \mathcal{P}$. This topology cleanly resolves all three apparent issues:

2 The "upper" prefix is to distinguish from the analogously defined *lower* Alexandrov topology, whose open sets are downard-closed sets. [82]

1. The Alexandrov topology clarifies the topological structure.

2. The open set structure on the Alexandrov topology on a poset $\mathcal{P}$ is order-reversing with respect to the partial order relation. That is, $x \leqslant y \implies (\uparrow x) \supseteq (\uparrow y)$. Given that a sheaf on a topological space X is a functor $\mathcal{F} : \mathbf{Open}(X)^{op} \to \mathcal{D}$ for some data category, this order-reversing property explains the absence of the opposite category construction.

3. Finally, assuming our data category $\mathcal{D}$ has enough limits and colimits, the functor category $[\mathcal{P}, \mathcal{D}]$, or equivalently the category of $\mathcal{D}$-valued cellular sheaves on $\mathcal{P}$, is categorically equivalent to the category $\mathbf{Shv}(\mathcal{P}, \mathcal{D})$ of $\mathcal{D}$ valued sheaves on $\mathcal{P}$ with respect to the Alexandrov topology [33, Theorem 4.2.10]. Since these sheaves must satisfy the sheaf condition, so do the cellular sheaves by pushing through the equivalence.

Remark 1.4.3. There are other natural topologies one can put on a posetal category. Such topologies and the sheaves on them have been studied by Lindenhovius [82] and Hemelaer [58].

Unfortunately, the Alexandrov topology fails to allow a cellular sheaf on an acyclic category $\mathcal{P}$ to be viewed as an honest sheaf on a topological space. By definition, a presheaf on a topological space X is a contravariant functor on the category of open sets $\mathbf{Open}(X)$, which is necessarily a thin category. Hence there is no way to accommodate multiple parallel arrows between objects in $\mathrm{Fc}(X)$.

This situation can be partially rectified through a Grothendieck topology. When a category $\mathcal{C}$ is endowed the **indiscrete** Grothendieck topology (also known as the **chaotic** topology), every contravariant functor $\mathcal{F} : \mathcal{C}^{op} \to \mathcal{D}$ is a sheaf on $\mathcal{C}$. This gives us a trivial way to view a cellular sheaf as a sheaf.

Proposition 1.4.4. *Let $\mathcal{F} : \mathcal{P} \to \mathcal{D}$ be a cellular sheaf on an acyclic category $\mathcal{P}$. $\mathcal{F}$ is a sheaf on $\mathcal{P}^{op}$ topologized by the indiscrete topology.*

While technically a sheaf, this is not a satisfying resolution. The cellular sheaf on $\mathcal{P}$ is a sheaf on $\mathcal{P}^{op}$, not $\mathcal{P}$. However, it is worth remarking that this really is no different than the Alexandrov topology on a poset. Given a poset $\mathcal{P}$, the indiscrete topology on $\mathcal{P}^{op}$ assigns a unique covering sieve to each object $x \in \mathcal{P}^{op}$ given by the overcategory $\mathcal{P}^{op}/x$. This exactly corresponds to the set of points $\{y \in \mathcal{P} : x \leqslant y\}$ with respect to the ordering on $\mathcal{P}$ — a basic open set in the upper Alexandrov topology. Hence while it is less conceptually and linguistically pleasing for a cellular sheaf on an acyclic category $\mathcal{P}$ to be a presheaf on $\mathcal{P}^{op}$, it is the natural generalization.

Remark 1.4.5. When $\mathcal{P} := \mathrm{Fc}(X)$ is the face category of a weakly-regular cell complex X, we can view a cellular sheaf $\mathcal{F} : \mathrm{Fc}(X) \to \mathcal{P}$ in a more satisfying way. Through the identifications $\mathrm{Fc}(X) \cong \mathbf{DExit}_X$ and $\mathbf{DEnt}_X \cong \mathbf{DExit}_X^{op}$, as discussed in Section 1.2.2.3, a cellular sheaf on $\mathrm{Fc}(X)$ is a sheaf on the discrete entrance path category $\mathbf{DEnt}_X$ with the indiscrete topology. This interpretation further links cellular sheaves to constructible sheaves, a connection first explored by Shepard [112].

1.4.3 *Sections of cellular sheaves*

The natural interpretation of a cellular sheaf $\mathcal{F} : \mathcal{P} \to \mathcal{D}$ is as a data structure. This is especially clear when the data category $\mathcal{D}$ has objects that are sets with certain additional structure, and morphisms that are structure preserving set maps (or more generally, when $\mathcal{D}$ is a concrete category). Objects in the data category $\mathcal{D}$ are interpreted as different spaces in which data can live. For each object $x \in \mathcal{P}$, the stalk $F(x) \in \mathcal{D}$ over x represents a choice of a space in which data can live. The restriction maps, on the other hand, provide local consistency conditions that may-or-may-not be satisfied by choices of data living over each point in $\mathcal{P}$. A consistent selection is a section.

Definition 1.4.6. Let $\mathcal{F} : \mathcal{P} \to \mathcal{D}$ be a cellular sheaf. Let $\mathcal{Z} \subseteq \mathcal{P}$ be a sub-category. The **space of sections** over $\mathcal{Z}$, denoted $\Gamma(\mathcal{F}; \mathcal{Z})$, is the limit

$$\Gamma(\mathcal{F}; \mathcal{Z}) := \lim \mathcal{F}\big|_{\mathcal{Z}}$$

in the category $\mathcal{D}$, when it exists. When $\mathcal{Z} = \mathcal{P}$, we call the sections $\Gamma(\mathcal{F}) := \Gamma(\mathcal{F}; \mathcal{P})$ the **global sections** of $\mathcal{F}$.

When the objects of $\mathcal{D}$ can be thought of as structured sets, such as vector spaces or R-modules, spaces of sections take on a definite meaning. Given a cellular sheaf $\mathcal{F} : \mathcal{P} \to \mathcal{D}$, and a subcategory $\mathcal{Z} \subseteq \mathcal{P}$, each point $\mathbf{x} \in \Gamma(\mathcal{F}; \mathcal{Z})$ can be thought of as a *locally consistent* choice of data living over each stalk $\sigma \in \mathrm{Ob}(\mathcal{Z})$. In particular, for each map $f : \sigma \to \tau$ in $\mathcal{Z}$, the corresponding restriction map $\mathcal{F}_f : \mathcal{F}(\sigma) \to \mathcal{F}(\tau)$ must map x_σ to x_τ.

1.4.4 *Sheaf operations*

When the data category $\mathcal{D}$ admits certain categorical operations like biproducts, tensor products, and pullbacks (among others), we obtain corresponding operations on cellular sheaves; these operations allow us to build new sheaves out of old ones. In particular, all six functors of Grothendieck's six functor formalism can be applied to cellular sheaves.

These operations generally lift pointwise from the data category to the sheaf category. For instance, given cellular sheaves $\mathcal{F}_1, \mathcal{F}_2 : \mathcal{P} \to \mathcal{D}$, their direct sum $\mathcal{F}_1 \oplus \mathcal{F}_2$ assigns to each cell σ the direct sum $\mathcal{F}_1(\sigma) \oplus \mathcal{F}_2$, with restriction maps acting componentwise. Similarly, pullback operations allow us to transfer sheaves across cellular maps $\phi : \mathcal{P} \to \mathcal{Q}$ while pushforward operations aggregate local data along such maps. The tensor product enables the construction of sheaves modeling coupled systems, where interactions between different data types must be captured within the sheaf structure. Access to these sheaf-building constructions is part of what makes the theory of cellular sheaves such a flexible framework for describing networked systems in applied mathematics.

The details of these operations for the adjoint pairs (pullback $\dashv$ pushforward) and (tensor $\dashv$ hom), as well as a few other operations like direct sums, are described in more detail for cellular sheaves of Hilbert spaces in Section 4.5. For more details on the general construction, as well as the missing "extraordinary" adjoint pair $((-)_! \dashv (-)^!)$, see [33].

1.4.5 *Cohomology*

Let $\mathcal{F} : \mathcal{P} \to \mathcal{A}$ be a cellular sheaf valued in an abelian category $\mathcal{A}$. In Definition 1.4.1 of a cellular sheaf, the requirement that the acyclic category $\mathcal{P}$ admit a grading $r : \mathcal{P} \to \mathbb{N}$ and a signed incidence structure $\epsilon : \mathrm{Mor}(\mathcal{P}) \to \{-1, 0, 1\}$ is exactly the structure necessary to define a cochain complex associated to $\mathcal{F}$.

Definition 1.4.7 (Associated cochain complex)**.** The **cochain complex associated to a cellular sheaf** $\mathcal{F} : \mathcal{P} \to \mathcal{A}$ is the cochain complex

$$\left(C^\bullet(\mathcal{P}, \mathcal{F}), \delta^\bullet\right) := C^0(\mathcal{P}; \mathcal{F}) \xrightarrow{\delta^0} C^1(\mathcal{P}; \mathcal{F}) \xrightarrow{\delta^1} C^2(\mathcal{P}; \mathcal{F}) \xrightarrow{\delta^3} \cdots$$

with k-**cochains** $C^k(\mathcal{P}; \mathcal{F})$ and k-**coboundary maps** $\delta^k : C^k(\mathcal{P}; \mathcal{F}) \to C^{k+1}(\mathcal{P}; \mathcal{F})$ defined by:

$$\begin{aligned} C^k(\mathcal{P}; \mathcal{F}) &:= \bigoplus_{r(\sigma)=k} \mathcal{F}(\sigma), \\ (\delta^k \mathbf{x})_\tau &:= \sum_{\substack{\sigma \lhd_1 \tau \\ f : \sigma \to \tau}} \epsilon(f) \mathcal{F}_f(\mathbf{x}_\sigma). \end{aligned}$$

The sum in the definition of δ^k is understood as being taken over covering morphisms into τ, and $\mathcal{F}_f := \mathcal{F}(f)$ is a shorthand for the image of the morphism f under the functor $\mathcal{F}$.

Proposition 1.4.8. $\left(C^\bullet(\mathcal{P}, \mathcal{F}), \delta^\bullet\right)$ *is a cochain complex.*

Proof. This follows straightforwardly from the coboundary condition in the definition of a signed incidence structure. Suppose that $\mathbf{x} \in C^k(\mathcal{P};\mathcal{F})$ is supported only on $F(\sigma)$ for a single point σ of grade k. One may compute:

$$
\begin{aligned}
(\delta^{k+1}\delta^k\mathbf{x})_\tau &= \sum_{\substack{\chi \lhd_1 \tau \\ g:\chi\to\tau}} \epsilon(g)\mathcal{F}_g((\delta^k\mathbf{x})_\chi) \\
&= \sum_{\substack{\chi \lhd_1 \tau \\ g:\chi\to\tau}} \mathcal{F}_g \left(\epsilon(g) \sum_{\substack{\sigma \lhd_1 \chi \\ f:\sigma\to\chi}} \epsilon(f)\mathcal{F}_f(\mathbf{x}_\sigma) \right) \\
&= \sum_{\substack{\chi \lhd_1 \tau \\ g:\chi\to\tau}} \sum_{\substack{\sigma \lhd_1 \chi \\ f:\sigma\to\chi}} \epsilon(g)\epsilon(f)\mathcal{F}_{g\circ f}(\mathbf{x}_\sigma) \\
&= \sum_{\substack{\sigma \lhd_2 \tau \\ h:\sigma\to\tau}} \sum_{(f,g)\in \mathrm{Cov}(h)} \epsilon(g)\epsilon(f)\mathcal{F}_h(\mathbf{x}_\sigma) \\
&= 0.
\end{aligned}
$$

Thus $\delta^{k+1} \circ \delta^k = 0$, and $\left(C^\bullet(\mathcal{P},\mathcal{F}), \delta^\bullet\right)$ is a cochain complex. ◻

Remark 1.4.9. This argument does not require that the cellular sheaf be valued in an abelian category. If $\mathcal{A}$ is merely additive, every cellular sheaf $\mathcal{F} : \mathcal{P} \to \mathcal{A}$ on a finite graded acyclic category will have an associated cochain complex.

Definition 1.4.10 (Cellular sheaf cohomology)**.** Let $\mathcal{F} : \mathcal{P} \to \mathcal{A}$ be a cellular sheaf valued in an abelian category $\mathcal{A}$, with associated cochain complex $\left(C^\bullet(\mathcal{P},\mathcal{F}), \delta^\bullet\right)$. The k^{th} **sheaf cohomology** of $\mathcal{P}$ with coefficients in $\mathcal{F}$ is the family of quotients

$$H^\bullet(\mathcal{P};\mathcal{F}) := \ker(\delta^{k+1})/\mathrm{im}(\delta^k)\,,$$

with the convention that $\mathrm{im}(\delta^{-1}) := 0$.

Remark 1.4.11. It is something of a misnomer to say "the" associated cochain complex and "the" sheaf cohomology of a cellular sheaf. The definition of the coboundary map $\delta^\bullet$ ultimately depends on a non-unique choice of a signed incidence structure ϵ on the GAC $\mathcal{P}$; different choices of ϵ lead to different definitions of δ. However, by a straightforward sign-flipping argument, it can be seen that the image and kernel of each coboundary map δ^k is invariant under different choices of signed incidence structure ϵ. Hence for most purposes (including cohomology) we may leave the specific choice of ϵ unspecified without creating ambiguity.

Remark 1.4.12. The 0^{th} cohomology, $H^0(\mathcal{P};\mathcal{F})$ is straightforward to interpret; $H^0(\mathcal{P};\mathcal{F})$ is isomorphic to the to the space of global sections of $\mathcal{F}$, when identified with a subspace of $C^0(\mathcal{P};\mathcal{F})$.

Higher cohomology can be interpreted as spaces of obstructions. We will describe this cohomology and its interpretation for cellular sheaves valued in the category of Hilbert spaces in more detail in Section 4.6. For more details on cellular sheaf cohomology in general, and its equivalence to the usual definition of sheaf cohomology in terms of derived categories, see [33].

1.4.6 *Cellular cosheaves*

As with everything [121], cellular sheaves admit a dual notion of *cellular cosheaves* by turning all the restriction maps around.

Definition 1.4.13. Let $\mathcal{P}$ be a graded acyclic category that admits a signed incidence structure. A **cellular cosheaf** on $\mathcal{P}$ is a contravariant functor $\mathcal{F} : \mathcal{P}^{op} \to \mathcal{D}$ for some data category $\mathcal{D}$. We call the $\mathcal{D}$-object $\mathcal{F}(\sigma) \in \mathcal{D}$ the **costalk** over σ for each $\sigma \in \mathcal{P}$. The $\mathcal{D}$-morphism $\mathcal{F}_f := \mathcal{F}(x \xrightarrow{f} y) : \mathcal{F}(y) \to \mathcal{F}(x)$ is called the **extension map** over $f : x \to y$.

When $\mathcal{G}$ is a multigraph (with self-loops allowed), viewed as a weakly-regular cell structure, a $\mathbf{Vect}_{\mathbb{F}}$-valued cellular cosheaf on $\mathcal{G}$ is exactly a **sheaf on** $\mathcal{G}$ in the sense of Friedman [42]. These sheaves on graphs were used by Friedman to prove the Hanna Neumann conjecture on finitely generated subgroups of free groups.

All the preceding constructions on cellular sheaves may themselves be dualized to cellular cosheaves. Sections become colimits of subcategories, the sheaf operators adapt easily, the associated cochain complex becomes an associated chain complex, and sheaf cohomology becomes cosheaf homology. While cellular cosheaves valued in abelian categories are somewhat harder to grasp than their cellular sheaf counterparts, cosheaves have proved useful for a variety of domains beyond the Hanna Neumann conjecture such as graphic statics [30, 31]. While the focus of this thesis is exclusively cellular sheaves of Hilbert spaces, we would be remiss not to mention this beautiful dual theory.

2

A BRIEF REVIEW OF BASIC FUNCTIONAL ANALYSIS

This chapter provides a concise review of the functional analytic foundations necessary for our subsequent development of cellular sheaves valued in Hilbert spaces. While the classical theory of weighted cellular sheaves operates within finite-dimensional vector spaces, the extension to infinite-dimensional Hilbert spaces requires careful attention to the underlying operator-theoretic machinery. We present these prerequisites in a self-contained manner, emphasizing those aspects most relevant to our later constructions. Proofs are largely omitted, but can be found in standard texts such as [38, 73, 100].

2.1 BANACH SPACES

Perhaps the most fundamental structure in classical functional analysis is the Banach space. A Banach space is a complete normed vector space: a vector space X equipped with a norm $\|\cdot\|$ such that every Cauchy sequence in X converges. This structure serves as a powerful generalization of the finite dimensional $\mathbb{R}^n$ and $\mathbb{C}^n$, and serves as the backbone of functional analysis. Unsurprisingly, prototypical examples of Banach spaces are $\mathbb{R}^n$ and $\mathbb{C}^n$ with their standard Euclidean norms, but the real utility emerges when working with infinite-dimensional spaces.

Definition 2.1.1. Let $\Bbbk$ be either $\mathbb{R}$ or $\mathbb{C}$. A $\Bbbk$-**Banach space** is a normed $\Bbbk$-vector space $(X, \|\cdot\|)$ that is complete with respect to the topology on X induced by $\|\cdot\|$.

Example 2.1.2. The following are examples of real Banach spaces.

1. The space $C[a, b]$ of continuous functions $f : [a, b] \to \mathbb{R}$ with the sup-norm $\|f\|_\infty := \sup\{|f(x)| : x \in [a, b]\}$.
2. For each $p \in [0, \infty]$, the space $\ell^p(\mathbb{N})$ of sequences (x_n) with finite p-norm $\|(x_n)\|_p := (\sum_{n\in\mathbb{N}} |x_n|^p)^{1/p}$.
3. The Lebesgue spaces $L^p[a, b]$ of Lebesgue measurable functions $f : [a, b] \to \mathbb{R}$ (modulo almost-everywhere agreement) with finite p-norm $\|f\|_p = \left(\int_a^b |f|^p \, dx\right)^{1/p}$.
4. More generally, the spaces $L^p(\Omega, \mathcal{F}, \mu ; \mathbb{R}^n)$ of $\mathcal{F}$-measurable functions $f : \Omega \to \mathbb{R}^n$ (modulo almost-everywhere) with finite p-norm, where $(\Omega, \mathcal{F}, \mu)$ is a measure space.

These real Banach spaces have complex Banach analogues. Most results we discuss will hold for both real and complex Banach spaces.

Remark 2.1.3. A linear subspace $Y \subseteq X$ of a Banach space X need not be a Banach space. While such a subspace inherits a norm by restriction, Y may not be complete with with respect to the norm, such as when Y is not closed in X. This is the only way a subspace can fail to be a Banach space.

Proposition 2.1.4. *Let X be a Banach space. A linear subspace $Y \subseteq X$ is a Banach space if and only if Y is topologically closed in X.*

2.2 HILBERT SPACES

Hilbert spaces form a special class of Banach spaces with well behaved geometry. In short, Hilbert spaces are Banach spaces in which angles between vectors can be defined. Consequently, Hilbert spaces admit stronger theorems and more geometric arguments that aren't available in a generic Banach space.

Definition 2.2.1. A $\mathbb{k}$-Banach space X is a $\mathbb{k}$-**Hilbert space** if the Banach space norm satisfies the **parallelogram identity**

$$\|x+y\|^2 + \|x-y\|^2 = 2\|x\|^2 + 2\|y\|^2.$$

Remark 2.2.2. This is not the most common definition of a Hilbert space. Usually, a Hilbert space is defined as an inner product space $(X, \langle -, - \rangle)$ which is complete with respect to the induced norm $\|x\| = \sqrt{\langle x, x \rangle}$. These definitions are equivalent; any such inner product space is easily seen to be a Banach space that satisfies the parallelogram identity. Conversely, in a real Banach space X that satisfies the parallelogram identity, the expression

$$\langle x, y \rangle := \frac{1}{4}\left(\|x+y\|^2 + \|x-y\|^2\right)$$

defines an inner product whose induced norm is exactly the Banach space norm $\| - \|$. A similar expression gives the inner product for a complex Hilbert space in terms of the norm. The additional structure of an inner product—and hence the ability to measure angles between vectors—makes Hilbert spaces particularly well-suited to problems in geometry, physics, and signal processing. However, when discussing maps between Hilbert spaces, especially in the context of category theory, we do not wish to require that maps respect the inner product structures directly. Hence we adopt a definition that treats Hilbert spaces as a special class of normed vector spaces. This choice also allows definitions on normed vector spaces and Banach spaces to immediately apply to Hilbert spaces without reference to the inner product.

Example 2.2.3. $\mathbb{R}^n$ and $\mathbb{C}^n$ are Hilbert spaces with their usual inner products. Several key examples of real infinite dimensional Hilbert spaces are specializations of their Banach counterparts.

1. The space $\ell^2(\mathbb{N})$ of square-summable sequences with inner product $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{n \in \mathbb{N}} x_n y_n$.
2. The space $L^2[a, b]$ of square-integrable real functions (modulo almost everywhere agreement with respect to Lebesgue measure) with inner product $\langle f, g \rangle = \int_a^b f(x)g(x)\, dx$.
3. More generally, the space $L^2(\Omega, \mathcal{F}, \mu)$ of square-integrable real-valued functions (modulo agreement μ-almost everywhere) on a measure space $(\Omega, \mathcal{F}, \mu)$.

These examples have complex Hilbert space analogs.

Remark 2.2.4. As a Banach space, every closed linear subspace of a Hilbert space X is itself a Hilbert space.

Every Hilbert space X admits a **Hilbert space basis** $\{e_\alpha : \alpha < \beta\}$ for some cardinal number β, such that for every vector $x \in X$, there is a unique sequence of coefficients $c_\alpha \in \mathbb{k}$, all but countably many $c_\alpha = 0$, such that x may be written as a limit of finite sums $\lim_{\alpha \to \beta} c_\alpha e_\alpha = x$. We abbreviate this limit to $x = \sum_\alpha c_\alpha e_\alpha$. Every Hilbert space basis for X has the same cardinality β; this invariant is the **dimension** of X. A $\mathbb{k}$-Hilbert space X is completely characterized its dimension. When a Hilbert space X admits a countable Hilbert space basis, X is said to be **separable**. Equivalently, X is separable if and only if X has a countable dense subset.

The geometric properties of Hilbert spaces admit a notion of orthogonality not present in a generic Banach space.

Definition 2.2.5. Let X be a Hilbert space. Two vectors $x, y \in X$ are **orthogonal**, denoted $x \perp y$, if $\langle x, y \rangle = 0$.

Notation 2.2.6. Let X be a Hilbert space, and $V \subseteq X$ a linear subspace. The **orthogonal complement** of V, denoted $V^\perp$, is the linear subspace

$$V^\perp := \{x \in X : x \perp v \text{ for all } v \in V\}.$$

Remark 2.2.7. The orthogonal complement $V^\perp$ is always a closed subspace of X—and thus a sub-Hilbert space—even when V is not closed. The *double complement* $(V^\perp)^\perp$, is the topological closure of V in X.

2.3 OPERATORS

Banach spaces are topological vector spaces. Therefore linear maps $A : X \to Y$ between Banach spaces can be categorized according to how they interact with the topologies of X and Y.

"Operator" serves as a catch-all term for partially defined linear maps, regardless of topological properties.

Definition 2.3.1. Let X and Y be vector spaces. A **(linear) operator** $A : X \to Y$ is a linear map defined on a linear subspace $\mathrm{Dom}(A) \subseteq X$. The domain of definition $\mathrm{Dom}(A)$ is the **domain** of A, and X is the **ambient space**.

Remark 2.3.2. Linear operators are not required to be globally defined. They need not even be defined on a sub-Banach space, as the linear subspace $\mathrm{Dom}(A)$ need not be closed. Therefore, when specifying a linear operator, the domain must be specified as well. This flexibility allows, for example, the analysis of differential operators on function spaces, which cannot be defined for *all functions*, but only those which are sufficiently differentiable.

Care must be taken when composing partially defined operators. Given an operator A, let $\mathcal{R}(A)$ denote its range. For a pair of composable unbounded operators $A : X \to Y$ and $B : Y \to Z$, it may be the case that $\mathcal{R}(A) \nsubseteq \mathrm{Dom}(B)$. In general, the domain of $B \circ A$ will be taken to be $\mathrm{Dom}(B \circ A) := \{x \in \mathrm{Dom}(A) \ : \ Ax \in \mathrm{Dom}(B)\} = A^{-1}(\mathrm{Dom}(B))$, unless otherwise stated.

There are many ways that a linear operator $A : X \to Y$ can interact with the topologies of X and Y. We now highlight a few of these topological properties a linear operator may posses, and the corresponding classes of operators.

2.3.1 *Bounded and unbounded operators*

Definition 2.3.3. Let $A : X \to Y$ be a linear operator between normed vector spaces. The **operator norm** of A is given by

$$\|A\|_{\mathrm{op}} := \sup_{x \in \mathrm{Dom}(A)} \frac{\|Ax\|_Y}{\|x\|_X}.$$

If $\|A\|_{\mathrm{op}}$ is finite, we say that A is a **bounded operator**. If $\|A\|_{\mathrm{op}} = \infty$, we say that A is an **unbounded operator**.

Remark 2.3.4. We adopt the convention that all bounded operators are assumed to be globally defined unless stated otherwise. That is, if $A : X \to Y$ is a bounded operator, the domain $\mathrm{Dom}(A)$ is assumed to be all of X, unless it is explicitly stated that A has a different domain. While unbounded operators can, in principle, be defined globally, some form of choice is required to exhibit such a map. Such pathological maps are rarely of interest, so we adopt the opposite convention for unbounded operators; an unbounded operator is assumed to be a partial operator unless otherwise stated.

Remark 2.3.5. The space $\mathcal{B}(X, Y)$, of bounded operators between Banach spaces X and Y is a vector space under pointwise addition and scaling. The operator norm $\| - \|_{\mathrm{op}}$ defines a Banach space structure on $\mathcal{B}(X, Y)$. Even when X and Y are Hilbert spaces, the space $\big(\mathcal{B}(X, Y), \|-\|_{op}\big)$ is a Banach space, unless X or Y is finite dimensional.

Example 2.3.6. Let $C^k([0,1];\mathbb{R})$ denote the real Banach space of k-times continuously differentiable real-valued functions supported on $[0,1]$, equipped with the sup-norm $\|-\|_\infty$. Let $D : C^0([0,1];\mathbb{R}) \to C^0([0,1];\mathbb{R})$ denote the derivative map $D(f) = \frac{d}{dx}f$ with domain $\mathrm{Dom}(D) = C^1([0,1];\mathbb{R}) \subseteq C^0([0,1];\mathbb{R})$. The operator D is unbounded; a continuous map $f : [0,1] \to \mathbb{R}$ with sup-nom $\|f\|_\infty = 1$ may have an arbitrarily large derivative at a point. It is not globally defined as not all continuous functions are continuously differentiable.

If $C^1([0,1];\mathbb{R})$ is instead equipped with the Banach space norm $\|f\| = \|f\|_\infty + \|Df\|_\infty$, the map $D : C^1([0,1];\mathbb{R}) \to C^0([0,1];\mathbb{R})$ is a globally defined bounded operator.

Boundedness is closely related to continuity, as evidenced by the following theorem.

Theorem 2.3.7. *Let $A : X \to Y$ be a Banach space operator with domain $\mathrm{Dom}(A)$. The following are equivalent.*

(i) A is bounded on its domain.

(ii) A is continuous on its domain.

(iii) A maps the unit ball $B_1(0) \cap \mathrm{Dom}(A)$ in A to a bounded set in Y.

Thus an operator $A : X \to Y$ with domain $\mathrm{Dom}(A)$ is bounded if and only if it is continuous on its domain. Similarly, unbounded operators are exactly those operators that are discontinuous on their domains. Hence a bounded operator is exactly a linear map that respects the topology of its domain as a normed vector space.

2.3.2 *Densely defined operators*

Definition 2.3.8. An operator $A : X \to Y$ is **densely defined** if $\mathrm{Dom}(A)$ is dense in X.

Remark 2.3.9. Since a linear subspace $V \subseteq X$ is dense in X *with respect to a chosen topology*, having a dense domain of definition is a topological property of an operator—not an algebraic property.

A Hilbert space operator $A : X \to Y$ may always be extended to a densely defined one. Suppose that $\overline{\mathrm{Dom}(A)} \subsetneq X$ is a proper subset. There is a linear operator $\hat{A} : X \to Y$ that extends A with domain $\mathrm{Dom}(\hat{A}) = \mathrm{Dom}(A) + \overline{\mathrm{Dom}(A)}^{\perp}$, where $\overline{\mathrm{Dom}(A)}^{\perp}$ denotes the orthogonal complement of $\overline{\mathrm{Dom}(A)}$ in the ambient space X, and "+" denotes the internal direct sum of linear subspaces. On elements $x \in \overline{\mathrm{Dom}(A)}^{\perp}$ we take $\hat{A}(x) = 0$.

2.3.3 *Closed operators*

Definition 2.3.10. Let $A : X \to Y$ be an unbounded Banach space operator with domain $\mathrm{Dom}(T) \subseteq X$. A is **closed** if its graph $\Gamma(A) := \{(x, Ax) \ : x \in \mathrm{Dom}(A)\} \subseteq X \oplus Y$ is topologically closed. A is **closable** if the closure $\overline{\Gamma(A)} \subseteq X \oplus Y$ is the graph of an operator $\overline{A} : X \to Y$. The extension $\overline{A}$ is called the **closure** of A.

A bounded operator is closed if and only if its domain is closed. Hence a globally-defined bounded operator is always closed. For unbounded operators, being closed is perhaps the most important way in which an unbounded operator can still be well behaved. Closedness (or closability) is the necessary attribute in order to carry out a variety of constructions, including those in spectral theory and semigroup theory, to be discussed later. The essence is the following equivalent definition of closedness: $A : X \to Y$ is closed if whenever x_n is a sequence in $\mathrm{Dom}(A)$ that converges to $x \in X$, and $A(x_n)$ converges to $y \in Y$, then $x \in \mathrm{Dom}(X)$ and $Ax = y$. Similarly, A is closable if and only if whenever $x_n \to 0$ in $\mathrm{Dom}(A)$ and $Ax_n \to y \in Y$, the limit $y = 0$. Hence we see that to be closed is to satisfy a weak form of continuity, as further evidenced by the following straightforward proposition.

Proposition 2.3.11. *Let* $A : X \to Y$ *be an unbounded operator. If* T *is closed, then* $\ker(T) \subseteq X$ *is closed.*

Closed operators enjoy another useful topological property. While an unbounded operator is discontinuous, a closed unbounded operator is continuous with respect to the graph norm.

Definition 2.3.12. Let $A : X \to Y$ be a closed Banach space operator with domain $\mathrm{Dom}(A)$. The **graph norm** on $\mathrm{Dom}(A)$ is the norm $\|x\|_{\Gamma(A)} := \left(\|x\|_X^2 + \|Ax\|_Y^2\right)^{1/2}$.

One may check that $\|-\|_{\Gamma(A)}$ is a well-defined norm on $\mathrm{Dom}(A)$. Moreover $\mathrm{Dom}(A)$ is complete with respect to $\|-\|_{\Gamma(A)}$, making $\left(\mathrm{Dom}(A), \|-\|_{\Gamma(A)}\right)$ a Banach space.

Proposition 2.3.13. *Let* $A : X \to Y$ *be a closed Banach space operator.* A *defines a bounded operator* $A : \mathrm{Dom}(A) \to Y$ *with respect to the graph norm on* A.

2.3.4 *Closed range operators*

Definition 2.3.14. Let $A : X \to Y$ be a Banach space operator. A has **closed range** if its range $\mathcal{R}(A) \subseteq Y$ is topologically closed.

To have closed range is a surprisingly restrictive property; even bounded operators usually fail to have closed range. A handful of equivalent characterizations to having closed range are offered by the **closed range theorem** [15, 78].

Theorem 2.3.15. *Let $A : X \to Y$ be a closed, densely defined Banach space operator. The following are equivalent.*

(i) A *has closed range.*

(ii) A^* *has closed range.*

(iii) 0 *is not an accumulation point in the spectrum* $\sigma(A^*A)$*.*

(iv) There is a constant C *such that* $\|x\| \leqslant C\|Ax\|$ *for all* $x \in \mathrm{Dom}(A) \cap \ker(A)^\perp$*.*

This theorem gives an indication of why many operators fail to have closed range. In order to have closed range, an operator cannot shrink inputs too quickly.

Remark 2.3.16. For an operator $A : X \to Y$, "closedness" and "closed range" are independent properties. There are Banach space operators with closed range that fail to be closed, and closed operators which fail to have closed range.

One useful property of a closed range operators between Hilbert spaces is that it admits a bounded Moore-Penrose pseudoinverse.

Proposition 2.3.17. *Let* $A : X \to Y$ *be a closed, densely defined Hilbert space operator. There is a unique closed, densely defined operator* $A^\dagger : Y \to X$ *with domain* $\mathrm{Dom}(A^\dagger) = \mathcal{R}(A) + \mathcal{R}(A)^\perp$*, called the* ***Moore-Penrose pseudoinverse*** *of* A*, which satisfies the following properties.*

(i) $\ker(A^\dagger) = \mathcal{R}(A)^\perp$.

(ii) $\mathcal{R}(A^\dagger) = \mathrm{Dom}(A) \cap \ker(A)^\perp$.

(iii) $A^\dagger A$ *is the orthogonal projection onto the closure* $\overline{\mathrm{Dom}(A) \cap \ker(A)^\perp}$*.*

(iv) $AA^\dagger$ *is the orthogonal projection onto the closure* $\overline{\mathcal{R}(A)}$*.*

Corollary 2.3.18. *When* $A : X \to Y$ *is a closed, densely defined Hilbert space operator with closed range, the pseudoinverse* $A^\dagger : Y \to X$ *is globally defined and bounded.*

2.3.5 *Structure preserving operators*

There are a variety of different ways in which a Banach space or Hilbert space operator can be "structure preserving." Four primary classes of structure preserving maps are isomorphisms, isometries, co-isometries, and unitary maps.

Definition 2.3.19. A globally-defined Banach space operator $A : X \to Y$ is an **isomorphism** if it is a bounded bijection with bounded inverse.

Isomorphisms are exactly linear homeomorphisms with respect to the topologies induced by the Banach space norms. By the open mapping theorem (Theorem 2.4.2), every bounded linear bijection has a bounded inverse.

While isomorphisms respect the topologies of Banach spaces, they do not respect the norms *per se*. In contrast, an isometry is an operator that preserves the norm structure exactly.

Definition 2.3.20. A globally-defined Banach space operator $A : X \to Y$ is an **isometry** if $\|Ax\|_Y = \|x\|_X$ for all $x \in \mathrm{Dom}(A)$.

All isometries are bounded by definition. A canonical example of an isometry is the right-shift operator on an ℓ^p-sequence space, given by $R(x_1, x_2, x_3, \cdots) = (0, x_1, x_2, x_3, \cdots)$. The Fourier transform on $L^2(\mathbb{R})$ provides another fundamental example, becoming an isometry after appropriate normalization.

While isometries respect the Banach space norms, there is no guarantee that an isometry is a bijection. When there is an isometric isomorphism between two Banach spaces, we say they are **isometrically isomorphic**. This is one of the strongest ways in which two Banach spaces can be equivalent to one another.

On Hilbert spaces, there are more designated classes of structure preserving maps. These maps are best characterized through the use of the linear adjoint, discussed in detail in Section 2.5. All the following classes of structure preserving maps are bounded, so one only needs the usual bounded linear adjoint.

Definition 2.3.21. Let $A : X \to Y$ be a Hilbert space operator. Let I_X and I_Y denote the identity maps on X and Y respectively.

- A is an **isometry** if $A^*A = I_X$.
- A is a **co-isometry** if $AA^* = I_Y$.
- A is **unitary** if A is both an isometry and a co-isometry.

Notation 2.3.22. Some sources reserve the term "unitary" for complex Hilbert spaces, and use "orthogonal" for the analogous property on real Hilbert spaces. We use "unitary" for real and complex Hilbert spaces for ease.

2.3.6 *Hilbert-Schmidt operators*

Definition 2.3.23. Let X and Y be Hilbert spaces, and $\{e_i\}_{i\in I}$ a Hilbert space basis of X. The **Hilbert-Schmidt norm** of a bounded operator $A : X \to Y$ is given by

$$\|A\|_{HS} := \left(\sum_{i\in I} \|Ae_i\|_X^2 \right)^{1/2} \in [0, \infty] \,.$$

A bounded operator $A : X \to Y$ is a **Hilbert-Schmidt operator** if $\|A\|_{HS} < \infty$.

The class of Hilbert-Schmidt operators between Hilbert spaces X and Y are closed under addition and scaling. Let $HS(X, Y) \subseteq \mathcal{B}(X, Y)$ denote the linear subspace of Hilbert-Schmidt operators. $HS(X, Y)$ forms a Hilbert space with inner product

$$\langle A, B\rangle_{HS} := \sum_{i\in I} \langle Ae_i, Be_i\rangle_Y \,.$$

This Hilbert-Schmidt inner product serves as an infinite dimensional analog of the familiar Frobenius inner product of matrices.

Hilbert-Schmidt operators are one of the most well-behaved classes of Hilbert space operators. Beyond being bounded, Hilbert-Schmidt operators are **compact**, meaning they map bounded sets to precompact sets. All finite rank operators are Hilbert-Schmidt.

Remark 2.3.24. The Hilbert-Schmidt norm on $HS(X, Y)$ is not equivalent to the operator norm on $HS(X, Y)$. However, they are related by the inequality $\| - \|_{op} \leqslant \| - \|_{HS}$.

2.4 FUNDAMENTAL THEOREMS

The theory of Banach spaces rests on four key theorems, sometimes called the *four pillars* of functional analysis [125]:

1. The Hahn-Banach theorem: bounded linear functionals can always be extended from subspaces while preserving their norm.

2. The open mapping theorem: every surjective continuous linear operator between Banach spaces is an open map.

3. The Banach-Steinhaus theorem: every pointwise-bounded family of continuous linear operators is uniformly bounded.

4. The closed graph theorem: a linear operator between Banach spaces is continuous if and only if its graph is closed.

These four theorems, while simple to state, capture deep properties of infinite-dimensional spaces and form the foundation upon which much of functional analysis is built. Proofs of these theorems may be found in any standard text on functional analysis, such as [15, 73, 100].

Theorem 2.4.1 (Hahn-Banach). *Let X be a normed $\Bbbk$-vector space, and $V \subseteq X$ a linear subspace. Every continuous linear functional $f : V \to \Bbbk$ may be extended to a continuous linear functional $\tilde{f} : X \to \Bbbk$ such that $\|\tilde{f}\|_{op} = \|f\|_{op}$.*

Theorem 2.4.2 (Open mapping theorem). *Let $A : X \to Y$ be a bounded surjective Banach space operator. For every open subset $U \subseteq X$, $A(U) \subseteq Y$ is open.*

Theorem 2.4.3 (Banach-Steinhaus). *Let X be a $\Bbbk$-Banach space and Y a normed $\Bbbk$-vector space, and $\mathcal{A} \subseteq \mathcal{B}(X, Y)$ a collection of bounded linear operators. If $\sup_{A \in \mathcal{A}} \|Ax\| < \infty$ for every $x \in X$, then $\sup_{A \in \mathcal{A}} \|A\|_{op} < \infty$.*

Theorem 2.4.4 (closed graph theorem). *Let $A : X \to Y$ be a globally-defined Banach space operator. A is bounded if and only if the graph $\Gamma(A)$ is closed.*

2.5 ADJOINTS

Definition 2.5.1. Let $A : X \to Y$ be a Hilbert space operator with domain $\mathrm{Dom}(A)$. An operator $B : Y \to X$ is an **adjoint** of A if $\langle Ax, y \rangle_Y = \langle x, A^* y \rangle_X$ for all $x \in \mathrm{Dom}(T)$ and $y \in \mathrm{Dom}(B)$, and $\mathrm{Dom}(B)$ is maximal with with respect to this property.

Remark 2.5.2. When $A : X \to Y$ is not densely defined, the adjoint A^* is not unique. As a trivial example, consider a Hilbert space X and the zero-operator $0 : X \to X$ with the one-point domain $\{0\}$. Any maximally-defined operator $A : X \to X$ will satisfy the definition of an adjoint of 0. However, when A is densely defined, the Hahn-Banach theorem and the Riesz representation theorem ensure the adjoint A^* is unique.

Proposition 2.5.3. *If $A : X \to Y$ is a densely defined Hilbert space operator, A has a unique adjoint, denoted $A^* : Y \to X$.*

Remark 2.5.4. The domains of A and A^* are intimately linked; extending the domain of T may require shrinking the domain of T^* to maintain the defining property of a linear adjoint.

The following proposition follows directly from the definition of the adjoint.

Proposition 2.5.5. *Let $A : X \to Y$ be a densely defined linear operator. The adjoint $A^* : Y \to X$ is a closed linear operator.*

Remark 2.5.6. The adjoints of bounded linear operators enjoy all the usual properties of adjoints in finite-dimensional linear algebra. For example, A^* is always globally defined, and the operator

$$(-)^* : \mathcal{B}(X, Y) \to \mathcal{B}(Y, X)$$

is an involutive isometric isomorphism, where $\mathcal{B}(-,-)$ is the Banach space of globally defined bounded operators equipped with the operator norm. Adjoints of unbounded operators do not enjoy these same properties. For example, given a pair of unbounded operators $A, B : X \to Y$, it is not in general the case that $(A + B)^* = A^* + B^*$.

Adjoints are closely linked to closedness of operators. Suppose that $A : X \to Y$ is a densely defined operator with adjoint $A^* : Y \to X$. If A^* itself is densely defined, then it has a unique closed adjoint $(A^*)^* : X \to Y$. Moreover, by the definition of the adjoint, this operator $(A^*)^*$ must extend the operator A; otherwise $(A^*)^*$ would not be maximally defined. It follows that A must be a closable operator, and $(A^*)^* = \overline{A}$ is its closure. This argument proves the following proposition.

Proposition 2.5.7. *Let $A : X \to Y$ be a densely defined Hilbert space operator. The following are equivalent.*

(i) A *is closable.*

(ii) A^* *is densely defined.*

For closed, densely defined operators, the adjoint maintains its geometry from finite dimensional linear algebra.

Proposition 2.5.8. *Let $A : X \to Y$ be a closed, densely defined Hilbert space operator. The following identities hold.*

(i) $\ker(A) = \mathcal{R}(A^*)^\perp$.

(ii) $\ker(A^*) = \mathcal{R}(A)^\perp$.

(iii) $\overline{\mathcal{R}(A)} = \ker(A^*)^\perp$.

(iv) $\overline{\mathcal{R}(A^*)} = \ker(A)^\perp$.

Remark 2.5.9. Observe that when A has closed range, one exactly recovers the familiar relationship $\mathcal{R}(A) = \ker(A^*)^\perp$ of finite dimensional linear algebra.

Definition 2.5.10. Let $T : X \to X$ be a densely defined unbounded Hilbert space operator. T is **symmetric** if $\langle Tx, y \rangle = \langle x, Ty \rangle$ for all $x, y \in \text{Dom}(T)$. The operator T is **self-adjoint** if $T = T^*$. Finally, T is **essentially self-adjoint** if T has a self-adjoint extension.

For bounded operators in $\mathcal{B}(X, X)$, "symmetric" and "self-adjoint" are coextensive. However, for a symmetric unbounded operator $A : X \to X$, it could be the case that $\mathrm{Dom}(A^*) \supsetneq \mathrm{Dom}(A)$, making A a symmetric operator that fails to be self-adjoint. In this case, A^* is a closed extension of A. If A is essentially self-adjoint, then $A^* = \overline{A}$ will be self-adjoint. Since A^* is always closed, every self-adjoint operator is closed.

We conclude this section with an essential theorem, sometimes referred to as **von Neumann's theorem**, due to its connection to von Neumann's seminal work on quantum mechanics [2]. Let $A : X \to Y$ be a Hilbert space operator with domain $\mathrm{Dom}(A)$. We say that A is **positive semidefinite** if $\langle Ax, x\rangle \geqslant 0$ for all $x \in \mathrm{Dom}(A)$. We say that A is **positive** if it is both positive semidefinite and self-adjoint.

Theorem 2.5.11 (von Neumann's theorem)**.** *Let $A : X \to Y$ be a closed, densely defined Hilbert space operator. Both A^*A and AA^* are positive operators.*

Unlike the other theorems of this section, the proof of von Neumann's theorem is not straightforward. A full proof may be found in [73, Theorem 3.24].

2.6 Quotients of Banach spaces

Taking the quotient of a Banach space X by a subspace V has some subtlety. Not every linear subspace $V \subseteq X$ is itself a Banach space. while the norm $\| - \|_X$ on X restricts to a norm on V, the subspace V need not be topologically complete with respect to this norm. For example, consider the linear subspace of all polynomials $P[a, b] \subseteq C[a, b]$. The linear subspace $P[a, b]$ is not itself a Banach space, but is merely a normed vector space. The Weierstrass approximation theorem ensures that every $f \in C[a, b]$ is a limit of Cauchy sequences in $P[a, b]$, so $P[a, b]$ cannot be complete with respect to the sup-norm. Fundamentally, the problem is that $P[a, b]$ is not closed. In general, a linear subspace V of a Banach space X is a Banach space with respect to its induced norm if and only if V is topologically closed in X.

Similar difficulties occur with quotients. Given a linear subspace V of a Banach space X, the quotient vector space X/V may be formed as usual as the set of cosets $\{x + V : x \in X\}$. When V is closed, the Banach space norm $\| - \|_X$ induces the following norm on X/V:

$$\|x + V\|_{X/V} = \inf\{\|x - v\|_X : v \in V\}.$$

This is a well-defined norm if and only if V is closed in X. Geometrically, this norm measures the distance from x to the closed subspace V. When V is closed, the quotient space X/V is complete with respect to $\| - \|_{X/M}$, and hence is a Banach space. Hence a quotient of *Banach spaces* is a Banach space, but the quotient of a Banach space by a linear space need not be Banach.

Proposition 2.6.1. *Let $A : X \to Y$ be a bounded Banach space operator. If the range $\mathcal{R}(A)$ is closed, there is an isomorphism of Banach spaces $X/\ker(A) \cong \mathcal{R}(A)$.*

Proof. Simply use the usual map $\phi : X/\ker(A) \to \overline{\mathcal{R}(A)}$ defined by $\phi(x+V) = Ax$. □

When $\mathcal{R}(A)$ isn't closed, this theorem fails. While there is still an isomorphism of vector spaces, the range $\mathcal{R}(A)$ will fail to be a Banach space, and hence ϕ will fail to be a Banach space isomorphism.

The quotient structure of Hilbert spaces is similar to that of Banach spaces. Let X be a Banach space. When $\| - \|_X$ satisfies the parallelogram identity, so does the induced norm $\| - \|_{X/V}$ on the quotient Banach space X/V for a closed subspace $V \subseteq X$. Hence a subspace quotient X/V in a Hilbert space is itself a Hilbert space exactly when V is closed.

The structure of a quotient Hilbert space X/V can be understood through orthogonal complements. Given a closed subspace V of a Hilbert space X, we can decompose X as an orthogonal direct sum, $X = V \oplus V^\perp$, where $V^\perp$ is the orthogonal complement of V in X. Applying the first isomorphism theorem (Proposition 2.6.1) to the orthogonal projection $P : X \to V^\perp$ gives a canonical isomorphism $X/V \cong V^\perp$, providing the quotient space with a natural Hilbert space structure inherited from X. Moreover, this isomorphism is unitary.

Using this isomorphism, we no longer need to take an infimum to define the quotient norm; the norm of an equivalence class $x + V$ in X/V is simply the norm of its unique representative in $V^\perp$. This additional structure makes Hilbert space quotients particularly easy to work with.

2.7 SEMIGROUPS

Definition 2.7.1. Let X be a Banach space. A **strongly continuous one-parameter semigroup** on X, or C_0**-semigroup** on X, is a map $T : [0,\infty) \to \mathcal{B}(X)$ that satisfies the following conditions.

(i) $T(0) = I$.

(ii) $T(s+t) = T(s)T(t)$ for all $s, t \geqslant 0$.

(iii) T is continuous in the strong operator topology on $\mathcal{B}(X)$.

If item (iii) is replaced with

(iii') T is continuous with respect to the operator norm on $\mathcal{B}(X)$,

then T is a **uniformly continuous semigroup**. If a C_0-semigroup has the additional property

(iv) $\|T(t)\|_{op} \leqslant 1$ for all $t \geqslant 0$,

then T is a **contraction semigroup**.

To every C_0-semigroup $T : [0, \infty) \to \mathcal{B}(X)$, there is an associated operator $A_T : X \to X$ called the **infinitesimal generator** defined by the limit

$$A_T(x) := \lim_{t \to 0^+} \frac{T(t)x - x}{t},$$

whenever the limit exists. The domain $\mathrm{Dom}(A_T)$ is exactly the set of $x \in X$ for which this limit converges. One may confirm by direct computation that the infinitesimal generator A_T is closed and densely defined. Moreover, there is a one-to-one correspondence between C_0-semigroups and their infinitesimal generators.

Remark 2.7.2. Even though the generator A_T may be an unbounded operator with a partial domain $\mathrm{Dom}(A_T) \subsetneq X$, the operator $T(t) : X \to X$ is bounded and globally defined for all $t \geqslant 0$.

When we wish to highlight the infinitesimal generator A of a semigroup T, we will often write e^{tA} or $\exp(tA)$ in place of $T(t)$. This evocative notation emphasizes the following relationship between C_0-semigroups ordinary differential equations on Banach spaces.

Proposition 2.7.3. *The map $x(t) = e^{tA}x_0$ is a **mild solution** to the initial value problem:*

$$\begin{aligned} \dot{x}(t) &= Ax(t) \\ x(0) &= x_0. \end{aligned}$$

That is, $x(t)$ satisfies the integral equation $x(t) = x(0) + \int_0^t Ax(s)\,ds$. When $x_0 \in \mathrm{Dom}(A)$, $x(t)$ is a strong solution to the initial value problem.

Not every closed, densely defined operator $A : X \to X$ generates a C_0-semigroup. There are a variety of useful and powerful theorems that characterize the operators that are infinitesimal generators.

Definition 2.7.4. Let $A : X \to X$ be a complex Banach space operator. A complex $\lambda \in \mathbb{C}$ is a **regular value** of A if the following conditions hold:

(i) The map $A_\lambda := A - \lambda I$ is injective.

(ii) The inverse $A_\lambda^{-1} : \mathcal{R}(A_\lambda) \to X$ is bounded.

(iii) $\mathcal{R}(A_\lambda)$ is dense in X.

The **resolvent set** of A, denoted $\rho(A)$, is the set of all regular values of A.

Remark 2.7.5. Under suitable conditions on $A : X \to Y$, the conditions defining a regular value of A simplify considerably. If A is a closed operator, then condition (iii) may be replaced with the requirement that A_λ is surjective. If A is bounded, all three conditions may be replaced with the single requirement that A_λ is a bounded linear isomorphism.

Remark 2.7.6. For a real Banach space X, the resolvent of an unbounded operator $A : X \to X$ is the resolvent of the complexification $A^{\mathbb{C}} : X^{\mathbb{C}} \to X^{\mathbb{C}}$. Through this complexification, the spectral theory of complex Banach spaces and Hilbert spaces may be adapted to real spaces.

Theorem 2.7.7 (Hille-Yosida)**.** *Let $A : X \to X$ be an unbounded linear operator, $r \in \mathbb{R}$, and $M > 0$. The map A is the generator of a C_0-semigroup T satisfying $\|T(t)\| \leqslant Me^{rt}$ if and only if the following conditions hold.*

(i) A *is closed and densely defined.*

(ii) Every real $\lambda > r$ is in the resolvent $\rho(A)$, and for all $n \in \mathbb{N}_{>0}$,

$$\|(\lambda I - A)^{-n}\| \leqslant \frac{M}{(\lambda - r)^n}.$$

Remark 2.7.8. Setting $M = 1$ and $r = 0$ gives necessary and sufficient criteria for A to generate a contraction semigroup. However, the Lumer-Phillips theorem is often more useful in practice.

Definition 2.7.9. An operator $A : X \to X$ is **dissipative** if for all $x \in \mathrm{Dom}(X)$ and $\lambda > 0$, we have

$$\|(\lambda I - A)x\| \geqslant \lambda\|x\|.$$

Similarly, A is **accretive** if for all $x \in \mathrm{Dom}(X)$ and $\lambda > 0$, we have

$$\|(\lambda I + A)x\| \geqslant \lambda\|x\|.$$

Theorem 2.7.10 (Lumer-Phillips)**.** *Let $A : X \to X$ be a Banach space operator. The map A is the generator of a contraction semigroup if and only if the following conditions hold.*

(i) A *is densely defined.*

(ii) A *is dissipative.*

(iii) $A - \lambda I$ *is surjective for some $\lambda > 0$.*

2.8 THE SPECTRAL THEOREM

Definition 2.8.1. Let $A : X \to X$ be an operator on a Banach space. The **spectrum** of A, denoted $\sigma(A)$, is the complement of the resolvent set of A in $\mathbb{C}$. That is, $\sigma(A) := \mathbb{C} \backslash \rho(A)$.

Proposition 2.8.2. *Let $A : X \to X$ be a Hilbert space operator. The spectrum $\sigma(A)$ and resolvent $\rho(A)$ satisfy the following properties.*

(i) $\sigma(A) \subseteq \mathbb{C}$ *is open and* $\rho(A) \subseteq \mathbb{C}$ *is closed.*

(ii) If A *is a bounded operator, then* $\sigma(A)$ *is a bounded subset of* $\mathbb{C}$*.*

(iii) If A *is self-adjoint, then* $\sigma(A)$ *is real.*

(iv) If A *is a positive operator, then* $\sigma(A)$ *is non-negative.*

Remark 2.8.3. Unlike finite dimensional vector spaces, not all points in the spectrum of an operator $A : X \to X$ are eigenvalues. That is, when $\lambda \in \sigma(A)$, there is no guarantee that there is a vector x such that $Ax = \lambda x$ (after possibly pushing through the complexification).

Even without eigenvectors, a self-adjoint operator $A : X \to X$ may be diagonalized by the spectral theorem.

Theorem 2.8.4 (spectral theorem)**.** *Let $A : X \to X$ be a self-adjoint operator on a separable $\mathbb{k}$-Hilbert space X. There is a finite measure space (Ω, μ), a unitary isomorphism $\Phi : X \to L^2(\Omega, \mu; \mathbb{k})$, and a real measurable μ-a.e. finite function $f : \Omega \to \mathbb{k}$ with corresponding multiplication operator M_f such that the following conditions hold.*

(i) $x \in \mathrm{Dom}(A)$ *if and only if* $M_f \Phi x = (f) \cdot (\Phi x)$ *is square integrable.*

(ii) $\Phi^{-1} M_f \Phi = A$.

(iii) The essential image of f *is exactly* $\sigma(A)$*.*

Remark 2.8.5. This is the multiplicative form of the spectral theorem; there are other equivalent formulations in terms of spectral measures and direct integrals. The spectral measure formulation will be useful for applying the Borel function calculus, which allows one to apply functions like $\cos(x)$, $\sin(x)/x$, or more generally any Borel function to Hilbert space operators. However, we utilize the multiplicative form whenever possible as it is the most intuitive formulation of the spectral theorem.

The spectral theorem for operators acts as an infinite-dimensional analogue to diagonalization. The multiplication operator M_f behaves like a diagonal operator on $L^2(\Omega, \mu; \mathbb{k})$, and the unitary operator Φ acts like a change of basis.

Part II

HILBERT SHEAVES

3

FOUNDATIONS FOR CELLULAR SHEAVES OF HILBERT SPACES

The functors of the classical theory of cellular sheaves, as developed in Chapter 1, are generally valued in abelian categories, such as the category of vector spaces. Hansen and Ghrist [54] introduced **weighted cellular sheaves**, valued in the category $\mathbf{FinHilb}_{\Bbbk}$ of finite-dimensional Hilbert spaces. Weighted cellular sheaves admit a rich spectral theory, akin to spectral graph theory, which has found application in opinion dynamics [55], neural networks [13, 54], clock synchronization [90], and beyond. However, some applications naturally lead to sheaves valued in infinite-dimensional Hilbert spaces with unbounded, partially-defined operators as restriction maps. This chapter develops the operator-theoretic foundations necessary to extend cellular sheaf theory to this infinite-dimensional setting.

The passage from finite to infinite dimensions introduces two fundamental complications. First, the category of Hilbert spaces and bounded operators, while extensively studied, lacks the abelian structure that makes the finite-dimensional theory so tractable. Second, when we further admit unbounded operators in our cellular sheaves, discontinuities are introduced which cause even basic categorical constructions to become delicate. Composition of morphisms requires careful attention to domains, cochain complexes associated to sheaves may fail to have well-defined cohomology groups, and certain categorical limits cease to exist.

This chapter outlines the addresses these challenges through three main developments. Section 3.1 introduces the necessary categories of Hilbert spaces, carefully distinguishing between bounded operators ($\mathbf{Hilb}_{\Bbbk}$) and unbounded operators ($\mathbf{Hilb}_{0,\Bbbk}$, $\mathbf{CoreHilb}_{0,\Bbbk}$). We demonstrate that while $\mathbf{Hilb}_{\Bbbk}$ retains the essential homological properties of $\mathbf{FinHilb}_{\Bbbk}$, the categories with unbounded operators require the framework of restriction categories to handle partially-defined morphisms. Section 3.2 reviews the theory of Hilbert complexes, due to Brüning and Lesch [17], to provide a suitable generalization of cochain complexes that accommodates unbounded coboundary operators. Finally, Section 3.3 analyzes block operators between direct sums of Hilbert spaces, establishing conditions under which these operators are closed or closable. Together, these developments provide the categorical, homological, and operator-theoretic infrastructure for the cellular sheaf theory developed in subsequent chapters.

To define cellular sheaves valued in Hilbert spaces, we first clarify the relevant categories of Hilbert spaces. The category of $\mathbb{k}$-Hilbert spaces and bounded operators, denoted $\mathbf{Hilb}_{\mathbb{k}}$, is a well studied category in the theories of categorical quantum mechanics [1, 10, 64], dagger categories [62, 109, 110], and operator categories [44, 123]. We will also need a category of Hilbert spaces with unbounded operators, $\mathbf{Hilb}_{0,\mathbb{k}}$, which is substantially less well-behaved. This category, with its partially-defined maps, may be fruitfully studied through the lens of restriction categories [26–28]. Other categories of Hilbert and Banach spaces have also been studied [79, 88], but are not suitable for cellular sheaf theory.

3.1.1 *The category of Hilbert spaces and bounded operators*

Definition 3.1.1. The category of **$\mathbb{k}$-Hilbert spaces and bounded operators**, denoted $\mathbf{Hilb}_{\mathbb{k}}$, consists of the following data.

- **Objects.** The objects of $\mathbf{Hilb}_{\mathbb{k}}$ are $\mathbb{k}$-Hilbert spaces.
- **Morphisms.** A morphism $A : X \to Y$ is a globally-defined $\mathbb{k}$-linear bounded operator.

Remark 3.1.2. The category $\mathbf{Hilb}_{\mathbb{k}}$ has been extensively studied. Moreover, the categorical structure has recently been axiomatized by Heunen and Kornell [63].

Remark 3.1.3. Morphisms in the category $\mathbf{Hilb}_{\mathbb{k}}$ are not required to respect the inner product structure on each object, but merely the topologies they induce. Continuity has no regard for orthogonality. Consequently, isomorphisms in $\mathbf{Hilb}_{\mathbb{k}}$ are simply bounded linear bijections, not unitary maps.

$\mathbf{Hilb}_{\mathbb{k}}$ is also a prototypical example of a **dagger category**, possibly first introduced by Burgin under the name "categories with involutions" [18]. A dagger category is a category $\mathcal{C}$ equipped with a **dagger**—a contravariant involutive functor $(-)^{\dagger} : \mathcal{C} \to \mathcal{C}^{op}$ which is the identity functor on objects. In $\mathbf{Hilb}_{\mathbb{k}}$, the dagger is given by the linear adjoint $(-)^*$. Dagger categories are used in the study of categorical quantum mechanics, where dagger compact categories form a general setting for studying the underlying operations of quantum theory [64].

The extra data of a Banach space norm makes $\mathbf{Hilb}_{\mathbb{k}}$ less well behaved than $\mathbf{Vect}_{\mathbb{k}}$, the category of $\mathbb{k}$-vector spaces with linear maps for morphisms. For example, $\mathbf{Hilb}_{\mathbb{k}}$ fails to be an abelian category. Perhaps the most straightforward way to see this is that for a morphism $A : X \to Y$, we have an isomorphism $\text{co-im}(A) \cong \text{im}(A)$ if and only if $\mathcal{R}(A)$ is closed in Y.

Nonetheless, $\mathbf{Hilb}_{\Bbbk}$ is *almost* abelian—it is a prototypical example of the recently defined concept of an R*-**category** [87]. In a few words, an R*-category is an additive dagger-category whose biproducts and kernels are dagger limits. They can be thought of through the following analogy; abelian categories are to abelian groups what R*-categories are to Hilbert spaces. As an R*-category, $\mathbf{Hilb}_{\Bbbk}$ is contained in several standard classes of weakly abelian categories. $\mathbf{Hilb}_{\Bbbk}$ is **quasi-abelian** in the sense of Schneiders [107, Definition 1.1.3], and hence is **homological** [103], satisfying the five, nine, and snake lemmas, as well as admitting the homology long exact sequence. Moreover, $\mathbf{Hilb}_{\Bbbk}$ is finitely complete.

Remark 3.1.4. One might object that morphisms in $\mathbf{Hilb}_{\Bbbk}$ failing to respect the inner products is an indication that $\mathbf{Hilb}_{\Bbbk}$ is not the "correct" category of Hilbert spaces to work with. If one wishes to avoid these problems, there are two clear options.

First, one could change the morphism from bounded linear maps to maps that respect the inner product structure directly, such as unitary operators, isometries, or partial isometries. While restricting to any of these classes of morphisms yields a valid category, it will not be adequate for cellular sheaf theory. These classes of maps are too restrictive to make for an interesting category of Hilbert spaces, and lack the necessary expressive power for applications of cellular sheaves.

Second, one could study $\mathbf{Hilb}_{\Bbbk}$ as a dagger category *per se*. Following *the way of the dagger* [72], to fully treat $\mathbf{Hilb}_{\Bbbk}$ as a dagger category is to preserve the structure of linear adjoints whenever possible. For example, while any continuous bijective linear map is an isomorphism in $\mathbf{Hilb}_{\Bbbk}$, a **dagger isomorphism** is an adjoint-preserving isomorphism, which is exactly a unitary map. While fruitful in other domains [64], treating $\mathbf{Hilb}_{\Bbbk}$ as a dagger category is too rigid for the theory of cellular sheaves; there are insufficiently many dagger-respecting limits (see Remark 4.2.3). Moreover, the dagger perspective cannot accommodate unbounded operators.

3.1.2 *Categories of Hilbert spaces and unbounded operators*

We will also utilize categories of Hilbert spaces with unbounded and partially-defined operators. Such categories are understudied in comparison to $\mathbf{Hilb}_{\Bbbk}$. We approach unbounded and partially defined operators through two key categories of Hilbert spaces.

- $\mathbf{Hilb}_{0,\Bbbk}$, the category of Hilbert spaces and partially-defined linear operators.
- $\mathbf{CoreHilb}_{0,\Bbbk}$, where unbounded partial operators are required to respect certain subspace containment relationships.

These categories of Hilbert spaces, which are best thought of as 2-categories, follow the approaches of Robinson and Rosolini's categories of partial maps [101], Carboni's bicategories of partial maps [19], and Cockett and Lack's restriction categories [26–28].

Notation 3.1.5. Let $A : X \to Y$ be an unbounded operator with $\mathrm{Dom}(A) = V \subseteq X$. To disambiguate the domain of A as an unbounded operator and the domain of A as a morphism in a category, we use **domain** exclusively to refer the operator domain V, and **source** to refer the object X.

Definition 3.1.6. The **category of $\Bbbk$-Hilbert spaces and partially-defined operators**, denoted $\mathbf{Hilb}_{0,\Bbbk}$, consists of the following data.

- **Objects.** The objects of $\mathbf{Hilb}_{0,\Bbbk}$ are $\Bbbk$-Hilbert spaces.
- **Morphisms.** A morphism $A : X \to Y$ is a $\Bbbk$-linear operator A with a specified domain $\mathrm{Dom}(A) \subseteq X$. Operators which map by the same rule on different partial domains constitute distinct morphisms in $\mathbf{Hilb}_{0,\Bbbk}$. When convenient, we will denote a morphism $A : X \to Y$ as a pair $(A, \mathrm{Dom}(A))$ to highlight this domain dependence.
- **Composition.** The composition of morphisms

  $$X \xrightarrow{A} Y \xrightarrow{B} Z$$

  is the usual composition $B \circ A : X \to Z$ with domain $\{x \in \mathrm{Dom}(A) \,:\, Ax \in \mathrm{Dom}(B)\}$.

Let A and B be a composable pair of closed densely-defined operators. Since $B \circ A$ need not be closed nor densely defined, we cannot restrict the class of morphisms to only closed densely-defined operators. This difficulty, at least for densely-defined operators, can be partially mitigated by including certain domain information in the objects and restricting the admissible ranges of operators. This motivates our second category of Hilbert spaces and unbounded operators.

Definition 3.1.7. The category of **cored $\Bbbk$-Hilbert spaces**, denoted $\mathbf{CoreHilb}_{0,\Bbbk}$, is the following category:

- **Objects.** Each object of $\mathbf{CoreHilb}_{0,\Bbbk}$ is a pair (X, V) of a $\Bbbk$-Hilbert space X and a linear subspace $V \subseteq X$. We call V the **core** of the object (X, V);
- **Morphisms.** A morphism $A : (X, V) \to (Y, W)$ is a $\Bbbk$-linear operator A with domain $\mathrm{Dom}(A) \supseteq V$ such that $A(V) \subseteq W$. When convenient, we will denote a morphism $A : X \to Y$ as a pair $(A, \mathrm{Dom}(A))$.
- **Composition.** The composition of morphisms

  $$(X, V) \xrightarrow{A} (Y, W) \xrightarrow{B} (Z, U)$$

  is the usual composition $B \circ A$ with domain $\{x \in \mathrm{Dom}(A) \,:\, Ax \in \mathrm{Dom}(B)\}$. By hypothesis, $V \subseteq \mathrm{Dom}(B \circ A)$.

Remark 3.1.8. We make a few remarks about these categories of Hilbert spaces and unbounded operators.

1. One may form a full subcategory $\mathbf{DenseCoreHilb}_{0,\Bbbk}$ induced by those objects (X, V) such that V is a dense linear subspace of X. Consequently, all morphisms (and their compositions) are densely-defined operators. This will be a convenient category to work in, as it will ensure maps out of the object may be assumed to have a common dense domain. Unfortunately, since the composition of closed densely-defined operators need not be closed, we cannot form a category of closed densely-defined operators with cores. Nonetheless, we use $\mathbf{ClDenseCoreHilb}_{0,\Bbbk}$ to refer to the *quiver* of densely-cored Hilbert spaces and closed densely-defined linear operators objects and morphisms.

2. $\mathbf{Hilb}_{\Bbbk}$ embeds into $\mathbf{Hilb}_{0,\Bbbk}$ by proper inclusion. Moreover, $\mathbf{Hilb}_{\Bbbk}$ embeds into the category $\mathbf{CoreHilb}_{0,\Bbbk}$ in two different ways:

$$(X \xrightarrow{A} Y) \mapsto ((X, 0) \xrightarrow{A} (Y, 0)),$$
$$(X \xrightarrow{A} Y) \mapsto ((X, X) \xrightarrow{A} (Y, Y)).$$

   Neither of these inclusions are full, and the second inclusion lands inside of the quiver $\mathbf{ClDenseCoreHilb}_{0,\Bbbk}$.

3. $\mathbf{Hilb}_{0,\Bbbk}$ itself embeds into $\mathbf{CoreHilb}_{0,\Bbbk}$ by:

$$\iota : (X \xrightarrow{A} Y) \mapsto ((X, 0) \xrightarrow{A} (Y, 0)).$$

   Similarly, there is a forgetful functor $\mathcal{U} : \mathbf{CoreHilb}_{0,\Bbbk} \to \mathbf{Hilb}_{0,\Bbbk}$ defined by

$$\mathcal{U} : \big((X, V) \xrightarrow{A} (Y, W)\big) \mapsto (X \xrightarrow{A} Y).$$

4. $\mathbf{Hilb}_{\Bbbk}$ is a subcategory of $\mathbf{TopVect}_{\Bbbk}$, the category of topological vector spaces with continuous linear maps. On the other hand, $\mathbf{Hilb}_{0,\Bbbk}$ and $\mathbf{CoreHilb}_{0,\Bbbk}$ are not subcategories of $\mathbf{TopVect}_{\Bbbk}$, as unbounded operators are discontinuous.

5. Taking adjoints is not a functorial operation on $\mathbf{Hilb}_{0,\Bbbk}$ or $\mathbf{CoreHilb}_{\Bbbk,0}$ as not every unbounded operator admits a uniquely defined adjoint. Similarly adjoints are not functorial on $\mathbf{DenseCoreHilb}_{0,\Bbbk}$. While every densely-defined morphism $A : (X, V) \to (Y, W)$ has an adjoint, it need not be densely defined (and hence not a morphism in $\mathbf{DenseCoreHilb}_{0,\Bbbk}$) when A isn't a closed operator. Therefore these categories are not dagger categories. However adjoints form a dagger-like structure on the *quiver* of Hilbert spaces and closed densely-defined operators.

Proposition 3.1.9. *The embedding* $\iota : \mathbf{Hilb}_{0,\Bbbk} \to \mathbf{CoreHilb}_{0,\Bbbk}$ *and the forgetful functor*

$$\mathcal{U} : \mathbf{CoreHilb}_{0,\Bbbk} \to \mathbf{Hilb}_{0,\Bbbk}$$

form an adjoint pair $\iota \dashv \mathcal{U}$.

Proof. Let $X \in \mathbf{Hilb}_{0,\Bbbk}$ and $(Y, W) \in \mathbf{CoreHilb}_{0,\Bbbk}$. We may define a natural bijection $\Phi_{X,(Y,W)}$: $\mathbf{CoreHilb}_{0,\Bbbk}(\iota(X), (Y, V)) \to \mathbf{Hilb}_{0,\Bbbk}(X, \mathcal{U}(Y, V))$ by

$$\Phi_{X,(Y,W)}((Y, 0) \xrightarrow{A} (X, V)) \mapsto (Y \xrightarrow{A} X).$$

It is straightforward to check this assignment in natural in X and (Y, W). □

Unlike $\mathbf{Hilb}_{\Bbbk}$, which is finitely complete, both $\mathbf{Hilb}_{0,\Bbbk}$ and $\mathbf{CoreHilb}_{0,\Bbbk}$ lack many finite limits. For example, $\mathbf{Hilb}_{0,\Bbbk}$ is missing certain pullbacks, as shown in the next example.

Example 3.1.10. Let $X \xrightarrow{A} Z \xleftarrow{B} Y$ be a cospan where $\mathrm{Dom}(A)$ and $\mathrm{Dom}(B)$ are proper subsets of X and Y. One may try to form a pullback by defining $K := \{(x, y) : Ax = By\} \subseteq X \oplus Y$, and form the following square with projections for legs.

$$\begin{array}{ccc} \overline{K} & \xrightarrow{\pi_Y} & Y \\ {\scriptstyle \pi_X}\downarrow & & \downarrow{\scriptstyle B} \\ X & \xrightarrow{A} & Z \end{array}$$

Consider another Hilbert space H and maps $\alpha : H \to X$ and $\beta : H \to Y$ such that $B\beta = A\alpha$. There is a map $\phi : H \to \overline{K}$ such that $\pi_X \phi = \alpha$ and $\pi_Y \phi = \beta$ if and only if $\mathrm{Dom}(\alpha) = \mathrm{Dom}(\beta)$. If this equality of domains does not hold, $\phi(h) = (\alpha(h), \beta(h))$ can only be defined on $\mathrm{Dom}(\alpha) \cap \mathrm{Dom}(\beta)$, and $\overline{K}$ fails to be a pullback.

Remark 3.1.11. A nearly identical argument can be used to show that the categories $\mathbf{CoreHilb}_{0,\Bbbk}$ and $\mathbf{DenseCoreHilb}_{0,\Bbbk}$ fail to have certain pullbacks as well.

This lack of limits poses a problem for defining cellular sheaves valued in Hilbert spaces with unbounded operators. For example, it becomes unclear how to define sections, which are usually defined as limits. The theory of restriction categories, as elucidated by Cockett and Lack [26–28], shows how to view certain limit-like constructions (including the proposed pullback square in Example 3.1.10) as a weaker notion of limit in a certain 2-category. The key is to add the extra structure of "restrictions," which model identity maps on partial domains.

Definition 3.1.12 (Restriction category [26, Definition 2.1.1]). A **restriction structure** on a category $\mathcal{C}$ is an assignment of a **restriction idempotent**[1] $\hat{f} : X \to X$ for each morphism $f : X \to Y$ that that satisfies the following axioms.

R.1 $f\hat{f} = f$ for all f.

R.2 $\hat{f}\hat{g} = \hat{g}\hat{f}$ for all arrows f, g with the same source.

R.3 $\widehat{g\hat{f}} = \hat{g}\hat{f}$ for all arrows f, g with the same source.

R.4 $\hat{g}f = f \circ \widehat{(gf)}$ for all composable pairs of arrows.

We call a pair $(\mathcal{C}, \widehat{(-)})$ of a category and a restriction structure a **restriction category**.

The intended interpretation of a restriction structure is that $f : X \to Y$ is a partial function with domain-of-definition $\mathrm{Dom}(f) \subseteq X$, and the restriction idempotent $\hat{f} : X \to X$ is the identity map $\mathrm{id}_X : x \to x$ with domain $\mathrm{Dom}(\mathrm{id}_X) = \mathrm{Dom}(f)$.

Proposition 3.1.13. *Let* $A := (A, \mathrm{Dom}(A)) : (X, V) \to (Y, W)$ *be a morphism in* $\mathbf{CoreHilb}_{0,\Bbbk}$. *The assignment* $\widehat{A} = (I_X, \mathrm{Dom}(A))$ *defines a restriction structure on* $\mathbf{CoreHilb}_{0,\Bbbk}$.

Proof. This is a straightforward computation of partial maps. For example, R.1 is confirmed by observing $(A, \mathrm{Dom}(A)) \circ (I_X, \mathrm{Dom}(A)) = (A, \mathrm{Dom}(A))$. The other properties are similar. □

Remark 3.1.14. $\mathbf{Hilb}_{0,\Bbbk}$ and $\mathbf{DenseCoreHilb}_{0,\Bbbk}$ are also restriction categories with the same restriction. This is witnessed by their inclusions into $\mathbf{CoreHilb}_{0,\Bbbk}$, and the fact that $\mathrm{Dom}(\hat{A}) = \mathrm{Dom}(A)$. In this section, we will mostly work with $\mathbf{CoreHilb}_{0,\Bbbk}$, as all results will similarly apply to $\mathbf{Hilb}_{0,\Bbbk}$ and $\mathbf{DenseCoreHilb}_{0,\Bbbk}$ by their inclusions into $\mathbf{CoreHilb}_{0,\Bbbk}$ as subcategories.

Definition 3.1.15. A morphism $f : X \to Y$ in a restriction category $\mathcal{C}$ is **total** if $\hat{f} = \mathrm{id}_X$.

In $\mathbf{CoreHilb}_{0,\Bbbk}$, the total morphisms are exactly those morphisms that are globally defined. The total maps in $\mathcal{C}$ define a wide subcategory $\mathbf{Total}(\mathcal{C}) \subseteq \mathcal{C}$.

Remark 3.1.16. Every restriction category $\mathcal{C}$ can be viewed as a 2-category in the following way. The 0-cells and 1-cells are exactly the objects and morphisms of $\mathcal{C}$. If $f, g : X \to Y$ are parallel morphisms in $\mathcal{C}$, there is a unique 2-cell $F : f \Rightarrow g$ if and only if $f = g\hat{f}$. That is, exactly when f is a restriction of g to a smaller partial domain.

Restriction categories come equipped with a notion of a **restriction limit**, which is a limit-like construction in the corresponding 2-categories.

1 The restriction idempotent of a morphism f is usually denoted by $\bar{f}$. We use $\hat{f}$ to disambiguate from the closure of an operator.

Definition 3.1.17 (Restriction limit [28, Section 4.4])**.** Let $\mathcal{C}$ be a restriction category, and $\mathcal{J}$ a diagram. The **restriction limit** of a functor $F : \mathcal{J} \to \mathcal{C}$ is a cone $p_J : L \to F(J)$ with total legs that satisfies the following universal property. If $q_J : M \to F(J)$ is a lax cone over F (meaning for all $f : I \to J$ in $\mathcal{J}$, $F(f) \circ q_I = q_J \circ (\widehat{F(f)\ q_I})$), there is a unique arrow $\phi : M \to L$ satisfying $p_J \circ \phi = q_J \circ e$, where e is the composite of all restriction idempotents $\widehat{q_J}$.

Notation 3.1.18. When working in a restriction category $\mathcal{C}$, we will use $\operatorname{reslim} F$ to denote the restriction limit of a functor $F : \mathcal{J} \to \mathcal{C}$.

Example 3.1.19. The pullback-like square constructed in Example 3.1.10 is a restriction limit in $\mathbf{Hilb}_{0,\Bbbk}$. In particular, the arrow $\phi : H \to \overline{K}$ defined by $\phi(h) = (\alpha(h), \beta(h))$ with domain $\mathrm{Dom}(\alpha) \cap \mathrm{Dom}(\beta)$ satisfies $\pi_Y \phi = \alpha\hat{\alpha}\hat{\beta}$ and $\pi_X \phi = \beta\hat{\alpha}\hat{\beta}$.

Remark 3.1.20. We make a few remarks about restriction limits.

1. Restriction limits admit a natural interpretation as 2-categorical limits, where the triangles

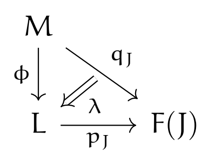

   commute up to a not-necessarily-invertible 2-cell λ.

2. By the universal property of restriction limits, when a functor $F : \mathcal{J} \to \mathcal{C}$ admits a restriction limit, the restriction limit is unique up to isomorphism.

3. When the functor $F : \mathcal{J} \to \mathcal{C}$ is valued in $\mathbf{Total}(\mathcal{C})$, the restriction limit of F in $\mathcal{C}$ is an honest limit in $\mathbf{Total}(\mathcal{C})$.

Theorem 3.1.21. $\mathbf{CoreHilb}_{0,\Bbbk}$ *admits all finite restriction limits.*

Proof. Every restriction idempotent $\hat{f}$ in $\mathbf{CoreHilb}_{0,\Bbbk}$ admits a section s and a retract r. Moreover, $\hat{f}$ can be uniquely determined from the data of the section s or the retract r. This makes $\mathbf{CoreHilb}_{0,\Bbbk}$ a **split retraction category** [26, Section 2.3.3].

Next, $\mathbf{Total}(\mathbf{CoreHilb}_{0,\Bbbk})$ has a terminal object given by $(0, 0)$, and admits all pullbacks by a similar construction to that in Example 3.1.10. Thus $\mathbf{Total}(\mathbf{CoreHilb}_{0,\Bbbk})$ admits all finite limits [84]. By [28, Proposition 4.12], $\mathbf{CoreHilb}_{0,\Bbbk}$ admits all finite restriction limits. □

Corollary 3.1.22. $\mathbf{Hilb}_{0,\Bbbk}$ *and* $\mathbf{DenseCoreHilb}_{0,\Bbbk}$ *admit all finite restriction limits.*

Remark 3.1.23. Finite restriction limits in $\mathbf{CoreHilb}_{0,\Bbbk}$ can be constructed in a manner similar to the restriction pullback in Example 3.1.10. Given a finite diagram $F : \mathcal{J} \to \mathbf{CoreHilb}_{0,\Bbbk}$, say that $\mathbf{x} = (x_J) \in \bigoplus_{J \in \mathcal{J}} F(J)$ is F-**admissible** if the following conditions hold.

(i) For all morphisms $f : J \to K$ in $\mathcal{J}$, the coordinate $x_J \in \mathrm{Dom}(F(f))$.

(ii) For all morphisms $f : J \to K$ in $\mathcal{J}$, we have $F(f)(x_J) = x_K$.

Let $L \subseteq \bigoplus_{J \in \mathcal{J}} F(J)$ denote the collection of all F-admissible points. The zero section is always F-admissible, so L is a linear subspace of the Hilbert space $\bigoplus_{J \in \mathcal{J}} F(J)$, and its closure $\overline{L}$ is a sub-Hilbert space. Moreover, L forms a cone over $F(\mathcal{J})$ with globally defined legs $p_J : (\overline{L}, L) \to F(J)$ given by coordinate projections. This is easily seen to be a restriction limit. Moreover, since L will be dense in $\overline{L}$, this construction of the restriction limit restricts to $\mathbf{DenseCoreHilb}_{0,\Bbbk}$. Pushing through the forgetful functor $\mathcal{U} : \mathbf{CoreHilb}_{0,\Bbbk} \to \mathbf{Hilb}_{0,\Bbbk}$ generates restriction limits in $\mathbf{Hilb}_{0,\Bbbk}$ as well.

3.1.3 *Additivity in restriction categories*

We now briefly discuss additivity in restriction categories. While $\mathbf{Hilb}_{0,\Bbbk}$ is not abelian (or even pre-additive category), unbounded operators have a clear additive structure sufficient for some homological algebra. Given parallel arrows $A, B : X \to Y$, we may define an operator sum $(A + B) : x \mapsto Ax + Bx$ with domain $\mathrm{Dom}(A) \cap \mathrm{Dom}(B)$. This sum operation gives the homset $\mathbf{Hilb}_{0,\Bbbk}(X, Y)$ the structure of a commutative monoid. Moreover, this monoid admits a weak notion of an additive inverse; for every operator A there is an operator $(-A)$ such that $A + (-A) = 0$ with domain $\mathrm{Dom}(A)$. This gives $\mathbf{Hilb}_{0,\Bbbk}(X, Y)$ the structure of an abelian **Clifford semigroup** [25], which form a subclass of **inverse semigroups** in the sense of Wagner [122] and Preston [98].

Definition 3.1.24. An inverse semigroup is a semigroup $(S, \cdot)$ such that for all $x \in S$, there is a unique $x^{-1} \in S$ such that $x = xx^{-1}x$ and $x^{-1} = x^{-1}xx^{-1}$. An inverse semigroup is a Clifford semigroup if every element $x \in S$ is in a subgroup of the semigroup.

Remark 3.1.25. Composition of linear operators does not give the homset $\mathbf{Hilb}_{0,\Bbbk}(X, Y)$ the structure of a Clifford semiring. In general, the operators $C(A + B)$ and $CA + CB$ do not have the same domain.

We now discuss additivity in a general restriction category $\mathcal{C}$. For an object $X \in \mathcal{C}$, let $\mathtt{D}(X)$ denote the set of restriction idempotents on X. $\mathtt{D}(X)$ admits the structure of a meet semilattice with order $e \leq f \iff ef = fe = e$ and meet $e \wedge f = ef = fe$.

Definition 3.1.26. Let $A : X \to Y$ be a morphism in a restriction category $\mathcal{C}$. A **corestriction** of A is a restriction idempotent $\check{A} \in \mathtt{D}(Y)$ such that $\check{A}A = A$, and if $e \in \mathtt{D}(Y)$ is a restriction idempotent such that $eA = A$, then $\check{A} \leq e$.

We characterize a few properties of the lattices $\mathtt{D}(X)$ and coresetrictions with the following lemma.

Lemma 3.1.27. *Let $\mathcal{C}$ be a restriction category. The following properties hold.*

(i) For a composable pair of arrows BA, $\widehat{BA} \leq \hat{A}$.

(ii) For a composable pair of arrows BA, if $\hat{B}A = A$, then $\widehat{BA} = \hat{A}$.

(iii) A morphism $A : X \to Y$ admits a corestriction if and only if $\bigwedge\{e \in \mathsf{D}(Y) \ : \ eA = A\}$ exists. In particular, if the semilattice $\mathsf{D}(X)$ is meet-complete for all $X \in \mathcal{C}$, then all morphisms admit corestrictions.

Proof. Let $\mathcal{C}$ be a restriction category.

(i) From the axioms of a restriction category, we compute $\widehat{BA}\hat{A} = \widehat{BA\hat{A}} = \widehat{BA}$. Therefore $\widehat{BA} \leq \hat{A}$.

(ii) $A = \hat{B}A = A\widehat{\hat{B}A}$, from which we conclude $\hat{A} = \widehat{A\widehat{\hat{B}A}} = \hat{A}\widehat{\hat{B}A}$. From (i), we recover $\widehat{BA} = \hat{A}$.

(iii) This follows directly from the definitions of meets and corestrictions.

□

Definition 3.1.28. Let $\mathcal{C}$ be a restriction category. The **domain category** of $\mathcal{C}$, denoted $\hat{\mathcal{C}}$, consists of the following data.

- **Objects.** Each object of $\hat{\mathcal{C}}$ is a pair (X, e) of an object $X \in \mathcal{C}$ and a restriction idempotent $e \in \mathsf{D}(X)$.
- **Morphisms.** A morphism $A : (X, e) \to (Y, f)$ is a $\mathcal{C}$-morphism $A : X \to Y$ such that $\hat{A} = e$ and $fA = A$.
- **Composition.** The composition $(X, e) \xrightarrow{A} (Y, f) \xrightarrow{B} (Z, g)$ is given by the composition BA in $\mathcal{C}$. This composition is well-defined by Lemma 3.1.27.

Remark 3.1.29. The domain category $\hat{\mathcal{C}}$ is not a restriction category. Instead, the partial domain data is packaged into the objects directly. Moreover, the structure of morphisms in $\hat{\mathcal{C}}$ allows us to only compose arrows that compose *properly*.

The presence of restriction limits in $\mathcal{C}$ guarantees the existence of certain classical limits and colimits in $\hat{\mathcal{C}}$. This correspondence is mediated by the following forgetful functor.

Definition 3.1.30. Let $\mathcal{U} : \hat{\mathcal{C}} \to \mathcal{C}$ denote the functor which acts by

$$\mathcal{U}\left((X, e) \xrightarrow{A} (Y, f)\right) := X \xrightarrow{A} Y.$$

Proposition 3.1.31. *Let $F : \mathcal{J} \to \hat{\mathcal{C}}$ be a finite diagram in $\hat{\mathcal{C}}$. If $\mathcal{U}F$ admits a restriction limit in $\mathcal{C}$, then F admits a limit in $\hat{\mathcal{C}}$.*

Proof. Let $p_J : L \to \mathcal{U}F(J)$ denote the restriction limiting cone in $\mathcal{C}$. Let $\ell \in \mathtt{D}(L)$ denote the restriction idempotent obtained by composing every restriction idempotent of the form $\widehat{\mathcal{U}F(\alpha) \circ p_J}$, where $\alpha : J \to J'$ is a $\mathcal{J}$-morphism. Consider the cone $q_J : (L, \ell) \to F(J)$ with legs $q_J = p_J \ell$. Given another cone $r_J : (M, m) \to F(J)$ in $\hat{\mathcal{C}}$, there is a is a lift to a cone $R_J : M \to \mathcal{U}F(J)$ in $\mathcal{C}$ with the same legs. There is a unique $\mathcal{C}$-morphism $\phi : M \to L$ such that $p_J \phi = R_J m$. The map $\ell\phi : (M, m) \to (L, \ell)$ witnesses that $r_J : (M, m) \to F(J)$ is a limiting cone. □

Corollary 3.1.32. *If $\mathcal{C}$ is finitely restriction complete, then $\hat{\mathcal{C}}$ is finitely complete. If $\mathcal{C}$ is restriction complete and each semilattice $\mathtt{D}(X)$ is meet-complete, then $\hat{\mathcal{C}}$ is complete.*

Example 3.1.33. The domain category $\widehat{\mathbf{Hilb}}_{0,\Bbbk}$ has pairs (X, V) of a Hilbert space X and a linear subspace $V \subseteq X$ for objects. A morphism $A : (X, V) \to (Y, W)$ in $\widehat{\mathbf{Hilb}}_{0,\Bbbk}$ is a linear map $A : X \to Y$ with domain $\mathrm{Dom}(A) = V$ and a range containment $\mathcal{R}(A) \subseteq W$. The domain category $\widehat{\mathbf{Hilb}}_{0,\Bbbk}$ is a wide subcategory of $\mathbf{CoreHilb}_{0,\Bbbk}$. Moreover, this category is easily seen to be categorically equivalent to the category of $\Bbbk$-vector spaces.

Definition 3.1.34. A **restriction pre-additive** category consists of the following data.

- A restriction category $\mathcal{C}$.
- All morphisms admit corestrictions.
- Each collection of restriction idempotents $\mathtt{D}(X)$ forms a lattice.
- Each homset $\mathcal{C}(X, Y)$ admits the structure of an abelian Clifford semigroup.

We require that the following conditions hold.

(i) For each pair of parallel morphisms $A, B : X \to Y$, the sum has restriction idempotent $\widehat{A+B} = \hat{A} \wedge \hat{B} = \hat{A}\hat{B}$ and corestriction $\widetilde{A+B} = \check{A} \vee \check{B}$.

(ii) For parallel morphisms $A, B : X \to Y$ and $C : Y \to Z$, if $\check{A}, \check{B} \leq \hat{C}$, then $C(A + B) = CA + CB$.

(iii) For parallel morphisms $A, B : Y \to Z$ and $C : X \to Y$, if $\check{C} \leq \hat{A} \wedge \hat{B}$, then $(A + B)C = AC + BC$.

Proposition 3.1.35. *If $\mathcal{C}$ is a restriction pre-additive category, the domain category $\hat{\mathcal{C}}$ is pre-additive.*

Proof. Since $A = A + (-A) + A$, condition (i) enforces that $\hat{A} = \hat{A} \wedge \widehat{(-A)}$, and therefore $\hat{A} = \widehat{(-A)}$. Similarly, $\check{A} = \widetilde{(-A)}$. In an abelian inverse semigroup, every additive idempotent α satisfies $\alpha = -\alpha$. Hence every additive idempotent can be written as $A + (-A)$ for some $A \in \mathcal{C}(X, Y)$. Given an idempotent α, it follows straightforwardly that $A + (-A) = \alpha$ if and only if $\hat{A} = \hat{\alpha}$.

The collection $G_\alpha := \{A \in \mathcal{C}(X,Y) \,:\, A + (-A) = \alpha\}$ forms an abelian subgroup of $\mathcal{C}(X,Y)$ with identity α. Moreover, we may write

$$\begin{aligned} G_\alpha &= \{A \in \mathcal{C}(X,Y) \,:\, \hat{A} = \hat{\alpha}\} \\ &= \hat{\mathcal{C}}\big((X,\hat{\alpha}),(Y,\mathrm{id}_Y)\big). \end{aligned}$$

Let $\hat{C}\big((X,e),(Y,f)\big)$ be a homset in the domain category. Given parallel morphisms $A, B : (X,e) \to (Y,f)$, we have $\widehat{A+B} = \hat{A} \wedge \hat{B} = e$, and $\widecheck{A+B} = \check{A} \vee \check{B} \leq f$. Therefore $\hat{C}\big((X,e),(Y,f)\big)$ is closed sums. Moreover, this homset is closed under inverses as $e = \widehat{(-A)}$ and $\widecheck{(-A)} \leq f$. Therefore $\hat{C}\big((X,e),(Y,f)\big)$ is an abelian subgroup of $\hat{C}\big((X,e),(Y,\mathrm{id}_Y)\big)$. The distributivity conditions (ii) and (iii) ensure that composition is bilinear. □

Definition 3.1.36. A **weak restriction zero-object** in a restriction category $\mathcal{C}$ is an object 0 with the following properties.

(i) 0 is **restriction terminal** [28, Section 4.1]; for every object A there is a total arrow $t_A : A \to 0$ such that $t_0 = \mathrm{id}_0$ and $t_B f = t_A \hat{f}$ for all $f : A \to B$.

(ii) 0 is initial in $\mathcal{C}$; for each X, there is a unique total morphism $!_X : 0 \to X$.

(iii) The **zero-morphism** $0_{XY} : X \to Y$ is the map $0_{X,Y} = !_Y t_X$. For each object X, we require $\widehat{0_{XX}} = 0_{XX}$.

Remark 3.1.37. Zero morphisms are not annihilative on the left. Given a map $g : X \to Y$, the composition $0_{YZ} g = 0_{XZ} \hat{g}$, which need not be total.

Definition 3.1.38. Let X and Y be objects in a restriction pre-additive category $\mathcal{C}$. A **binary restriction biproduct** of X and Y is an object $X \oplus Y$ and a collection of four total maps

$$\iota_X : X \to X \oplus Y, \quad \iota_Y : Y \to X \oplus Y, \quad \pi_X : X \oplus Y \to X, \quad \pi_Y : X \oplus Y \to Y$$

such that the following conditions hold.

(i) **Restriction product**. $X \oplus Y$, equipped with its projection maps π_X and π_Y, is a **restriction product**; for maps $A : Z \to X$ and $B : Z \to Y$ there is a unique map $\langle A, B \rangle : Z \to X \oplus Y$ such that $A\hat{B} = \pi_X \langle A, B \rangle$, $B\hat{A} = \pi_Y \langle A, B \rangle$, and $\widehat{\langle A, B \rangle} = \hat{A}\hat{B}$.

(ii) **Weak restriction coproduct**. $X \oplus Y$, equipped with its inclusion maps ι_X and ι_Y, is a **weak restriction coproduct**[2]; for any maps $A : X \to Z$ and $B : Y \to Z$, there is a unique morphism

2 We say "weak" to differentiate from Cockett and Lack's notion of a **restriction coproduct** [28], which is simply a coproduct.

$[A, B] : X \oplus Y \to Z$ such that $[A, B]\iota_X = A$ and $[A, B]\iota_Y = B$, and the restriction $[A, B]$ is controlled componentwise:

$$\widehat{[A, B]}\iota_X = \iota_X \hat{A}, \quad \widehat{[A, B]}\iota_Y = \iota_Y \hat{B}.$$

(iii) **Biproduct equations.** We require that the following equations hold:

$$\pi_X \iota_X = \mathrm{id}_X, \quad \pi_Y \iota_Y = \mathrm{id}_Y, \quad \pi_X \iota_Y = 0_{YX}, \quad \pi_Y \iota_X = 0_{XY}.$$

We additionally require that $\iota_X \pi_X + \iota_Y \pi_Y = \mathrm{id}_{X \oplus Y}$.

Theorem 3.1.39. *Let $\mathcal{C}$ be a restriction pre-additive category. If $\mathcal{C}$ has a weak restriction zero-object and admits all binary restriction biproducts, then $\hat{\mathcal{C}}$ is additive.*

Proof. First, we show that $\hat{\mathcal{C}}$ admits a zero object. Let 0 be the restriction zero-object in $\mathcal{C}$. Since 0 is initial in $\mathcal{C}$, $\mathrm{id}_0 : 0 \to 0$ is the unique restriction idempotent in $\mathtt{D}(0)$. For each $X \in \mathcal{C}$ and restriction idempotent $e \in \mathtt{D}(X)$, $e \circ ! = !$, where $! : 0 \to X$ is the unique morphism from 0 to X. Therefore there is a unique morphism $! : (0, \mathrm{id}_0) \to (X, e)$ in $\hat{\mathcal{C}}$, and $(0, \mathrm{id}_0)$ is initial. Meanwhile, Proposition 3.1.31 ensures that $(0, \mathrm{id}_0)$ is terminal in $\hat{\mathcal{C}}$.

Next, we show that $\hat{\mathcal{C}}$ admits binary biproducts. Let (X, e) and (Y, f) be objects in $\hat{\mathcal{C}}$. The morphism $e \oplus f := \iota_X e \pi_X + \iota_Y f \pi_Y$ is restriction idempotent on $X \oplus Y$. We claim $(X \oplus Y, e \oplus f)$ is a biproduct in $\hat{\mathcal{C}}$. By the Proposition 3.1.31, $(X \oplus Y, e \oplus f)$ is a product. To check that $(X \oplus Y, e \oplus f)$ is a coproduct, let $A : (X, e) \to (Z, g)$ and $B : (Y, f) \to (Z, g)$. Any map $\phi : (X \oplus Y, e \oplus f) \to (Z, g)$ such that $\phi\iota_X = A$ and $\phi\iota_Y = B$ lifts to a morphism $\phi : X \oplus Y \to Z$ in $\mathcal{C}$ that satisfies the universal property of the weak restriction coproduct. Moreover, the map $[A, B] : X \oplus Y \to Z$ defines a map $[A, B] : (X \oplus Y, e \oplus f) \to (Z, g)$ that satisfies the universal property of the coproduct. Therefore $(X \oplus Y, e \oplus f)$ is a coproduct. Finally, the biproduct equations hold by the construction of $e \oplus f$. Therefore $\hat{\mathcal{C}}$ admits all finite biproducts.

Since $\hat{\mathcal{C}}$ is pre-additive by Proposition 3.1.35 and admits all finite biproducts, $\hat{\mathcal{C}}$ is additive. □

3.2 HILBERT COMPLEXES

As seen in section 1.4.5, a cellular sheaf valued in an Abelian category naturally induces a cochain complex. Since $\mathbf{Hilb}_{\Bbbk}$ is quasi-abelian, it is straightforward to port results to cellular sheaves valued in Hilbert spaces and bounded operators. However, to generalize to unbounded partially defined operators, which do not form a quasi-abelian category, requires more care. We now discuss the theory of Hilbert complexes: a class of sufficiently well behaved cochain complexes of Hilbert spaces with unbounded coboundary maps. Hilbert complexes were introduced by

Brüning and Lesch [17] to study elliptic complexes and de Rahm cohomology. Since their introduction, Hilbert complexes have proved invaluable for studying partial differential equations [95, 96], finite-element exterior calculus [8, 65], and more.

Definition 3.2.1. A **Hilbert complex** $(X^\bullet, V^\bullet, \delta^\bullet)$ is the data of a family of Hilbert space operators $\{\delta^j : X^j \to X^{j+1}\}_{j \in \mathbb{N}}$ with $\mathrm{Dom}(\delta^j) = V^j$, such that the following conditions hold.

(i) Each operator δ^j is closed and densely defined.

(ii) $\mathcal{R}(\delta^j) \subseteq V^{j+1}$ for all j.

(iii) $\delta^{j+1}\delta^j(x) = 0$ for all j and all $x \in V^j$.

Notation 3.2.2. Despite the domain conditions, we will frequently abbreviate the Hilbert complex $(X^\bullet, V^\bullet, \delta^\bullet)$ to $X^\bullet$, and draw the Hilbert complex as

$$X^0 \xrightarrow{\delta^0} X^1 \xrightarrow{\delta^1} \cdots .$$

We say a Hilbert complex is **finite** if $X^j = 0$ for all j sufficiently large. We say that a Hilbert complex is **bounded** if each map δ^j is a globally-defined bounded operator. For clarity, we adopt the convention that $X^{-1} = 0$, and δ^{-1} is the zero-map.

3.2.1 *The Hodge decomposition*

While not a cochain complex in an Abelian category, Hilbert complexes give enough structure to recover analogues of key results like Hodge decompositions, cohomology groups, and Laplacians. However, care must be taken to accommodate the fact that domains $V^\bullet$ and ranges $\mathcal{R}^\bullet := \mathcal{R}(\delta^\bullet)$ need not be closed linear subspaces.

Definition 3.2.3. Let $(X^\bullet, V^\bullet, \delta^\bullet)$ be a Hilbert complex. A k-**cocycle** is a point in the kernel $x \in \ker(\delta^k)$. We use $\mathfrak{Z}^k(X^\bullet) := \ker(\delta^k)$ to refer to the Hilbert space of k-cocycles. Similarly, a k-**coboundary** is a point in the image $y \in \mathcal{R}(\delta^{k-1})$. We use $\mathfrak{B}^k(X^\bullet)$ to refer to the space of k-coboundaries. Finally, let $\mathfrak{B}^{k\perp}$ denote the orthogonal complement of $\mathfrak{B}^k$ in X^k. A point $x \in X^k$ is k-**harmonic** if $x \in \mathfrak{B}^{k\perp} \cap \mathfrak{Z}^k$. We let $\mathfrak{H}^k(X^\bullet)$ denote the k-harmonic space.

Notation 3.2.4. When the Hilbert complex is clear from context, we reduce the notation to $\mathfrak{Z}^k$, $\mathfrak{B}^k$, and $\mathfrak{H}^k$.

Remark 3.2.5. Since the kernel of a closed operator is closed, $\mathfrak{Z}^k$ is closed, and hence a Hilbert space. Similarly, $\mathfrak{H}^k$ is a Hilbert space as the intersection of two closed subspaces. $\mathfrak{B}^k$, on the other hand, is only a linear subspace in general. When $\mathfrak{B}^k$ is closed for all k, we say that the Hilbert complex $(X^\bullet, V^\bullet, \delta^\bullet)$ is **closed**.

In a Hilbert complex, we get a decomposition of each X^k in terms of the spaces of coboundaries, cocycles, and harmonic points. In particular, we may use the properties of direct sums and orthogonal complements to compute:

$$\begin{aligned} X^k &= \mathfrak{Z}^k \oplus \mathfrak{Z}^{k\perp} \\ &= \left[\mathfrak{Z}^k \cap (\overline{\mathfrak{B}^k} \oplus \mathfrak{B}^{k\perp})\right] \oplus \mathfrak{Z}^{k\perp} \\ &= (\mathfrak{Z}^k \cap \overline{\mathfrak{B}^k}) \oplus (\mathfrak{Z}^k \cap \mathfrak{B}^{k\perp}) \oplus \mathfrak{Z}^{k\perp} \\ &= \overline{\mathfrak{B}^k} \oplus \mathfrak{H}^k \oplus \mathfrak{Z}^{k\perp}. \end{aligned}$$

This resulting decomposition $X^k = \overline{\mathfrak{B}^k} \oplus \mathfrak{H}^k \oplus \mathfrak{Z}^{k\perp}$ is known as the **weak Hodge decomposition** of X^k. When $\mathfrak{B}^k$ is closed, we recover the **strong Hodge decomposition** $X^k = \mathfrak{B}^k \oplus \mathfrak{H}^k \oplus \mathfrak{Z}^{k\perp}$.

3.2.2 *Morphisms and the domain complex*

Care must be taken with domains to define morphisms of Hilbert complexes as chain maps. In order to commute, we require that the chain maps respect the domains of the coboundary operators of the Hilbert complexes.

Definition 3.2.6. Let $(X^\bullet, V^\bullet, \delta_X^\bullet)$ and $(Y^\bullet, W^\bullet, \delta_Y^\bullet)$ be Hilbert complexes. A **morphism**

$$f^\bullet : (X^\bullet, V^\bullet, \delta_X^\bullet) \to (Y^\bullet, W^\bullet, \delta_Y^\bullet)$$

is a family of bounded, globally defined operators $f^k : X^k \to Y^k$ that satisfy the following conditions.

(i) $f^k(\mathrm{Dom}(\delta_X^k)) \subseteq \mathrm{Dom}(\delta_Y^k)$ for all k.

(ii) $\delta_Y^k f^k x = f^{k+1} \delta_X^{k+1} x$ for all $x \in \mathrm{Dom}(\delta_X^k)$.

Hilbert complexes and Hilbert complex morphisms form a category, $\mathbf{HilbComp}_{\mathbb{k}}$. For some homological algebra, Hilbert complex morphisms are not adequately structured. Diagram chases, like those of the five and snake lemmas, cannot be performed; lifts of elements by the components of a Hilbert complex morphism may fail to lie inside the domain of the coboundary operators. To mitigate these difficulties, we introduce the domain complex.

Definition 3.2.7. The **domain complex** of a Hilbert complex $(X^\bullet, V^\bullet, \delta_X^\bullet)$ is the Hilbert complex $(V^\bullet, V^\bullet, \delta^\bullet)$, where the domain V^k is a Hilbert space with the graph inner product

$$\langle x, y \rangle_{\Gamma(\delta^k)} = \langle x, y \rangle_{X^k} + \langle \delta^k x, \delta^k y \rangle_{X^{k+1}}.$$

It is straightforward to verify that the domain complex of a Hilbert complex is a well-defined bounded Hilbert complex. Moreover, a Hilbert complex morphism $f^\bullet : (X^\bullet, V^\bullet, \delta_X^\bullet) \to (Y^\bullet, W^\bullet, \delta_Y^\bullet)$ induces a morphism of the domain complexes $f^\bullet|_{V^\bullet} : (V^\bullet, V^\bullet, \delta_X^\bullet) \to (W^\bullet, W^\bullet, \delta_Y^\bullet)$. By imposing conditions on both Hilbert complex morphisms and their induced morphisms between domain complexes, key results of homological algebra may be recovered through the following adaptation of "exactness".

Definition 3.2.8. Let $(X, V_X) \xrightarrow{f} (Y, V_Y) \xrightarrow{g} (Z, V_Z)$ denote a sequence of Hilbert spaces X, Y, Z containing linear subspaces V_X, V_Y, V_Z respectively, and bounded globally defined operators $f : X \to Y$ and $g : Y \to Z$ such that $f(V_X) \subseteq V_Y$, and $g(V_Y) \subseteq V_Z$. This sequence of maps is **exact** at Y if $\mathcal{R}(f) = \ker(g)$. The sequence is **pair exact** at (Y, V_Y) if it is exact at Y, and $\mathcal{R}(f|_{V_X}) = \ker(g|_{V_Y})$. That is, the pair (f, g) is algebraically exact on both the full Hilbert spaces, and on the underlying linear subspaces.

This definition of exactness may be lifted to morphisms of Hilbert complexes by requiring exactness gradewise.

Definition 3.2.9. Let $(X^\bullet, V_X^\bullet, \delta_X^\bullet) \xrightarrow{f^\bullet} (Y^\bullet, V_Y^\bullet, \delta_Y^\bullet) \xrightarrow{g^\bullet} (Z^\bullet, V_Z^\bullet, \delta_Z^\bullet)$ be a composable pair of Hilbert complex morphisms. This sequence of maps is **exact** at $(Y^\bullet, V_Y^\bullet, \delta_Y^\bullet)$ if $(X^k, V_X^k) \xrightarrow{f^k} (Y^k, V_Y^k) \xrightarrow{g^k} (Z^k, V_Z^k)$ is pair exact in each grade k.

3.2.3 *Cohomology of Hilbert complexes*

Hilbert complexes also come equipped with a notion of cohomology.

Definition 3.2.10. Let $(X^\bullet, V^\bullet, \delta^\bullet)$ be a Hilbert complex. The k^{th}-**reduced cohomology** of the Hilbert complex is the Hilbert space quotient $H^k := \mathfrak{Z}^k/\overline{\mathfrak{B}^k}$. When $\mathfrak{B}^k$ is closed (δ^{k-1} has closed range), we call $\mathfrak{Z}^k/\overline{\mathfrak{B}^k} = \mathfrak{Z}^k/\mathfrak{B}^k$ the k^{th}-**cohomology**.

The k^{th}-reduced cohomology always carries a Hilbert space structure as a quotient of Hilbert spaces. Moreover, we have the following relationship between the k^{th}-reduced cohomology and the k^{th}-harmonic space.

Theorem 3.2.11. *Let $(X^\bullet, V^\bullet, \delta^\bullet)$ be a Hilbert complex. There is a natural unitary equivalence $H^k \cong \mathfrak{H}^k$.*

Proof. $\mathfrak{H}^k = \mathfrak{Z}^k \cap \mathfrak{B}^{k\perp}$ is a closed subspace of $\mathfrak{Z}^k$. Let $P : \mathfrak{Z}^k \to \mathfrak{H}^k$ denote the orthogonal projection onto the k-harmonic space. P induces a unitary isomorphism $\hat{P} : H^k \to \mathfrak{H}^k$. □

Unfortunately, one cannot easily derive an "honest" cohomology theory for an arbitrary Hilbert complex. As discussed in Section 3.1, the category $\mathbf{Hilb}_{\Bbbk}$ is not an abelian category. In particular, if $\mathfrak{B}^k$ isn't closed, the k^{th}-cohomology $\mathfrak{Z}^k/\mathfrak{B}^k$ will not be a Hilbert space, but merely a vector space. This vector space is well defined, of course, but ceases to hold an obvious interpretation for the Hilbert complex. Further, even if each $\mathfrak{B}^k$ is closed and all cohomologies exist as Hilbert spaces, it can still be difficult to interpret. For example, since each H^k may be an infinite dimensional Hilbert space, standard interpretations in terms of Betti numbers are not applicable.

One approach to interpreting the cohomology is through obstructions via exact sequences. While the homological algebra of Hilbert complexes is well-studied for the de Rahm complex [7, 36], the general homological algebra of Hilbert is comparatively understudied [49].

Lemma 3.2.12 (Snake lemma for Hilbert complexes)**.** *Suppose we have a commutative diagram of Hilbert spaces*

$$\begin{array}{ccccccccc}
 & & \ker(A) & \longrightarrow & \ker(B) & \longrightarrow & \ker(C) & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & X & \xrightarrow{f} & Y & \xrightarrow{g} & Z & \longrightarrow & 0 \\
 & & \downarrow A & & \downarrow B & & \downarrow C & & \\
0 & \longrightarrow & X' & \xrightarrow{f'} & Y' & \xrightarrow{g'} & Z' & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
 & & X'/\overline{\mathcal{R}(A)} & \longrightarrow & Y'/\overline{\mathcal{R}(B)} & \longrightarrow & Z'/\overline{\mathcal{R}(C)} & &
\end{array}$$

with morphisms in the top row given by restriction of f *and* g *to the kernels of* A *and* B*, and morphisms in the bottom row given by* $[x'] \mapsto [f'x']$ *and* $[y'] \mapsto [g'y']$*. Further suppose that the rows* $(X, \mathrm{Dom}(A)) \xrightarrow{f} (Y, \mathrm{Dom}(B)) \xrightarrow{g} (Z, \mathrm{Dom}(C)) \xrightarrow{0} (0, 0)$ *and* $0 \xrightarrow{0} X' \xrightarrow{f'} Y' \xrightarrow{g'} Z'$ *are pair exact and exact respectively, and each operator* A, B, C *is closed and densely-defined. There is a bounded globally-defined connecting morphism* d *forming a bounded Hilbert complex*

$$\ker(A) \longrightarrow \ker(B) \longrightarrow \ker(C) \xrightarrow{d} X'/\overline{\mathcal{R}(A)} \longrightarrow Y'/\overline{\mathcal{R}(B)} \longrightarrow Z'/\overline{\mathcal{R}(C)} \; .$$

This sequence is always exact at $\ker(B)$*. Moreover, the following hold.*

(i) If A *has closed range, the sequence is exact at* $\ker(C)$*.*

(ii) If B *has closed range, the sequence is exact at* $X'/\overline{\mathcal{R}(B)}$*.*

(iii) If C *has closed range, the sequence is exact at* $Y'/\overline{\mathcal{R}(B)}$*.*

Proof. The maps of the top and bottom row are all bounded and globally defined, and exactness at $\ker(B)$ is automatic. To construct the map d, we perform the usual diagram chase of the snake lemma; for $z \in \ker(C)$, since $g|_{\mathrm{Dom}(B)}$ is a surjection, there is a $y \in \mathrm{Dom}(B)$ such that $gy = z$. Since $z \in \ker(C)$, we recover that $By \in \ker(g')$, and hence there is a unique $x' \in X'$ such that $f'x' = By'$.

Denote this point x' by $d_0(y)$, as it only depended on our choice of y. We define $d(z) := [x']$. To check d is well defined, notice that d only depended on our choice of $y \in \mathrm{Dom}(B)$ such that $gy = z$. Take $\tilde{y} \in \mathrm{Dom}(B)$ to be a second such choice, and $\tilde{x}' = d_0(\tilde{y})$ to be the corresponding point in X'. Consider the difference $w = y - \tilde{y}$. The point Bw is still in the kernel of g', so $x' - \tilde{x}'$ is the unique point in X' mapped to Bw by f'. w is in the linear space $\ker(g) \cap \mathrm{Dom}(B)$, so there is an $\alpha \in \mathrm{Dom}(A)$ such that $f\alpha = w$. By the injectivity of f', we recover $A\alpha = x' - \tilde{x}'$, proving $x - x' \in \mathcal{R}(A) \subseteq \overline{\mathcal{R}(A)}$. Therefore $d : \ker(C) \to X'/\overline{\mathcal{R}(A)}$ is globally well-defined.

Now we check that d gives the structure of a Hilbert complex. Suppose $k \in \ker(B)$. To check that $gk \in \ker(d)$, it suffices to notice that $d_0(k) = 0$. Additionally, for $z \in \ker(C)$ and a choice of $y \in \mathrm{Dom}(B)$ such such that $gy = z$, we may check that $[f(d_0(y))] = [By] = [0]$ in $Y'/\overline{\mathcal{R}(B)}$.

To confirm that d is bounded, recognize that $g : \mathrm{Dom}(B) \to \ker(C)$ admits a bounded right inverse when $\mathrm{Dom}(B)$ is topologized with the graph norm of B, and $f' : X' \to \mathcal{R}(f') = \ker(g')$ admits a bounded inverse when $\mathcal{R}(f')$ is topologized as a sub-Hilbert space of Y'. The map $d : \ker(C) \to X'/\overline{\mathcal{R}(A')}$ may be written as the composition of bounded operators

$$\ker(C) \xrightarrow{g^{-1}} \mathrm{Dom}(B) \xrightarrow{B\big|_{g^{-1}(\ker(C))}} \ker(g') \xrightarrow{(f')^{-1}} X' \xrightarrow{\pi} X'/\overline{\mathcal{R}(A')} \ ,$$

which proves d is bounded.

Finally, in the case that $\mathcal{R}(A)$, $\mathcal{R}(B)$ and $\mathcal{R}(C)$ are respectively closed, these maps align with the usual snake lemma in vector spaces, yielding exactness at the corresponding locations. □

Remark 3.2.13. This version of the snake lemma for Hilbert complexes is stronger than the version from the abelian category structure of the domain category $\widehat{\mathbf{Hilb}}_{0,\Bbbk}$. In the domain category, the *topological quotient* $X'/\overline{\mathcal{R}(A)}$ lacks a categorical interpretation. However, when all maps have closed range, the algebraic and topological quotients agree, and the snake lemma provides the same exact sequence.

We may now use the snake lemma to provide an interpretation for higher cohomologies like H^1 for closed Hilbert complexes. Suppressing the domains and coboundary maps for notational ease, let $X^\bullet, Y^\bullet$, and $Z^\bullet$ be closed Hilbert complexes. Suppose we have a short exact sequence of complexes

$$0 \to X^\bullet \xrightarrow{f^\bullet} Y^\bullet \xrightarrow{g^\bullet} Z^\bullet \to 0.$$

Applying the snake lemma gives a long exact sequence

$$H^0(X^\bullet) \longrightarrow H^0(Y^\bullet) \longrightarrow H^0(Z^\bullet) \longrightarrow H^1(X^\bullet) \longrightarrow H^2(Y^\bullet) \longrightarrow \cdots .$$

The connecting homomorphism $d : H^0(Z^\bullet) \to H^1(Y^\bullet)$ can now be interpreted as the obstruction to lifting a 0-cocycle in $Z^\bullet$ to a 0-cocycle in $Y^\bullet$. In particular, every 0-cocycle $z^0 \in Z^0$ has a preimage $y^0 \in Y^0$ under g^0. In general, y^0 will not be a cocycle, but its image $y^1 = \delta_Y^0 y^0$ will live in the kernel $\ker(g^1)$. By exactness, there is a unique $x^1 \in X^1$ such that $f^1 x^1 = y^1$. Moreover, $d[z^0] = [x^1] \in H^1(X^\bullet)$, showing that the obstruction to y^0 being a cocycle lives in $H^1(X^\bullet)$. The other cohomology groups $H^k(X^\bullet)$ can be interpreted in a similar manner.

3.2.4 *Subcomplexes and relative cohomology*

Definition 3.2.14. Recall that a morphism of Hilbert complexes $f^\bullet : X^\bullet \to Y^\bullet$ is a collection of globally-defined bounded Hilbert space morphisms $f^k : X^k \to Y^k$ such that $f^k(V^k) \subseteq W^k$ and the following diagram commutes (with respect to the domain of $\delta_X^\bullet$).

$$\begin{array}{ccccccc} \cdots & \longrightarrow & X^k & \xrightarrow{\delta_X^k} & X^{k+1} & \longrightarrow & \cdots \\ & & \downarrow{f^k} & & \downarrow{f^{k+1}} & & \\ \cdots & \longrightarrow & Y^k & \xrightarrow{\delta_Y^k} & Y^{k+1} & \longrightarrow & \cdots \end{array}$$

When each component ι^k of a Hilbert complex morphism $\iota^\bullet : X^\bullet \to Y^\bullet$ is an inclusion of a sub-Hilbert space, we say that $(X^\bullet, V^\bullet, \delta_X^\bullet)$ is a **subcomplex** of $(Y^\bullet, W^\bullet, \delta_Y^\bullet)$.

Subcomplexes allow us to define the relative cohomology of the Hilbert complex. Suppose $(X^\bullet, V^\bullet, \delta_X^\bullet)$ is a subcomplex of $(Y^\bullet, W^\bullet, \delta_Y^\bullet)$. Each X^k is a sub-Hilbert space of Y^k, and we may form a Hilbert space quotient Y^k/X^k. The correspondence $\delta_{Y/X}^k : Y^k/X^k \to Y^{k+1}/X^{k+1}$ defined by $\delta_{Y/X}^k : [y] \mapsto [\delta_Y^k y]$ has domain $D^k = \{[y] \in Y^k/X^k \;:\; y \in W_k\}$ induces the **relative Hilbert complex**

$$\cdots \longrightarrow Y^k/X^k \xrightarrow{\delta_{Y/X}^k} Y^{k+1}/X^{k+1} \longrightarrow \cdots .$$

The **relative (reduced) homology** of $Y^\bullet$ with respect to $X^\bullet$ is the (reduced) homology of this complex, denoted $H^\bullet(Y^\bullet, X^\bullet)$.

By Theorem 3.2.11, there is a unitary equivalence between the k^{th}-reduced homology of the relative Hilbert complex and the k^{th}-**relative harmonic space** $H^k(Y^\bullet, X^\bullet) \cong \mathfrak{H}^k(Y^\bullet, X^\bullet)$. By the identification of $Y^k/X^k \cong X^{k\perp}$, the relative harmonic space $\mathfrak{H}^k(Y^\bullet, X^\bullet)$ can be identified with the harmonic points $y \in \mathfrak{H}^k(Y^\bullet)$ such that $y \in X^{k\perp}$. That is, $\mathfrak{H}^k(Y^\bullet, X^\bullet)$ is the collection of harmonic points in Y^k that vanish on X^k.

As a cochain complex, a Hilbert complex admits a Hodge Laplacian. This Laplacian (and its associated Dirichlet problem) has been well-studied [8, 65], and has recently been linked to a variety of coupling problems in physics [14]. To define the Hodge Laplacian, we begin with the dual complex of a Hilbert complex.

Definition 3.2.15. The **dual complex** of a Hilbert complex is chain complex obtained by reversing arrows via adjoints. We denote the dual complex by $(X_\bullet, V^*_\bullet, d_\bullet)$, where $X_k := X^k$, and $d_k := (\delta^{k-1})^*$ with domain $V^*_k \subseteq X_k$. Hence the dual complex is a Hilbert complex (with descending indices)

$$X_0 \xleftarrow{d_1} X_1 \xleftarrow{d_2} \cdots .$$

Remark 3.2.16. Due to the decreasing indexing, the dual complex to a Hilbert complex is a chain complex instead of cochain complex. For finite Hilbert complexes, this difference is inconsequential, but the change is significant for infinite complexes. Moreover, the dual complex to a bounded Hilbert complex is always bounded, and the dual complex to a closed complex is always closed by the closed range theorem (Theorem 2.3.15).

Definition 3.2.17. Let $(X^\bullet, V^\bullet, \delta^\bullet)$ be a Hilbert complex. The **Hodge Laplacian** of $X^\bullet$ is the chain operator $\mathcal{L} := \delta d + d\delta$ with graded components given by

$$\mathcal{L}^k := \delta^{k-1} d_k + d_{k+1}\delta^k.$$

This is a (generically) unbounded operator with domain

$$\mathrm{Dom}(\mathcal{L}^k) = \{x \in V^k \cap V^*_k \,:\, d_k x \in V^{k-1} \text{ and } \delta^k x \in V^*_{k+1}\}.$$

Notation 3.2.18. We write the Hodge Laplacian the sum $\mathcal{L} = \mathcal{L}_+ + \mathcal{L}_-$ of the **up Laplacian** $\mathcal{L}^k_+ := d_{k+1}\delta^k$ and **down Laplacian** $\mathcal{L}^k_- := \delta^{k-1} d_k$.

Theorem 3.2.19. *The kernel of the Hodge Laplacian is the harmonic space. That is,* $\ker(\mathcal{L}^k) = \mathfrak{H}^k$.

Proof. Each coboundary map is closed, and we may identify $\mathfrak{B}^{k\perp} = \ker(d_k)$. Consequently $\mathfrak{H}^k = \ker(d_k) \cap \ker(\delta^k)$, from which it follows that $\mathfrak{H}^k \subseteq \ker(\mathcal{L}^k)$.

To prove the reverse inclusion, consider the up Laplacian and the down Laplacian separately. We may check that both $\mathcal{R}(d_k) \perp \ker(\delta^{k-1})$ and $\mathcal{R}(\delta^k) \perp \ker(d_{k+1})$, which forces $\ker(\mathcal{L}^k_-) = \ker(d_k)$ and $\ker(\mathcal{L}^k_+) = \ker(\delta^k)$. Moreover, the orthogonality of the weak Hodge decomposition witnesses that $\mathcal{R}(\delta^{k-1}) \perp \mathcal{R}(d_{k+1})$, and consequently that $\mathcal{R}(\mathcal{L}^k_-) \perp \mathcal{R}(\mathcal{L}^k_+)$. Therefore the kernel of $\mathcal{L}^k = \mathcal{L}^k_- + \mathcal{L}^k_+$ is exactly the intersection $\ker(d_k) \cap \ker(\delta^k) = \mathfrak{H}^k$. □

We immediately obtain the following corollary for the grade-zero Laplacian.

Corollary 3.2.20. *The kernel of $\mathcal{L}^0$ is exactly the kernel of δ^0.*

Theorem 3.2.21. *The Hodge Laplacian $\mathcal{L}^k$ is a positive operator.*

Proof. Consider the block operator $\begin{bmatrix} \delta^{k-1} & d_{k+1} \end{bmatrix} : X^{k-1} \oplus X_{k+1} \to X^k$. This block operator is closed and has dense domain $V^{k-1} \oplus V^*_{k+1}$. To verify this, suppose that $(x_n, y_n) \in V^{k-1} \oplus V^*_{k+1}$ is a sequence of points in the domain such that $(x_n, y_n) \to (x, y) \in X^{k-1} \oplus X_{k+1}$, and $\delta^{k-1}x_n + d_{k+1}y_n \to z$ in X^k. Since $\mathcal{R}(\delta^{k-1}) \subseteq \ker(\delta^k)$, and $\ker(\delta^k) \perp \mathcal{R}(d_{k+1})$, it follows that $\delta^{k-1}x_n \to z_1$ and $d_{k+1}y_n \to z_2$ separately in X^k. δ^{k-1} and d_{k+1} are closed operators, so $(x, y) \in V^{k-1} \oplus V^*_{k+1}$ and $\delta^{k-1}x + d_{k+1}y = z_1 + z_2 = z$. Hence $\begin{bmatrix} \delta^{k-1} & d_{k+1} \end{bmatrix}$ is closed.

Next, the following computation demonstrates that

$$\begin{bmatrix} \delta^{k-1} & d_{k+1} \end{bmatrix}^* = \begin{bmatrix} d_k \\ \delta^k \end{bmatrix}$$

with domain $V^*_k \cap V^k$; for $(x, y) \in V^{k-1} \oplus V^*_{k+1}$ and $z \in V^*_k \cap V^k$:

$$\begin{aligned} \left\langle \begin{bmatrix} \delta^{k-1} & d_{k+1} \end{bmatrix} (x, y), z \right\rangle &= \langle \delta^{k-1}x, z \rangle + \langle d_{k+1}y, z \rangle \\ &= \langle x, d_k z \rangle + \langle y, \delta_k z \rangle \\ &= \left\langle (x, y), \begin{bmatrix} d_k \\ \delta^k \end{bmatrix} z \right\rangle. \end{aligned}$$

By von Neumann's Theorem (2.5.11), it follows that $\mathcal{L}^k = \begin{bmatrix} \delta^{k-1} & d_{k+1} \end{bmatrix} \begin{bmatrix} \delta^{k-1} & d_{k+1} \end{bmatrix}^*$ is a positive operator. □

3.3 BLOCK OPERATORS OF HILBERT SPACES

A fundamental aspect of finite-dimensional linear algebra is the ability to represent a linear map between direct sums of vector spaces by a block matrix—a matrix whose entries are themselves matrices encoding maps between the vector space summands. While block operators generalize straightforwardly for bounded Hilbert space operators, the theory is more subtle for unbounded operators [92, 117]. We highlight and develop a few key results for bounded and unbound block operators.

Let $\{X_j\}_{j\in J}$ be a family of Hilbert spaces with a possibly infinite index set J. The **direct sum** $\mathcal{X} := \bigoplus_{j\in J} X_j$ is the Hilbert space

$$\left\{ \mathbf{x} = (x_j)_{j\in J} \ : \ \sum_{j\in J} \|x_j\|^2_{X_j} < \infty \right\}$$

equipped with the inner product

$$\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{X}} := \sum_{j\in J} \langle x_j, y_j \rangle.$$

For finite index sets J, the direct product is a categorical biproduct in $\mathbf{Hilb}_{\mathbb{k}}$. However this is not the case for infinite index sets; the direct sum fails to be either a product or a coproduct. Nonetheless, for any summand X_j, one may factor the identity map $I_j : X_j \to X_j$ as $\pi_j \circ \iota_j$, where $\iota_j : X_j \hookrightarrow \mathcal{X}$ is the inclusion of the j'th summand X_j as $X_j \oplus \bigoplus_{j'\neq j} 0$, and $\pi_j : \mathcal{X} \to X_j$ is the projection $\mathbf{x} \mapsto x_j$.

Let $\{X_j\}_{j\in J}$ and $\{Y_i\}_{i\in I}$ be Hilbert spaces with finite index sets I, J, and let $\mathcal{X} := \bigoplus_{j\in J} X_j$ and $\mathcal{Y} := \bigoplus_{i\in I} Y_i$ denote the direct sums. Given an operator $\mathbf{A} : \mathcal{X} \to \mathcal{Y}$ with domain $\mathbf{V} := \mathrm{Dom}(\mathbf{A})$, if the domain $\mathbf{V}$ splits as $\mathbf{V} = \bigoplus_j V_j$ with $V_j \subseteq X_j$, there is an induced operator $A_{ij} : X_j \to Y_i$ for each pair (i, j) such that the following diagram commutes.

$$\begin{array}{ccc} \mathcal{X} & \xrightarrow{\mathbf{A}} & \mathcal{Y} \\ {\scriptstyle \iota_j}\uparrow & & \downarrow{\scriptstyle \pi_i} \\ X_j & \xrightarrow{A_{ij}} & Y_i \end{array}$$

The domain of A_{ij} may be taken to be $V_{ij} := V_j$. The collection of operators A_{ij} for $i \in I$ and $j \in J$ constitutes a representation of $\mathbf{A}$ as a **block operator**

$$[A_{ij}]_{i\in I, j\in J} : \mathcal{X} \to \mathcal{Y}.$$

This block operator acts like matrix multiplication; given an input $\mathbf{x} \in \mathcal{X}$, the image of $[A_{ij}]$ on i-component of $\mathcal{Y}$ is

$$(\mathbf{A}\mathbf{x})_i = \sum_{j\in J} A_{ij} x_j.$$

The condition that the domain of $\mathbf{A}$ split over the summands of $\mathcal{X}$ is not always satisfied. Consequently, not every unbounded operator $\mathbf{A} : \mathcal{X} \to \mathcal{Y}$ admits a representation as a block operator. The converse process of assembling a collection of operators between summands into a block operator also has has domain subtleties. Given Hilbert space operators $A_{ij} : X_j \to Y_i$, one may always form a block operator $[A_{ij}]$. However, the domain on which $[A_{ij}]$ may be defined can be arbitrarily small—possibly only $\{0\}$. This can even be the case for unbounded operators with finite index sets I and J. Even when $[A_{ij}]$ is a well-defined operator with a dense domain, it may fail to have essential properties like being closed. This section is dedicated to exploring when such block operators are well-defined Hilbert space operators, and when they are well behaved.

3.3.1 *Finite block operators*

We restrict our attention to **finite block operators** where I and J are finite index sets. Fix $n, m \in \mathbb{N}_{\geqslant 1}$, Hilbert spaces $\mathcal{X} = X_1 \oplus \cdots \oplus X_m$ and $\mathcal{Y} = Y_1 \oplus \cdots \oplus Y_n$, and operators $A_{ij} : X_j \to Y_i$ with dense domain $V_{ij} \subseteq X_j$. The theory of 2x2 block operators, their properties, and their spectra have been well developed. For example, see [21, 48, 117], the last of which has applications to mathematical physics, fluid dynamics, and quantum mechanics. Several of the results in this section are generalizations of the work of Tretter [117].

When each block operator A_{ij} is bounded and globally defined, the induced finite block operator is well-behaved.

Proposition 3.3.1. *Suppose each $A_{ij} : X_j \to Y_i$ is a globally-defined bounded Hilbert space operator. The finite block operator $\mathbf{A} := [A_{ij}]$ has the following properties.*

(i) $\mathbf{A}$ is globally-defined and bounded (and hence closed).

(ii) The linear adjoint of $\mathbf{A}$ is given by $[A_{ij}]^* = [A_{ji}^*]$.

(iii) The composition of such bounded finite block operators is given by matrix multiplication.

Proof. Using the Cauchy-Schwartz inequality and the triangle inequality, it can be shown directly that the operator norm of $\mathbf{A}$ is bounded above by

$$\|\mathbf{A}\|_{\text{op}} \leqslant \sqrt{\sum_{i=1}^{n} \sum_{j=1}^{m} \|A_{ij}\|^2},$$

proving that $\mathbf{A}$ is bounded. Once working with a globally-defined bounded operator the other results are direct computations. □

When the A_{ij} maps are unbounded and partially defined, analysis of block operators becomes substantially more complex. We first investigate the domain of definition. For a fixed $j \in \{1, \ldots, m\}$, the **common core** of the operators $A_{1j}, \ldots, A_{nj}$ is the intersection $\tilde{V}_j := \bigcap_{i=1}^{n} V_{ij}$. The maximal domain of definition for the block operator $\mathbf{A} = [A_{ij}]$ is the direct sum $\bigoplus_{j \in J} \tilde{V}_j \subseteq \mathcal{X}$. Going forward, we will always take this to be the domain of a block operator $[A_{ij}]$ over finite index sets I, J unless otherwise stated. We get the following proposition for free.

Proposition 3.3.2. *The finite block operator* $[A_{ij}]$ *is densely defined if and only if the collection of operators* $\{A_{1j}, \dots, A_{nj}\}$ *has a densely-defined common core for each* $j \in J$.

Remark 3.3.3. For globally-defined operators A_{ij} (bounded or unbounded), the resulting finite block operator $[A_{ij}]$ is globally defined.

For the remainder of this section, we will always assume that $\mathbf{A} = [A_{ij}]$ is densely defined. Even when $[A_{ij}]$ is known to be densely defined, unbounded block operators have subtleties that must be navigated. We enumerate a few of those subtleties here.

1. If $\mathbf{A}$ is closable, the closure $\overline{\mathbf{A}}$ need not admit a block operator structure.
2. Moreover, when $\mathbf{A}$ is closable, the adjoint $\mathbf{A}^* \neq [A_{ji}^*]$ in general. The operator $[A_{ji}^*]$, known as the formal adjoint, need not even be densely defined when $\mathbf{A}$ is closable [91, Example 6.5].
3. Given block operators $\mathbf{A} = [A_{ij}] : \mathcal{X} \to \mathcal{Y}$ and $\mathbf{B}[B_{ki}] : \mathcal{Y} \to \mathcal{X}$, one may form a block operator $\mathbf{C} = [C_{kj}]$ with $C_{kj} = \sum_i B_{ki} A_{ij}$. This operator, defined with it's maximal domain is an extension $\mathbf{BA} \subseteq \mathbf{C}$. In general, these domains will not agree, and the containment will be strict.

3.3.2 *Closable and closed block operators*

To determine when $[A_{ij}]$ is closed or closable is more challenging. There is not a clear way to provide necessary and sufficient conditions in terms of properties of the underlying A_{ij} maps. In particular, it is not the case that $[A_{ij}]$ is closed (resp. closable) whenever each block A_{ij} is closed (resp. closable), as illustrated by the following example.

Example 3.3.4. Let $A : \ell^2(\mathbb{N};\mathbb{R}) \to \ell^2(\mathbb{N};\mathbb{R})$ be the operator defined by $A\left(\sum_j x_j e_j\right) = (\sum_j j x_j e_j)$, where e_j is the j^{th} standard basis element, with domain $\text{Dom}(A) = \{\mathbf{x} : A\mathbf{x} \in \ell^2(\mathbb{N};\mathbb{R})\}$. A is a diagonal operator, and hence is closed. Consider the block operator

$$\mathbf{A} := \begin{bmatrix} A & -A \end{bmatrix} : \ell^2(\mathbb{N};\mathbb{R}) \oplus \ell^2(\mathbb{N};\mathbb{R}) \to \ell^2(\mathbb{N};\mathbb{R}),$$

with domain $\text{Dom}(A) \oplus \text{Dom}(A)$. Fix $\mathbf{x} = \sum_j \frac{1}{j} e_j \in \ell^2(\mathbb{N};\mathbb{R})$ where e_j is the j^{th} standard basis element, , and let $\mathbf{x}^n = \sum_{j=1}^n \frac{1}{j} e_j$. Notice that $(\mathbf{x}^n, \mathbf{x}^n)^T \to (\mathbf{x}, \mathbf{x})^T$ in $\ell^2(\mathbb{N};\mathbb{R}) \oplus \ell^2(\mathbb{N};\mathbb{R})$, and $\mathbf{A}(\mathbf{x}^n, \mathbf{x}^n) = 0$ for all n, forming a constant sequence. The limit $(\mathbf{x}, \mathbf{x})^T$ is not in the domain of $\mathbf{A}$, showing that $\mathbf{A}$ is not closed.

The remainder of this section is dedicated to determining when a finite block operator is closed or closable. We begin with two general purpose lemmas. A first way to force a finite block operator $\mathbf{A}$ be closed is by controlling the graph norm.

Lemma 3.3.5. *Suppose each* A_{ij} *is a closed (resp. closable) operator. If there is a* $C > 0$ *such that*

$$\sum_{i=1}^{n}\sum_{j=1}^{m}\|A_{ij}x_j\|^2 \leqslant C^2\|(\mathbf{x},\mathbf{A}\mathbf{x})\|^2_{\Gamma(\mathbf{A})}$$

for all $\mathbf{x} \in \mathrm{Dom}(\mathbf{A})$*, then* $\mathbf{A} = [A_{ij}]$ *is closed (resp. closable).*

Proof. We prove the "closed" version of the lemma; the "closable" version is similar. Suppose $(\mathbf{x}^{(n)}, \mathbf{A}\mathbf{x}^{(n)})$ is a Cauchy sequence in the graph $\Gamma(\mathbf{A})$. Both $\mathbf{x}^{(n)}$ and $\mathbf{A}\mathbf{x}^{(n)}$ are Cauchy in $\mathcal{X}$ and $\mathcal{Y}$ respectively, and thus have limits $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}$ respectively. By hypothesis, for each $i \in I$ and $j \in J$, the sequence $A_{ij}x_j^{(n)}$ is Cauchy in n, and converges to a $y_{ij} \in Y_i$. Moreover, $x_j = \lim_{n\to\infty} x_j^{(n)}$ is in the domain V_{ij} and satisfies $A_{ij}x_j = y_{ij}$ since A_{ij} is closed. It follows that $\mathbf{x} \in \mathrm{Dom}(\mathbf{A})$ and $\mathbf{A}\mathbf{x} = \mathbf{y}$. Hence $\Gamma(\mathbf{A})$ is complete and $\mathbf{A}$ is a closed operator. □

Our second general lemma simplifies the problem by allowing us to check closedness row by row.

Lemma 3.3.6. *Let* $\mathbf{A} = [A_{ij}] : \mathcal{X} \to \mathcal{Y}$ *be a finite block operator. Let* $R_i = \begin{bmatrix} A_{i1} & \cdots & A_{im} \end{bmatrix}$ *denote the* i^{th} *row of* $\mathbf{A}$*. If each* $R_i : \mathcal{X} \to Y_i$ *with domain* $\mathrm{Dom}(R_i) = \mathrm{Dom}(A_{i1}) \oplus \cdots \oplus \mathrm{Dom}(A_{im})$ *is closed (resp. closable), then* $\mathbf{A}$ *is closed (resp. closable).*

Proof. Note that $\mathrm{Dom}(\mathbf{A}) = \bigcap_i \mathrm{Dom}(R_i)$ is contained in the domain of each row. Let $\mathbf{x}^n \to 0$ denote a convergent sequence sequence in $\mathrm{Dom}(\mathbf{A})$ such that $\mathbf{A}\mathbf{x}^n \to \mathbf{y}$. If R_i is closable, then the i^{th} coordinate $\mathbf{y}_i = 0$. Therefore if each R_i is closable, then $\mathbf{A}$ is closable.

Similarly, let $\mathbf{x}^n \to \mathbf{x}$ denote a sequence in $\mathrm{Dom}(\mathbf{A})$ that converges to a point $\mathbf{x} \in \mathcal{X}$, and $\mathbf{A}\mathbf{x}^n = \mathbf{y}$. If R_i is closed, then $\mathbf{x} \in \mathrm{Dom}(R_i)$, and $R_i\mathbf{x} = \mathbf{y}_i$. Consequently, if each R_i is closed, then $\mathbf{A}$ is closed.

□

Remark 3.3.7. Note that this results also holds if we take the domain of R_i to be $\mathrm{Dom}(\mathbf{A})$.

Remark 3.3.8. If R_i is closable, then each block A_{ij} of R_i is closable [91, Corollary 3.4]. It follows that in order to apply this lemma, it is necessary (but not sufficient) that every entry of $\mathbf{A}$ is closable. However, it is not the case that all blocks of a closable block operator are closable [91, Remark 4.4].

While these lemmas provide practical criteria to check if a block operator is closed or closable, they do not provide any insight into how to *pick blocks* A_{ij} such that $[A_{ij}]$ is closed. Looking at blocks, it is not obvious when the summability condition, nor row-wise closability are satisfied. Thankfully, there are a variety of block-wise conditions under which we can ensure that the finite block operator $\mathbf{A}$ is closed. One of the most simple involves having enough bounded operators in the block operator.

Proposition 3.3.9. *Let $\mathbf{A} = [A_{ij}]$ be a finite block operator such that each row contains at most one unbounded operator—all other entries are bounded and globally defined. If each unbounded operator is closed (resp. closable), then $\mathbf{A}$ is closed (resp. closable).*

Proof. Let $\mathbf{x}^{(n)}$ be a sequence in $\mathrm{Dom}(\mathbf{A})$ such that $\mathbf{x}^{(n)} \to \mathbf{x} \in \mathcal{X}$. Suppose $\mathbf{y}^{(n)} := \mathbf{A}\mathbf{x}^{(n)} \to \mathbf{y} \in \mathcal{Y}$. For each row index i, let j_i denote the column-index of the entry in the i^{th} row that is unbounded. For a fixed i we may write

$$y_i^{(n)} - \sum_{j \neq j_i} A_{ij} x_j^{(n)} = A_{ij_i} x_{j_i}^{(n)}.$$

Taking the limit as $n \to \infty$ on the left hand side yields $y_i - \sum_{j \neq j_i} A_{ij} x_j$. On the right hand side, since A_{ij_i} is closed, we recover $x_{j_i} \in V_{ij_i}$. It follows that $y_i = \sum_j A_{ij} x_j$. Repeating this argument for each i shows that $\mathbf{A}$ is closed. The proof for the closable case is similar. □

Another way to enforce that a finite block operator $[A_{ij}]$ is closed is to impose conditions of the ranges of the underlying blocks.

Proposition 3.3.10. *Let $\mathbf{A} = [A_{ij}]$ be a finite block operator with each A_{ij} a closed operator. Let $R_{ij} := \overline{\mathcal{R}(A_{ij})}$ denote the closure of the range of A_{ij}. If the following conditions hold for each i, then $\mathbf{A}$ is a closed operator.*

(i) The internal direct sum $R_{i1} + \cdots + R_{im} \subseteq Y_i$ is a closed subspace.

(ii) $R_{ij} \cap R_{ik} = \{0\}$ whenever $j \neq k$.

In particular, these conditions hold when the range-closures $\{R_{i1}, \ldots, R_{im}\}$ are pairwise orthogonal.

Proof. Let $\mathbf{x}^{(n)}$ be a sequence in $\mathrm{Dom}(\mathbf{A})$ such that $\mathbf{x}^{(n)} \to \mathbf{x} \in \mathcal{X}$. Suppose $\mathbf{y}^{(n)} := \mathbf{A}\mathbf{x}^{(n)} \to \mathbf{y} \in \mathcal{Y}$. Let $P_{ij} : \bigoplus_j R_{ij} \to R_{ij}$ denote the orthogonal projection operator onto the R_{ij} component, and let $\mathbf{A}_i : \mathcal{X} \to Y_i$ denote the block operator defined by the i^{th} row of $\mathbf{A}$. Since P_{ij} is bounded, in the i^{th} row we may compute:

$$\begin{aligned} P_{ij} y_i &= \lim_{n \to \infty} P_{ij} \mathbf{A}_i \mathbf{x}^{(n)} \\ &= \lim_{n \to \infty} A_{ij} x_j^{(n)}. \end{aligned}$$

Since A_{ij} is closed, $x_j \in V_{ij}$, and $P_{ij}y_i = A_{ij}x_j$. Applying this argument to each row individually shows $\mathbf{x} \in \mathrm{Dom}(\mathbf{A})$ and $\mathbf{Ax} = \mathbf{y}$, proving that $\mathbf{A}$ is closed. □

Remark 3.3.11. Under the hypotheses of Proposition 3.3.10, if each operator A_{ij} has closed range, then $\mathbf{A}$ has closed range as well.

Definition 3.3.12. Let X, Y, Z be Banach spaces, and let $A : X \to Y$ and $B : X \to Z$ be operators. A is **relatively bounded** with respect to B (or simply B-**bounded**) if $\mathrm{Dom}(B) \subseteq \mathrm{Dom}(A)$, and there are constants $a, b \geqslant 0$ such that

$$\|Ax\|_Y \leqslant a\|x\|_X + b\|Bx\|_Z$$

for all $x \in \mathrm{Dom}(B)$. The infimum

$$\delta := \inf\{b \geqslant 0 \,:\, \exists a \geqslant 0 \text{ witnessing the relative bound}\}$$

is called the **relative bound** of A with respect to B.

Definition 3.3.13. Let $\mathbf{A} = [A_{ij}]$ be an $n \times m$ finite block operator with $n \geqslant m$. $\mathbf{A}$ is **diagonally dominant** of order r if for all $i \leqslant n$ and $j \neq i$, the operator A_{ij} is A_{jj}-bounded with relative bound $\delta_{ij} < r$. That is, each entry of each column is relatively bounded (with relative bound $< r$) by the diagonal entry in the column.

Proposition 3.3.14. *Let $\mathbf{A} = [A_{ij}]$ be an $(m + k) \times m$ block operator with $k \geqslant 0$. If $\mathbf{A}$ is diagonally dominant of order $\frac{1}{2}(mk + (m-1)^2)^{-1}$ and A_{ii} is a closed (resp. closable) operator for each $1 \leqslant i \leqslant n$, then $\mathbf{A}$ is a closed (resp. closable) operator.*

Proof. Write $\mathbf{A} = \mathbf{D} + \mathbf{B}$ where $\mathbf{D}$ and $\mathbf{B}$ are the $(m + k) \times m$ block operators consisting of the diagonal and non-diagonal entries of $\mathbf{A}$ respectively. By the closed operator form of the Kato-Rellich theorem [73, Theorem IV.1.16], to prove that $\mathbf{A}$ is closed, it will suffice to prove that $\mathbf{D}$ is closed and that $\mathbf{B}$ is $\mathbf{D}$-bounded with relative bound $\delta < 1$.

$\mathbf{D}$ is closed immediately by Proposition 3.3.10. To show that $\mathbf{B}$ is $\mathbf{D}$-bounded with relative bound $\delta < 1$, we use the relative bounding hypothesis and a remark in section V.4.1 in [73] to pick a pair of constants $a, b \geqslant 0$ with $b < (mk + (m-1)^2)^{-1}$ such that for each $1 \leqslant j \leqslant m$ and $x_j \in \mathrm{Dom}(A_{jj})$,

$$\|A_{ij}x_j\|^2 \leqslant a\|x_j\|^2 + b\|A_{jj}x_j\|^2.$$

Next, write $\mathbf{B}$ in blocks

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_0 \\ \mathbf{B}_1 \end{bmatrix}$$

where $\mathbf{B}_0$ is an $m \times m$ block operator with zero diagonal, and $\mathbf{B}_1$ is $k \times m$. For $\mathbf{B}_0$, compute:

$$\begin{aligned}
\|\mathbf{B}_0\mathbf{x}\|^2 &= \sum_{i=1}^{m} \left\| \sum_{j \neq i} A_{ij} x_j \right\|^2 \\
&\leqslant \sum_{i=1}^{m} \sum_{j \neq i} (m-1) \left\| A_{ij} x_j \right\|^2 \\
&\leqslant \sum_{i=1}^{m} \sum_{j \neq i} (m-1) \left(a\|x_j\|^2 + b\|A_{jj}x_j\|^2 \right) \\
&= \sum_{i=1}^{m} (m-1)^2 \left(a\|x_i\|^2 + b\|A_{ii}x_i\|^2 \right) \\
&= (m-1)^2 a\|\mathbf{x}\|^2 + (m-1)^2 b\|\mathbf{D}\mathbf{x}\|^2 .
\end{aligned}$$

Working with $\mathbf{B}_1$, we similarly compute:

$$\begin{aligned}
\|\mathbf{B}_1\mathbf{x}\|^2 &= \sum_{i=m+1}^{m+k} \left\| \sum_{j=1}^{m} A_{ij} x_j \right\|^2 \\
&\leqslant \sum_{i=m+1}^{m+k} \sum_{j=1}^{m} m \left\| A_{ij} x_j \right\|^2 \\
&\leqslant \sum_{i=m+1}^{m+k} \sum_{j=1}^{m} m \left(a\|x_j\|^2 + b\|A_{jj}x_j\|^2 \right) \\
&= \sum_{i=1}^{m} mk \left(a\|x_i\|^2 + b\|A_{ii}x_i\|^2 \right) \\
&= mka\|\mathbf{x}\|^2 + mkb\|\mathbf{D}\mathbf{x}\|^2 .
\end{aligned}$$

Letting $C := \left(mk + (m-1)^2\right)$, combining these computations yields:

$$\begin{aligned}
\|\mathbf{B}\mathbf{x}\|^2 &= \|\mathbf{B}_0\mathbf{x}\|^2 + \|\mathbf{B}_1\mathbf{x}\|^2 \\
&\leqslant Ca\|\mathbf{x}\|^2 + Cb\|\mathbf{D}\mathbf{x}\|^2.
\end{aligned}$$

Since $bC < 1$ by hypothesis, applying the Kato-Rellich theorem proves that $\mathbf{A} = \mathbf{D} + \mathbf{B}$ is a closed operator. □

Remark 3.3.15. The requisite bound can be improved by observing that the relative bounds within each column trade off against each other. That is, if the operators in column j are A_{jj}-

bounded with a smaller relative bound, then the operators in column $j' \neq j$ are allowed to have a larger relative bound. See [117, Theorem 2.2.8] for details.

Remark 3.3.16. Diagonal dominance, while essential for the previous argument, is a more restrictive condition than is needed. By rearranging the order of the summands in $\mathcal{X}$ and $\mathcal{Y}$, any block operator where every column has a dominant operator (with the required relative bounds), each living in a distinct row, can be rearranged into a diagonally dominant operator.

Remark 3.3.17. Given an $n \times (n+k)$ block operator $\mathbf{A} = [A_{ij}]$ with more columns than rows, we may still apply Proposition 3.3.14 through a padding argument. Extend $\mathbf{A}$ to a square block operator

$$\hat{\mathbf{A}} := \begin{bmatrix} \mathbf{A} \\ \mathbf{0} \end{bmatrix}.$$

This extended operator is closed (resp. closable) if and only if $\mathbf{A}$ is closed (resp. closable). However, $\hat{\mathbf{A}}$ can only be diagonally dominant when the entries in the columns of index $j > n$ are all bounded operators (and hence are relatively bounded by 0).

Yet another way to force a finite block operator to be closed is to essentially combine the hypotheses of Proposition 3.3.10 and Proposition 3.3.14.

Definition 3.3.18. Let $\mathbf{A} = [A_{ij}]$ be an $n \times m$ finite block operator. The i_0^{th} row of $\mathbf{A}$ is **row dominant** of order r if for all i, j, the operator A_{ij} is $A_{i_0 j}$-bounded with relative bound $\delta_{ij} < r$. That is, each entry of each column is relatively bounded (with relative bound $< r$) by the entry in row i_0.

Proposition 3.3.19. *Let $\mathbf{A} = [A_{ij}]$ be an $n \times m$ block operator. Let i_0 be a row index, and let $R_j := \overline{\mathcal{R}(A_{i_0 j})}$ denote the closure of the range of $A_{i_0 j}$. The finite block operator is closed (resp. closable) if the following conditions hold.*

(i) Row i_0 is row dominant of order $1/m$.

(ii) Each $A_{i_0 j}$ is a closed (resp. closable) operator.

(iii) The closed ranges $\{R_j\}_j : 1 \leqslant j \leqslant m\}$ are pairwise orthogonal in Y_{i_0}.

Proof. Without loss of generality, suppose i_0 is the top row of $\mathbf{A}$. Write $\mathbf{A}$ as $\mathbf{A} = \mathbf{A}_0 + \mathbf{B}$ where $\mathbf{A}_0$ is the top row of $\mathbf{A}$ with all other entries zeroed out, and $\mathbf{B}$ is $\mathbf{A}$ with the top row zeroed out. By Proposition 3.3.10, $\mathbf{A}_0$ is a closed operator.

Again, pick $a, b \geqslant 0$ with $b < 1/m$ such that for each A_{ij} with $i > 1$,

$$\|A_{ij} x_j\|^2 \leqslant a\|x_j\|^2 + b\|A_{1j} x_j\|^2.$$

We now compute:

$$
\begin{aligned}
\|\mathbf{B}\mathbf{x}\|^2 &= \sum_{i=2}^{n} \left\| \sum_{j=1}^{m} A_{ij} x_j \right\|^2 \\
&\leqslant \sum_{i=2}^{n} \sum_{j=1}^{m} m \|A_{ij} x_j\|^2 \\
&\leqslant \sum_{i=2}^{n} \sum_{j=1}^{m} m(a\|x_j\|^2 + b\|A_{1j} x_j\|^2) \\
&= ma\|\mathbf{x}\|^2 + mb\|\mathbf{A}_0 \mathbf{x}\|^2.
\end{aligned}
$$

Since $mb < 1$ by hypothesis, the Kato-Rellich theorem again proves that $\mathbf{A} = \mathbf{A}_0 + \mathbf{B}$ is closed. □

To conclude this section, we summarize a few ways to enforce that a block operator be closed or closable.

- Ensure there is at most one unbounded operator per-row.
- Ensure the images of the operators in each row are orthogonal.
- Check for a dominant operator in each column, each living in a distinct row.
- Check for a dominant row containing operators with orthogonal images.

4
CELLULAR SHEAVES OF HILBERT SPACES

Having established the necessary categorical and operator-theoretic foundations in Chapter 3, we now turn to the central construction of this thesis: cellular sheaves valued in Hilbert spaces. This chapter develops a systematic theory of cellular sheaves valued in the category $\mathbf{Hilb}_{0,\Bbbk}$ of Hilbert spaces and unbounded operators. The passage from finite to infinite dimensions introduces fundamental complications that require careful treatment. When restriction maps are unbounded and partially defined, the composition of morphisms requires precise domain bookkeeping, the associated cochain complexes may fail to satisfy the hypotheses of a Hilbert complex, and even basic sheaf operations such as the sheaf hom become problematic.

We address these challenges through a two-stage approach. Section 4.1 introduces pre-Hilbert sheaves as the most general cellular sheaves valued in $\mathbf{Hilb}_{0,\Bbbk}$. Section 4.2 and Section 4.3 describe the sections and associated cochain complexes of pre-Hilbert sheaves. While pre-Hilbert sheaves generalize weighted cellular sheaves directly, they may exhibit pathological behavior: their coboundary operators need not be closable, their cohomology groups may fail to exist, and their spectral theory may be ill-defined. Section 4.4 identifies the additional hypotheses necessary to obtain well-behaved objects, leading to the definition of Hilbert sheaves proper. A pre-Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ qualifies as a Hilbert sheaf when its associated coboundary operators $\mathring{\delta}^k$ are closable, ensuring that the cochain complex forms a genuine Hilbert complex in the sense of Brüning and Lesch [17].

The remainder of the chapter develops the fundamental constructions and properties of Hilbert sheaves. We introduce two distinguished classes of Hilbert sheaves that merit special attention: bounded Hilbert sheaves, where all restriction maps are bounded operators, and closed Hilbert sheaves, where all coboundary operators have closed range. These classes exhibit particularly favorable properties; bounded sheaves admit normalization under suitable conditions, while closed sheaves possess honest (rather than reduced) cohomology groups. Throughout, we illustrate the theory with concrete examples, including sheaves arising from differential operators on manifolds and block operator constructions.

4.1 PRE-HILBERT SHEAVES

We now extend the scope of weighted cellular sheaves from finite dimensional Hilbert spaces to arbitrary Hilbert spaces. The most direct generalization yields what we call pre-Hilbert sheaves—functors on GACs valued in the the category $\mathbf{Hilb}_{0,\Bbbk}$ of Hilbert spaces and unbounded operators. While these objects arise naturally and many fundamental constructions from weighted cellular sheaf theory carry over, the unboundedness introduces complications that will necessitate additional hypotheses in subsequent sections.

Definition 4.1.1. A **cellular pre-Hilbert sheaf** (or simply a **pre-Hilbert sheaf**) is a cellular sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ (as per Definition 1.4.1) with grade-wise finite domain $\mathcal{P}$, valued in the category of Hilbert spaces and operators. In particular, it consists of the following data.

- A grade-wise finite graded acyclic category $\mathcal{P}$ that admits a signed incidence structure ϵ.
- For each object $\sigma \in \mathcal{P}$, a Hilbert space $\mathcal{F}(\sigma)$ called the **stalk** over σ.
- For each indecomposable morphism $f : \sigma \to \tau$ in $\mathcal{P}$, a $\mathbf{Hilb}_{0,\Bbbk}$-morphism $\mathcal{F}_f : \mathcal{F}(\sigma) \to \mathcal{F}(\tau)$ called the **restriction map** over f.
- All other morphisms in the image of $\mathcal{F}$ are determined by composition.

Remark 4.1.2. Pre-Hilbert sheaves directly generalize weighted cellular sheaves, which are valued in the category $\mathbf{FinHilb}_{\Bbbk}$ of finite dimensional Hilbert spaces [54, Section 3.1]. In $\mathbf{FinHilb}_{\Bbbk}$, all operators are bounded and globally defined.

Remark 4.1.3. The "pre" prefix is not meant to invoke pre-Hilbert spaces; each stalk of a pre-Hilbert sheaf is an honest Hilbert space. We will soon see that extra hypotheses are needed to ensure that a pre-Hilbert sheaf behaves in a manner similar to a weighted cellular sheaf. It is only under those additional hypotheses that we will call a cellular sheaf a "Hilbert sheaf". In short, "pre-Hilbert sheaf" should be parsed as "pre-(Hilbert sheaf)", not as "(pre-Hilbert) sheaf."

Several essential concepts and constructions on weighted cellular sheaves can be straightforwardly adapted to pre-Hilbert sheaves.

4.2 SECTIONS

Definition 4.2.1. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a pre-Hilbert sheaf, and let $\mathcal{Z} \subseteq \mathcal{P}$ be a subcategory of $\mathcal{P}$. The **space of sections** over $\mathcal{Z}$ is the restriction limit

$$\Gamma(\mathcal{Z}; \mathcal{F}) := \operatorname{reslim} \mathcal{F}\big|_{\mathcal{Z}}$$

where $\mathcal{F}|_{\mathcal{Z}}$ is the restriction of $\mathcal{F}$ to the sub-category $\mathcal{Z}$. When $\mathcal{Z} = \mathcal{P}$, we recover the **space of global sections** $\Gamma(\mathcal{F}) := \Gamma(\mathcal{P}; \mathcal{F})$.

Since $\mathbf{Hilb}_{0,\Bbbk}$ admits all finite restriction limits, spaces of sections are always well-defined. Moreover, since finite dimensional subspaces of Hilbert spaces are always closed, when $\mathcal{F}$ is valued in the subcategory $\mathbf{FinHilb}_{\Bbbk} \subseteq \mathbf{Hilb}_{0,\Bbbk}$, these restriction-limit sections exactly recover the usual spaces of sections of a weighted cellular sheaf.

Like weighted cellular sheaves, the spaces of sections of a pre-Hilbert sheaf admit a nice characterization through the construction in Remark 3.1.23. Given a sub-category $\mathcal{Z} \subseteq \mathcal{P}$, an $\mathcal{F}|_{\mathcal{Z}}$-admissible point $\mathbf{x}$ is a choice $x_\sigma \in F(\sigma)$ for each object $\sigma \in \mathcal{Z}$ such that whenever $f : \sigma \to \tau$ is a $\mathcal{Z}$-morphism, then $\mathcal{F}_f(x_\sigma) = x_\tau$. An $\mathcal{F}|_{\mathcal{Z}}$-admissible point can thus be interpreted as a locally consistent choice of data in each stalk. The space of sections $\Gamma(\mathcal{Z}; \mathcal{F})$ is the closure of space of $\mathcal{F}|_{\mathcal{Z}}$-admissible points.

Remark 4.2.2. When $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{\Bbbk} \subseteq \mathbf{Hilb}_{0,\Bbbk}$ has all bounded and globally-defined restriction maps, all restriction limits agree with the usual limits in $\mathbf{Hilb}_{\Bbbk}$.

The space of global sections of a pre-Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is the restriction limit $\operatorname{reslim} \mathcal{F}$. That is, $\Gamma(\mathcal{F})$ is the closure of the space of $\mathcal{F}$-admissible points; choices $x_\sigma \in \mathcal{F}(\sigma)$ for each object $\sigma \in \mathcal{P}$ such that whenever $f : \sigma \to \tau$ is a $\mathcal{P}$-morphism, then $\mathcal{F}_f(x_\sigma) = x_\tau$. There may be global sections which merely are limits of sequences of $\mathcal{F}$-admissible points, but are not themselves $\mathcal{F}$-admissible.

As an object in $\mathbf{Hilb}_{0,\Bbbk}$, the space of global sections $\Gamma(\mathcal{F})$ comes equipped with a Hilbert space inner product. However, this inner product is non-canonical; the structure is merely defined up to isomorphisms, or equivalently, up to the size of a Hilbert space basis. In particular, the space of global sections is not defined uniquely up to unitary isomorphism.

As constructed via restriction limit, $\Gamma(\mathcal{F})$ inherits its inner product structure as a subspace

$$\Gamma(\mathcal{F}) \subseteq \bigoplus_{\sigma \in \mathcal{P}} \mathcal{F}(\sigma).$$

While this is a legitimate construction of $\Gamma(\mathcal{F})$ as a Hilbert space, there is a preferable construction. We present this construction for global sections, but a similar construction holds for all spaces of sections. Every $\mathcal{F}$-admissible point $\mathbf{x} \in \bigoplus_{\sigma \in \mathcal{P}} \mathcal{F}(\sigma)$ is uniquely determined by a choice of a point $x_\sigma \in \mathcal{F}(\sigma)$ for each object $\sigma \in \mathcal{P}$ of rank $r(\sigma) = 0$; all other values x_τ can be determined by applying restriction maps. We may pare off this superfluous information, and instead identify $\mathbf{x}$ as a point in $C^0(\mathcal{P}, \mathcal{F}) = \bigoplus_{r(\sigma)=0} \mathcal{F}(\sigma)$. The closure of this set of $\mathcal{F}$-admissible points (with different cone-legs) also satisfies the universal property of the restriction limit $\operatorname{reslim} \mathcal{F}$, and can be identified as the space of global sections. The difference between these two constructions of

$\Gamma(\mathcal{F})$ is entirely analogous to the construction of a pullback as a cone over a cospan vs. as a commutative square.

Going forward, we will always identify $\Gamma(\mathcal{F})$ with the second construction, and the Hilbert space structure it receives as a subspace of $C^0(\mathcal{P};\mathcal{F})$. This identification is advantageous for at least two reasons. First, as we will see, this aligns better with the Hodge theory of Hilbert complexes. Second, when necessary, it will allow us to safely consider pre-Hilbert sheaves whose underlying categories $\mathcal{P}$ have infinitely many objects. Provided $\mathcal{P}$ is locally finite and has at most finitely many objects of each rank, $C^0(\mathcal{P};\mathcal{F})$ will be a Hilbert space and all arguments will go through without extra work.

Remark 4.2.3. It has been observed [52, 54] that it would be preferable to use the technology of *dagger limits* [62]—which define limits in dagger categories up to unitary isomorphism—to define spaces of sections. This is, unfortunately, not possible in the dagger category $\mathbf{Hilb}_{\mathbb{k}}$. As shown by Heunen and Karvonen, $\mathbf{Hilb}_{\mathbb{k}}$ does not admit all dagger limits; it does not even admit all *dagger pullbacks* [62]. Moreover, By the construction of dagger limits, when a pre-Hilbert sheaf $\mathcal{F}$ valued in $\mathbf{Hilb}_{\mathbb{k}}$ does admit the construction of $\Gamma(\mathcal{F})$ as a dagger limit (such as when all restriction maps are dagger-monomorphism), this dagger limit does not agree with either proposed construction of $\Gamma(\mathcal{F})$. However, the inclusion of $\Gamma(\mathcal{F}) \hookrightarrow C^0(\mathcal{P};\mathcal{F})$ is a *dagger kernel* in the sense of Heunen and Jacobs [61].

4.3 THE ASSOCIATED COCHAIN COMPLEX

Given a pre-Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$, we may form an associated cochain complex of Hilbert spaces

$$\left(C^\bullet(\mathcal{P},\mathcal{F})\,,\,\mathring{\delta}^\bullet\right) := C^0(\mathcal{P};\mathcal{F}) \xrightarrow{\mathring{\delta}^0} C^1(\mathcal{P};\mathcal{F}) \xrightarrow{\mathring{\delta}^1} C^2(\mathcal{P};\mathcal{F}) \xrightarrow{\mathring{\delta}^3} \cdots$$

with k-**cochains** $C^k(\mathcal{P};\mathcal{F})$ and k-**coboundary maps** $\mathring{\delta}^k : C^k(\mathcal{P};\mathcal{F}) \to C^{k+1}(\mathcal{P};\mathcal{F})$ defined by:

$$\begin{aligned} C^k(\mathcal{P};\mathcal{F}) &:= \bigoplus_{r(\sigma)=k} F(\sigma)\,, \\ (\mathring{\delta}^k \mathbf{x})_\tau &:= \sum_{\substack{\sigma \triangleleft_1 \tau \\ f:\sigma\to\tau}} \epsilon(f)\mathcal{F}_f(\mathbf{x}_\sigma)\,. \end{aligned}$$

The operator $\mathring{\delta}^k$ is the block operator $[\mathring{\delta}^k_{\tau,\sigma}]$ where σ and τ range over objects of $\mathcal{P}$ of rank k and $k+1$ respectively. The block $\mathring{\delta}^k_{\tau,\sigma}$ is given by

$$\mathring{\delta}^k_{\tau,\sigma} := \sum_{f:\sigma\to\tau} \epsilon(f)\mathcal{F}_f\,,$$

and has domain $\mathrm{Dom}(\mathring{\delta}^k_{\tau,\sigma}) = \bigcap_{f:\sigma\to\tau} \mathrm{Dom}(\mathcal{F}_f)$, where the empty sum is understood to be the globally defined zero-operator. The domain of the block operator $\mathring{\delta}^k$ is therefore given by $\bigoplus_{r(\sigma)=k} \left(\bigcap_\tau \mathrm{Dom}(\mathring{\delta}^k_{\tau,\sigma})\right) \subseteq \bigoplus_{r(\sigma)=k} \mathcal{F}(\sigma)$.

This is a "cochain complex" in the sense that $\mathring{\delta}^{k+1}\mathring{\delta}^k$ is the zero-operator on its domain $\{\mathbf{x} \in \mathrm{Dom}(\mathring{\delta}^k) \, : \, \mathring{\delta}^k(\mathbf{x}) \in \mathrm{Dom}(\mathring{\delta}^{k+1})\}$. This can be checked by repeating the argument in the proof Proposition 1.4.8. Many important properties from weighted cellular sheaves are preserved.

Proposition 4.3.1. *The kernel* $\ker(\mathring{\delta}^0)$ *is exactly the set of* $\mathcal{F}$*-admissible points. Moreover, the closure of the kernel is the set of global sections* $\Gamma(\mathcal{F})$*.*

In general, the cochain complex $\left(C^\bullet(\mathcal{P},\mathcal{F}), \mathring{\delta}^\bullet\right)$ will not be a Hilbert complex. There is no guarantee that any of the following conditions hold.

1. $\mathring{\delta}^k$ is densely-defined for all k.
2. $\mathring{\delta}^k$ is closed for all k.
3. $\mathcal{R}(\mathring{\delta}^k) \subseteq \mathrm{Dom}(\mathring{\delta}^{k-1})$.

Unless we have these properties, the cochain complex associated to a pre-Hilbert sheaf will not be a Hilbert complex. Consequently, we will not have access to the (weak) Hodge decomposition and (reduced) cohomology. Moreover, since $\mathring{\delta}^k$ need not be a closed densely-defined operator, the adjoint $(\mathring{\delta}^k)^*$ may not be densely-defined (or even exist!) in general, preventing the existence of a Hodge Laplacian. We must restrict to a better behaved collection of cellular sheaves of Hilbert spaces to ensure these qualities. This leads to the following definition.

Definition 4.3.2. A **Hilbert sheaf** is a pre-Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ whose associated cochain complex $\left(C^\bullet(\mathcal{P},\mathcal{F}), \mathring{\delta}^\bullet\right)$ has the following properties.

(i) $\mathring{\delta}^k$ is densely-defined and closable for all k.

(ii) $\mathcal{R}(\mathring{\delta}^k) \subseteq \mathrm{Dom}(\mathring{\delta}^{k+1})$ for all k.

This definition is synthetic in the sense that a Hilbert sheaf is defined to be a pre-Hilbert sheaf with exactly the properties we desire; the required properties are not of the underlying functor $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$, but of the associated cochain complex. However, the definition gives no concrete method for determining if a pre-Hilbert sheaf is a Hilbert sheaf. This is partly out of necessity; it is very difficult to give clear and concise collections of necessary or sufficient conditions for a pre-Hilbert sheaf to be a Hilbert sheaf. Moreover, most pre-Hilbert sheaves (with "most" understood in an informal sense) are not Hilbert sheaves. As witnessed in Section 3.3, it is very difficult to ensure that a block operator is closable.

A Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ has two distinct associated cochain complexes; first there is the usual associated cochain complex as constructed in the previous subsection. As with a pre-Hilbert sheaf, this may-or-may-not be a Hilbert complex. Second, there is an **associated Hilbert complex** $\left(C^\bullet(\mathcal{P},\mathcal{F}), \delta^\bullet\right)$, where $\delta^k := \overline{\mathring{\delta}^k}$ is the closure of the coboundary operator $\mathring{\delta}^k$. When the associated cochain complex is itself a Hilbert complex, we say the Hilbert sheaf is **proper**.

The Hilbert complex $\left(C^\bullet(\mathcal{P},\mathcal{F}), \delta^\bullet\right)$ associated to a Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is the better analog of the usual cochain complex of a weighted cellular sheaf. In particular, since $\ker(\delta^0) = \overline{\ker(\mathring{\delta}^0)}$, the kernel of the coboundary operator δ^0 is exactly the space of global sections.

There are two special classes of Hilbert sheaves worth extra attention. First, a Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_\Bbbk \subseteq \mathbf{Hilb}_{0,\Bbbk}$ where every restriction map is globally defined and bounded will generally be better behaved than its unbounded cousins. We call such a Hilbert sheaf **bounded**. On the other hand, when each coboundary operator δ^k has closed range (hence yielding honest cohomology instead of merely reduced cohomology), we say that a Hilbert sheaf is **closed**.

4.4.1 *The category of Hilbert sheaves*

We may form a category $\mathbf{HilbShv}_\Bbbk(\mathcal{P})$ of Hilbert sheaves on a GAC $\mathcal{P}$ as a subcategory of the functor category $[\mathcal{P}, \mathbf{Hilb}_{0,\Bbbk}]$.

Definition 4.4.1. The category $\mathbf{HilbShv}_\Bbbk(\mathcal{P})$ consists of the following data.

- **Objects.** Each object of $\mathbf{HilbShv}_\Bbbk(\mathcal{P})$ is a Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$.
- **Morphisms.** A **Morphism of Hilbert sheaves** is a natural transformation $\phi : \mathcal{F} \Rightarrow \mathcal{G}$ between Hilbert sheaves such that each component ϕ_σ for $\sigma \in \mathcal{P}$ is a bounded globally-defined linear operator satisfying $\phi_\sigma(\mathrm{Dom}(\mathcal{F}_f)) \subseteq \mathrm{Dom}(\mathcal{G}_f)$ for each covering morphism f with domain σ in $\mathcal{P}$.
- **Composition.** Composition is the usual composition of natural transformations.

Note that the underlying operators in a morphism of Hilbert sheaves must be globally defined and bounded. This condition ensures that Hilbert sheaf morphisms induce a morphism of Hilbert complexes. Suppose $\phi : \mathcal{F} \to \mathcal{G}$ is a morphism in $\mathbf{HilbShv}_{\mathbb{k}}(\mathcal{P})$. There is an induced map of Hilbert complexes

$$\begin{array}{ccccccc}
\cdots \longrightarrow & C^k(\mathcal{P};\mathcal{F}) & \xrightarrow{\delta^k(\mathcal{P};\mathcal{F})} & C^{k+1}(\mathcal{P};\mathcal{F}) & \longrightarrow \cdots \\
& \phi^k \big\downarrow & & \phi^{k+1} \big\downarrow & \\
\cdots \longrightarrow & C^k(\mathcal{P};\mathcal{G}) & \xrightarrow{\delta^k(\mathcal{P};\mathcal{G})} & C^{k+1}(\mathcal{P};\mathcal{G}) & \longrightarrow \cdots
\end{array}$$

where the map ϕ^k is the direct sum

$$\phi^k = \bigoplus_{r(\sigma)=k} \phi_\sigma.$$

Since $\phi_\tau \mathcal{F}_f = \mathcal{G}_f \phi_\sigma$ for each covering morphism $f : \sigma \to \tau$ in $\mathcal{P}$, the chain map $\phi^\bullet$ is a morphism of Hilbert complexes.

Notation 4.4.2. We let $\mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{P})$ denote the full sub-category of Hilbert sheaves where all restriction maps are bounded and globally defined. We call such Hilbert sheaves **bounded** Hilbert sheaf.

4.4.2 *Finding Hilbert sheaves*

We now turn to the problem of when a Pre-Hilbert sheaf is a Hilbert sheaf. We start with a few useful examples.

Proposition 4.4.3. *Every pre-Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{\mathbb{k}} \subseteq \mathbf{Hilb}_{0,\mathbb{k}}$ with bounded restriction maps is a proper Hilbert sheaf.*

Proof. When each restriction map $\mathcal{F}_f$ is bounded and globally defined, each coboundary map δ^k is also bounded and globally defined. The associated cochain complex is therefore a Hilbert complex. □

Remark 4.4.4. It follows from this proposition that $\mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{P})$ is exactly the functor category $[\mathcal{P}, \mathbf{Hilb}_{\mathbb{k}}]$.

Example 4.4.5. Consider a pre-Hilbert sheaf $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{0,\mathbb{k}}$ where $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a finite undirected multi-graph without self loops, viewed as a posetal category. Suppose there is a closed densely-defined operator $A : X \to Y$ such that $\mathcal{F}_f = A$ for every restriction map. Each edge e has exactly two bounding vertices u, v, and the coboundary map $\mathring{\delta} : C^0(\mathcal{G};\mathcal{F}) \to C^1(\mathcal{G};\mathcal{F})$ maps into $\mathcal{F}(e)$ by

$$(\mathring{\delta}\mathbf{x})_e = A \oplus (-A)$$

up to a choice of orientation.

In general, this coboundary operation is not closed (see Example 3.3.4). However, it is closable, as witnessed by a straightforward computation. Hence this is an *improper* Hilbert sheaf.

As a concrete example, consider the one-edge graph, and the following Hilbert sheaf

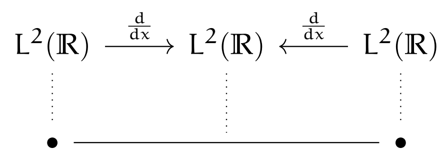

where the domain of the derivative operator is $\{f \in L^2(\mathbb{R}) : f' \in L^2(\mathbb{R})\}$. This derivative operator is easily seen to be closed. The associated cochain complex is given by

$$L^2(\mathbb{R}) \oplus L^2(\mathbb{R}) \xrightarrow{\begin{bmatrix} \frac{d}{dx} & -\frac{d}{dx} \end{bmatrix}} L^2(\mathbb{R}) \quad ,$$

with coboundary map $\mathring{\delta} = \begin{bmatrix} \frac{d}{dx} & (-\frac{d}{dx}) \end{bmatrix}$. This is not a closed operator. The closure, δ, has domain $\mathrm{Dom}(\delta) = \{(f,g)^\mathsf{T} \in L^2(\mathbb{R}) \oplus L^2(\mathbb{R}) : \frac{d}{dx}(f-g) \in L^2\}$, and acts by *weak differentiation* of the difference, $\delta(f,g)^\mathsf{T} = \frac{d}{dx}(f-g)$, yielding the associated Hilbert complex

$$L^2(\mathbb{R}) \oplus L^2(\mathbb{R}) \xrightarrow{\quad \delta \quad} L^2(\mathbb{R}) \ .$$

Remark 4.4.6. As this example shows, the closure of the block operator $\mathring{\delta}^k$ need not be a block operator itself with respect to the same decomposition.

Remark 4.4.7. More generally, any Hilbert sheaf where every restriction map at a given rank is the same closable operator will be a (usually improper) Hilbert sheaf.

It is, in general, difficult to write theorems for checking if a pre-Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is a Hilbert sheaf. To ensure that the coboundary operator $\mathring{\delta}^k$ is densely-defined and closable is often a matter of checking properties of the block operator directly through ad hoc methods. However, there are a few handy tricks for building a Hilbert sheaf from the ground up.

Suppose we have a GAC $\mathcal{P}$ on which we would like to define a Hilbert sheaf. We start by building a functor $\mathcal{F}_0 : \mathcal{P} \to \mathbf{DenseCoreHilb}_{0,\Bbbk}$ valued in cored Hilbert spaces. Further, assume that stalks $\mathcal{F}_0(\sigma) = (X_\sigma, V_\sigma)$ have already been assigned. For each covering map $f : \sigma \to \tau$, take $\mathcal{F}_0(f) : (X_\sigma, V_\sigma) \to (X_\tau, V_\tau)$ to be an operator such that $\mathcal{R}(\mathcal{F}_0(f)) \subseteq V_\tau$. After pushing through the forgetful functor $\mathcal{U} : \mathbf{DenseCoreHilb}_{0,\Bbbk} \to \mathbf{Hilb}_{0,\Bbbk}$, we arrive at a pre-Hilbert sheaf $\mathcal{F} = \mathcal{U}\mathcal{F}_0$ whose coboundary operators $\mathring{\delta}^k$ are densely defined and satisfy $\mathcal{R}(\mathring{\delta}^k) \subseteq \mathrm{Dom}(\mathring{\delta}^{k+1})$. To further ensure that each $\mathring{\delta}^k$ is closable, the results in Section 3.3 are helpful.

We end this section with a large class of Hilbert sheaves on a network $\mathcal{G}$. We start by briefly recalling a few definitions from differential geometry. We assume that all manifolds are real and second countable, and that all vector bundles have finite rank.

Notation 4.4.8. Let M be a smooth manifold, and $\pi : E \to M$ a smooth vector bundle over M. We denote the fiber over $p \in M$ by $E_p := \pi^{-1}(p)$. We further denote the space of sections of E by $\Gamma(E)$, and the space of smooth sections by $\Gamma^\infty(E)$.

Notation 4.4.9. We use the following standard notation for multi-indices. Given an n-tuple of non-negative integers $\alpha = (\alpha_1, \ldots, \alpha_n)$, let $|\alpha| = \sum_j \alpha_j$. For an $\mathbf{x} \in \mathbb{R}^n$, let $\mathbf{x}^\alpha := \prod_{j=1}^n x_j^{\alpha_j}$. Finally, let $\frac{\partial^{|\alpha|}}{\partial \mathbf{x}^\alpha} = \frac{\partial^{|\alpha|}}{\partial x_1^{\alpha_1} \cdots \partial x_n^{\alpha_n}}$.

Definition 4.4.10. Let M be a real smooth manifold, and $E \to M$ as smooth vector bundle over M. A **metric** on E is an assignment of an inner product $h_p : E_p \times E_p \to \mathbb{R}$ to each fiber E_p of E. We further require that for any pair of smooth sections $s_1, s_2 \in \Gamma^\infty(E)$, that the map $\phi(p) : M \to \mathbb{R}$ defined by $\phi(p) := h_p(s_1(p), s_2(p))$ is smooth.

Definition 4.4.11. Let M be a smooth manifold, $E \to M$ a smooth vector bundle equipped with a metric h, and μ a volume form on M. A section $s \in \Gamma(E)$ is an L^2-**section** if

$$\int_M h_p(s(p), s(p))\, d\mu(p) < \infty.$$

We denote the space of L^2-sections by $L^2(E)$.

Remark 4.4.12. The space $L^2(E)$ is a real Hilbert space with respect to pointwise addition, pointwise scaling, and the inner product

$$\langle s_1, s_2 \rangle_{L^2} = \int_M h_p(s_1(p), s_2(p))\, d\mu(p).$$

Moreover, the space $\Gamma_c^\infty(E)$ of compactly supported smooth sections is dense in $L^2(E)$. This standard result may be proven using a mollifier argument and a partition of unity. See [79] for details.

Definition 4.4.13. Let M be a smooth manifold, and E, F a pair of smooth vector bundles over M, with a shared collection of trivializing neighborhoods $\{U_j\}_{j \in J}$. Let $\Gamma^\infty(E)$ and $\Gamma^\infty(F)$ denote the spaces of smooth sections of E and F. A linear map $D : \Gamma^\infty(E) \to \Gamma^\infty(F)$ is a **differential operator** if there is an integer $m \geqslant$ such that in each neighborhood U_j with local coordinates $\mathbf{x}$ we may write

$$D\big|_{U_j}(\mathbf{x}) = \sum_{|\alpha| \leqslant m} A_\alpha^{(i)}(\mathbf{x}) \frac{\partial^{|\alpha|}}{\partial \mathbf{x}^\alpha}$$

where $A^{(i)}$ is a $q \times p$ matrix of smooth functions.

Remark 4.4.14. Let M be a smooth manifold with a volume form μ, and E, F a pair of smooth vector bundles over M, both equipped with metrics. A differential operator $D : \Gamma^\infty(E) \to \Gamma^\infty(F)$ defines an unbounded operator $D : L^2(E) \to L^2(F)$ with domain $\mathrm{Dom}(D) := \Gamma_c^\infty(E)$.

Lemma 4.4.15. *Let $D : L^2(E) \to L^2(F)$ be a differential operator with domain $\Gamma_c^\infty(E)$. D is closable.*

Proof. This is a standard fact about differential operators. The essence of the argument is that any such $D : \Gamma^\infty(E) \to \Gamma^\infty(F)$, due to its smooth coefficients, admits a **formal adjoint** $D^\dagger : \Gamma^\infty(F) \to \Gamma^\infty(E)$, such that the integration by parts formula

$$\int_M h_p(D^\dagger t(p), s(p))\, d\mu = \int_M h_p(t(p), Ds(p))\, d\mu$$

holds for every pair of sections $s \in \Gamma^\infty(E)$ and $t \in \Gamma^\infty(F)$ such that $\mathrm{supp}(s) \cap \mathrm{supp}(t)$ is compact. For a full proof, see [108, Lemma 12.8]. From this formal adjoint, we see that $D : L^2(E) \to L^2(F)$ admits a densely defined adjoint D^* as an unbounded Hilbert space operator, from which we may conclude that D is closable. □

Lemma 4.4.16. *Let $D = [D_{ij}] : \bigoplus_{j=1}^m \Gamma^\infty(E^{(j)}) \to \bigoplus_{i=1}^n \Gamma^\infty(F^{(i)})$ be a finite block operator with each component D_{ij} a differential operator. D is a differential operator.*

Proof. Let $\mathbf{s} = (s_1, \dots, s_m)^T \in \bigoplus_{j=1}^m \Gamma^\infty(E^{(j)})$ be a tuple of smooth sections. $\mathbf{s}$ is a smooth section of the direct sum bundle $E^{(1)} \oplus \cdots \oplus E^{(m)} \to M$. In local coordinates, each term D_{ij} can be written as a smooth-function weighted sum of derivatives. Consequently, $D\mathbf{s} \in \bigoplus_{i=1}^n \Gamma^\infty(F^{(i)}$ itself is a smooth-function weighted sum of derivatives, and hence a differential operator. □

With these lemmas in hand, we can define a broad class of Hilbert sheaves.

Theorem 4.4.17. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite network, and $\mathcal{F} : \mathcal{G} \to$ $\mathbf{Hilb}_{0,\mathbb{R}}$ be a pre-Hilbert sheaf consisting of the following data.*

- *For each $\sigma \in \mathcal{V} \amalg \mathcal{E}$, the stalk $\mathcal{F}(\sigma) := L^2(E^\sigma)$, where E^σ is a smooth vector bundle, equipped with a metric, over a smooth manifold M^σ with volume form μ^σ.*
- *For each covering morphism $f : v \to e$, the restriction map $\mathcal{F}_f : L^2(E^v) \to L^2(E^e)$ is a differential operator.*

$\mathcal{F}$ is a Hilbert sheaf.

Proof. The coboundary operator $\mathring{\delta} : C^0(\mathcal{G}; \mathcal{F}) \to C^1(\mathcal{G}; \mathcal{F})$ is a finite block operator with differential operators for blocks. Lemma 4.4.16 ensures that δ is itself a differential operator, which is closable by Lemma 4.4.15. Therefore $\mathcal{F}$ is a Hilbert sheaf. □

Remark 4.4.18. The closure δ corresponds to extending the domain of $\mathring{\delta}$ to sufficiently weakly-differentiable sections with square integrable derivatives, which corresponds to the Sobolev space $H^k(\bigoplus_{\nu \in \mathcal{V}} E^\nu) \subseteq L^2(\bigoplus_{\nu \in \mathcal{V}} E^\nu) = C^0(\mathcal{G}; \mathcal{F})$, where k is the order of δ as a differential operator.

4.5 SHEAF OPERATIONS

We now return to the sheaf operations briefly discussed in Section 1.4.4. While most sheaf operations, when applied to a pre-Hilbert sheaf, will yield a new pre-Hilbert sheaf, care must be taken with Hilbert sheaves. Not all sheaf operations respect the closability of the coboundary operator without additional assumptions. To that end, we state operations in terms of pre-Hilbert sheaves, and provide conditions under which following each operation with a proposition that states when the operation yields a Hilbert sheaf. These criteria should not be considered exhaustive.

Remark 4.5.1. Since bounded pre-Hilbert sheaves are always Hilbert sheaves, and all sheaf operations preserve boundedness of restriction maps, it follows that all sheaf operations respect bounded Hilbert sheaves.

4.5.1 *Direct sum*

Definition 4.5.2. Let $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be pre-Hilbert sheaves. The **direct sum** of $\mathcal{F}$ and $\mathcal{G}$ is the pre-Hilbert sheaf $\mathcal{F} \oplus \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ with stalks $(\mathcal{F} \oplus \mathcal{G})(\sigma) = \mathcal{F}(\sigma) \oplus \mathcal{G}(\sigma)$ and restriction maps given by $(\mathcal{F} \oplus \mathcal{G})_f = \mathcal{F}(f) \oplus \mathcal{G}(f)$.

Proposition 4.5.3. *If $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ are both Hilbert sheaves, the direct sum $\mathcal{F} \oplus \mathcal{G}$ is a Hilbert sheaf as well.*

Proof. The space of k-cochains of $\mathcal{F} \oplus \mathcal{G}$ decomposes as $C^k(\mathcal{P}; \mathcal{F} \oplus \mathcal{G}) = C^k(\mathcal{P}; \mathcal{F}) \oplus C^k(\mathcal{P}; \mathcal{G})$. Hence the coboundary operator $\mathring{\delta}^k_{\mathcal{F} \oplus \mathcal{G}}$ itself can be written as a block operator

$$\mathring{\delta}^k_{\mathcal{F} \oplus \mathcal{G}} = \begin{bmatrix} \mathring{\delta}^k_{\mathcal{F}} & 0 \\ 0 & \mathring{\delta}^k_{\mathcal{G}} \end{bmatrix}$$

with respect to this decomposition, with domain $\mathrm{Dom}(\mathring{\delta}^k_{\mathcal{F}}) \oplus \mathrm{Dom}(\mathring{\delta}^k_{\mathcal{G}})$. It easily follows that $\mathcal{F} \oplus \mathcal{G}$ is a Hilbert sheaf. □

4.5.2 *Tensor product*

Definition 4.5.4. The **tensor product** of two pre-Hilbert sheaves $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is the pre-Hilbert sheaf $\mathcal{F} \otimes \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ with stalks $(\mathcal{F} \otimes \mathcal{G})(\sigma) = \mathcal{F}(\sigma) \otimes \mathcal{G}(\sigma)$ and restriction maps given by $(\mathcal{F} \otimes \mathcal{G})_f = \mathcal{F}(f) \otimes \mathcal{G}(f)$.

Proposition 4.5.5. *Let $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be Hilbert sheaves whose coboundary operators $\mathring{\delta}^k_{\mathcal{F}}$ and $\mathring{\delta}^k_{\mathcal{G}}$ have closable rows. The tensor product $\mathcal{F} \otimes \mathcal{G}$ is also a Hilbert sheaf.*

Proof. Let $\tau \in \mathcal{P}$ be a $(k+1)$-cell, and let R_τ denote the row of the k-coboundary operator $\mathring{\delta}^k$ of $\mathcal{F} \otimes \mathcal{G}$ that maps onto $(\mathcal{F} \otimes \mathcal{G})(\tau)$ To prove that R_τ is closable, it suffices to show that the adjoint map R^*_τ has dense domain. There are inclusions

$$\bigcap_{\sigma \lhd_1 \tau} \bigcap_{f:\sigma \to \tau} \mathrm{Dom}\left((\mathcal{F}_f \otimes \mathcal{G}_f)^*\right) \subseteq \mathrm{Dom}(R^*_\tau) \subseteq (\mathcal{F} \otimes \mathcal{G})(\tau).$$

Since $\mathcal{F}$ and $\mathcal{G}$ each row of the coboundary operators $\mathring{\delta}^k_{\mathcal{F}}$ and $\mathring{\delta}^k_{\mathcal{G}}$ are closable, there are dense linear subspaces:

$$X_{\mathcal{F}} := \bigcap_{\sigma \lhd_1 \tau} \bigcap_{f:\sigma \to \tau} \mathrm{Dom}(\mathcal{F}^*_f) \subseteq \mathcal{F}(\tau)\,,$$
$$X_{\mathcal{G}} := \bigcap_{\sigma \lhd_1 \tau} \bigcap_{f:\sigma \to \tau} \mathrm{Dom}(\mathcal{G}^*_f) \subseteq \mathcal{G}(\tau)\,.$$

There is a sequence of dense subspaces

$$X_{\mathcal{F}} \otimes_{\mathrm{alg}} X_{\mathcal{G}} \subseteq \mathcal{F}(\tau) \otimes_{\mathrm{alg}} \mathcal{G}(\tau) \subseteq \mathcal{F}(\tau) \otimes \mathcal{G}(\tau) = (\mathcal{F} \otimes \mathcal{G})(\tau)$$

where $\otimes_{\mathrm{alg}}$ denotes the algebraic tensor product. Since $\mathrm{Dom}(A^* \otimes B^*) \subseteq \mathrm{Dom}\left((A \otimes B)^*\right)$ for any pair of Hilbert space operators, $X_{\mathcal{F}} \otimes_{\mathrm{alg}} X_{\mathcal{G}} \subseteq \mathrm{Dom}(R^*_\tau)$. Therefore R^*_τ is densely defined, making R_τ closable. By Lemma 3.3.6, $\mathring{\delta}^k$ is closable, and $\mathcal{F} \otimes \mathcal{G}$ is a Hilbert sheaf. □

4.5.3 *Sheaf hom*

Unlike for weighted cellular sheaves, the sheaf hom cannot always be defined for pre-Hilbert sheaves. The essential problem is that for infinite dimensional Hilbert spaces X and Y, the space of bounded linear operators $\mathbf{Hilb}_{\Bbbk}(X, Y)$ is a Banach space, but not a Hilbert space. This obstruction to the sheaf hom construction also applies to bounded Hilbert sheaves. However, by restricting to Hilbert sheaf morphisms to those whose components are Hilbert Schmidt operators.

Notation 4.5.6. Let $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be pre-Hilbert sheaves. For $\sigma \in \mathcal{P}$, let $HS_\sigma(\mathcal{F}, \mathcal{G})$ denote the collection of natural transformations $\phi : \mathcal{F}\big|_{st(\sigma)} \Rightarrow \mathcal{G}\big|_{st(\sigma)}$ that satisfy the following conditions.

(i) Each component $\phi_\tau : \mathcal{F}(\tau) \to \mathcal{G}(\tau)$ is a bounded, globally defined Hilbert Schmidt operator.

(ii) For every morphism f of $st(\sigma)$ with source τ, there is an inclusion $\phi_\tau(\mathrm{Dom}(\mathcal{F}_f)) \subseteq \mathrm{Dom}(\mathcal{G}_f)$.

Lemma 4.5.7. *Let $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be pre-Hilbert sheaves. If the restriction maps of $\mathcal{G}$ are closed, the space $HS_\sigma(\mathcal{F}, \mathcal{G})$ is a Hilbert space for each $\sigma \in \mathcal{P}$.*

Proof. We may identify $HS_\sigma(\mathcal{F}, \mathcal{G})$ as a linear subspace of the Hilbert space[1]

$$\mathcal{H} := \bigoplus_{\tau \in st(\sigma)} HS(\mathcal{F}(\sigma), \mathcal{G}(\sigma)),$$

where $HS(\mathcal{F}(\sigma), \mathcal{G}(\sigma))$ is the Hilbert space of Hilbert Schmidt operators between $\mathcal{F}(\sigma)$ and $\mathcal{G}(\sigma)$. We now must verify that $HS_\sigma(\mathcal{F}, \mathcal{G})$ is closed as a subspace of $\mathcal{H}$.

Let ϕ^n denote a sequence of natural transformations in $HS_\sigma(\mathcal{F}, \mathcal{G})$ such that $\phi^n \to \phi \in \mathcal{H}$. Each component ϕ^n_τ converges to ϕ_τ in the Hilbert-Schmidt norm, and hence in the operator norm as well.

Let $f : \tau \to \gamma$ be a morphism of $\mathcal{P}$. Let $x \in \mathrm{Dom}(\mathcal{F}_f)$. For each $n \in \mathbb{N}$, we have that $\mathcal{G}_f \phi^n_\tau x = \phi^n_\gamma \mathcal{F}_f x$. By the observed operator norm convergence, we have strong convergence, so $\phi^n_\tau x \to \phi_\tau x$ and $\phi^n_\gamma \mathcal{F}_f x \to \phi_\gamma \mathcal{F}_f x$. Since $\mathcal{G}_f$ is closed, $G_f \phi^n_\tau x \to G_f \phi_\tau x = \phi_\gamma \mathcal{F}_f x$, proving that $\phi \in HS_\sigma(\mathcal{F}, \mathcal{G})$. □

Remark 4.5.8. Since every natural transformation in $HS_\sigma(\mathcal{F}, \mathcal{G})$ contains the data of a natural transformation $HS_\tau(\mathcal{F}, \mathcal{G})$ for every $\tau \geqslant \sigma$, we may identify $HS_\tau(\mathcal{F}, \mathcal{G}) \subseteq HS_\sigma(\mathcal{F}, \mathcal{G})$ as a sub-Hilbert space.

With this lemma, we may define a Hilbert-Schmidt analog to the sheaf hom.

Definition 4.5.9. Let $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be pre-Hilbert sheaves, with $\mathcal{G}$ having closed restriction maps. The **Hilbert Schmidt sheaf hom** from $\mathcal{F}$ to $\mathcal{G}$ is the pre-Hilbert sheaf $HS(\mathcal{F}, \mathcal{G})$ with stalks $HS(\mathcal{F}, \mathcal{G})(\sigma) = HS_\sigma(\mathcal{F}, \mathcal{G})$, and all restriction maps given by orthogonal projection.

Since every restriction map in $HS(\mathcal{F}, \mathcal{G})$ is an orthogonal projection, and hence bounded, we immediately get the following result by Proposition 4.4.3.

Proposition 4.5.10. *Let $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be pre-Hilbert sheaves, with $\mathcal{G}$ having closed restriction maps. The Hilbert Schmidt sheaf hom $HS(\mathcal{F}, \mathcal{G})$ is a Hilbert sheaf.*

1 Note that when $st(\sigma)$ contains infinitely many object, the maps in $\mathcal{H}$ must have square-summable Hilbert Schmidt norms.

Remark 4.5.11. Unsurprisingly, there is no longer a tensor $\dashv$ hom adjunction for Hilbert sheaves. This is an immediate consequence of the lack of a tensor $\dashv$ hom adjunction for Hilbert spaces in general.

4.5.4 *Kernel*

Definition 4.5.12. Let $\eta : \mathcal{F} \to \mathcal{G}$ be a morphism of Hilbert sheaves. The **kernel** of η is the Pre-Hilbert sheaf $\ker(\eta) : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ with stalks $\ker(\eta)(\sigma) = \ker(\eta_\sigma) \subseteq \mathcal{F}(\sigma)$, and restriction maps $\ker(\eta)_f = \mathcal{F}_f|_{\ker(\eta_\sigma)}$ for $f : \sigma \to \tau$.

Proposition 4.5.13. *The kernel of a Hilbert sheaf morphism is a Hilbert sheaf.*

Proof. Let $\mathring{\delta}^k$ denote the k-coboundary operator of $\mathcal{F}$. The k-coboundary operator of $\ker(\eta)$ is the restriction of $\mathring{\delta}^k$ to a sub-Hilbert space $C^k(\mathcal{P}; \ker(\eta)) \subseteq C^k(\mathcal{P}; \mathcal{F})$. Since $\mathring{\delta}^k$ is closable, this restriction is closable as well. □

4.5.5 *Pullback*

Definition 4.5.14. Let $\phi : \mathcal{P} \to \mathcal{Q}$ be a cellular map. The **pullback** of a pre-Hilbert sheaf $\mathcal{F} : \mathcal{Q} \to \mathbf{Hilb}_{0,\Bbbk}$ by ϕ is the pre-Hilbert sheaf $\phi^*\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ with stalks $\phi^*\mathcal{F}(\sigma) = \mathcal{F}(\phi(\sigma))$ and restriction maps $\phi^*\mathcal{F}_f = \mathcal{F}_{\phi(f)}$.

Proposition 4.5.15. *Let $\mathcal{P}, \mathcal{Q}$ be finite GACs which admit signed incidence structures, $\phi : \mathcal{P} \to \mathcal{Q}$ a covering map, and $\mathcal{F} : \mathcal{Q} \to \mathbf{Hilb}_{0,\Bbbk}$ a Hilbert sheaf. The pullback $\phi^*\mathcal{F}$ is a Hilbert sheaf.*

Proof. It is straightforward to verify that there is an $n \in \mathbb{N}$ such that for each object $\sigma \in \mathbb{Q}$, $|\phi^{-1}(\sigma)| = n$. Therefore we may write

$$C^k(\mathcal{P}; \phi^*\mathcal{F}) = \bigoplus^n C^k(\mathcal{P}; \mathcal{F}),$$

where $\bigoplus^n$ denotes the n-fold direct sum. With respect to this direct sum decomposition, we may write $\mathring{\delta}^k_{\phi^*\mathcal{F}} = \bigoplus^n \mathring{\delta}^k_{\phi^*\mathcal{F}}$ since ϕ is an isomorphism star-wise. It follows that $\mathring{\delta}^k_{\phi^*\mathcal{F}}$ is closable. □

Remark 4.5.16. In general, the pullback of a Hilbert sheaf need not be a Hilbert sheaf. The essential problem is that an arbitrary pullback may duplicate blocks in the coboundary matrix in a manner that breaks closability. Without additional structure to the cellular map ϕ to control the preimages of cells, closability cannot be guaranteed.

4.5.6 *Pushforward*

Defining the pushforward of a pre-Hilbert sheaf requires the use of restriction limits. Let $\phi : \mathcal{P} \to \mathcal{Q}$ be a cellular map. For each object $\sigma \in \mathcal{Q}$, let $\mathcal{P}_{\phi,\sigma}$ denote the full subcategory of $\mathcal{P}$ generated by the set of objects $\{\sigma' \in \mathcal{P} : \phi(\sigma') \geqslant \sigma\}$ in the underlying poset of $\mathcal{Q}$. By domain restriction, each such category $\mathcal{P}_{\phi,\sigma}$ defines a functor $\mathcal{F}|_{\mathcal{P}_{\phi,\sigma}} : \mathcal{P}_{\phi,\sigma} \to \mathbf{Hilb}_{0,\Bbbk}$. This functor has a restriction limit

$$\phi_*\mathcal{F}(\sigma) := \operatorname{reslim} \mathcal{F}|_{\mathcal{P}_{\phi,\sigma}}.$$

When $\sigma \leqslant \tau$ in the poset underling $\mathcal{Q}$, the restriction-limiting cone over $\mathcal{F}|_{\mathcal{P}_{\phi,\sigma}}$ with apex $\phi_*\mathcal{F}(\sigma)$ defines a cone over $\mathcal{F}|_{\mathcal{P}_{\phi,\tau}}$. Since both of these cones have total legs, there is a unique operator

$$\phi_*\mathcal{F}_f : \phi_*\mathcal{F}(\sigma) \to \phi_*\mathcal{F}(\tau)$$

such that $q_\gamma = p_\gamma \phi_*\mathcal{F}_f$, where p_γ and q_γ are the legs of the cones with apex $\phi_*\mathcal{F}(\tau)$ and $\phi_*\mathcal{F}(\sigma)$ respectively.

Definition 4.5.17. Let $\phi : \mathcal{P} \to \mathcal{Q}$ be a cellular map. The **pushforward** of a pre-Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ by ϕ is the pre-Hilbert sheaf $\phi_*\mathcal{F} : \mathcal{Q} \to \mathbf{Hilb}_{0,\Bbbk}$ with stalks $\phi_*\mathcal{F}(\sigma)$ and restriction maps $\phi^*\mathcal{F}_f$.

Remark 4.5.18. It is, in general, very difficult to ensure that the pushforward of a Hilbert sheaf $\mathcal{F}$ is a Hilbert sheaf outside of the bounded case. The fundamental challenge is that when the restriction limits used to define the stalks of $\phi_*\mathcal{F}$ are not honest limits, there is no clear relationship between the coboundary operators $\mathring{\delta}^k_{\mathcal{F}}$ and $\mathring{\delta}^k_{\phi_*\mathcal{F}}$.

Proposition 4.5.19. *Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{\Bbbk}$ be a bounded Hilbert sheaf, and $\phi : \mathcal{P} \to \mathcal{Q}$ a cellular map. The pushforward $\phi_*\mathcal{F}$ is a Hilbert sheaf.*

4.6 COHOMOLOGY

Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf with associated Hilbert complex $\left(C^\bullet(\mathcal{P},\mathcal{F}), \delta^\bullet\right)$. As a Hilbert complex, we get a corresponding **reduced sheaf cohomology**

$$H^k(\mathcal{P};\mathcal{F}) := \ker(\delta^k)/\overline{\mathcal{R}(\delta^{k-1})}.$$

When $\mathcal{R}(\delta^{k-1})$ is closed, we call H^k the k^{th}-**sheaf cohomology**.

In degree zero, the (reduced) sheaf cohomology is easy to interpret. By Theorem 3.2.11, there is a unitary isomorphism $\Gamma(\mathcal{F}) = \ker(\delta^0) \cong H^k(\mathcal{P};\mathcal{F})$ between the reduced sheaf cohomology and the space of global sections. Higher cohomology groups for closed Hilbert sheaves can be interpreted as obstructions to lifting cocyclic sections in the manner of Section 3.2.3.

We may also work with relative cohomology of Hilbert sheaves. Say that a full subcategory $\mathcal{B} \subseteq \mathcal{P}$ is a **subcomplex**[2] of $\mathcal{P}$ if whenever $y \in \mathcal{B}$ and $x \leqslant y$ in the underlying poset, we have $x \in \mathcal{B}$ as well. When $\mathcal{B}$ is a subcomplex of $\mathcal{P}$, every signed incidence structure ϵ on $\mathcal{P}$ restricts to a signed incidence structure on $\mathcal{P}$. The restriction $\mathcal{F}\big|_{\mathcal{B}} : \mathcal{B} \to \mathbf{Hilb}_{0,\mathbb{k}}$ is itself a Hilbert sheaf.

We may view the restriction $\mathcal{F}\big|_{\mathcal{B}} : \mathcal{B} \to \mathbf{Hilb}_{0,\mathbb{k}}$ as a Hilbert sheaf on $\mathcal{P}$ in the following way. Let $\iota : \mathcal{B} \to \mathcal{P}$ denote the inclusion of $\mathcal{B}$ as a subcategory. Both ι^* and ι_* preserve Hilbert sheaves, so we may apply both to $\mathcal{F}$ and obtain a Hilbert sheaf $\iota_*\iota^*\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$. This Hilbert sheaf essentially looks look $\mathcal{F}$, except the stalk $\mathcal{F}(\sigma) = 0$ (with the restriction maps into and out of $\mathcal{F}(\sigma)$ adjusted accordingly) whenever $\sigma \notin \mathcal{B}$. There is a natural transformation $\mathcal{F} \to \iota_*\iota^*\mathcal{F}$ given by the identity map I on stalks over $\sigma \in \mathcal{B}$, and 0 on the stalks over $\sigma \notin \mathcal{B}$.

We may also form the Hilbert sheaf $\ker(\iota_*\iota^*) : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$, which has stalks

$$\ker(\iota_*\iota^*)(\sigma) = \begin{cases} \mathcal{F}(\sigma) & \text{if } \sigma \notin \mathcal{B} \\ 0 & \text{if } \sigma \in \mathcal{B} \end{cases}$$

and restriction maps given by domain restriction. There is a morphism of Hilbert sheaves from $\ker(\iota_*\iota^*)$ to $\mathcal{F}$ whose components are stalk-wise inclusion maps.

Note that $\iota_*\iota^*\mathcal{F}$ and $\ker(\iota_*\iota^*)$ are non-zero on complementary stalks. Consequently, the Hilbert complex (reduced) cohomology of the Hilbert complex associated to $\ker(\iota_*\iota^*)$ may be identified with the relative cohomology of $\mathcal{F}$ with respect to $\iota_*\iota^*\mathcal{F}$. We thus use the notation $H^\bullet\big(\mathcal{P}, \mathcal{B}; \mathcal{F}\big) := H^\bullet\big(\mathcal{P}; \ker(\iota_*\iota^*)\big)$.

The identified natural transformations yields a short exact sequence of Hilbert sheaves

$$0 \to \ker(\iota_*\iota^*) \to \mathcal{F} \to \iota_*\iota^*\mathcal{F} \to 0\,,$$

2 We adopt the term "subcomplex" to evoke a subcomplex of a regular cell complex, which play an analogous role when a cellular sheaf is defined on a face poset.

which functorially induces an exact sequence of the associated Hilbert complexes:

$$\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \downarrow & & \downarrow & & \downarrow & & \\
0 & \longrightarrow & C^k(\mathcal{P};\ker(\iota_*\iota^*)) & \longrightarrow & C^k(\mathcal{P};\mathcal{F}) & \longrightarrow & C^k(\mathcal{P};\iota_*\iota^*\mathcal{F}) & \longrightarrow & 0\ \cdot \\
 & & \downarrow{\scriptstyle \delta^k_{\ker(\iota_*\iota^*)}} & & \downarrow{\scriptstyle \delta^k_{\mathcal{F}}} & & \downarrow{\scriptstyle \delta^k_{\iota_*\iota^*\mathcal{F}}} & & \\
 & & \vdots & & \vdots & & \vdots & &
\end{array}$$

The snake lemma (Lemma 3.2.12) provides the following Hilbert complex.

$$\begin{array}{ccccc}
H^0(\mathcal{P},;\mathcal{F}) & \longrightarrow & H^0(\mathcal{P};\mathcal{F}) & \longrightarrow & H^0(\mathcal{P};\iota_*\iota^*\mathcal{F}) \\
 & & \swarrow{\scriptstyle d^0} & & \\
H^1(\mathcal{P},;\mathcal{F}) & \longrightarrow & H^1(\mathcal{P};\mathcal{F}) & \longrightarrow & H^1(\mathcal{P};\iota_*\iota^*\mathcal{F}) \\
 & & \swarrow{\scriptstyle d^1} & & \\
\dots & & & &
\end{array}$$

When $\mathcal{F}$ is a closed Hilbert sheaf, the restriction $\iota_*\iota^*\mathcal{F}$ is closed as well. It follows that this sequence is exact at all locations except $H^k(\mathcal{P};\iota_*\iota^*\mathcal{F})$.

Remark 4.6.1. Note that when $\mathcal{F}$ is a closed Hilbert sheaf, it does not necessarily follow that $\ker(\iota_*\iota^*\mathcal{F})$ is closed.

4.7 RANGE AND NORMALIZATION

Beyond having an interpretable cohomology, closed Hilbert sheaves, whose coboundary operators δ^k have closed range, have desirable properties that generic Hilbert sheaves lack. The following proposition justifies the emphasis on closed Hilbert sheaves.

Proposition 4.7.1. *Let $A : X \to Y$ be a closed, densely defined Hilbert space operator. If A has closed range, then $0 \notin \sigma(A)$ or 0 is an isolated eigenvalue.*

Proof. This follows straightforwardly from the closed range theorem (Theorem 2.3.15). Since A has closed range, the restriction $A_0 := A\big|_{\ker(A)^\perp}$ is bounded below; there is a constant $C > 0$ such that $\|A_0x\| \geqslant C\|x\|$ for all $x \in \mathrm{Dom}(A_0)$. A_0 is injective, and admits an inverse map A_0^{-1} : $Y \to \ker(A)^\perp$ with domain $\mathrm{Dom}(A_0^{-1}) = \mathcal{R}(A_0)$. A_0^{-1} is a bounded operator with operator norm $\|A_0^{-1}\|_{op} \leqslant \frac{1}{C}$. The spectral radius of a bounded operator is bounded above by the operator norm, so $|\lambda| \leqslant \frac{1}{C}$ for all $\lambda \in \sigma(A_0^{-1})$. By the spectral inversion formula, $\sigma(A)\backslash\{0\} = \sigma(A_0)\backslash\{0\} = \{\lambda \in \mathbb{C} : \frac{1}{\lambda} \in \sigma(A_0^{-1})\backslash\{0\}\}$ is bounded away from zero. □

When $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is a closed Hilbert sheaf, each coboundary operator $\delta^k : C^k(\mathcal{P};\mathcal{F}) \to C^{k+1}(\mathcal{P};\mathcal{F})$ has a "gap" in its spectrum surrounding 0. As we will observe in the following chapters, the closed range condition also impacts the Hodge Laplacian of the corresponding Hilbert complex, causing the spectral theory and dynamics to more closely mirror that of the finite dimensional case. Other closed range conditions also allow the recovery of finite dimensional results, like sheaf normalization.

Definition 4.7.2. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf. $\mathcal{F}$ is **normalized** if for all $\sigma \in \mathcal{P}$ and $x, y \in \mathcal{F}(\sigma) \cap \ker(\delta)^{\perp}$, we have $\langle \delta x, \delta y\rangle = \langle x, y\rangle$.

Remark 4.7.3. Note that we are defining normalization with respect to the *Hilbert complex* associated to the cellular sheaf.

Remark 4.7.4. It follows directly from the definition that if $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is a normalized sheaf, the coboundary operator $\delta^\bullet$ is a bounded operator in each grade. Under some additional mild hypotheses on the structure of $\mathcal{P}$, it follows that $\mathcal{F}$ is a bounded Hilbert sheaf.

At times, it is convenient to normalize a cellular sheaf—replace a cellular sheaf $\mathcal{F}$ with an isomorphic normalized sheaf $\hat{\mathcal{F}}$.

Theorem 4.7.5. *Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf such that $\mathcal{P}$ has a maximal rank N. $\mathcal{F}$ can be normalized if and only if the associated Hilbert complex $(C^\bullet(\mathcal{P};\mathcal{F}), \delta^\bullet)$ is bounded and each column of δ^k has closed range.*

Proof. First, suppose there is a k such that $\delta^k : C^k(\mathcal{P};\mathcal{F}) \to C^{k+1}(\mathcal{P};\mathcal{F})$ is unbounded. Then δ^k as a linear function is unbounded for all pairs of Hilbert space norms on the vector spaces $C^k(\mathcal{P};\mathcal{F})$ and $C^{k+1}(\mathcal{P};\mathcal{F})$. It follows that the sheaf $\mathcal{F}$ cannot be normalized. Hence boundedness of each δ^k is required.

Next, we prove that when each δ^k is bounded, $\mathcal{F}$ can be normalized if and only if each δ^k has closed range. By the closed graph theorem, each δ^k is necessarily globally defined. Starting from the highest rank, we modify the inner product structure on each stalk to achieve normalization. Suppose all cells of rank $k+1$ have already been normalized; that is, for all $\tau \in \mathcal{P}$ with $r(\tau) \geqslant k+1$ and all $x, y \in \mathcal{F}(\tau) \cap \ker(\delta^{k+1})^{\perp}$, we have $\langle \delta^{k+1}x, \delta^{k+1}y\rangle = \langle x, y\rangle$. Fix a cell $\sigma \in \mathcal{P}$ of rank k. To normalize the stalk $\mathcal{F}(\sigma)$, take the orthogonal projections $P, Q : \mathcal{F}(\sigma) \to \mathcal{F}(\sigma)$ onto $\ker(\delta^{k+1})$ and $\ker(\delta^{k+1})^{\perp}$ respectively. We may take these projections as $\ker(\delta^{k+1}) \cap \mathcal{F}(\sigma)$ is topologically closed. Define a new inner product on $\mathcal{F}(\sigma)$ by

$$\langle x, y\rangle_{*\sigma} := \langle \delta^k Qx, \delta^k Qy\rangle_{*C^{k+1}} + \langle Px, Py\rangle_\sigma\,,$$

where $\langle -, - \rangle_\sigma$ is the un-normalized inner product on $\mathcal{F}(\sigma)$, and $\langle -, - \rangle_{*C^{k+1}}$ is the normalized inner product on C^{k+1}. The stalk $\mathcal{F}(\sigma)$ is normalized with respect to this new inner product. Repeating this process inductively normalizes the sheaf.

One may check that this new inner product also induces a Hilbert space structure on $\mathcal{F}(\sigma)$ if and only if the column of δ^k that acts on $\mathcal{F}(\sigma)$ has closed range. When this inner product does induce a Hilbert space structure, the two Hilbert space structures must have equivalent norms. Let $\mathcal{F}' : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$ denote the normalized sheaf obtained at the end of this inductive procedure. The natural transformation $\eta : \mathcal{F} \Rightarrow \mathcal{F}'$ given by $\eta_\sigma : x \mapsto x$ witnesses that these two Hilbert sheaves are isomorphic. $\square$

Remark 4.7.6. This "every column of δ^k has closed range" condition is quite restrictive; this will not hold in general, and neither implies nor is implied by $\mathcal{F}$ being a closed Hilbert sheaf.

5

HILBERT SHEAF LAPLACIANS

This chapter investigates the spectral properties of Laplacians associated with cellular sheaves valued in Hilbert spaces. The sheaf Laplacian, defined as the Hodge Laplacian of the Hilbert complex from Chapter 3, encodes local consistency of stalk-wise data and global harmonic structure. Building on the finite-dimensional theory, the infinite-dimensional setting reveals new phenomena, such as the delicate relationship between closed-range conditions, spectral gaps, and the solvability of harmonic extension problems. This chapter establishes when the finite-dimensional results of [54] and [52] generalize.

Section 5.1 introduces the Hilbert sheaf Laplacian $\mathcal{L}^k = (\delta^k)^*\delta^k + \delta^{k-1}(\delta^{k-1})^*$ as a closed, densely defined, positive operator on the space of k-cochains. Its fundamental properties follow from the general theory of Hodge Laplacians. The kernel consists precisely of harmonic cochains, which are unitarily isomorphic to reduced cohomology classes. Section 5.2 examines the harmonic extension problem, wherein one seeks to extend data over a subset of cells, thought of as a boundary constraint, to a harmonic cochain on the complement of the boundary cells. For bounded Hilbert sheaves, we establish existence through block operator analysis; for unbounded sheaves, we employ the theory of shorted operators to characterize when solutions exist. The analysis reveals that solvability depends crucially on whether boundary data lies in the domain of a certain quadratic form. Section 5.3 develops the spectral theory of sheaf Laplacians. The interaction of sheaf morphisms with the Laplacian spectra is explored, establishing spectral containment relationships. When coboundary operators have closed range, zero becomes an isolated eigenvalue, and we establish interlacing results for eigenvalues under sheaf morphisms. The chapter concludes by examining how sheaf operations, particularly direct sums and pullbacks, interact with Laplacian spectra.

5.1 THE HILBERT SHEAF LAPLACIAN

In Section 3.2, we observed that every Hilbert complex has an associated Hodge Laplacian. The Hodge Laplacian $\mathcal{L}^\bullet$ forms a positive operator on each grade of the Hilbert complex. We begin this chapter investigating the corresponding Laplacian $\mathcal{L}^k(\mathcal{P};\mathcal{F})$ that arises from the Hilbert complex $(C^\bullet, \delta^\bullet)$ associated to a Hilbert sheaf $F : \mathcal{P} \to \mathbf{Hilb}_{\mathbb{k}}$.

Definition 5.1.1. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf with corresponding Hilbert complex $(C^\bullet, \delta^\bullet)$. The **sheaf Laplacian** for $\mathcal{F}$ is the chain map $\mathcal{L}^\bullet : C^\bullet \to C^\bullet$ defined by

$$\mathcal{L}^k := (\delta^k)^* \delta^k + \delta^{k-1}(\delta^{k-1})^*$$

on k-cochains.

Remark 5.1.2. Since the coboundary map itself depends on the choice of signed incidence relation on $\mathcal{P}$, the Laplacian does as well.

As the Hodge Laplacian of a Hilbert complex, Theorem 3.2.21 and Theorem 3.2.19 apply to Hilbert sheaf Laplacians, and give the following properties.

Proposition 5.1.3. *Let $\mathcal{L}^\bullet$ be the Hilbert sheaf Laplacian of a Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$. $\mathcal{L}^\bullet$ has the following properties.*

(i) $\mathcal{L}^k$ *is a closed, densely defined, positive operator for each* $k \in \mathbb{N}$.

(ii) The kernel $\ker(\mathcal{L}^k) \subseteq C^k(\mathcal{P}; \mathcal{F})$ *is exactly the* k*-harmonic space* $\mathfrak{H}^k$*, and is unitarily isomorphic to the* k^{th}*-reduced cohomology* $H^k(\mathcal{P}; \mathcal{F}) = \ker(\delta^k)/\mathcal{R}(\delta^{k-1})$.

(iii) In particular, the kernel of $\mathcal{L}^0$ *is exactly the kernel of* δ^0*, which is unitarily isomorphic to the space of global sections* $\ker(\delta^0)$.

Example 5.1.4. We return to the simple Hilbert sheaf

$$\begin{array}{ccccc} L^2(\mathbb{R}) & \xrightarrow{\frac{d}{dx}} & L^2(\mathbb{R}) & \xleftarrow{\frac{d}{dx}} & L^2(\mathbb{R}) \\ \vdots & & \vdots & & \vdots \\ \bullet & \rule{2em}{0.4pt} & \rule{2em}{0.4pt} & \rule{2em}{0.4pt} & \bullet \end{array}$$

from Example 4.4.5. The degree-zero Laplacian $\mathcal{L} : L^2(\mathbb{R}) \oplus L^2(\mathbb{R}) \to L^2(\mathbb{R}) \oplus L^2(\mathbb{R})$ is given by

$$\mathcal{L}\begin{pmatrix} f \\ g \end{pmatrix} = \begin{pmatrix} \frac{d^2}{dx^2}(g - f) \\ \frac{d^2}{dx^2}(f - g) \end{pmatrix}$$

with domain $\mathrm{Dom}(\mathcal{L}) = \{(f, g)^T \; : \; \frac{d^2}{dx^2}(f - g) \in L^2(\mathbb{R})\}$.

Definition 5.2.1. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a cellular sheaf of Hilbert spaces, and let $\mathcal{U} \subseteq \text{Ob}(\mathcal{P})$ be a collection of objects of rank k in $\mathcal{P}$. We say that a k-cochain $\mathbf{x} \in C^k(\mathcal{P};\mathcal{F})$ is **harmonic on** $\mathcal{U}$ if $\mathcal{L}^k\mathbf{x}\big|_{\mathcal{U}} = 0$, where $\mathbf{x}\big|_{\mathcal{U}}$ is the orthogonal projection of $C^k(\mathcal{P};\mathcal{F})$ onto the subspace $\bigoplus_{u\in\mathcal{U}} \mathcal{F}(u)$.

We have already observed that $\mathfrak{H}^k$, the space of k-harmonic cochains, is exactly the kernel of the sheaf Laplacian $\mathcal{L}^k$, and that the 0-harmonic cochains may be identified with the space of global sections (Proposition 5.1.3). This identification carries particular significance in the context of cellular sheaves, where the coboundary operator $\delta^0 : C^0(\mathcal{P};\mathcal{F}) \to C^1(\mathcal{P};\mathcal{F})$ encodes the compatibility conditions between local sections. The harmonic cochains represent a natural generalization of classical harmonic functions to the sheaf-theoretic setting. In the finite-dimensional case, these cochains admit an explicit characterization through the Hodge decomposition. However, for Hilbert sheaves with unbounded operators, the relationship between harmonic cochains and cohomology classes becomes more subtle, particularly when the coboundary maps fail to have closed range. This motivates the study of harmonic extension problems, wherein we seek to extend partial data defined on a subcomplex to a harmonic cochain on the entire domain.

The harmonic extension problem naturally arises in several contexts within applied topology. For instance, when modeling distributed systems or sensor networks via cellular sheaves, one often possesses measurements or constraints on a subset of nodes and seeks to infer consistent values throughout the network. The existence and uniqueness of such extensions depend crucially on the spectral properties of the restricted Laplacian operators, as we shall demonstrate.

Definition 5.2.2. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf, $\mathcal{B}$ be a subcomplex of $\mathcal{P}$, and $\mathcal{U} = \text{Ob}(\mathcal{P})\backslash\text{Ob}(\mathcal{B})$. A **harmonic extension** of a k-cochain $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$ is an $\mathbf{x} \in C^k(\mathcal{P};\mathcal{F})$ such that $(\mathcal{L}^k\mathbf{x})\big|_{\mathcal{U}} = 0$ and $\mathbf{x}\big|_{\mathcal{B}} = \mathbf{y}$. The **harmonic extension problem** is to find a harmonic extension $\mathbf{x} \in C^k(\mathcal{P};\mathcal{F})$ of a specified cochain $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$.

Remark 5.2.3. When $\mathcal{B}$ is empty, the harmonic extension problem reduces to the problem of finding a harmonic cochain.

5.2.1 *Bounded harmonic extensions*

Analysis of the harmonic extension problem for a bounded Hilbert sheaf is straightforward. In a slight abuse of notation, write $C^k(\mathcal{P};\mathcal{F}) = C^k(\mathcal{U};\mathcal{F}) \oplus C^k(\mathcal{B};\mathcal{F})$ for each k. Write the Laplacian $\mathcal{L}^k : C^k(\mathcal{U};\mathcal{F}) \oplus C^k(\mathcal{B};\mathcal{F}) \to C^{k+1}(\mathcal{U};\mathcal{F}) \oplus C^{k+1}(\mathcal{B};\mathcal{F})$ as a block operator

$$\mathcal{L}^k = \begin{bmatrix} \mathcal{L}^k_{\mathcal{U},\mathcal{U}} & \mathcal{L}^k_{\mathcal{U},\mathcal{B}} \\ \mathcal{L}^k_{\mathcal{B},\mathcal{U}} & \mathcal{L}^k_{\mathcal{B},\mathcal{B}} \end{bmatrix}.$$

The harmonic extension problem asks us to find $\mathbf{w} \in C^k(\mathcal{U};\mathcal{F})$ and $\mathbf{z} \in C^k(\mathcal{B};\mathcal{F})$ such that

$$\begin{bmatrix} \mathcal{L}^k_{\mathcal{U},\mathcal{U}} & \mathcal{L}^k_{\mathcal{U},\mathcal{B}} \\ \mathcal{L}^k_{\mathcal{B},\mathcal{U}} & \mathcal{L}^k_{\mathcal{B},\mathcal{B}} \end{bmatrix} \begin{bmatrix} \mathbf{w} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{z} \end{bmatrix}.$$

It suffices to find a $\mathbf{w}$ such that $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}\mathbf{w} = -\mathcal{L}^k_{\mathcal{B},\mathcal{U}}\mathbf{y}$. In other words, to show the harmonic extension problem has a solution for all $\mathbf{y}$, it suffices to show that $\mathcal{R}(\mathcal{L}^k_{\mathcal{U},\mathcal{B}}) \subseteq \mathcal{R}(\mathcal{L}^k_{\mathcal{U},\mathcal{U}})$. With respect to this same decomposition of k-cochain spaces, we may write

$$\delta^k = \begin{bmatrix} \delta^k_{\mathcal{U},\mathcal{U}} & \delta^k_{\mathcal{U},\mathcal{B}} \\ 0 & \delta^k_{\mathcal{B},\mathcal{B}} \end{bmatrix},$$

where $\delta^k_{\mathcal{B},\mathcal{U}} = 0$ by the closure property of a subcomplex. We may express the k-Laplacian as

$$\begin{aligned} \mathcal{L}^k &= \begin{bmatrix} \mathcal{L}^k_{\mathcal{U},\mathcal{U}} & \mathcal{L}^k_{\mathcal{U},\mathcal{B}} \\ \mathcal{L}^k_{\mathcal{B},\mathcal{U}} & \mathcal{L}^k_{\mathcal{B},\mathcal{B}} \end{bmatrix} \\ &= \begin{bmatrix} (\delta^k_{\mathcal{U},\mathcal{U}})^*\delta^k_{\mathcal{U},\mathcal{U}} & (\delta^k_{\mathcal{U},\mathcal{U}})^*\delta^k_{\mathcal{U},\mathcal{B}} \\ (\delta^k_{\mathcal{U},\mathcal{B}})^*\delta^k_{\mathcal{U},\mathcal{U}} & (\delta^k_{\mathcal{U},\mathcal{B}})^*\delta^k_{\mathcal{U},\mathcal{B}} + (\delta^k_{\mathcal{B},\mathcal{B}})^*\delta^k_{\mathcal{B},\mathcal{B}} \end{bmatrix} \\ &\quad + \begin{bmatrix} \delta^{k-1}_{\mathcal{U},\mathcal{U}}(\delta^{k-1}_{\mathcal{U},\mathcal{U}})^* + \delta^{k-1}_{\mathcal{U},\mathcal{B}}(\delta^{k-1}_{\mathcal{U},\mathcal{B}})^* & \delta^{k-1}_{\mathcal{U},\mathcal{B}}(\delta^{k-1}_{\mathcal{B},\mathcal{B}})^* \\ \delta^{k-1}_{\mathcal{B},\mathcal{B}}(\delta^{k-1}_{\mathcal{U},\mathcal{B}})^* & \delta^{k-1}_{\mathcal{B},\mathcal{B}}(\delta^{k-1}_{\mathcal{B},\mathcal{B}})^* \end{bmatrix}. \end{aligned}$$

Writing $T = \begin{bmatrix} (\delta^k_{\mathcal{U},\mathcal{U}})^* & \delta^{k-1}_{\mathcal{U},\mathcal{U}} & \delta^{k-1}_{\mathcal{U},\mathcal{B}} \end{bmatrix}$ and $S = \begin{bmatrix} (\delta^k_{\mathcal{U},\mathcal{B}})^* & 0 & \delta^{k-1}_{\mathcal{B},\mathcal{B}} \end{bmatrix}$, we may write $\mathcal{L}^k_{\mathcal{U},\mathcal{U}} = TT^*$ and $\mathcal{L}^k_{\mathcal{U},\mathcal{B}} = TS^*$. We have a range inclusion $\mathcal{R}(TS^*) \subseteq \mathcal{R}(T)$. Finally, $\mathcal{R}(T) = \mathcal{R}(TT^*)$ if and only if T has closed range. We now present the following theorem.

Theorem 5.2.4. *Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$ be a bounded Hilbert sheaf and $\mathcal{B}$ a subcomplex of $\mathcal{P}$. The harmonic extension problem has a solution for every $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$ if and only if $\delta^k_{\mathcal{U},\mathcal{U}}$, $\delta^{k-1}_{\mathcal{U},\mathcal{U}}$, and $\delta^{k-1}_{\mathcal{U},\mathcal{B}}$ have closed* *ranges. When* δ^k *itself has closed range as well, the solution is unique if and only if the map* $H^k(\mathcal{P}, \mathcal{B}; \mathcal{F}) \to H^k(\mathcal{P}; F)$ *is the zero map.*

Proof. We have already established the existence of solutions to the harmonic extension problem. To establish uniqueness, we must determine when $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}$ is injective. We relate the decomposition $C^k(\mathcal{P}; \mathcal{F}) = C^k(\mathcal{B}; \mathcal{F}) \oplus C^k(\mathcal{U}; \mathcal{F})$ to the relative cohomology with respect to $\mathcal{B}$. Since $C^k(\mathcal{U}; \mathcal{F}) = C^k(\mathcal{B}; \mathcal{F})^\perp$, we may identify $C^k(\mathcal{U}; \mathcal{F})$ with the space of relative k-cochains $C^k(\mathcal{P}, \mathcal{B}; \mathcal{F})$ that vanish on $\mathcal{B}$. The coboundary operator of the relative Hilbert complex is exactly the block $\delta^k_{\mathcal{U},\mathcal{U}}$ in the block representation of δ^k; we recover the relative Hilbert complex

$$\cdots \longrightarrow C^k(\mathcal{P}, \mathcal{B}; \mathcal{F}) \xrightarrow{\delta^k_{\mathcal{U},\mathcal{U}}} C^{k+1}(\mathcal{P}, \mathcal{B}; \mathcal{F}) \longrightarrow \cdots$$

and its Hodge Laplacian $\Delta^k = (\delta^k_{\mathcal{U},\mathcal{U}})^* \delta^k_{\mathcal{U},\mathcal{U}} + \delta^{k-1}_{\mathcal{U},\mathcal{U}} (\delta^{k-1}_{\mathcal{U},\mathcal{U}})^*$. Looking at the block decomposition of $\mathcal{L}^k$, we may write $\mathcal{L}^k_{\mathcal{U},\mathcal{U}} = \Delta^k + \delta^{k-1}_{\mathcal{U},\mathcal{B}} (\delta^{k-1}_{\mathcal{U},\mathcal{B}})^*$. The kernel of a sum of positive operators is the intersection of the kernels, so $\ker(\mathcal{L}^k_{\mathcal{U},\mathcal{U}}) = \ker(\Delta^k) \cap \ker((\delta^{k-1}_{\mathcal{U},\mathcal{B}})^*)$. Now, consider the map

$$d : \mathfrak{H}^{k-1}(\mathcal{B}; \mathcal{F}) \to \mathfrak{H}^k(\mathcal{P}, \mathcal{B}; \mathcal{F})$$

from the cohomology long exact sequence on Hodge representatives. The map is given by the restriction $d := \delta^{k-1}_{\mathcal{B},\mathcal{U}}\big|_{\mathfrak{H}^{k-1}(\mathcal{B};\mathcal{F})}$. Therefore $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}$ is injective if and only if $\mathcal{R}(d)$ is dense in $\mathfrak{H}^k(\mathcal{P}, \mathcal{B}; \mathcal{F})$. Since δ^k has closed range, the relative cohomology Hilbert complex is exact at $\mathfrak{H}^k(\mathcal{P}, \mathcal{B}; \mathcal{F})$, so $\mathcal{R}(d)$ is dense in $\mathfrak{H}^k(\mathcal{P}, \mathcal{B}; \mathcal{F})$ if and only if $j : \mathfrak{H}^k(\mathcal{P}, \mathcal{B}; \mathcal{F}) \to \mathfrak{H}^k(\mathcal{P}; \mathcal{F})$ is the zero map. □

5.2.2 *Unbounded harmonic extension*

The preceding analysis of the harmonic extension problem for a bounded Hilbert sheaf does not directly apply to an unbounded Hilbert sheaf. There are two key difficulties that must be overcome.

1. **Domain splitting.** When $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is an unbounded Hilbert sheaf, while the operator $\mathring{\delta}^k$ splits over the decompositions $C^\bullet(\mathcal{P}; \mathcal{F}) \cong C^\bullet(\mathcal{P}, \mathcal{B}; \mathcal{F}) \oplus C^\bullet(\mathcal{B}; \mathcal{F})$, there is no guarantee that the Laplacian $\mathcal{L}^k$ does. The restriction $(\mathcal{L}^k \mathbf{x})\big|_{\mathcal{U}}$ may not be well defined.
2. **Well-posedness.** When the Laplacian does split, Harmonic extension becomes trivially impossible when $\mathbf{y} \in C^($\mathcal{B}; \mathcal{F})$ falls outside the shared domain $\mathrm{Dom}(\mathcal{L}^k_{\mathcal{U},\mathcal{B}}) \cap \mathrm{Dom}(\mathcal{L}^k_{\mathcal{B},\mathcal{B}})$.

Nonetheless, well-posed harmonic extension problems for unbounded Hilbert sheaves may have interpretable solutions.

Example 5.2.5. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be the Hilbert sheaf

$$\begin{array}{ccccc} L^2([0,1]) & \xrightarrow{\frac{d}{dx}} & L^2([0,1]) & \xleftarrow{I} & L^2([0,1]) \\ \vdots & & \vdots & & \vdots \\ \bullet & \text{———} & \text{———} & \text{———} & \star \end{array}$$

with $\mathrm{Dom}(\frac{d}{dx}) = \{f \in L^2([0,1]) \,:\, f' \in L^2([0,1])\}$. Let $\{\star\}$ be a one-point subcomplex of $\mathcal{P}$. $\mathcal{F}$ is a proper Hilbert sheaf with Laplacian

$$\mathcal{L}\begin{pmatrix} f \\ g \end{pmatrix} = \begin{pmatrix} (g - f')' \\ g - f' \end{pmatrix}$$

with domain $\mathrm{Dom}(\mathcal{L}) = \{(f,g)^T \;:\; f' \in L^2 \text{ and } (f'-g)' \in L^2\}$. Since these derivatives are weak, this Laplacian can be written in block-matrix form over the decomposition $C^0(\mathcal{P},\{\star\};\mathcal{F}) \oplus C^0(\{\star\};\mathcal{F}) \cong L^2([0,1]) \oplus L^2([0,1])$ as

$$\mathcal{L} = \begin{bmatrix} -\frac{d^2}{dx^2} & \frac{d}{dx} \\ -\frac{d}{dx} & I \end{bmatrix}.$$

The range of $\frac{d}{dx}$ is a superset of than the range of $\frac{d^2}{dx^2}$, so the Harmonic extension problem may fail to have a solution. For a fixed $g \in C^0(\{\star\};\mathcal{F}) \cong L^2([0,1])$ the harmonic extension problem asks us to find an $f \in L^2([0,1])$ such that $f'' - g' = 0$. This is satisfied by $f = cx + \int g$, where $\int g$ is any anti-derivative of g. As expected, the harmonic extension problem has a solution exactly when $g \in \mathcal{R}(\frac{d}{dx})$. Moreover, the solution when $c = 0$ is a harmonic section and f is a solution to the)(weak) differential equation $f' = g$.

Following the approach of Arlinskiĭ [6] for linear relations, we analyze the harmonic extension problem through the theory of shorted operators. The following analysis also applies to bounded Hilbert sheaves.

Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf with associated cochain complex $(C^\bullet(\mathcal{P};\mathcal{F}), \delta^\bullet)$. Let $\mathcal{B}$ be a subcomplex of $\mathcal{P}$, and let P and Q denote the orthogonal projections from $C^k(\mathcal{P};\mathcal{F})$ onto $C^k(\mathcal{B};\mathcal{F})$ and $C^k(\mathcal{U};\mathcal{F}) = C^k(\mathcal{B};\mathcal{F})^\perp$ respectively. Let $\mathcal{D} := \mathrm{Dom}(\delta^k) \cap \mathrm{Dom}((\delta^{k-1})^*)$. The coboundary operators δ^k define a a quadratic form

$$\mathfrak{a}(u,v) := \langle \delta^k u, \delta^k v\rangle + \langle (\delta^{k-1})^* u, (\delta^{k-1})^* v\rangle$$

on $C^k(\mathcal{P};\mathcal{F}) \times C^k(\mathcal{P};\mathcal{F})$ with domain $\mathcal{D} \times \mathcal{D}$. When $u \in \mathrm{Dom}(\mathcal{L}^k)$ and $v \in \mathcal{D}$, we have $\mathfrak{a}(u,v) = \langle \mathcal{L}^k u, v\rangle$.

For $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$, let $\mathcal{A}_\mathbf{y} := \{\mathbf{x} \in \mathcal{D} \ : \ P\mathbf{x} = \mathbf{y}\}$. Note that $\mathcal{A}_\mathbf{y}$ may be empty. We may rephrase the harmonic extension problem in terms for the forms $\mathfrak{a}_\mathcal{B}$ and $\mathfrak{a}$.

Proposition 5.2.6. *Let* $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$. $\mathbf{u} \in C^k(\mathcal{P};\mathcal{F})$ *is a solution to the harmonic extension problem for* $\mathbf{y}$ *if and only if* $\mathbf{u}$ *is a minimizer of* $\mathfrak{q}(-) := \mathfrak{a}(-,-)$ *over* $\mathcal{A}_\mathbf{y}$.

Proof. We prove the real case; the complex case is similar. Let $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$. A solution to the harmonic extension problem for $\mathbf{y}$ is an $\mathbf{x} \in C^k(\mathcal{B};\mathcal{F})^\perp$ such that $\mathcal{L}^k(\mathbf{x}+\mathbf{y}) = \mathbf{z} \in C^k(\mathcal{B};\mathcal{F})$. Set $\mathbf{u} := \mathbf{x}+\mathbf{y}$, and let $\mathbf{x}' \in C^k(\mathcal{B};\mathcal{F})^\perp \cap \mathcal{D}$, and consider the quantity

$$\mathfrak{q}(\mathbf{u}+\mathbf{x}') = \|(\mathcal{L}^k)^{1/2}\mathbf{u}\|^2 + \|(\mathcal{L}^k)^{1/2}\mathbf{x}\|^2 + 2\langle \mathcal{L}^k\mathbf{u}, \mathbf{x}'\rangle.$$

Since $\mathcal{L}^k\mathbf{u} \perp \mathbf{x}'$, it follows that $\mathfrak{q}(\mathbf{u}+\mathbf{x}') \geqslant \mathfrak{q}(\mathbf{u})$, and $\mathbf{u}$ minimizes $\mathfrak{q}(\mathbf{u})$.

Conversely, suppose that $\mathbf{u} = \mathbf{x}+\mathbf{y}$ minimizes $\mathfrak{q}$ on $\mathcal{A}_\mathbf{y}$. For any $\mathbf{x}' \in C^k(\mathcal{B};\mathcal{F})^\perp \cap \mathcal{D}$ and $t \in \mathbb{R}$, we have $\mathfrak{q}(\mathbf{u}+t\mathbf{x}') \geqslant \mathfrak{q}(\mathbf{u})$. Expanding and rearranging yields

$$2t\langle \mathcal{L}^k\mathbf{u}, \mathbf{x}'\rangle + t^2\mathfrak{q}(\mathbf{x}') \geqslant 0.$$

This quadratic in t can only be positive if the linear term $2t\langle \mathcal{L}^k\mathbf{u}, \mathbf{x}'\rangle = 0$, from which we conclude that $P\mathcal{L}^k\mathbf{u} = 0$, making $\mathbf{u}$ a solution to the harmonic extension problem. ☐

Remark 5.2.7. We define a new functional

$$\mathfrak{a}_\mathcal{B}(\mathbf{y}) := \inf\{\mathfrak{a}(\mathbf{x},\mathbf{x}) \ : \ \mathbf{x} \in \mathcal{A}_\mathbf{y}\}$$

with $\mathrm{Dom}(\mathfrak{a}_\mathcal{B}) := \{\mathbf{y} \in C^k(\mathcal{B};\mathcal{F}) \ : \ \mathcal{A}_\mathbf{y} \neq \varnothing\}$. As a minimizer of $\mathfrak{q}$, a harmonic extension $\mathbf{u}$ realizes the infimum $\mathfrak{a}_\mathcal{B}(\mathbf{y})$. The quadratic form $\mathfrak{a}_\mathcal{B}(\mathbf{y})$ defines an energy functional, which is minimized by a harmonic extension.

Theorem 5.2.8 ([6, Theorem 3.1])**.** *Let* $A : X \to X$ *be a positive Hilbert space operator, and* Y *a subspace of* X. *The set*

$$\Xi(A,Y) := \left\{\tilde{A} \textit{ a positive operator with } \tilde{A} \leqslant A \textit{ and } \mathcal{R}(\tilde{A}) \subseteq Y\right\}$$

has a unique maximal element A_Y, *called the* ***shortening*** *of* A. *The map* A_Y *has the following properties.*

(i) $Y^\perp \subseteq \ker A_Y$.

(ii) $A_Y\big|_Y$ *is self-adjoint in* Y.

(iii) $\mathcal{R}(A_Y^{1/2}) = \mathcal{R}(A^{1/2}) \cap Y$.

Let $\mathcal{L}^k_\mathcal{B}$ denote the shortening of the sheaf Laplacian $\mathcal{L}^k$ with respect to the subspace $C^k(\mathcal{B};\mathcal{F}) \subseteq C^k(\mathcal{P};\mathcal{F})$. We may relate the shortening $\mathcal{L}^k_\mathcal{B}$ to the quadratic operator $\mathfrak{a}_\mathcal{B}$ by the following proposition, adapted from [77].

Proposition 5.2.9. *There is an equality* $\mathfrak{a}_{\mathcal{B}}(\mathbf{y}) = \langle \mathcal{L}^k_{\mathcal{B}}\mathbf{y}, \mathbf{y}\rangle$.

Proof. For each $t \geqslant 0$, let $B_t := (\mathcal{L}^k : tP)$ denote the parallel sum of positive operators [6, Definition 4.1]. Since tP is bounded, each parallel sum B_t is a bounded, positive operator. The B_t operators have strong resolvent limit

$$\text{SR-}\lim_{t\to\infty} B_t = \mathcal{L}^k_{\mathcal{B}}$$

by [6, Theorem 4.3]. Moreover, by the idempotence of shorting and [6, Proposition 3.2], it follows that $\text{SR-}\lim_{t\to\infty}(B_t)_{\mathcal{B}} := (\mathcal{L}^k_{\mathcal{B}})_{\mathcal{B}} = \mathcal{L}^k_{\mathcal{B}}$. Since each $(B_t)_{\mathcal{B}}$ is bounded and positive, they generate a globally defined, non-negative quadratic form $\mathfrak{b}_t(\mathbf{x}) = \langle (B_t)_{\mathcal{B}}\mathbf{x}, \mathbf{x}\rangle$. Moreover, this family is bounded below by zero, and is monotonically increasing in t. For each $\mathbf{x} \in \text{Dom}(\mathcal{L}^k_{\mathcal{B}})$, Kato's monotone convergence theorem for quadratic forms ([73, Theorem VIII.3.13a]) ensures $\lim_{t\to\infty} \mathfrak{b}_t(\mathbf{x}) = \langle \mathcal{L}^k_{\mathcal{B}}\mathbf{x}, \mathbf{x}\rangle$.

Conversely, since each operator $(B_t)_{\mathcal{B}}$ is bounded, we may apply Kreĭn's variational identity [77]:

$$\mathfrak{b}_t(\mathbf{x}) = \inf_{\mathbf{y}\in\mathcal{B}^\perp} \left\{\langle (B_t)_{\mathcal{B}}(\mathbf{x}+\mathbf{y}), \mathbf{x}+\mathbf{y}\rangle\right\} .$$

The quantity on the right hand side is monotonically increasing in t for each $\mathbf{x}$, from which we may conclude that $\lim_{t\to\infty} \mathfrak{b}_t(\mathbf{x}) = \mathfrak{a}_{\mathcal{B}}(\mathbf{x})$ for all $\mathbf{x} \in \text{Dom}(\mathcal{L}^k_{\mathcal{B}})$. This identity may be weakly extended to $\mathbf{x} \in \text{Dom}\left((\mathcal{L}^k_{\mathcal{B}})^{1/2}\right)$. ☐

Corollary 5.2.10. *The harmonic extension problem has a solution at* $\mathbf{y}$ *if and only if* $\mathbf{y} \in \text{Dom}(\mathfrak{a}_{\mathcal{B}})$. *Moreover,* $\mathbf{y} \in \text{Dom}(\mathcal{L}^k_{\mathcal{B}})$ *if and only if the energy minimizer* $\mathbf{u} = \mathbf{y} + \mathbf{x}$ *belongs to* $\text{Dom}(\mathcal{L}^k)$, *and* $\mathbf{z} := P\mathcal{L}^k\mathbf{u} = \mathcal{L}^k_{\mathcal{B}}\mathbf{y}$ *is determined uniquely.*

5.3 Laplacian spectra

We begin by summarizing some properties of the domains and ranges of the up and down Laplacians respectively.

Proposition 5.3.1. *Let* $C^k = \overline{\mathfrak{B}^k} \oplus \mathfrak{H}^k \oplus \mathfrak{Z}^{k\perp}$ *denote the weak Hodge decomposition for the space of* k*-cochains. The following hold.*

(i) $\overline{\mathcal{R}(\mathcal{L}^k_-)} \subseteq \overline{\mathfrak{B}^k} \subseteq \ker(\mathcal{L}^k_+) = \ker(\delta^k)$.

(ii) $\overline{\mathcal{R}(\mathcal{L}^k_+)} \subseteq \mathfrak{Z}^{k\perp} \subseteq \ker(\mathcal{L}^k_-) = \ker((\delta^{k-1})^*)$.

(iii) $\mathfrak{H}^k = \ker(\mathcal{L}^k_+) \cap \ker(\mathcal{L}^k_-) = \ker(\mathcal{L}^k)$.

Proof. Since all four operators δ^k and δ^{k-1} are closed operators (and hence closed kernels closed kernels) with $\mathcal{R}(\delta^{k-1}) \subseteq \ker(\delta^k)$, we immediately recover $\overline{\mathfrak{B}^k} \subseteq \ker(\delta^k)$. A similar argument with the adjoints $(\delta^k)^*$ and $(\delta^{k-1})^*$ (plus a use of the identity $\ker(A)^\perp = \overline{\mathcal{R}(A^*)}$ for a closed operator A) yields the second item. The third item is the main content of the proof of Theorem 3.2.21. □

Corollary 5.3.2. *The up and down Laplacians $\mathcal{L}_+^k$ and $\mathcal{L}_-^k$ restrict to maps:*

$$\mathcal{L}_+^k|_{\mathfrak{Z}^{k\perp}} : \mathfrak{Z}^{k\perp} \to \mathfrak{Z}^{k\perp},$$
$$\mathcal{L}_-^k|_{\overline{\mathfrak{B}^k}} : \overline{\mathfrak{B}^k} \to \overline{\mathfrak{B}^k}.$$

When clear from context, we will use $\mathcal{L}_+^k$ and $\mathcal{L}_-^k$ to refer to both the up/down-Laplacian and its restriction. We are now able to begin analyzing the spectrum of the sheaf Laplacian $\mathcal{L}^k$.

Proposition 5.3.3. *0 is not an eigenvalue of $\mathcal{L}_+^k|_{\mathfrak{Z}^{k\perp}}$ nor $\mathcal{L}_-^k|_{\overline{\mathfrak{B}^k}}$. If δ^k (resp. δ^{k-1}) has closed range, then 0 is not in the spectrum $\sigma(\mathcal{L}_+^k|_{\mathfrak{Z}^{k\perp}})$ (resp. $\sigma(\mathcal{L}_-^k|_{\overline{\mathfrak{B}^k}})$).*

Proof. We prove the result for the restricted up-Laplacian $\mathcal{L}_+^k : \mathfrak{Z}^{k\perp} \to \mathfrak{Z}^{k\perp}$; the corresponding argument for the down-Laplacian is essentially identical. Since $\mathfrak{Z}^{k\perp} \perp \ker(\mathcal{L}_+^k) \cap \mathfrak{Z}^{k\perp}$, zero cannot be an eigenvalue. If $\mathcal{R}(\delta^k)$ is closed, the closed range theorem guarantees that δ^k is bounded below on $\mathfrak{Z}^{k\perp}$; there is a $c > 0$ such that $\|\delta^k x\| \geqslant c\|x\|$ for all $x \in \mathfrak{Z}^{k\perp} \cap \mathrm{Dom}(\delta^k)$. The up-Laplacian on $\mathfrak{Z}^{k\perp} \cap \mathrm{Dom}(\mathcal{L}_+^k)$ is bounded below by $c^2 I$, as $\langle \mathcal{L}_k^+ x, x\rangle = \|\delta^k x\|^2 \geqslant c^2\|x\|^2$. It follows that $\mathcal{L}_k^+ - c^2 I$ is a positive operator. $0 \notin \sigma(\mathcal{L}_k^+)$ since $\sigma(\mathcal{L}_k^+ - c^2 I) = \sigma(\mathcal{L}_k^+) + c^2$. □

Proposition 5.3.4. *The spectrum of $\mathcal{L}^k$ is given by $\sigma(\mathcal{L}^k) = \{0\} \cup \sigma(\mathcal{L}_+^k) \cup \sigma(\mathcal{L}_-^k)$. Moreover, if $\mathcal{R}(\delta^{k-1})$ and $\mathcal{R}(\delta^k)$ are closed, then $\sigma(\mathcal{L}^k)\backslash\{0\} = \sigma(\mathcal{L}_+^k|_{\mathfrak{Z}^{k\perp}}) \cup \sigma(\mathcal{L}_-^k|_{\overline{\mathfrak{B}^k}})$.*

Proof. The weak Hodge decomposition $C^k = \overline{\mathfrak{B}^k} \oplus \mathfrak{H}^k \oplus \mathfrak{Z}^{k\perp}$ can equivalently be written as $C^k = \overline{\mathcal{R}(\delta^{k-1})} \oplus \ker(\mathcal{L}^k) \oplus \overline{\mathcal{R}((\delta^k)^*)}$. Utilizing the identities of Proposition 5.3.1, the operator $\mathcal{L}^k : C^k \to C^k$ can be written as a block-diagonal operator

$$\mathcal{L}^k = \begin{bmatrix} \mathcal{L}_-^k|_{\overline{\mathfrak{B}^k}} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \mathcal{L}_+^k|_{\mathfrak{Z}^{k\perp}} \end{bmatrix}.$$

It easily follows that the spectrum of $\mathcal{L}^k$ is the union of the spectra of the blocks: $\sigma(\mathcal{L}^k) = \{0\} \cup \sigma(\mathcal{L}_+^k) \cup \sigma(\mathcal{L}_-^k)$. When δ^{k-1} and δ^k have closed range, Proposition 5.3.3 ensures that zero is not in the spectra of the restricted up/down-Laplacians.

□

Corollary 5.3.5. *Suppose that δ^k and δ^{k-1} both have closed range. Then 0 is an isolated eigenvalue of $\mathcal{L}^k$.*

Proof. In the proof of Proposition 5.3.3, we saw that when δ^k and δ^{k-1} have closed ranges, the spectra of the restricted up/down-Laplacians are bounded away from zero. Hence Proposition 5.3.4 proves that $0 \in \sigma(\mathcal{L}^k)$ is an isolated eigenvalue. □

Proposition 5.3.6. $\sigma(\mathcal{L}^k_+|_{\mathfrak{Z}^{k\perp}}) = \sigma(\mathcal{L}^{k-1}_-|_{\overline{\mathfrak{B}^{k-1}}})$.

Proof. Using the polar decomposition [100, Theorem VIII.32], we may write $\delta^k = U\sqrt{\mathcal{L}^k_+}$, where $U : C^k \to C^{k+1}$ is a partial isometry with initial space $\ker(\delta)^\perp$ and final space $\overline{\mathcal{R}(\delta^k)}$, and $\sqrt{\mathcal{L}^k_+}$ is the unique self-adjoint operator such that $\sqrt{\mathcal{L}^k_+} \circ \sqrt{\mathcal{L}^k_+} = \mathcal{L}^k_+$. We may now compute:

$$\begin{aligned}
\mathcal{L}^{k+1}_- &= \delta^k(\delta^k)^* \\
&= \left(U\sqrt{\mathcal{L}^k_+}\right)\left(\sqrt{\mathcal{L}^k_+}U^*\right) \\
&= U\mathcal{L}^k_+U^*.
\end{aligned}$$

U restricts to a unitary map $U : \mathfrak{Z}^{k\perp} \to \overline{\mathfrak{B}^{k+1}}$, so $\sigma(\mathcal{L}^k_+|_{\mathfrak{Z}^{k\perp}}) = \sigma(\mathcal{L}^{k-1}_-|_{\overline{\mathfrak{B}^{k-1}}})$ have the same spectra. □

Remark 5.3.7. It also follows from the same argument that $\mathcal{L}^k_+|_{\mathfrak{Z}^{k\perp}}$ and $\mathcal{L}^{k-1}_-|_{\overline{\mathfrak{B}^{k-1}}}$ have the same eigenvalues. Moreover, the eigenvectors may be related by the following chain of implications:

$$\begin{aligned}
x \in \mathfrak{Z}^{k\perp} \cap \mathrm{Dom}(\mathcal{L}^k_+) \text{ is an eigenvector of } \mathcal{L}^k_+ &\iff (\delta^k)^*\delta^k x = \lambda x \\
&\implies \delta^k(\delta^k)^*\delta^k x = \lambda\delta^k x \\
&\iff \mathcal{L}^{k+1}_-\delta^k x = \lambda\delta^k x.
\end{aligned}$$

These eigenvectors x and $\delta^k x$ have the same eigenvalue.

5.3.1 *Morphisms*

Let $\mathcal{F}$ and $\mathcal{G}$ be Hilbert sheaves on the same GAC $\mathcal{P}$. The existence of a spectrally well-behaved Hilbert sheaf morphism $\phi : \mathcal{F} \to \mathcal{G}$ enforces a relationship between the spectra of Hilbert sheaf Laplacian. We catalog a few results that relate the spectra $\sigma(\mathcal{L}^k_{\mathcal{F}})$ and $\sigma(\mathcal{L}^k_{\mathcal{G}})$ in the presence of morphisms with different properties.

Definition 5.3.8. Let $\phi^\bullet : (X^\bullet, V_X^\bullet, \delta_X^\bullet) \to (Y^\bullet, V_Y^\bullet, \delta_Y^\bullet)$ be a Hilbert complex morphism. $\phi^\bullet$ is a **bimorphism** if $\phi^\bullet$ is also a Hilbert complex morphism of the dual complexes $\phi^\bullet : (X_\bullet, V^*_{X,\bullet}, d_{X,\bullet}) \to (Y_\bullet, V^*_{Y,\bullet}, d_{Y,\bullet})$.

Bimorphisms constitute a class of Hilbert space morphisms that respect the sheaf Laplacian and its up/down components.

Lemma 5.3.9. *Suppose $\phi : \mathcal{F} \to \mathcal{G}$ is a Hilbert complex morphism whose induced Hilbert complex morphism $\phi^\bullet : (C^\bullet(\mathcal{P};\mathcal{F}), \delta^\bullet_{\mathcal{F}}) \to (C^\bullet(\mathcal{P};\mathcal{G}), \delta^\bullet_{\mathcal{G}})$ is a bimorphism. Then for each k, $\phi^k \mathcal{L}^k_{\mathcal{F},+} = \mathcal{L}^k_{\mathcal{F},+}\phi^k$, $\phi^k \mathcal{L}^k_{\mathcal{F},-} = \mathcal{L}^k_{\mathcal{F},-}\phi^k$, and $\phi^k\mathcal{L}^k_{\mathcal{F}} = \mathcal{L}^k_{\mathcal{F}}\phi^k$.*

Proof. Using the bimorphism property of $\phi^\bullet$, we may check:

$$\begin{aligned}\phi^k \mathcal{L}^k_{\mathcal{F},+} &= (\delta^k_{\mathcal{G}})^* \phi^{k+1} \delta^k_{\mathcal{F}} \\ &= \mathcal{L}^k_{\mathcal{G},+}\phi^k.\end{aligned}$$

Similarly, $\phi^k \mathcal{L}^k_{\mathcal{F},-} = \mathcal{L}^k_{\mathcal{G},-}\phi^k$. By linearity, $\phi^k\mathcal{L}^k_{\mathcal{F}} = \mathcal{L}^k_{\mathcal{G}}\phi^k$. □

Proposition 5.3.10. *Let $\phi : \mathcal{F} \to \mathcal{G}$ be a Hilbert complex morphism whose induced chain map $\phi^\bullet$ is a bimorphism. If ϕ^k is an isometry, then $\sigma(\mathcal{L}^k_{\mathcal{F}}) \subseteq \sigma(\mathcal{L}^k_{\mathcal{G}})$. If ϕ^k is a co-isometry, then $\sigma(\mathcal{L}^k_{\mathcal{G}}) \subseteq \sigma(\mathcal{L}^k_{\mathcal{F}})$.*

Proof. As a bimorphism, we have $\phi^k\mathcal{L}^k_{\mathcal{F}} = \mathcal{L}^k_{\mathcal{G}}\phi^k$. If ϕ^k is an isometry, precomposing with $(\phi^k)^*$ yields $\mathcal{L}^k_{\mathcal{F}} = (\phi^k)^*\mathcal{L}^k_{\mathcal{G}}\phi^k$. The closed image $M := \mathcal{R}(\phi^k)$ is an invariant subspace of $\mathcal{L}^k_{\mathcal{G}}$, in the sense that if $\mathbf{x} \in \mathcal{R}(\phi^k) \cap \mathrm{Dom}(\mathcal{L}^k_{\mathcal{G}})$, then $\mathcal{L}^k_{\mathcal{G}}\mathbf{x} \in \mathcal{R}(\phi^k)$. It follows that $(\phi^k)^*\mathcal{L}^k_{\mathcal{G}}\phi^k = (\phi^k)^*\left(\mathcal{L}^k_{\mathcal{G}}|_M\right)\phi^k$. The map ϕ^k is unitary onto its image, and $\sigma(\mathcal{L}^k_{\mathcal{F}}) = \sigma(\mathcal{L}^k_{\mathcal{G}}|_M) \subseteq \sigma(\mathcal{L}^k_{\mathcal{G}})$. The proof for a co-isometric ϕ^k is similar. □

Corollary 5.3.11. *Let $\phi : \mathcal{F} \to \mathcal{G}$ be a Hilbert complex morphism whose induced chain map $\phi^\bullet$ is a bimorphism. If ϕ^k is unitary, then $\sigma(\mathcal{L}^k_{\mathcal{F}}) = \sigma(\mathcal{L}^k_{\mathcal{G}})$.*

5.3.2 *Eigenvalues*

Hilbert sheaf morphisms between Hilbert sheaves also offer control over the eigenvalues of the sheaf Laplacians. The control offered is especially meaningful when the Hilbert sheaf Laplacian has discrete spectrum consisting of primarily eigenvalues.

Example 5.3.12. Let $S^1 := \{x \in \mathbb{R}^2 : \|x\| = 1\}$ denote the unit circle in $\mathbb{R}^2$, and consider the Hilbert sheaf

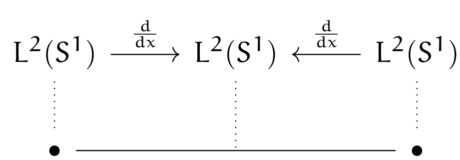

with Lebesgue measure. This is a Hilbert sheaf by Theorem 4.4.17. A similar computation to Example 6.1.8 shows the Laplacian $\mathcal{L} : L^2(S^1) \oplus L^2(S^1) \to L^2(S^1) \oplus L^2(S^1)$ is given by

$$\mathcal{L}\begin{pmatrix} f \\ g \end{pmatrix} = \begin{pmatrix} \frac{d^2}{dx^2}(g-f) \\ \frac{d^2}{dx^2}(f-g) \end{pmatrix}$$

where $\frac{d^2}{dx^2}$ denotes weak differentiation on $L^2(S^1)$. The domain of $\mathcal{L}$ is the Sobolev space $H^2(S^1)$. Under the unitary change of variables $(s,t)^T = (\frac{f+g}{\sqrt{2}}, \frac{g-f}{\sqrt{2}})^T$, the Laplacian $\mathcal{L}$ takes the form

$$\mathcal{L} = \begin{bmatrix} 0 & 0 \\ 0 & -2\frac{d^2}{dx^2} \end{bmatrix}.$$

We may now analyze the spectrum of $\mathcal{L}$ component-wise. The first diagonal block $0 : L^2(S^1) \to L^2(S^1)$ has spectrum $\sigma(0) = \{0\}$: an eigenvalue of infinite multiplicity. The second diagonal block $-2\frac{d^2}{dx^2} : L^2(S^1) \to L^2(S^1)$ is a scaling of the usual Laplacian on the unit circle, and has spectrum of all eigenvalues $\sigma(-2\frac{d^2}{dx^2}) = \{2k^2 : k \geqslant 0\}$, where 0 has multiplicity 1, and all other eigenvalues have multiplicity 2. Therefore, the spectrum of $\mathcal{L}$ is given by

$$\sigma(\mathcal{L}) = \{0\} \cup \{2k^2 \,:\, k \geqslant 1\}.$$

$\sigma_{\text{ess}}(\mathcal{L}) = \{0\}$, and each positive value is an eigenvalue of multiplicity 2.

Remark 5.3.13. More generally, consider a network sheaf $\mathcal{F}$ of differential operators as in Theorem 4.4.17. The Hilbert sheaf Laplacian $\mathcal{L}$ is a (weak) differential operator on the direct sum bundle of the vertex stalks. If the manifold underlying the smooth vector bundle on each vertex stalk is compact, and the Laplacian $\mathcal{L}$ is elliptic on $\ker(\mathcal{L})^\perp$, Then the spectrum $\sigma(\mathcal{L})$ will have the form

$$\sigma(\mathcal{L}) = \{0\} \cup \{\lambda_k \,:\, k \geqslant 1\}$$

where $\sigma_{\text{ess}}(\mathcal{L}) = 0$, and each λ_k is an isolated positive eigenvalue of finite multiplicity.

Notation 5.3.14. Let $A : X \to X$ be a Hilbert space operator. Let $M(A, \lambda)$ denote the (cardinal) multiplicity of $\lambda \in \mathbb{C}$ as an eigenvalue in the spectrum $\sigma(A)$.

Proposition 5.3.15. *Let $\phi : \mathcal{F} \to \mathcal{G}$ be a Hilbert sheaf morphism with induced chain map $\phi^\bullet : C^\bullet(\mathcal{P};\mathcal{F}) \to C^\bullet(\mathcal{P};\mathcal{G})$. If ϕ^k is an isometry for each k, then $M(\mathcal{L}^k_{\mathcal{F},+}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{G},+}, 0)$. If $\phi^\bullet$ is an isometric bimorphism, then additionally $M(\mathcal{L}^k_{\mathcal{F},-}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{G},-}, 0) M(\mathcal{L}^k_{\mathcal{F}}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{G}}, 0)$. Similarly, if ϕ^k is a co-isometry for each k, then $M(\mathcal{L}^k_{\mathcal{G},+}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{F},+}, 0)$. If ϕ^k is a co-isometric bimorphism, then additionally $M(\mathcal{L}^k_{\mathcal{G},-}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{F},-}, 0)$, and $M(\mathcal{L}^k_{\mathcal{G}}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{F}}, 0)$.*

Proof. ϕ^k maps cochains in $\ker(\mathcal{L}^k_{\mathcal{F},+})$ into $\ker(\mathcal{L}^k_{\mathcal{G},+})$. If ϕ^k is an isometry, then $\dim\ker(\mathcal{L}^k_{\mathcal{F},+})$ is bounded above by $\dim\ker(\mathcal{L}^k_{\mathcal{G},+})$, and $M(\mathcal{L}^k_{\mathcal{F},+}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{G},+}, 0)$. When $\phi^\bullet$ is a bimorphism, ϕ^k must map the space of harmonic cochains $\mathfrak{H}^k_{\mathcal{F}}$ into $\mathfrak{H}^k_{Y}$. When ϕ^k is an isometry, it follows that $M(\mathcal{L}^k_{\mathcal{F}}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{G}}, 0)$. The argument for co-isometries is similar. □

Remark 5.3.16. If ϕ is merely an injection on the k^{th} grade, then $M(\mathcal{L}^k_{\mathcal{F},+}, 0) \leqslant M(\mathcal{L}^k_{\mathcal{G},+}, 0)$. If ϕ is merely a surjection on the k^{th} grade, then $M(\mathcal{L}^k_{\mathcal{F},+}, 0) \geqslant M(\mathcal{L}^k_{\mathcal{G},+}, 0)$

This result is not, in general, useful for understanding the Hilbert sheaf Laplacian. In general, unless the images of restriction maps intersect in finite dimensional subspaces, the multiplicity of zero as an eigenvalue of the sheaf Laplacian will be infinite. A more useful theorem for the analysis of eigenvalues is the Courant-Fischer theorem [106, Theorem 12.1].

Theorem 5.3.17 (Courant-Fischer theorem)**.** *Let $A : X \to X$ be a positive operator on a Hilbert space X. The eigenvalues of A (counted with multiplicity) lying below the essential spectrum $\sigma_{ess}(A)$ may be enumerated in increasing order $\lambda_1 \leqslant \lambda_2 \leqslant \cdots$. The j^{th} eigenvalue λ_j is given by the formula*

$$\lambda_j = \inf_{\substack{V \subseteq X \\ \dim(V)=j}} \sup_{\substack{x \in V \cap \mathrm{Dom}(A) \\ x \neq 0}} \frac{\langle Ax, x\rangle}{\|x\|^2}.$$

Remark 5.3.18. For any compact interval $I \subseteq \mathbb{R}$ lying below the essential spectrum $\sigma_{\text{ess}}(A)$, there can be at most finitely many eigenvalues of A lying in I.

Proposition 5.3.19. *Suppose $\phi : \mathcal{F} \to \mathcal{G}$ is a Hilbert sheaf morphism whose induced chain map $\phi^\bullet$ is an isometric bimorphism. Let $\tilde{\mathcal{L}}^k_{\mathcal{F}}$ and $\tilde{\mathcal{L}}^k_{\mathcal{G}}$ denote the restrictions of $\mathcal{L}^k_{\mathcal{F}}$ and $\mathcal{L}^k_{\mathcal{G}}$ to the orthogonal complements $\ker(\mathcal{L}^k_{\mathcal{F}})^\perp$ and $\ker(\mathcal{L}^k_{\mathcal{G}})^\perp$ respectively. Order the eigenvalues $\{\lambda_j\}$ of $\tilde{\mathcal{L}}^k_{\mathcal{F}}$ and $\{\mu_j\}$ of $\tilde{\mathcal{L}}^k_{\mathcal{G}}$, below the essential spectra, in increasing order. The following hold.*

(i) *$\{\lambda_j\}$ and $\{\mu_j\}$ can be put in one-to-one correspondence.*

(ii) $\mu_j \leqslant \lambda_j$.

(iii) *If $\phi^k(\ker(\mathcal{L}^k_{\mathcal{F}})^\perp)$ has finite codimension $r < \infty$ in $\ker(\mathcal{L}^k_{\mathcal{G}})^\perp$, then $\mu_{j+r} \geqslant \lambda_j$.*

Proof. Let $\tilde{\mathcal{L}}^k_{\mathcal{F}}$ and $\tilde{\mathcal{L}}^k_{\mathcal{G}}$ denote the restrictions of $\mathcal{L}^k_{\mathcal{F}}$ and $\mathcal{L}^k_{\mathcal{G}}$ to the orthogonal complements $\ker(\mathcal{L}^k_{\mathcal{F}})^\perp$ and $\ker(\mathcal{L}^k_{\mathcal{G}})^\perp$ respectively. Enumerate the eigenvalues of $\tilde{\mathcal{L}}^k_{\mathcal{F}}$ and $\tilde{\mathcal{L}}^k_{\mathcal{G}}$ below the essential spectra in increasing order by $\{\lambda_j\}_j$ and $\{\mu_i\}_i$ respectively.

As a bimorphism, ϕ^k and $(\phi^k)^*$ commutes with each grade of the Laplacian: $\phi^k \mathcal{L}^k_{\mathcal{F}} = \mathcal{L}^k_{\mathcal{G}} \phi^k$ and $\mathcal{L}^k \mathcal{F}(\phi^k)^* = (\phi^k)^* \mathcal{L}^k_{\mathcal{G}}$. For $\mathbf{y} \in \ker(\mathcal{L}^k_{\mathcal{G}})$ we have $(\phi^k)^* \mathbf{y} \in \ker(\mathcal{L}^k_{\mathcal{F}})$. It follows that for $\mathbf{x} \in \mathrm{Dom}(\tilde{\mathcal{L}}^k_{\mathcal{F}})$ and $\mathbf{y} \in \ker(\mathcal{L}^k_{\mathcal{G}})$, that $\phi^k(\mathbf{x}) \perp \mathbf{y}$, and hence $\phi^k(\mathrm{Dom}(\tilde{\mathcal{L}}^k_{\mathcal{F}})) \subseteq \mathrm{Dom}(\tilde{\mathcal{L}}^k_{\mathcal{G}})$. $M :=$ $\phi^k(\ker(\mathcal{L}^k_{\mathcal{F}})^\perp)$ is a stable subspace of $\tilde{\mathcal{L}}^k_{\mathcal{G}}$.

Using injectivity, each eigenvector ψ of $\mathcal{L}^k_{\mathcal{F}}$ is mapped to a unique eigenvector $\phi^k(\psi)$ of $c\tilde{a}lL^k_{\mathcal{G}}$. Therefore the eigenvalues $\{\lambda_j\}_j$ and $\{\mu_j\}_j$ can be placed in one-to-one correspondence. Further, since ϕ^k is unitary onto its image, the eigenvalues of $\tilde{\mathcal{L}}^k_{\mathcal{F}}$ are exactly the eigenvalues of $\tilde{\mathcal{L}}^k_{\mathcal{G}}|_M$. It follows from the Courant-Fischer theorem that $\mu_j \leqslant \lambda_j$, for all j. In particular, every subspace $V \subseteq \ker(\mathcal{L}^k_{\mathcal{F}})^\perp$ of dimension j is mapped by ϕ^k to a subspace $\phi^k(V) \subseteq \ker(\mathcal{L}^k_{\mathcal{G}})^\perp$ of dimension j. Further, $V \cap \mathrm{Dom}(\mathcal{L}^k_{\mathcal{F}})$ is mapped by ϕ^k into $\phi^k(V) \cap \mathrm{Dom}(\mathcal{L}^k_{\mathcal{G}})$. There is an equality

$$\sup_{\substack{x \in V \cap \mathrm{Dom}(\mathcal{L}^k_F) \\ x \neq 0}} \frac{\langle \mathcal{L}^k_{\mathcal{F}} x, x \rangle}{\|x\|^2} = \sup_{\substack{y \in \phi^k(V) \cap \mathrm{Dom}(\mathcal{L}^k_G) \\ y \neq 0}} \frac{\langle \mathcal{L}^k_{\mathcal{G}} y, y \rangle}{\|y\|^2}.$$

It immediately follows that $\mu_j \leqslant \lambda_j$.

Finally, suppose that M has finite codimension $r < 0$ in $\ker(\mathcal{L}^k_{\mathcal{G}})^\perp$. It directly follows from the min-max formulation that if $\lambda_j = \mu_j$, there can be at most r eigenvalues μ_k between λ_j and λ_{j+1}. That is, $\mu_{j+r} \geqslant \lambda_j$ for all j. □

5.3.3 *Sheaf operations*

Several of the sheaf operations (Section 4.5) interact with the Laplacian spectra in controlled ways. The direct sum is the most straightforward.

Proposition 5.3.20. *Let $\mathcal{F}, \mathcal{G} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be Hilbert sheaves. The sheaf Laplacian of the direct sum $\mathcal{F} \oplus \mathcal{G}$ satisfies*

$$\sigma\left(\mathcal{L}^k_{\mathcal{F}\oplus\mathcal{G}}\right) = \sigma\left(\mathcal{L}^k_{\mathcal{F}}\right) \cup \sigma\left(\mathcal{L}^k_{\mathcal{G}}\right).$$

Proof. As observed in Proposition 4.5.3, the coboundary operator of $\mathcal{F} \oplus \mathcal{G}$ decomposes as

$$\mathring{\delta}^k_{\mathcal{F}\oplus\mathcal{G}} = \mathrm{diag}(\mathring{\delta}^k_{\mathcal{F}}, \mathring{\delta}^k_{\mathcal{G}}) : C^k(\mathcal{P};\mathcal{F}) \oplus C^k(\mathcal{P};\mathcal{G}) \longrightarrow C^{k+1}(\mathcal{P};\mathcal{F}) \oplus C^{k+1}(\mathcal{P};\mathcal{G}).$$

This diagonal operator acts on $C^k(\mathcal{P};\mathcal{F})$ and $C^k(\mathcal{P};\mathcal{G})$ independently, so $\delta^k_{\mathcal{F}\oplus\mathcal{G}} = \mathrm{diag}(\delta^k_{\mathcal{F}}, \delta^k_{\mathcal{G}})$ and $\mathcal{L}^k_{\mathcal{F}\oplus\mathcal{G}} = \mathrm{diag}(\mathcal{L}^k_{\mathcal{F}}, \mathcal{L}^k_{\mathcal{G}})$. Therefore $\sigma\left(\mathcal{L}^k_{\mathcal{F}\oplus\mathcal{G}}\right) = \sigma\left(\mathcal{L}^k_{\mathcal{F}}\right) \cup \sigma\left(\mathcal{L}^k_{\mathcal{G}}\right)$. □

Pullbacks by covering maps also have well-behaved spectral properties for Bounded Hilbert sheaves.

Proposition 5.3.21. *Let $\phi : \mathcal{P} \to \mathcal{Q}$ be a covering morphism of GACs, and $\mathcal{F} : \mathcal{Q} \to \mathbf{Hilb}_{\Bbbk}$ a bounded Hilbert sheaf. There is an inclusion of spectra $\sigma\left(\mathcal{L}^k_{\mathcal{F}}\right) \subseteq \sigma\left(\mathcal{L}^k_{\phi^*\mathcal{F}}\right)$.*

Proof. The map ϕ induces a cochain map $(\Phi^*)^\bullet : C^\bullet(\mathcal{Q};\mathcal{F}) \to C^\bullet(\mathcal{P};\phi^*\mathcal{F})$ defined by

$$\left((\Phi^*)^k(\mathbf{x})\right)_\tau = \mathbf{x}_{\phi(\tau)}$$

for each $\tau \in C^k(\mathcal{Q};\mathcal{F})$. We first confirm that Φ^* is a bimorphism of Hilbert complexes. By the definition of a covering map, one may check that for each $\tau \in C^{k+1}(\mathcal{P};\phi^*\mathcal{F})$ and $\mathbf{x} \in C^k(\mathcal{Q};\mathcal{F})$,

$$\begin{aligned}
\left(\delta^k_{\phi^*\mathcal{F}}(\Phi^*)^k\mathbf{x}\right)_\tau &= \sum_{\substack{\sigma' \lhd_1 \tau \\ g:\sigma'\to\tau}} \epsilon_{\phi^*\mathcal{F}}(g)\mathcal{F}_{\phi(g)}(\mathbf{x}_{\phi(\sigma')}) \\
&= \sum_{\substack{\sigma \lhd_1 \phi(\tau) \\ f:\sigma\to\phi(\tau)}} \epsilon_{\mathcal{F}}(f)\mathcal{F}_f(\mathbf{x}_\sigma) \\
&= \left((\Phi^*)^{k+1}\delta^k_{\mathcal{F}}\mathbf{x}\right)_\tau,
\end{aligned}$$

which demonstrates that $\delta^k_{\phi^*\mathcal{F}}(\Phi^*)^k = (\Phi^*)^{k+1}\delta^k_{\mathcal{F}}$. A similar computation utilizing the pre-image property of covering maps shows that $(\delta^k_{\phi^*\mathcal{F}})^*(\Phi^*)^{k+1} = (\Phi^*)^k(\delta^k_{\mathcal{F}})^*$, making Φ^* a Hilbert space bimorphism.

Next, observe that $(\Phi^*)^k$ is a scaled isometry; identifying $C^k(\mathcal{P};\phi^*\mathcal{F}) \cong \bigoplus^n C^k(\mathcal{Q};\mathcal{F})$, the map $(\Phi^*)^k(\mathbf{x}) = \bigoplus^n \mathbf{x}$. It directly follows that $\frac{1}{\sqrt{n}}(\Phi^*)^k : C^k(\mathcal{Q};\mathcal{F}) \to C^k(\mathcal{P};\phi^*\mathcal{F})$ defines an isometric bimorphism. By Proposition 5.3.10, $\sigma\left(\mathcal{L}^k_{\mathcal{F}}\right) \subseteq \sigma\left(\mathcal{L}^k_{\phi^*\mathcal{F}}\right)$. □

Part III

DYNAMICS

6

LINEAR DYNAMICS

This chapter investigates dynamical systems on cellular sheaves valued in Hilbert spaces, focusing on heat flow and wave propagation. The sheaf Laplacian introduced in Chapter 5 naturally generates a heat equation and a wave equation on cochain spaces, whose solutions exhibit asymptotic behavior controlled by the spectral properties of the Laplacian. Beyond the classical heat flow and wave propagation on cochains, we explore two dynamical systems unique to the sheaf-theoretic setting: relative heat flows that solve harmonic extension problems dynamically, and restriction map diffusion that evolves the sheaf structure itself toward a collection of restriction maps admitting a prescribed global section. These linear dynamics provide distributed algorithms for solving networked systems of equations.

Section 6.1 establishes that the negative sheaf Laplacian $-\mathcal{L}^k$ generates a contraction semigroup via the Lumer-Phillips theorem, yielding well-posed heat flows for all initial cochains. The long-term behavior of these flows is characterized through spectral decomposition: heat flows converge to orthogonal projections onto harmonic cochains, with convergence in the strong operator topology for general Hilbert sheaves and in operator norm for closed sheaves with spectral gaps. Section 6.1.2 adapts the heat flow to solve harmonic extension problems dynamically. Given boundary data on a subcomplex, we construct relative heat flows that either converge to harmonic extensions when they exist, or diverge when the problem is ill-posed. Section 6.2 introduces a dynamical system that evolves restriction maps rather than cochains. For network sheaves with fixed stalks, we derive gradient flows on spaces of bounded operators. The analysis distinguishes between general Banach space settings and the geometric case of Hilbert-Schmidt operators, the latter of which can be identified as a heat flow of a different sheaf Laplacian. Section 6.2.2 examines joint dynamics where cochains and restriction maps evolve simultaneously as a coupled system. The chapter concludes with an analysis of the wave equation on a Hilbert sheaf. Unlike heat flow, wave propagation does not converge, instead exhibiting oscillatory behavior, the time-average of which is a harmonic cochain.

6.1 HEAT FLOW

Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf with associated Hilbert complex $\left(C^\bullet(\mathcal{P};\mathcal{F}), \delta^\bullet\right)$. Each grade of the sheaf Laplacian $\mathcal{L}^k$ is a positive operator (Theorem 3.2.21). Using the spectral theorem and the Lumer-Phillips theorem, one may show that $-\mathcal{L}^k$ is the infinitesimal generator of a contraction semigroup on $C^k(\mathcal{P};\mathcal{F})$. We prove this as a general lemma.

Lemma 6.1.1. *Let $A : X \to X$ be an operator on a separable Hilbert space. If A is negative, then A is the infinitesimal generator of a contraction semi-group.*

Proof. We present the proof for real Hilbert spaces, as the argument can be straightforwardly adapted to complex Hilbert spaces. Let (Ω, μ), $\Phi : X \to L^2(\Omega;\mu)$, and $f : \Omega \to \mathbb{R}$ satisfy the conditions of the spectral theorem for A. If M_f generates a contraction semigroup T on $L^2(\Omega, \mu; \mathbb{R})$, then $\Phi T \Phi^{-1}$ is a contraction semigroup on X with generator A. Hence it suffices to check that M_f satisfies the conditions of the Lumer-Phillips theorem.

M_f is self-adjoint, and hence is densely defined. Since $f(\omega) \leqslant 0$ for μ-almost every $\omega \in \Omega$, M_f is dissipative. Finally, every positive $\lambda \in \mathbb{R}$ is in the resolvent $\rho(M_f)$. Since M_f is closed, it follows that $M_f - \lambda I$ is a surjection for all $\lambda > 0$. Therefore M_f generates a contraction semigroup on $L^2(\Omega, \mu; \mathbb{R})$. □

Applying Lemma 6.1.1 to $-\mathcal{L}^k$ shows that $-\mathcal{L}^k$ is the infinitesimal generator of a contraction semigroup on $C^k(\mathcal{P};\mathcal{F})$.

Definition 6.1.2. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$ be a Hilbert sheaf with sheaf Laplacian $\mathcal{L}^k$. The k^{th}-**heat semigroup** of $\mathcal{F}$ is the semigroup $\exp(-t\mathcal{L}^k)$ on $C^k(\mathcal{P};\mathcal{F})$. Given a k-cochain $\mathbf{x}_0 \in C^k(\mathcal{P};\mathcal{F})$, the **heat flow** of $\mathbf{x}_0$ is the path

$$\mathbf{x}_t = \exp(-t\mathcal{L}^k)\mathbf{x}_0.$$

As a C_0-semigroup, the heat flow $\mathbf{x}_t = \exp(-t\mathcal{L}^k)\mathbf{x}_0$ is a mild solution to initial value problem:

$$\begin{aligned} \dot{\mathbf{x}} &= -\mathcal{L}^k\mathbf{x}\,, \\ \mathbf{x}(0) &= \mathbf{x}_0\,. \end{aligned}$$

In particular, when $\mathbf{x}_0 \in \mathrm{Dom}(\mathcal{L}^k)$, the heat flow is a classical solution to the initial value problem.

6.1.1 *Long term behavior of the heat flow*

Since the heat semigroups of a cellular sheaf $\mathcal{F}$ are generated by a negative operator, we get an additional property: controlled dynamics as $t \to \infty$. We start with a general lemma.

Lemma 6.1.3. *Let e^{tA} be a contraction semigroup on a separable Hilbert space X generated by a negative operator $A : X \to X$. Let $P : X \to X$ denote the orthogonal projection onto the kernel $\ker(A)$. As $t \to \infty$, the semigroup e^{tA} converges to P in the strong operator topology.*

Proof. Let (Ω, μ), $\Phi : X \to L^2(\Omega, \mu; \mathbb{R})$ and $f : \Omega \to \mathbb{R}$ satisfy the conditions of the spectral theorem for A. Notice that $T_0(t) := M_{\exp(tf)}$ defines a contraction semigroup on $L^2(\Omega, \mu; \mathbb{R})$ with generator M_f. For any $k \in \ker(M_f)$, the support $\operatorname{supp}(k)$ must be contained in the set $\{\omega \in \Omega : f(\omega) = 0\}$, up to a set of a measure zero. It follows that $M_{\exp(tf)}k = k$ for all $t \geqslant 0$. Now for any $g \in \ker(M_f)^\perp \subseteq L^2(\Omega, \mu; \mathbb{R})$, the function $G_t = \exp(2tf)g^2$ converges pointwise to 0 almost everywhere as $t \to \infty$ since the essential range of f is non-positive. Moreover, G_t is dominated by g^2, which is integrable. The dominated convergence theorem now ensures that

$$\lim_{t\to\infty} \|M_{\exp(tf)}g\|^2 = \lim_{t\to\infty} \int G_t \, d\mu = 0.$$

Therefore $M_{\exp(tf)}$ converges to the orthogonal projection onto the kernel of M_f in the strong operator topology. Identifying $e^{tA} = \Phi^{-1} M_{\exp(tf)} \Phi$ proves that $e^{tA} \to P$ in the strong operator topology as well. □

Lemma 6.1.4. *The result of Proposition Lemma 6.1.3 holds when X is not separable as well.*

Proof. Suppose X is not separable. We use Zorn's lemma to prove that X can be decomposed into a (possibly uncountable) direct sum $X = \bigoplus_{i\in I} W_i$ where each W_i is separable and A is stable on each W_i.

Let $\mathcal{D}$ denote the collection of all decompositions $X = (\bigoplus_{i\in I} W_i) \oplus Z$, that satisfies the following conditions.

1. Each W_i is non-empty and separable.
2. A is invariant on Z and on each W_i.
3. Either $Z = 0$ or Z is not separable.

We may partially order $\mathcal{D}$ according to the following rule:

$$\left[\left(\bigoplus_{i\in I} W_i\right) \oplus Z \leqslant \left(\bigoplus_{j\in J} W'_j\right) \oplus Z'\right] \iff \left[\exists \text{ an injective map } f : I \hookrightarrow J \text{ s.t. } W_i = W'_{f(i)}\right].$$

That is, the J-decomposition further peels off more separable Hilbert spaces from Z. Since every chain in $\mathcal{D}$ is bounded, we may use Zorn's lemma to take a maximal decomposition in $\mathcal{D}$.

Now suppose $(\bigoplus_{i\in I} W_i) \oplus Z$ is a decomposition and $Z \neq 0$. Take $v \in Z$ to be any non-zero vector, and $\lambda \in \rho(A)$ be any value in the resolvent set with corresponding operator $R_\lambda := (A - \lambda I)^{-1}$. Let $[v] := \{v, R_\lambda v, R_\lambda^2 v, \ldots\}$ denote the orbit of v under R_λ, and set

$$Z_v := \overline{\operatorname{span}[v]}.$$

The space Z_v is closed, A-invariant, and has a countable basis given by $[v]$, making it an invariant separable sub-Hilbert space of Z. Moreover $\dim(Z_v) \geqslant 1$. Consequently,

$$\left(\bigoplus_{i\in I} W_i\right) \oplus Z \leqslant \left(\bigoplus_{i\in I} W_i\right) \oplus Z_v \oplus (Z_v^\perp \cap Z).$$

Thus if $Z \neq 0$, the decomposition cannot be maximal. Therefore the maximal decomposition from Zorn's lemma is a decomposition of X into separable, stable, subspaces.

Applying the separable case on each component proves the general result. Let $X = \bigoplus_{i\in I} W_i$ be a decomposition of X into separable sub-Hilbert spaces such that $A(W_i) \subseteq W_i$ for each $i \in I$. For $x \in X$, there is a countable sub-index set $J \subseteq I$ such that $x_0 = \sum_{j\in J}(x_0)_j$, where each $(x_0)_j$ is the W_j component of x_0. By the separable case, each $e^{tA}(x_0)_j \to 0$ as $t \to \infty$. It follows that $e^{tA}x_0 \to 0$ as well. □

We now return to the Hilbert sheaf Laplacian $\mathcal{L}^k : C^k(\mathcal{P};\mathcal{F}) \to C^k(\mathcal{P};\mathcal{F})$. The application of Lemma 6.1.4 to the k^{th} heat semigroup yields the following theorem.

Theorem 6.1.5. *The k^{th} heat semigroup $\exp(-t\mathcal{L}^k)$ converges in the strong operator topology as $t \to \infty$ to the orthogonal projection operator $P : C^k(\mathcal{P};\mathcal{F}) \to C^k(\mathcal{P};\mathcal{F})$ onto the space $\mathfrak{H}^k$ of k-harmonic cochain. In particular, for all $\mathbf{x}_0 \in C^k(\mathcal{P};\mathcal{F})$, the heat flow $\mathbf{x}_t$ converges to the nearest harmonic cochain as $t \to \infty$.*

Corollary 6.1.6. *A heat flow in $C^0(\mathcal{P};\mathcal{F})$ with initial value $\mathbf{x}_0$ converges to the nearest global section to $\mathbf{x}_0$.*

Remark 6.1.7. When the Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is not proper, there is no guarantee that the kernel of the block operator $\mathring{\delta}^k$ is closed; hence $\ker(\mathring{\delta}^k)$ may be merely a dense linear subspace of $\ker(\delta^k) = \mathfrak{H}^k$, and the limit $\lim_{t\to\infty} \mathbf{x}_t$ of a heat flow on $C^k(\mathcal{P};\mathcal{F})$ may not be in the kernel of the block operator $\mathring{\delta}^k$. This represents a genuine distinction from the finite dimensional theory of dynamics on weighted sheaves. However, when $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_\Bbbk$ is a bounded sheaf, convergence to a point in the kernel of the block operator $\mathring{\delta}^k$ is guaranteed.

Example 6.1.8. We return again to the simple Hilbert sheaf

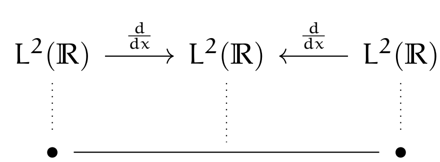

from Example 4.4.5 and Example 5.1.4. The heat flow $(f_t, g_t)^T \in L^2(\mathbb{R}) \oplus L^2(\mathbb{R})$ with initial value $(f_0, g_0)^T$ obeys the differential equation

$$\begin{pmatrix} \dot{f} \\ \dot{g} \end{pmatrix} = \begin{pmatrix} \frac{d^2}{dx^2}(f-g) \\ \frac{d^2}{dx^2}(g-f) \end{pmatrix}.$$

Making the change of variables

$$\begin{aligned} u_t &= f_t + g_t \\ v_t &= f_t - g_t, \end{aligned}$$

we may compute $\dot{u} = 0$ and $\dot{v} = -2\frac{d^2}{dx^2}v$. That is, v_t obeys the heat equation on $\mathbb{R}$, and dissipates to 0 as $t \to \infty$. Hence taking limits, we see that $f_\infty + g_\infty = f_0 + g_0$ and $f_\infty - g_\infty = 0$. It follows that $f_\infty = g_\infty = \frac{f_0+g_0}{2}$.

Example 6.1.9. Let M be a compact real Riemannian manifold with canonical volume form μ. Let $L^2(M;\mu)$ denote the space of square integrable real-valued functions on M. The **Laplace-Beltrami operator** Δ^{LB} is the operator $\Delta = d^*d$, where $d : L^2(M;\mu) \to L^2(T^*M)$ is the gradient operator. The domain of Δ^{LB} is taken to be the maximal domain. We get a Hilbert sheaf

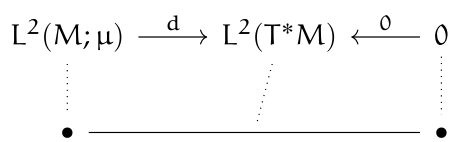

with coboundary operator $\delta = \begin{bmatrix} -d & 0 \end{bmatrix}$. This is sheaf is proper, with corresponding sheaf Laplacian

$$\mathcal{L} = \begin{bmatrix} -\Delta^{LB} & 0 \\ 0 & 0 \end{bmatrix}.$$

The induced heat flow $f_t = \exp(t\Delta^{LB})$ is the usual Laplace-Beltrami based heat flow on $f \in L^2(M;\mu)$, and f_t converges as $t \to \infty$ to the nearest harmonic function in $L^2(M;\mu)$. Hence the usual Laplace-Beltrami heat flow may be recognized as a special case of heat flow on a Hilbert sheaf.

Example 6.1.10. Consider the Hilbert sheaf

$$
\begin{array}{ccccc}
L^2([0,1]) & \xrightarrow{\begin{bmatrix} \frac{d}{dx} \\ I \end{bmatrix}} & L^2([0,1]) \oplus L^2([0,1]) & \xleftarrow{\begin{bmatrix} I \\ I \end{bmatrix}} & L^2([0,1]) \\
\vdots & & \vdots & & \vdots \\
\bullet & \text{———————} & \bullet & \text{———————} & \bullet
\end{array}
$$

with corresponding sheaf Laplacian

$$
\mathcal{L} = \begin{bmatrix} I - \frac{d^2}{dx^2} & \frac{d}{dx} - I \\ -\frac{d}{dx} - I & I \end{bmatrix}
$$

on domain $H^2([0,1]) \oplus H^1([0,1]) \subseteq L^2([0,1]) \oplus L^2([0,1])$. The induced heat flow $(f_t, g_t)^T$ converges to the nearest point $(f_\infty, g_\infty)^T$ such that $f'_\infty - g_\infty = 0$ and $f_\infty - g_\infty = 0$. That is, f_∞ must be a weak solution to the ODE $f'_\infty = f_\infty$, so $f_\infty(x) = Ce^x$ almost everywhere for some $C \in \mathbb{R}$.

Example 6.1.11. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite network and $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{0,\mathbb{k}}$ be a Hilbert sheaf of differential operators, as in Theorem 4.4.17. Let $B_\sigma \to M_\sigma$ denote the smooth vector bundle on each object $\sigma \in \mathcal{E} \cup \mathcal{V}$. The coboundary operator $\delta : C^0(\mathcal{G}; \mathcal{F}) \to C^1(\mathcal{G}; \mathcal{F})$ decomposes as a block operator $\delta = [\delta_{ev}]$ over vertex-stalks and edge-stalks; for v a bounding vertex of an edge e, the block $\delta_{ev} : \mathcal{F}(v) \to \mathcal{F}(e)$ is given by

$$
\delta_{ev} = \sum_{f : v \to e} \epsilon(f) \overline{\mathcal{F}_f} \,.
$$

This is a (weak) differential operator and the coboundary operator encodes a networked system of homogeneous linear differential equations. A global section $\mathbf{x}$ of $\mathcal{F}$ exactly corresponds to a solution to this networked system; for each edge e with distinct bounding vertices u, v, $\mathbf{x}_u$ and $\mathbf{x}_v$ are smooth sections of B_u and B_v respectively such that $\delta_{eu}\mathbf{x}_u = \delta_{ev}\mathbf{x}_v$. For an edge e with a unique bounding vertex v, we instead recover $\delta_{ev}\mathbf{x}_v = 0$. The heat flow of $\mathcal{F}$ exactly converges to such a solution to the networked differential equation.

The rate of convergence of heat flow is controlled by the spectrum of the sheaf Laplacian $\mathcal{L}^k$ through the following theorem.

Proposition 6.1.12. *Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$ be a Hilbert sheaf. Suppose either $0 \notin \sigma(\mathcal{L}^k)$, or 0 is an isolated eigenvalue in $\sigma(\mathcal{L}^k)$. The heat semigroup $\exp(-t\mathcal{L}^k)$ converges in the operator norm to the orthogonal projection operator $P : C^k(\mathcal{P}; \mathcal{F}) \to C^k(\mathcal{P}; \mathcal{F})$ onto $\mathfrak{H}^k$.*

Proof. Let $\tilde{\mathcal{L}}^k : \ker(\mathcal{L}^k)^\perp \to \ker(\mathcal{L}^k)^\perp$ denote the restriction of $\mathcal{L}^k$ to the orthogonal complement of its kernel. Since 0 is an isolated eigenvalue of $\mathcal{L}^k$, there is an $r > 0$ such that $\lambda > r$ for all $\lambda \in \sigma(\tilde{\mathcal{L}}^k)$. By the spectral theorem, we may bound $\|\exp(-t\tilde{\mathcal{L}}^k)\|_{op} \leqslant e^{-tr}$, from which it follows that $\exp(-t\tilde{\mathcal{L}}^k) \to 0$ in the operator norm. It directly follows that $\exp(-t\mathcal{L}^k)$ converges to P in the operator norm. □

Remark 6.1.13. In the preceding proof, the larger the value of $r > 0$, the faster the convergence of $\exp(-t\mathcal{L}^k) \xrightarrow{op} P$. Hence we see that the rate of convergence is controlled by $\inf \sigma(\tilde{\mathcal{L}}^k)$. When $\inf \sigma(\tilde{\mathcal{L}}^k) = 0$, the convergence is merely in the strong topology, and distance bounds for cochains cannot be given uniformly in norm.

The following corollary follows since closed Hilbert sheaves have spectral gaps (Corollary 5.3.5).

Corollary 6.1.14. *When $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ is a closed Hilbert sheaf, each heat semigroup $\exp(-t\mathcal{L}^k)$ converges in the operator norm to the orthogonal projection P^k onto k-harmonic cochains.*

6.1.2 *Relative heat flows and harmonic extension*

Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf, and $\mathcal{B}$ be a subcomplex of $\mathcal{P}$. Suppose the domain of the sheaf Laplacian $\mathcal{L}$ splits over the decomposition $C^k(\mathcal{P};\mathcal{F}) \cong C^k(\mathcal{P},\mathcal{B};\mathcal{F}) \oplus C^k(\mathcal{B};\mathcal{F})$, and write

$$\mathcal{L} = \begin{bmatrix} \mathcal{L}^k_{\mathcal{U},\mathcal{U}} & \mathcal{L}^k_{\mathcal{U},\mathcal{B}} \\ \mathcal{L}^k_{\mathcal{B},\mathcal{U}} & \mathcal{L}^k_{\mathcal{B},\mathcal{B}} \end{bmatrix}.$$

Given $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$, we may try to form a heat flow for $\mathbf{x} \in C^k(\mathcal{P},\mathcal{B};\mathcal{F})$ that converges to a harmonic extension of $\mathbf{y}$. The dynamics of $\mathbf{x}$ should obey

$$\dot{\mathbf{x}} = -\mathcal{L}^k_{\mathcal{U},\mathcal{U}}\mathbf{x} - \mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y}.$$

We may again use semigroup theory to tackle this problem. Note that $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}$ can be extended to a positive operator; it is positive symmetric and densely defined, and has a self-adjoint Friedrichs extension. Hence $-\mathcal{L}^k_{\mathcal{U},\mathcal{U}}$ is the infinitesimal generator of a contraction semigroup $\exp(-t\mathcal{L}^k_{\mathcal{U},\mathcal{U}})$ on $C^k(\mathcal{P},\mathcal{B};\mathcal{F})$. As observed in Section 5.2, $\mathbf{y}$ will have a harmonic extension exactly when $\mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y} \in \mathcal{R}(\mathcal{L}^k_{\mathcal{U},\mathcal{U}})$. When this is the case, take $\mathbf{w}$ such that $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}\mathbf{w} = \mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y}$, and define the **relative heat flow** of $\mathbf{x}_0 \in C^k(\mathcal{P},\mathcal{B};\mathcal{F})$ to be the path:

$$\mathbf{x}_t = \exp(-t\mathcal{L}^k_{\mathcal{U},\mathcal{U}})(\mathbf{x}_0 + \mathbf{w}) - \mathbf{w}. \tag{1}$$

One may check that this is a mild solution to the initial value problem:

$$\dot{\mathbf{x}} = -\mathcal{L}^k_{\mathcal{U},\mathcal{U}}\mathbf{x} - \mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y},$$
$$\mathbf{x}(0) = \mathbf{x}_0.$$

The path $\mathbf{x}_t$ is well-behaved asymptotically. By Lemma 6.1.4, $\exp(-t\mathcal{L}^k_{\mathcal{U},\mathcal{U}})$ converges in the strong operator topology to the orthogonal projection operator P onto the kernel $\ker(\mathcal{L}^k_{\mathcal{U},\mathcal{U}})$ as $t \to \infty$. It follows by some simple geometry that $\mathbf{x}_\infty = P(\mathbf{x}_0 + \mathbf{w}) - \mathbf{w}$ is the closest point to $\mathbf{x}_0$ in $C^k(\mathcal{P}, \mathcal{B}; \mathcal{F})$ that solves the harmonic extension problem for $\mathbf{y}$. We have proven the following theorem.

Theorem 6.1.15. *Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf, and $\mathcal{B}$ be a subcomplex of $\mathcal{P}$. Suppose the domain of the sheaf Laplacian $\mathcal{L}$ splits over the decomposition $C^k(\mathcal{P};\mathcal{F}) \cong C^k(\mathcal{P},\mathcal{B};\mathcal{F}) \oplus C^k(\mathcal{B};\mathcal{F})$. Let $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$ and $\mathbf{w} \in C^k(\mathcal{P},\mathcal{B};\mathcal{F})$ such that $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}\mathbf{w} = \mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y}$. For all $\mathbf{x}_0 \in C^k(\mathcal{P},\mathcal{B};\mathcal{F})$, the relative heat flow $\mathbf{x}_t = \exp(-t\mathcal{L}^k_{\mathcal{U},\mathcal{U}})(\mathbf{x}_0 + \mathbf{w}) - \mathbf{w}$ converges to the closest point $\mathbf{x}_\infty \in C^k(\mathcal{P},\mathcal{B};\mathcal{F})$ such that $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}\mathbf{x}_\infty + \mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y} = 0$.*

For an ill-posed harmonic extension problem, the relative heat flow diverges. For any $\mathbf{y} \in C^k(\mathcal{B};\mathcal{F})$, we may define a relative heat flow that acts as a mild solution to $\dot{\mathbf{x}} = -\mathcal{L}^k_{\mathcal{U},\mathcal{U}}\mathbf{x} - \mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y}$ using variation of constants. In particular, set $T(t) = \exp(-t\mathcal{L}^k_{\mathcal{U},\mathcal{U}})$ and $\mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y} = \mathbf{b}_{\ker} + \mathbf{b}_{\ker^\perp}$ where $\mathbf{b}_{\ker} \in \ker(\mathcal{L}^k_{\mathcal{U},\mathcal{U}})$ and $\mathbf{b}_{\ker^\perp} \in \ker(\mathcal{L}^k_{\mathcal{U},\mathcal{U}})^\perp$, and define the relative heat flow[1]

$$\begin{aligned}\mathbf{x}_t &:= T(t)\mathbf{x}_0 + \int_0^t T(t-s)\mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y}\, ds \\ &= T(t)\mathbf{x}_0 + \int_0^t T(t-s)\mathbf{b}_{\ker}\, ds + \int_0^t T(t-s)\mathbf{b}_{\ker^\perp}\, ds.\end{aligned}$$

We now analyze these three pieces separately. As $t \to \infty$, $T(t)\mathbf{x}_0$ converges to $P\mathbf{x}_0$, where P is the orthogonal projection onto the kernel of $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}$. Since $T(t)\mathbf{b}_{\ker} = \mathbf{b}_{\ker}$ for all $t \geqslant 0$, we may write $\int_0^t T(t-s)\mathbf{b}_{\ker}\, ds = t\mathbf{b}_{\ker}$. Finally, it can be seen through the spectral theorem that $\int_0^t T(t-s)\mathbf{b}_{\ker^\perp}\, ds$ either converges to a point (if $\mathbf{b}_{\ker^\perp} \in \mathcal{R}(\mathcal{L}^k_{\mathcal{U},\mathcal{U}})$, putting us in the previous case) or grows sub-linearly. In either case, the linear growth of the $t\mathbf{b}_{\ker}$ dominates, and the relative heat flow diverges to infinity.

The relative harmonic flow can be used to detect when a harmonic extension problem has a solution. Simply check if the relative heat flow converges; if it does, the limiting value is a solution to the harmonic extension problem. If the relative heat flow diverges, there is no solution.

1 When there is a $\mathbf{w} \in C^k(\mathcal{P},\mathcal{B};\mathcal{F})$ such that $\mathcal{L}^k_{\mathcal{U},\mathcal{U}}\mathbf{w} = \mathcal{L}^k_{\mathcal{U},\mathcal{B}}\mathbf{y}$, this expression for the relative heat flow reduces to equation (1).

We now turn our attention away from heat dynamics on spaces of cochains of a Hilbert sheaf, and consider evolving restriction maps according to a heat flow. We work over a network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with a given choice of vertex and edge stalks. Instead of driving a cochain $\mathbf{x}_0 \in C^0$ toward a global section, we designate a fixed cochain $\mathbf{x}$ and drive the restriction maps themselves toward a sheaf for which $\mathbf{x}_0$ is a global section.

Definition 6.2.1. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite multigraph (allowing self-loops). A **network Hilbert sheaf** on $\mathcal{G}$ is a Hilbert sheaf $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{0,\mathbb{k}}$, where $\mathcal{G}$ is viewed as a weakly regular cell structure.

For the entirety of this section, fix a finite multigraph $\mathcal{G}$, and a *function* $F : \mathrm{Ob}(\mathcal{G}) \to \mathrm{Ob}(\mathbf{Hilb}_{\mathbb{k}})$. We think of F as a choice of a stalk for each vertex and edge of $\mathcal{G}$, without a corresponding assignment of restriction maps. The function F may be extended (non-uniquely) to a network Hilbert sheaf $\mathcal{F}$ by specifying restriction maps. We let $\mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{G}; F)$ denote the set of all such bounded extensions. Since there are no commutativity constraints for restriction maps that must be satisfied on a network Hilbert sheaf, there is a one-to-one correspondence between sheaves in $\mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{G}; F)$ and a choice of a bounded linear operator for each morphism of $\mathcal{G}$.

Given a pair of Hilbert spaces X and Y, let $\mathcal{B}(X, Y) := \mathbf{Hilb}_{\mathbb{k}}(X, Y)$ denote the space of bounded linear operators $X \to Y$. The operator norm $\| - \|_{\mathrm{op}}$ gives $\mathcal{B}(X, Y)$ the structure of a Banach space. Consequently, we may identify $\mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{G}; F)$ with the Banach space

$$\mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{G}; F) \equiv \bigoplus_{f : v \to e} \mathcal{B}\big(F(v), F(e)\big) \subseteq \mathcal{B}\big(C^0(\mathcal{G}; F), C^1(\mathcal{G}; F)\big)$$

where the direct sum is taken over all non-identity morphisms in $\mathcal{G}$.

Fix a bounded network Hilbert sheaf $\mathcal{F} \in \mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{G}; F)$ and a zero-cochain $\mathbf{x} \in C^0(\mathcal{G}; F)$. Let e be an edge of $\mathcal{G}$, with incoming maps $F : v \to e$ and $g : u \to e$. Consider the evolution

$$\frac{d}{dt}\mathcal{F}_f = -(\mathcal{F}_f(x_v) - \mathcal{F}_g(x_u))x_v^*, \tag{2}$$

where $x_v^* := \langle -, x_v \rangle$ is the bounded linear functional on $F(v)$ induced by x_v. Applying these dynamics to each restriction map $\mathcal{F}_f$ yields a coupled system, and a corresponding first order autonomous linear dynamical system

$$\dot{\mathcal{F}} = \Psi(\mathcal{F}) \tag{3}$$

on the Banach space $\mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{G}; F)$. We call this system **restriction map diffusion**.

To analyze the long-term dynamics of restriction map diffusion, we use the following lemma.

Lemma 6.2.2. *Let X, Y be Hilbert spaces, $x \neq 0$ a point in X, and $A : \mathcal{B}(X, Y) \to \mathcal{B}(X, Y)$ the bounded linear map $A\phi := \phi(x)x^*$. The initial value problem*

$$\begin{aligned} \dot{\phi} &= -A\phi \\ \phi(0) &= \phi_0 \end{aligned}$$

has a unique solution ϕ_t in $\mathcal{B}(X, Y)$ that converges to a well defined limit $\phi_\infty = \lim_{t\to\infty} \phi_t$ which is the closest point to ϕ_0 in $\ker(A)$.

Proof. The map A has a non-trivial kernel consisting of exactly those $\phi \in \mathcal{B}(X, Y)$ that vanish on x_0. For any ϕ, we may compute $A^2\phi = \|x\|^2 A\phi$, witnessing that the range $\mathcal{R}(A)$ of A is an eigenspace of A with eigenvalue $\|x\|^2$. Writing

$$\phi = \left(\phi - \frac{1}{\|x\|^2} A\phi\right) + \frac{1}{\|x\|^2} A\phi$$

witnesses a direct sum decomposition $\mathcal{B}(X, Y) = \ker(A) \oplus \mathcal{R}(A)$. The spectrum of A is now given by $\sigma(A) = \{0, \|x\|^2\}$, which is real, non-negative, and has an isolate eigenvalue at 0. Chapter 2 section 2 of [34] now ensures the dynamical system

$$\begin{aligned} \dot{\phi} &= -A\phi \\ \phi(0) &= \phi_0 \end{aligned}$$

has a unique solution ϕ_t for each initial value ϕ_0, that converges to $\phi_\infty = \phi_0 - \frac{1}{\|x\|^2} A\phi_0$.

Consider the projection map $P : X \to \ker(A)$ by $P(\phi) = \phi - \frac{1}{\|x\|^2} A\phi$. This map has operator norm $\|P\|_{op} = 1$, so $P(\phi)$ is the nearest point to ϕ in $\ker(A)$ [114]. □

To use this lemma to analyze restriction map diffusion, notice that the pairs of restriction maps into each edge of $\mathcal{G}$ evolve independently from one another. Consequently, each edge-component δ_e of the coboundary map δ evolves independently. Let e be an edge of $\mathcal{G}$, with incoming covering maps $f : v \to e$ and $g : u \to e$, where it is possible that $u = v$. The coboundary component δ_e, viewed as a linear map $\delta_e : F(v) \oplus F(u) \to F(e)$ evolves in $\mathcal{B}\big(F(v) \oplus F(u)\,,\, F(e)\big)$ according to

$$\dot{\delta_e} = -\delta_e \begin{pmatrix} x_v \\ x_u \end{pmatrix} \begin{pmatrix} x_v \\ x_u \end{pmatrix}^* . \tag{4}$$

Notice that these dynamics on δ_e are equivalent to the dynamics on the restriction maps into e, as

$$\dot{\delta}_e = \begin{bmatrix} -\dot{\mathcal{F}}_f & \dot{\mathcal{F}}_g \end{bmatrix} = (\mathcal{F}_g x_u - \mathcal{F}_u x_v) \begin{bmatrix} x_v^* & x_u^* \end{bmatrix},$$

after possibly expanding the domain of δ_e to $F(u) \oplus F(u)$ when $u = v$. Applying Lemma 6.2.2 initialized at ${}^0\delta_e$ to this system yields a solution ${}^t\delta_e$, which converges to the nearest operator ${}^\infty\delta_e$ such that $(x_v, x_u)^T$ is in the kernel of δ_e. Taking a block-operator representation for ${}^t\delta_e$ gives evolutions ${}^t\mathcal{F}_f$ and ${}^t\mathcal{F}_g$.

Applying these dynamics to each edge simultaneously yields dynamics ${}^t\mathcal{F}_f$ for every restriction map f in $\mathcal{G}$; combining these paths in the Banach space $\mathbf{BHilbShv}_{\mathbb{k}}(\mathcal{G}; F)$ gives an evolution of sheaves ${}^t\mathcal{F}$ starting from the initial bounded Hilbert sheaf ${}^0\mathcal{F}$. The limiting sheaf ${}^\infty\mathcal{F} = \lim_{t\to\infty} {}^t\mathcal{F}$ is the closest sheaf to ${}^0\mathcal{F}$ for which $\mathbf{x}$ is a global section. We have proved the following theorem.

Theorem 6.2.3. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a network and $F : Ob(\mathcal{G}) \to Ob(\mathbf{Hilb}_{\mathbb{k}})$ be a choice of a Hilbert space for each vertex and edge. Let ${}^0\mathcal{F}$ be an initial choice of a bounded Hilbert sheaf extending F, and $\mathbf{x} \in C^0(\mathcal{G}; F)$ a fixed choice of a zero-cochain. Restriction map diffusion on $\mathbf{BHilbShv}(\mathcal{G}; F)$ has a solution ${}^t\mathcal{F}$, which converges as $t \to \infty$ to the nearest network sheaf of Hilbert spaces for which $\mathbf{x}$ is a global section.*

6.2.1 *Hilbert-Schmidt restriction map diffusion*

Under additional hypotheses on the admissible restriction maps, we may view these dynamics as taking place in a Hilbert space and converging to an orthogonal projection in sheaf-space. Say that a Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{\mathbb{k}}$ is a **Hilbert-Schmidt sheaf** if each restriction map is Hilbert-Schmidt. For such a Hilbert sheaf, all coboundary operators $\delta^\bullet$ and sheaf Laplacians $\mathcal{L}^k$ are also Hilbert-Schmidt operators.

The operator $A : HS(X, Y) \to HS(X, Y)$ that acts by $A(\phi) = \phi(x_0)x_0^*$ for some fixed $x_0 \neq 0$ is a rank-one operator, and consequently is Hilbert-Schmidt. It follows that restriction map diffusion in $\mathcal{B}(X, Y)$, when started from a Hilbert Schmidt operator stays inside of $HS(X, Y) \subseteq \mathcal{B}(X, Y)$ (with the derivative taken with respect to the operator norm). Moreover, the bound $\| - \|_{op} \leqslant \| - \|_{HS}$, ensure that the derivative in $HS(X, Y)$ with respect to $\| - \|_{op}$ agrees with the derivative with respect to $\| - \|_{HS}$, whenever the $\| - \|_{HS}$-derivative exists. Applying this to restriction map diffusion yields the following corollary to Theorem 6.2.3.

Theorem 6.2.4. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a network and $F : Ob(\mathcal{G}) \to Ob(\mathbf{Hilb}_{\mathbb{k}})$ be a choice of a Hilbert space for each vertex and edge. Let ${}^0\mathcal{F}$ be an initial choice of a Hilbert-Schmidt sheaf extending F, and $\mathbf{x} \in C^0(\mathcal{G}; F)$ a fixed choice of a zero-cochain. Restriction map diffusion on the Hilbert space of Hilbert-Schmidt sheaves has a solution ${}^t\mathcal{F}$, which converges as $t \to \infty$ to the nearest Hilbert-Schmidt sheaf for which $\mathbf{x}$ is a global section.*

The projection onto the nearest sheaf for which $\mathbf{x}$ is a global section is now properly an *orthogonal projection* with respect to the Hilbert space structure of Hilbert-Schmidt operators.

Remark 6.2.5. The Hilbert-Schmidt setting reveals that restriction map diffusion is itself a heat flow. Let X, Y, Z be Hilbert spaces, and take non-zero fixed points $x \in X$ and $y \in Y$. Let $\Phi_X : HS(X, Z) \to Z$ denote the evaluation map $\Phi_X A = Ax$, which has adjoint $\Phi^* z = -zx^*$. Define Φ_Y similarly. The heat dynamics for the Hilbert sheaf

$$\begin{array}{ccccc} HS(X,Z) & \xrightarrow{\Phi_X} & Z & \xleftarrow{\Phi_Y} & HS(Y,Z) \\ \vdots & & \vdots & & \vdots \\ \bullet & \multicolumn{3}{c}{\rule[0.5ex]{6em}{0.4pt}} & \bullet \end{array}$$

evolves $A \in HS(X, Z)$ and $B \in HS(Y, Z)$ by

$$\begin{pmatrix} \dot{A} \\ \dot{B} \end{pmatrix} = -\begin{pmatrix} (Ax - By)x^* \\ (By - Ax)y^* \end{pmatrix}.$$

These are essentially the dynamics of restriction map diffusion. Indeed, given a Hilbert-Schmidt network sheaf $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{\mathbb{k}}$, and a zero-cochain $\mathbf{x} \in C^0(\mathcal{G}; \mathcal{F})$, we may build a new network sheaf on $\mathrm{Map}(\mathcal{F}, \mathbf{x}) : \mathcal{G} \to \mathbf{Hilb}_{\mathbb{k}}$ with the following data.

- **Stalks.** For each vertex v, let E_v denote the collection of edges bounded by v. The vertex stalk $\mathrm{Map}(\mathcal{F}, \mathbf{x})(v)$ is the direct sum

  $$\mathrm{Map}(\mathcal{F}, \mathbf{x})(v) = \bigoplus_{e \in E_v} HS\big(\mathcal{F}(v), \mathcal{F}(e)\big).$$

  The edge stalks $\mathrm{Map}(\mathcal{F}, \mathbf{x})(e) = \mathcal{F}(e)$ are unchanged from $\mathcal{F}$.

- **Restriction maps.** For each map $f : v \to e$, the restriction map $\mathrm{Map}(\mathcal{F}, \mathbf{x})_f$ is given by

  $$\mathrm{Map}(\mathcal{F}, \mathbf{x})_f(A) = A_e(\mathbf{x}_v)$$

  where A_e is the e-component of $A \in \bigoplus_{e \in E_v} HS\big(\mathcal{F}(v), \mathcal{F}(e)\big)$ and $\mathbf{x}_v$ is the v-component of $\mathbf{x}$.

By the previous analysis, heat flows on the Hilbert sheaf $\mathrm{Map}(\mathcal{F}, \mathbf{x})$ exactly encodes restriction map diffusion.

6.2.2 *Joint dynamics*

On a Hilbert-Schmidt network sheaf $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{0,\Bbbk}$ with a choice of a cochain $\mathbf{x}_0 \in C^0(\mathcal{G}; \mathcal{F})$, one may diffuse both the cochain $\mathbf{x}_t$ and the sheaf ${}^t\mathcal{F}$ simultaneously. These joint dynamics evolve according to

$$\begin{aligned} \dot{\mathbf{x}} &= -\alpha \delta^* \delta \mathbf{x} \\ \dot{\delta}_e &= -\beta \delta_e \mathbf{x}_e \mathbf{x}_e^*, \end{aligned}$$

where $\mathbf{x}_e$ is the pair of components of $\mathbf{x}$ in the stalks over the vertices bounding e. Additionally, $\alpha, \beta \geqslant 0$ are real parameters that control the comparative speeds of the cochain diffusion and the restriction map diffusion.

The first step toward understanding joint dynamics is to recognize this system as gradient descent. Let $HS(C^0, C^1)$ denote the Hilbert space of Hilbert-Schmidt operators between zero-cochains and one-cochains, and $\mathbf{HSCbdry} \subseteq HS(C^0, C^1)$ the space of coboundary operators, having the correct sparsity pattern as block operators. We work in the space

$$Z := C^0(\mathcal{G}; F) \oplus \mathbf{HSCbdry}$$

as the dynamics on each component δ_e induces dynamics for the whole coboundary matrix by $\dot{\delta} = -P(\delta \mathbf{x})\mathbf{x}^*$. To recover the α and β coefficients, we renormalize our stalks by premultiplying by the diagonal block operator $M : Z \to Z$ that scales the cochain component of Z by $\alpha \geqslant 0$ and the coboundary map component of Z by β. This renormalization induces a new inner product $\langle \mathbf{z}_1, \mathbf{z}_2 \rangle_M := \langle \mathbf{z}_1, M\mathbf{z}_2 \rangle_Z$ on Z. The induced norm of $\langle -, - \rangle_M$ is equivalent to the norm induced by $\langle -, - \rangle_Z$ as a direct sum. Consider the potential function

$$V \begin{pmatrix} \mathbf{x} \\ \delta \end{pmatrix} = \frac{1}{2} \|\delta \mathbf{x}\|^2 .$$

The gradient of the potential function $V : Z \to \mathbb{R}$ with respect to $\langle -, - \rangle_M$ is given by

$$\nabla_M V \begin{pmatrix} \mathbf{x} \\ \delta \end{pmatrix} = \begin{pmatrix} \alpha \delta^* \delta \mathbf{x} \\ \beta P(\delta \mathbf{x}) \mathbf{x}^* \end{pmatrix},$$

and consequently we recover our dynamics in Z as the gradient descent

$$\begin{pmatrix} \dot{\mathbf{x}}_t \\ \dot{\delta}_t \end{pmatrix} = -\nabla V \begin{pmatrix} \mathbf{x}_t \\ \delta_t \end{pmatrix}.$$

Proposition 6.2.6. *The initial value problem*

$$\begin{aligned} \dot{\mathbf{z}} &= -\nabla V \mathbf{z} \\ \mathbf{z}(0) &= \mathbf{z}_0 \end{aligned}$$

has unique global solutions for all $\mathbf{z}_0$*.*

Proof. $\nabla_M V$ is locally Lipschitz, so the Picard-Lindelöf theorem [15] promises unique local solutions. We check that

$$\begin{aligned} \frac{d}{dt}\|\mathbf{z}_t\|_M^2 &= \langle \mathbf{z}_t, \dot{\mathbf{z}}_t \rangle_M \\ &= -\alpha \|\delta_t \mathbf{x}_t\|_M^2 - 2\beta \sum_{e \in \mathcal{E}} \|\delta_e \mathbf{x}_e\|_M^2 \|x_e\|_M^2 \\ &\leqslant 0. \end{aligned}$$

Therefore trajectories are bounded, and the unique local solutions can be globally extended. □

Convergence of joint dynamics to a pair $(\mathbf{x}_\infty, \delta_\infty)^T$ with $\delta_\infty \mathbf{x}_\infty = 0$ is difficult to guarantee. One would like to use Lyapunov theory, but the infinite dimensional setting has caveats. Consider the following version of LaSalle's invariance principle, due to Hale [51, Theorem 1]

Lemma 6.2.7. *Suppose* u *is a dynamical system on a Banach space* X*. If* V *is a Lyapunov function on* $A \subseteq X$ *and an orbit* $\mathbf{x}_t$ *of* u *belongs to* A *and is precompact, then* $\mathbf{x}_t$ *converges to a point in the largest invariant set in* $S = \{\mathbf{y} \in A : \dot{V}(\mathbf{y}) = 0\}$*.*

The requirement of precompact trajectories is quite restrictive; without precompactness, a trajectory $\mathbf{x}_t$ can spiral through infinite dimensions, dissipating the potential energy $V(\mathbf{x}_t)$ to zero, yet not converging. There are a dearth of tools to find precompact trajectories for joint dynamics. In particular, the linearization of the second derivative of the potential function V does not have a finite dimensional kernel, preventing the use of a Łojasiewicz–Simon inequality [23, 56, 113]. Nonetheless we can prove some conditions under which the dynamics converge. For ease of notation, we work in the case that $\alpha = \beta = 1$, as though we have already renormalized via the operator M.

To begin, we define a collection of quantities associated to the evolution of $(\mathbf{x}_t, \delta_t)^T$. We set:

$$\begin{aligned} p_t &:= \|\mathbf{x}_t\| \\ q_t &:= \|\delta_t\| \\ r_t &:= p_t^2 - q_t^2 \\ s_t &:= \|\delta_t \mathbf{x}_t\| \\ V_t &:= V((\mathbf{x}_t, \delta_t)^T). \end{aligned}$$

Lemma 6.2.8. *The quantities* p_t, q_t, r_t, s_t, V_t *satisfy the following.*

(i) $\frac{d}{dt} p_t^2 = -4V_t \leqslant 0$.

(ii) $\frac{d}{dt} q_t^2 = -4V_t \leqslant 0$.

(iii) $\frac{d}{dt} r_t = 0$.

(iv) $\frac{d}{dt} s_t^2 = -2(\|\delta_t^* \delta_t \mathbf{x}_t\|^2 + \|P(\delta_t \mathbf{x}_t \mathbf{x}_t^*)\|^2)$.

(v) $\frac{d}{dt} V_t = -\left(\|\delta_t^* \delta_t \mathbf{x}_t\|^2 + \|P\delta_t \mathbf{x}_t \mathbf{x}_t^*\|^2\right)$.

Proposition 6.2.9. *If* $(\inf_{t>0} p_t) > 0$*, then* $(\mathbf{x}_t, \delta_t)$ *converges to an equilibrium point of* V*. In particular, this condition is satisfied when* $r > 0$*.*

Proof. The arc-length of the trajectory $(\mathbf{x}_t, \delta_t)^T$ in Z is given by the integral

$$\int_0^\infty \left\| \begin{pmatrix} \dot{\mathbf{x}}_t \\ \dot{\delta}_t \end{pmatrix} \right\| dt = \int_0^\infty \sqrt{(-\dot{V}_t)}\, dt.$$

Hence the trajectory has finite arclength if and only if $\sqrt{(-\dot{V}_t)}$ is integrable over the interval $[0, \infty)$. A sufficient criterion for this integrability is to find a bound of the form $\dot{V}_t \leqslant -kV_t$ for a positive constant k; this implies an upper bound:

$$\begin{aligned} -\dot{V}_t &= \|\delta_t^* \delta_t \mathbf{x}_t\|^2 + \|P\delta_t \mathbf{x}_t \mathbf{x}_t^*\|^2 \\ &\leqslant 2(q_t^2 + p_t^2)V_t \\ &\leqslant 2(q_0^2 + p_0^2)V_0 e^{-kt}, \end{aligned}$$

from which we conclude $\sqrt{-\dot{V}_t}$ is integrable. To find such a bound, we work with the following upper bound on $\dot{V}_t$:

$$\dot{V}_t = -\left(\|\delta_t^* \delta_t \mathbf{x}_t\|^2 + \|P\delta_t \mathbf{x}_t \mathbf{x}_t^*\|^2\right)$$

$$\leqslant -2p_t^2 V_t$$
$$\leqslant 2(\inf_t p_t)V_t\,.$$

By hypothesis, $\inf_t p_t > 0$, giving the required bound. Every path of finite length is precompact, so LaSalle's invariance principle (Lemma 6.2.7), the trajectory $(\mathbf{x}_t, \delta_t)$ converges to an equilibrium point $(\mathbf{x}_\infty, \delta_\infty)^T$ such that $\delta_\infty \mathbf{x}_\infty = 0$. □

Remark 6.2.10. This proposition effectively says that so long as $\mathbf{x}_t$ does not converge to 0, then the joint dynamics $(\mathbf{x}_t, \delta_t)^T$ must converge to some equilibrium point. Therefore if we can bound p_t below, joint dynamics converge nicely. For example, when $r = p_t^2 - q_t^2 > 0$, we know $p_t > \sqrt{r} > 0$ for all t. Moreover, while the coboundary component δ_t (and hence the sheaf ${}^t\mathcal{F}$) need not converge, joint dynamics always converges in the zero-cochain component.

For another criterion, we may adapt proposition 7.2.3 of [52] by the exact same argument.

Corollary 6.2.11. *Let* $K_t := \delta_t^* \delta_t - \mathbf{x}_t \mathbf{x}_t^*$ *be an operator from* $C^0(\mathcal{G}; F) \to C^0(\mathcal{G}; F)$. *If there is a vertex* v *such that* K_0 *is not positive semidefinite on the stalk* $F(v)$, *then joint dynamics converge.*

For yet another criterion, we may compare $\|\mathbf{x}_0\|$ to the norm of its initial image $\|\delta_0 \mathbf{x}_0\|$.

Corollary 6.2.12. *If* $\|\mathbf{x}_0\| > \|\delta_0 \mathbf{x}_0\|$, *then joint dynamics converge.*

Proof. Consider the quantity $\Phi(t) = \frac{s_t^2}{p_t^2} = \frac{\|\delta_t \mathbf{x}_t\|^2}{\|\mathbf{x}_t\|^2}$. We may compute the derivative:

$$\frac{d}{dt}\Phi(t) = \frac{p_t^2 (s_t^2)' - s_t^2 (p_t^2)'}{p_t^4}$$
$$\leqslant \frac{2(\|\delta_t \mathbf{x}_t\|^4 - \|\mathbf{x}_t\|^2 \|\delta_t^* \delta_t \mathbf{x}_t\|^2)}{\|\mathbf{x}_t\|^4}\,.$$

By the Cauchy-Schwartz inequality, we may bound $\|\delta_t \mathbf{x}_t\|^4 \leqslant \|\mathbf{x}_t\|^2 \|\delta_t^* \delta_t \mathbf{x}_t\|^2$, from which we conclude $\frac{d}{dt}\Phi(t) \leqslant 0$. Therefore if $\|\mathbf{x}_0\|^2 \geqslant \|\delta_0 \mathbf{x}_0\|^2$, then $\|\mathbf{x}_t\|^2 \geqslant \|\delta_t \mathbf{x}_t\|^2$ for all t. If $\|\mathbf{x}_t\| \to 0$, then $\|\delta_t \mathbf{x}_t\| \to 0$ as well, and $(\mathbf{x}_t, \delta_t)^T \to (0, 0)^T$, proving that joint dynamics always converge for such a starting point. □

6.3 Wave Propagation

While heat flows dissipate energy and converge to equilibrium states, wave dynamics preserve energy and exhibit oscillatory behavior. This section develops the theory of wave propagation on Hilbert sheaves, extending the classical wave equation to the sheaf-theoretic setting. On a suitable energy space $\mathcal{E}^k$, the wave operator W^k generates a unitary semigroup whose orbits solve the second-order wave equation $\ddot{\mathbf{x}} = -\tilde{\mathcal{L}}^k\mathbf{x}$. While individual trajectories oscillate indefinitely, their time averages exhibit convergence properties that enable distributed computation of harmonic cochains, in line with the finite-dimensional case on a weighted cellular sheaf [115].

Definition 6.3.1. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$ be a Hilbert sheaf. Let $\mathcal{L}^k : C^k(\mathcal{P};\mathcal{F}) \to C^k(\mathcal{P};\mathcal{F})$ be the k^{th}-sheaf Laplacian as usual, and let $\tilde{\mathcal{L}}^k : \ker(\mathcal{L}^k)^\perp \to \ker(\mathcal{L}^k)^\perp$ denote its restriction to the stable subspace $\ker(\mathcal{L}^k)^\perp$. Let $(\tilde{\mathcal{L}}^k)^{1/2}$ denote the unique positive square root of $\tilde{\mathcal{L}}^k$. Define the k^{th}-**energy space** of $\mathcal{F}$ to be the Hilbert space $\mathcal{E}^k := \operatorname{Dom}\left((\tilde{\mathcal{L}}^k)^{1/2}\right) \oplus \ker(\mathcal{L}^k)^\perp$ with the inner product:

$$\left\langle \begin{pmatrix} x_1 \\ y_1 \end{pmatrix}, \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} \right\rangle_{\mathcal{E}} := \left\langle (\tilde{\mathcal{L}}^k)^{1/2}x_1, (\tilde{\mathcal{L}}^k)^{1/2}x_2 \right\rangle + \langle y_1, y_2 \rangle.$$

Remark 6.3.2. When passing from $C^k(\mathcal{P};\mathcal{F})$ to the energy space $\mathcal{E}^k$, we ignore the components of cochains in the kernel of $(\mathcal{L}^k)^{1/2}$. Essentially, cochains in $(\mathcal{L}^k)^{1/2}$ carry no energy, and are implicitly fixed by wave dynamics.

Definition 6.3.3. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\mathbb{k}}$ be a Hilbert sheaf. The k^{th}-**wave operator** $W^k : \mathcal{E}^k \to \mathcal{E}^k$ is the block operator

$$W^k = \begin{bmatrix} 0 & I \\ -\tilde{\mathcal{L}}^k & 0 \end{bmatrix},$$

with domain $\operatorname{Dom}(W^k) = (\operatorname{Dom}(\tilde{\mathcal{L}}^k) \oplus \operatorname{Dom}((\tilde{\mathcal{L}}^k)^{1/2}) \subseteq \mathcal{E}^k$.

The operator W^k is a closed, densely defined, and can be seen to be skew-adjoint; W^k satisfies $(W^k)^* = -W^k$. To confirm this, we first check that W^k is skew-symmetric. For $(x_1, y_1)^\mathsf{T}$ and $(x_2, y_2)^\mathsf{T}$ in the domain $\operatorname{Dom}(W^k)$, we compute:

$$\begin{aligned}
\left\langle W^k \begin{pmatrix} x_1 \\ y_1 \end{pmatrix}, \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} \right\rangle_{\mathcal{E}} &= \left\langle \begin{pmatrix} y_1 \\ -\tilde{\mathcal{L}}x_1 \end{pmatrix}, \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} \right\rangle_{\mathcal{E}} \\
&= \left\langle (\tilde{\mathcal{L}}^k)^{1/2}y_1, (\tilde{\mathcal{L}}^k)^{1/2}x_2 \right\rangle_{C^k} - \left\langle \tilde{\mathcal{L}}^k x_1, y_2 \right\rangle_{C^k} \\
&= -\left\langle (\tilde{\mathcal{L}}^k)^{1/2}x_1, (\tilde{\mathcal{L}}^k)^{1/2}y_2 \right\rangle_{C^k} + \left\langle y_1, \tilde{\mathcal{L}}^k x_2 \right\rangle_{C^k} \\
&= \left\langle \begin{pmatrix} x_1 \\ y_1 \end{pmatrix}, -W^k \begin{pmatrix} x_2 \\ y_2 \end{pmatrix} \right\rangle_{\mathcal{E}}.
\end{aligned}$$

This confirms that W^k is skew-symmetric. It is straightforward to check that $\operatorname{Dom}(W^k)$ is the maximal domain on which the adjoint can be defined, confirming that W^k is skew-adjoint. By

Stone's theorem [100], W^k is the infinitesimal generator of a strongly continuous one-parameter unitary group $U(t) := e^{tW^k}$ on $\mathcal{E}^k$.

Definition 6.3.4. Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf with sheaf Laplacian $\mathcal{L}^k$. The k^{th}-**wave group** of $\mathcal{F}$ is the strongly continuous one parameter unitary group $\exp(tW^k)$ on the energy space $\mathcal{E}^k$, where W^k is k^{th}-wave operator. Given an initial point $(\mathbf{x}_0, \mathbf{y}_0)^\top \in \mathcal{E}^k$, the **wave propagation** of $(\mathbf{x}_0, \mathbf{y}_0)^\top$ is the path

$$\begin{pmatrix} \mathbf{x}_t \\ \mathbf{y}_t \end{pmatrix} = \exp(tW^k) \begin{pmatrix} \mathbf{x}_0 \\ \mathbf{y}_0 \end{pmatrix}$$

in $\mathcal{E}^k$.

Remark 6.3.5. Let $(\mathbf{x}_t, \mathbf{y}_t)^\top$ be the wave propagation of an initial value $(\mathbf{x}_t, \mathbf{y}_t)^\top$. Using the properties of semigroups, we have that $(\dot{\mathbf{x}}_t, \dot{\mathbf{y}}_t)^\top = W(\mathbf{x}_t, \mathbf{y}_t)^\top = (\mathbf{y}_t, -\mathcal{L}\mathbf{x}_t)^\top$. It follows that the first component $\mathbf{x}_t$ is a solution to the **wave equation**

$$\ddot{\mathbf{x}} = -\tilde{\mathcal{L}}\mathbf{x} \tag{5}$$

subject to initial conditions $\mathbf{x}(0) = \mathbf{x}_0$ and $\dot{\mathbf{x}}(0) = \mathbf{y}_0$. Accordingly, we will often denote the wave propagation $(\mathbf{x}_t, \dot{\mathbf{x}}_t)$.

Remark 6.3.6. Since the wave group $e^{tW} : \mathcal{E}^k \to \mathcal{E}^k$ is a globally defined unitary operator for each t, the energy norm $\|(\mathbf{x}_t, \dot{\mathbf{x}}_t)\|_{\mathcal{E}}$ is constant. This may also be confirmed through direct computation. Even though $\|(\mathbf{x}_t, \dot{\mathbf{x}}_t)\|_{\mathcal{E}}$ is constant, it is possible that $\|\mathbf{x}_t\|$ diverges to infinity in $C^k(\mathcal{P}; \mathcal{F})$.

6.3.1 *Solutions of the wave equation*

Proposition 6.3.7. *Let $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ be a Hilbert sheaf, and let W^k be the k^{th}-wave operator. The spectrum of W^k is purely imaginary and given by*

$$\sigma(W^k) = \left\{ \pm\lambda i \, : \lambda \in \sigma(\tilde{\mathcal{L}}^k) \right\}.$$

Proof. This is a general fact about block operators of the form of W^k. We handle the case where $\mathcal{F}$ is a complex Hilbert sheaf—the real case follows via complexification. A value $\lambda \in \mathbb{C}$ is in the spectrum of W^k if and only if the operator $(\lambda I - W)$ fails to be boundedly invertible onto its image. Solving the equation $(\lambda I - W)(x, y)^\top = (w, z)^\top$ yields the system:

$$(\tilde{\mathcal{L}}^k + \lambda^2 I)x = z + \lambda w$$

$$y = \lambda x - w.$$

Hence $(\lambda I - W)$ has a bounded inverse if and only if $(\tilde{\mathcal{L}}^k + \lambda^2 I)$ has a bounded inverse. This is equivalent to the statement that $-\lambda^2 \in \sigma(\tilde{\mathcal{L}}^k)$, from which we derive that

$$\sigma(W^k) = \big\{ \pm \lambda i : \lambda \in \sigma(\tilde{\mathcal{L}}^k) \big\}.$$

Since $\tilde{\mathcal{L}}^k$ is a self-adjoint operator (and hence has a real spectrum), the spectrum of W is purely imaginary. □

More can be said by restricting our attention to eigenvalues. For simplicity, assume we're working with a complex Hilbert sheaf—the following results can be adapted for real Hilbert sheaves through complexification. When $(\lambda, \mathbf{u})$ is an eigenvalue-eigenvector pair for $\tilde{\mathcal{L}}^k$, we get a corresponding duo of eigenvectors for W^k, $\mathbf{z}^+ = (\mathbf{u}, i\omega\mathbf{u})^T$ and $\mathbf{z}^- = (\mathbf{u}, -i\omega\mathbf{u})^T$ where $\omega = \sqrt{\lambda}$, with corresponding eigenvalues $i\omega$ and $-i\omega$ respectively. We call the eigenvectors $\mathbf{z}^+$ and $\mathbf{z}^-$ **normal modes** of W. Normal modes represent purely-oscillatory solutions to the wave equation.

When $\tilde{\mathcal{L}}^k$ (and hence W^k) have pure point spectrum, every wave propagation can be expressed as an infinite sum of normal modes. In particular, we may write $(\mathbf{x}_0, \dot{\mathbf{x}}_0)^T = \sum_{\lambda \in \sigma(\tilde{\mathcal{L}}^k)} a_\lambda^+ \mathbf{z}_\lambda^+ + a_\lambda^- \mathbf{z}_\lambda^-$, where $a_\lambda^\pm \in \mathbb{C}$ is a scalar. We may now express the wave propagation $(\mathbf{x}_t, \dot{\mathbf{x}}_t)^T$ as

$$(\mathbf{x}_t, \dot{\mathbf{x}}_t)^T = \sum_{\lambda \in \sigma(\tilde{\mathcal{L}})} a_\lambda^+ e^{i\sqrt{\lambda}t} \mathbf{z}_\lambda^+ + a_\lambda^- e^{-i\sqrt{\lambda}t} \mathbf{z}_\lambda^- .$$

However, the sheaf Laplacian will generically fail to have pure point spectrum. We may nonetheless arrive at a similar representation through spectral calculus. We start with the following theorem.

Theorem 6.3.8. *The wave propagation of* $(\mathbf{x}_0, \dot{\mathbf{x}}_0)^T$ *can be represented as*

$$\begin{pmatrix} \mathbf{x}_t \\ \dot{\mathbf{x}}_t \end{pmatrix} = \begin{pmatrix} \cos\big(t(\tilde{\mathcal{L}}^k)^{1/2}\big)\mathbf{x}_0 + \big(\tilde{\mathcal{L}}^k\big)^{-1/2} \sin\big(t(\tilde{\mathcal{L}}^k)^{1/2}\big)\dot{\mathbf{x}}_0 \\ -\big(\tilde{\mathcal{L}}^k\big)^{1/2} \sin\big(t(\tilde{\mathcal{L}}^k)^{1/2}\big)\mathbf{x}_0 + \cos\big(t(\tilde{\mathcal{L}}^k)^{1/2}\big)\dot{\mathbf{x}}_0 \end{pmatrix}.$$

Proof. Let $U(t) := \exp(tW^k)$ denote the wave group. We may write $U(t) = C(t) + S(t)$, where $C(t) = \frac{1}{2}\big(U(t) + U(-t)\big)$ and $S(t) = \frac{1}{2}\big(U(t) - U(-t)\big)$. Analyzing the strongly continuous group $C(t)$, we may derive the following identities:

$$C(0) = I$$
$$\dot{C}(0) = 0$$

$$\ddot{C}(t) = -\begin{bmatrix} \tilde{\mathcal{L}} & 0 \\ 0 & \tilde{\mathcal{L}} \end{bmatrix} C(t)\,.$$

$C(t)$ is a solution to the initial value problem $\ddot{T} + \begin{bmatrix} \tilde{\mathcal{L}} & 0 \\ 0 & \tilde{\mathcal{L}} \end{bmatrix} T = 0$, subject to the initial conditions $T(0) = I$ and $\dot{T}(0) = 0$ in the Banach space $\mathcal{E}^k$. The unique mild solution to this differential equation is given by

$$C(t) = \begin{bmatrix} \cos\left((\tilde{\mathcal{L}}^k)^{1/2} t\right) & 0 \\ 0 & \cos\left((\tilde{\mathcal{L}}^k)^{1/2} t\right) \end{bmatrix},$$

where $\cos\left((\tilde{\mathcal{L}}^k)^{1/2} t\right)$ is the bounded operator defined by the Borel function calculus. A similar analysis of $S(t)$ yields[2]

$$S(t) = \begin{bmatrix} 0 & \left(\tilde{\mathcal{L}}^k\right)^{-1/2} \sin\left(t(\tilde{\mathcal{L}}^k)^{1/2}\right) \\ \left(\tilde{\mathcal{L}}^k\right)^{1/2} \sin\left(t(\tilde{\mathcal{L}}^k)^{1/2}\right) & 0 \end{bmatrix}.$$

The solution follows by addition. □

Remark 6.3.9. This solution is analogous to the d'Alembert solution to the wave equation in one dimension.

6.3.2 *Long term behavior of wave propagation*

Theorem 6.3.10. *Let $(\mathbf{x}_t, \dot{\mathbf{x}}_t)^T$ be a wave propagation with initial point $(\mathbf{x}_0, \dot{\mathbf{x}}_0)^T$. The time-average position of $\mathbf{x}_t$ is zero. That is,*

$$\lim_{T\to\infty} \frac{1}{T} \int_0^T \mathbf{x}_t \, dt = 0\,.$$

Proof. Using Theorem 6.3.8, write the cochain

$$\mathbf{x}_t = \cos\left(t(\tilde{\mathcal{L}}^k)^{1/2}\right)\mathbf{x}_0 + \left(\tilde{\mathcal{L}}^k\right)^{-1/2} \sin\left(t(\tilde{\mathcal{L}}^k)^{1/2}\right)\dot{\mathbf{x}}_0\,.$$

We analyze these two terms in the sum separately.

The cosine term. Let $\Phi(t) := \left(\tilde{\mathcal{L}}^k\right)^{-1/2} \sin\left(t(\tilde{\mathcal{L}}^k)^{1/2}\right)$. The function $f(t, x) = \sin(tx)/x$, appropriately extended to a total function on $\mathbb{R}^2$, is a bounded Borel function for each t and $\frac{\partial f}{\partial t}$ is a

2 The operator $\left(\tilde{\mathcal{L}}^k\right)^{-1/2} \sin\left(t(\tilde{\mathcal{L}}^k)^{1/2}\right)$ is the operator obtained by applying the bounded Borel function $f(x) = \frac{\sin(tx)}{x}$ to the operator $(\tilde{\mathcal{L}}^k)^{1/2}$; the notation is not intended to imply that $(\tilde{\mathcal{L}}^k)^{1/2}$ is invertible.

bounded Borel function for all t. The Borel function calculus therefore allows us to differentiate $\Phi(t) = f(t, (\tilde{\mathcal{L}}^k)^{1/2})$ and obtain

$$\begin{aligned}\frac{d}{dt}\Phi(t) &= \frac{\partial f}{\partial t}(t, (\tilde{\mathcal{L}}^k)^{1/2}) \\ &= \cos(t(\tilde{\mathcal{L}}^k)^{1/2}).\end{aligned}$$

The integral of the cosine term can now be evaluated as

$$\frac{1}{T}\int_0^T \cos\left(t(\tilde{\mathcal{L}}^k)^{1/2}\right)\mathbf{x}_0 \, dt = \frac{\left(\tilde{\mathcal{L}}^k\right)^{-1/2}\sin\left(T(\tilde{\mathcal{L}}^k)^{1/2}\right)\mathbf{x}_0}{T}.$$

Since $\mathbf{x}_0 \in \ker(\tilde{\mathcal{L}}^k)^{\perp}$, the spectral measure form of the spectral theorem guarantees

$$\left\|\frac{\left(\tilde{\mathcal{L}}^k\right)^{-1/2}\sin\left(T(\tilde{\mathcal{L}}^k)^{1/2}\right)\mathbf{x}_0}{T}\right\|^2 = \int_{\lambda\in\sigma((\tilde{\mathcal{L}}^k)^{1/2})}\left(\frac{\sin(T\lambda)}{T\lambda}\right)^2 d\mu_{x_0}(\lambda).$$

Let $m > 0$ be a small real number. On the interval $(0, m)$, the integral may be bounded above by $\mu_{x_0}((0, m))$ as the integrand is bounded above by 1. On the interval (m, ∞), the integral may be bounded above by $\frac{\|x_0\|}{T^2m^2}$ as the integrand is bounded above by $\frac{1}{T^2\lambda^2}$. Combining these bounds yields

$$\left\|\frac{\left(\tilde{\mathcal{L}}^k\right)^{-1/2}\sin\left(T(\tilde{\mathcal{L}}^k)^{1/2}\right)\mathbf{x}_0}{T}\right\|_X^2 \leqslant \mu_{x_0}((0, m)) + \frac{\|\mathbf{x}_0\|^2}{T^2m^2}.$$

The upper bound converges to $\mu_{x_0}((0, m))$ as $T \to \infty$. Since m may be chosen arbitrarily small, this quantity goes to zero. It follows that the cosine term satisfies

$$\lim_{T\to\infty}\frac{1}{T}\int_0^T \cos\left(T(\tilde{\mathcal{L}}^k)^{1/2}\right)\mathbf{x}_0 \, dt = 0.$$

The sine term. A similar analysis to the cosine term gives the following identity:

$$\int_0^T (\tilde{\mathcal{L}}^k)^{-1/2}\sin(t(\tilde{\mathcal{L}}^k)^{1/2})\frac{1}{T}\dot{\mathbf{x}}_0 \, dt = \frac{(\tilde{\mathcal{L}}^k)^{-1}(I - \cos(T(\tilde{\mathcal{L}}^k)^{1/2}))\dot{\mathbf{x}}_0}{T}.$$

Again using the spectral theorem, since $\dot{\mathbf{x}}_0$ is orthogonal to the kernel of $\tilde{\mathcal{L}}^k$, this term goes to zero as $T \to \infty$.

Conclusion. Combining these results, when $\dot{\mathbf{x}}_0 \perp \ker(\tilde{\mathcal{L}}^k)$, the time-average position $\tilde{\mathbf{x}} = \lim_{T\to\infty}\frac{1}{T}\int_0^T \mathbf{x}_t \, dt$ converges to zero. □

This argument may be adapted to give a procedure for finding harmonic cochains in $C^k(\mathcal{P};\mathcal{F})$. Consider the space $\mathcal{H}^k := \mathrm{Dom}\left((\mathcal{L}^k)^{1/2}\right) \oplus \ker(\mathcal{L}^k)^\perp$, which strictly contains the energy space $\mathcal{E}^k$. The energy inner product $\langle -, - \rangle_{\mathcal{E}}$ defines a semi-inner product on $\mathcal{H}^k$; there are non-zero points $x \in \mathcal{H}^k$ such that $\langle x, x \rangle = 0$. Moreover, one may extend W^k to $\mathcal{H}^k$ by

$$\mathcal{W}^k = \begin{bmatrix} 0 & I \\ -\mathcal{L}^k & 0 \end{bmatrix}.$$

This operator is still well-defined as $\mathcal{R}(\mathcal{L}^k) \subseteq \ker(\mathcal{L}^k)^\perp$. For any $(\mathbf{x}_0, \dot{\mathbf{x}}_0)^\top \in \mathcal{H}^k$, writing $\mathbf{x}_0 = \mathbf{y}_0 + \mathbf{k}_0$ where $\mathbf{y}_0 \in \ker((\mathcal{L}^k)^{1/2})^\perp$ and $\mathbf{k}_0 \in \ker((\mathcal{L}^k)^{1/2})$, one may compute that

$$\mathcal{W}^k \begin{pmatrix} \mathbf{x}_0 \\ \dot{\mathbf{x}}_0 \end{pmatrix} = W^k \begin{pmatrix} \mathbf{y}_0 \\ \dot{\mathbf{x}}_0 \end{pmatrix}.$$

This allows us to extend the wave propagation dynamics from $\mathcal{E}^k$ to $\mathcal{H}^k$, yielding

$$\begin{pmatrix} \mathbf{x}_t \\ \dot{\mathbf{x}}_t \end{pmatrix} = \begin{pmatrix} \mathbf{k}_0 \\ 0 \end{pmatrix} + e^{tW^k} \begin{pmatrix} \mathbf{y}_0 \\ \dot{\mathbf{x}}_0 \end{pmatrix}.$$

Corollary 6.3.11. *Let $(\mathbf{x}_t, \dot{\mathbf{x}}_t)^\top$ be a wave propagation on $\mathcal{H}^k$ with initial point $(\mathbf{x}_0, 0)^\top$. The time-average position*

$$\tilde{\mathbf{x}} = \lim_{T\to\infty} \frac{1}{T} \int_0^T \mathbf{x}_t \, dt$$

is the nearest k-harmonic cochain to $\mathbf{x}_0$.

Proof. Write $\mathbf{x}_0 = \mathbf{y}_0 + \mathbf{k}_0$ where $\mathbf{y}_0 \in \ker((\mathcal{L}^k)^{1/2})^\perp$ and $\mathbf{k}_0 \in \ker((\mathcal{L}^k)^{1/2})$. Since $(\mathcal{L}^k)^{1/2}$ and $\mathcal{L}^k$ have the same kernel, $\mathbf{k}_0$ is the nearest harmonic cochain in $C^k(\mathcal{P};\mathcal{F})$. The wave dynamics on $\mathcal{H}^k$ decompose as

$$\frac{1}{T} \int_0^T \mathbf{x}_t \, dt = \mathbf{k}_0 + \frac{1}{T} \int_0^T \mathbf{y}_t \, dt$$

Taking limits and applying Theorem 6.3.10 yields the desired result. □

Remark 6.3.12. This corollary essentially allows the use of wave propagation for distributed computation of a harmonic cochain. If each k-cell represents an agent, the dynamics can be computed locally at each k-cell utilizing only the information from neighboring k-cells. By continuously evolving local data via wave propagation and recording the signal, taking an average stalk-wise yields a harmonic cochain.

7

NONLINEAR DYNAMICS

In Chapter 6, we observed that the sheaf Laplacian $\mathcal{L}^k$ of a Hilbert sheaf $\mathcal{F} : \mathcal{P} \to \mathbf{Hilb}_{0,\Bbbk}$ defines a heat equation $\dot{\mathbf{x}} = -\mathcal{L}^k \mathbf{x}$ on the space of k-cochains $C^k(\mathcal{P}; \mathcal{F})$. Crucially, this is a first order linear differential equation which can be solved via semigroup theory. In this chapter, we explore two distinct nonlinear analogues to this heat flow. We restrict our attention to the case of network Hilbert sheaves for simplicity, though many results can be adapted for the k-Laplacian $\mathcal{L}^k : C^k(\mathcal{P}; \mathcal{F}) \to C^k(\mathcal{P}; \mathcal{F})$ of an arbitrary Hilbert sheaf. Both approaches to nonlinear heat flow come from the following observation.

Proposition 7.0.1. *Let $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{\Bbbk}$ be a bounded network Hilbert sheaf with Laplacian $\mathcal{L} = \delta^* \delta$. Every heat flow $\mathbf{x}_t$ on the space of zero-cochains $C^0(\mathcal{G}; \mathcal{F})$ is a gradient descent trajectory of the quadratic potential function*

$$V(\mathbf{x}) := \frac{1}{2} \|\delta \mathbf{x}\|^2.$$

In particular, $\nabla V = \mathcal{L}$. This suggests an identification of linear heat dynamics on a Hilbert sheaf with gradient descent with respect to a quadratic distance function $g(-) = \| - \|^2$. We may adapt this framework in two different ways.

1. First, we may consider replacing the quadratic $g(-) = \| - \|^2$ with a different function for measuring "distance" in $C^1(\mathcal{G}; \mathcal{F})$. Changing the distance function g will result in different, nonlinear sheaf Laplacian $\mathcal{L}^g : C^0(\mathcal{G}; \mathcal{F}) \to C^0(\mathcal{G}; \mathcal{F})$, with different dynamics. This approach, which we call a **one-cochain nonlinearity**, was first explored for weighted cellular sheaves in [55, Section 10], and expanded in [52].

2. Second, we may take a truly "dynamics first" approach by replacing the coboundary operator δ with a *nonlinear* map. Such a nonlinear coboundary map may be obtained by allowing restriction maps to themselves be nonlinear. We call this approach a **zero-cochain nonlinearity**, as the coboundary map itself is nonlinear on $C^0(\mathcal{G}; \mathcal{F})$.

The first approach, explored in Section 7.1, generalizes the quadratic potential to incorporate edge-wise nonlinearities. While this framework has been studied in the finite-dimensional setting, we briefly review the construction and note the additional care required when extending to infinite-dimensional Hilbert sheaves, particularly regarding compactness assumptions for LaSalle invariance arguments.

The remainder of the chapter focuses on the second approach, which forms the primary contribution of this chapter. Section 7.2 develops the general theory of C^0-nonlinear Hilbert sheaves, where restriction maps need not be linear. The challenge here lies in defining an appropriate notion of sheaf Laplacian when the coboundary operator δ loses linearity. We address this by employing various generalized gradients (Fréchet, Clarke, and convex subdifferential) depending on the regularity of the potential function $V_{\mathcal{F}}(\mathbf{x}) = \frac{1}{2}\|\delta\mathbf{x}\|^2$. The section concludes by extending the framework to Riemannian network sheaves, where stalks are smooth Riemannian manifolds rather than Hilbert spaces.

Section 7.3 and Section 7.4 examine two special cases of C^0-nonlinear Hilbert sheaves that admit tractable analysis. Section 7.3 studies affine network sheaves, where restriction maps take the form $\mathcal{F}_f(x) = A_f x + b_f$. We demonstrate that heat flows on such sheaves converge to ordinary least squares solutions of potentially inconsistent inhomogeneous linear systems. Moreover, we establish a cohomological interpretation through the language of torsors, following the work of [46]. Specifically, we show that the cohomology of the linear structure sheaf with restrictions maps $\mathcal{F}_f(x) = A_f$ encodes obstructions to the existence of global sections for affine Hilbert sheaves with restriction maps $\mathcal{F}_f(x) = A_f x + b_f$.

Section 7.4 investigates continuous piecewise affine (CPWA) Hilbert sheaves, where restriction are continuous piecewise affine maps. Maps of this class arise naturally in applications to neural networks with ReLU activation functions. The analysis requires tools from non-smooth analysis as the resulting dynamics constitute state-dependent switched affine systems. We establish global existence of Filippov solutions and prove that fast heat flows—those that minimize time spent in sliding modes—are bounded and converge to generalized critical points of the potential function. The polyhedral structure underlying CPWA maps provides sufficient geometric control to ensure well-behaved long-term dynamics despite the lack of smoothness.

Throughout, we maintain focus on the interplay between the algebraic structure of sheaves and the analytic properties of their associated dynamical systems, demonstrating how nonlinearity in restriction maps enriches both the theoretical framework and potential applications of cellular sheaf theory.

7.1 One-cochain nonlinearities

Let $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{0,\Bbbk}$ be a network sheaf of Hilbert spaces on a finite graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with coboundary map δ. We briefly discuss nonlinear dynamics of the form $\dot{\mathbf{x}} = \frac{1}{2}g(\delta\mathbf{x})$ for different choices of g on bounded Hilbert sheaves. This approach to nonlinear sheaf Laplacians in the finite dimensional setting is well-trod ground [52, 54]. We first recast the usual heat flow on space of zero-cochains $C^0(\mathcal{G};\mathcal{F})$ as subgradient descent.

7.1.1 *Heat flow as subgradient descent*

Definition 7.1.1. Let $V : X \to \overline{\mathbb{R}}$ be a partially-defined extended real-valued functional on a Hilbert space X. V is **proper** if V has non-empty domain, is not uniformly infinite, and never takes on the value of $-\infty$.

Let $A : X \to Y$ be a closed, densely defined operator with $D := \mathrm{Dom}(A)$. We may define a potential function

$$V(x) := \begin{cases} \frac{1}{2}\|Ax\|^2 & \text{if } x \in D \\ \infty & \text{else.} \end{cases}$$

The potential function V is not continuous in general, and cannot be differentiated. However, V is lower-semicontinuous, convex, and proper. These properties allow us to compute the subdifferential of V as

$$\partial V(x) := \{g \in X : V(y) \geqslant V(x) + \langle g, y - x\rangle \quad \forall y \in X\}.$$

The subdifferential $\partial V(x) = \{A^*Ax\}$ for all $x \in \mathrm{Dom}(A^*A)$ [15]. Moreover, when A is bounded (and V is Fréchet differentiable), this subdifferential exactly agrees with the usual gradient of V.

Applying this analysis to the (potentially unbounded) coboundary operator δ of the network Hilbert sheaf $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{0,\Bbbk}$ yields the identification of the heat flow $\dot{\mathbf{x}} = -\mathcal{L}\mathbf{x}$ with (sub)gradient descent with respect to the functional $V(\mathbf{x}) = \frac{1}{2}\|\delta\mathbf{x}\|^2$.

7.1.2 *Edgewise nonlinearities*

In the previous analysis of the potential function $V(\mathbf{x}) = \frac{1}{2}\|\delta\mathbf{x}\|^2$, the potential function decomposes edgewise as

$$V(\mathbf{x}) = \frac{1}{2}\sum_{e\in\mathcal{E}} \|(\delta\mathbf{x})_e\|^2,$$

where $\mathcal{E}$ is the set of edges of the network $\mathcal{G}$. The analysis via subgradients may easily be repeated for unbounded network Hilbert sheaves when the distance function $\|-\|^2$ is replaced with a different convex function g on each edge, such as the p-norm $\|-\|^p$ for any choice of $p \geqslant 1$. However, to consider non-convex g, it is useful to restrict attention to bounded Hilbert sheaves.

Definition 7.1.2. Let $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{\Bbbk}$ be a bounded network Hilbert sheaf on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. A C^1**-nonlinearity** is a family of globally defined maps $\{\phi_e : \mathcal{F}(e) \to \mathbb{R} \mid e \in \mathcal{E}\}$. The data of a C^1-nonlinearity induces a block map $\Phi : C^1(\mathcal{G};\mathcal{F}) \to \mathbb{R}$ by

$$\Phi(\mathbf{y}) := \sum_{e\in\mathcal{E}} \phi_e \mathbf{y}_e.$$

We affectionately call the map Φ the **middle map** of the C^1-nonlinearity due to its role in defining the C^1-nonlinear Laplacian.

Definition 7.1.3. Let $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{\Bbbk}$ be a bounded network Hilbert sheaf on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with C^1-nonlinearity $\{\phi_e\}_{e\in\mathcal{E}}$. The **potential function** associated to the C^1-nonlinearity is the map $V^\Phi : C^0(\mathcal{G};\mathcal{F}) \to \mathbb{R}$ defined by

$$V^\Phi(\mathbf{x}) := \frac{1}{2}\Phi(\delta\mathbf{x}) = \frac{1}{2}\sum_{e\in\mathcal{E}} \phi_e\big((\delta\mathbf{x})_e\big) .$$

When Φ is continuously differentiable in a neighborhood of a point $\delta\mathbf{x} \in C^1(\mathcal{G};\mathcal{F})$, the potential function V^Φ is differentiable at $\mathbf{x}$, and has gradient

$$\nabla V^\Phi = \frac{1}{2}\delta^*(\nabla\Phi)\delta .$$

Definition 7.1.4. Let $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{0,\Bbbk}$ be a network Hilbert sheaf on a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with C^1-nonlinearity $\{\phi_e\}_{e\in\mathcal{E}}$. The C^1**-nonlinear Laplacian** is the map $\mathcal{L}^\Phi : C^0(\mathcal{G};\mathcal{F}) \to C^0(\mathcal{G};\mathcal{F})$ defined by

$$\mathcal{L}^\Phi := \frac{1}{2}\delta^*(\nabla\Phi)\delta.$$

The C^1-nonlinear Laplacian $\mathcal{L}^\Phi$ induces a smooth dynamical system on $C^0(\mathcal{G};\mathcal{F})$, wherever the composition $\Phi\delta(x)$ is differentiable. The resulting dynamical system

$$\dot{\mathbf{x}} = -\mathcal{L}^\Phi\mathbf{x} \tag{6}$$

is called the C^1**-nonlinear heat flow** with respect to the C^1-nonlinearity Φ. The Picard-Lindelöf theorem ensures the local existence of C^1-nonlinear heat flows for all initial values $\mathbf{x}_0 \in \mathrm{Dom}(\mathcal{L}^\Phi)$.

Remark 7.1.5. In [55, Section 7.3], these C^1-nonlinear dynamics are explored for finite dimensional weighted cellular sheaves. Many of the results can be straightforwardly adapted for the infinite dimensional setting. However, care must be taken with arguments based on LaSalle invariance, which require additional assumptions or argumentation to ensure precompactness.

7.2 Zero-cochain nonlinearities

Our second approach to defining a nonlinear sheaf Laplacian is to directly adapt the definition of a network Hilbert sheaf to allow for nonlinear restriction maps.

Definition 7.2.1. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite graph, viewed as a weakly-regular cell structure. A C^0**-nonlinear Hilbert sheaf** $\mathcal{F}$ on $\mathcal{G}$ consists of the following data.

- For each $\sigma \in \mathcal{V} \amalg \mathcal{E}$, the choice of a Hilbert space **stalk** $\mathcal{F}(\sigma)$.
- For each covering morphism $f : v \to e$ in $\mathcal{G}$, a **restriction map** $\mathcal{F}_f : \mathcal{F}(v) \to \mathcal{F}(e)$ with domain $\mathrm{Dom}(f)$.

From these data, we define the spaces of zero-cochains $C^0(\mathcal{G};\mathcal{F})$ and one-cochains $C^1(\mathcal{G};\mathcal{F})$ as before; namely by the direct sum of stalks over vertices and edges respectively. We also define a **coboundary map** $\mathring{\delta} : C^0(\mathcal{G};\mathcal{F}) \to C^1(\mathcal{G};\mathcal{F})$ with respect to a choice of an orientation of each edge by

$$(\mathring{\delta}\mathbf{x})_e = \mathcal{F}_g x_v - \mathcal{F}_f x_u\,,$$

where f and g are the covering maps corresponding to the source and target of the edge e. We further require that $\mathring{\delta}$ be **closable**, meaning that the closure of the graph $\Gamma(\mathring{\delta}) \subseteq C^0(\mathcal{G};\mathcal{F}) \times C^1(\mathcal{G};\mathcal{F})$ is the graph of a function, which we denote by $\delta : C^0(\mathcal{G};\mathcal{F}) \to C^1(\mathcal{G};\mathcal{F})$. In an abuse of notation, we also call the closure δ the **coboundary** map.

Remark 7.2.2. Like Hilbert sheaves with unbounded linear operators, the restriction map $\mathcal{F}_f : \mathcal{F}(v) \to \mathcal{F}(e)$ is not required to be globally defined nor continuous.

Consider the potential function $V_{\mathcal{F}}(\mathbf{x}) := \frac{1}{2}\|\delta\mathbf{x}\|^2$ on $C^0(\mathcal{G};\mathcal{F})$. When V admits a generalized gradient $\partial V_{\mathcal{F}}(x)$ for all $x \in \mathrm{Dom}(\delta)$, we may define a C^0**-nonlinear sheaf Laplacian**

$$\mathcal{L} := \partial V_{\mathcal{F}}(x)$$

and a corresponding **heat flow**

$$\dot{\mathbf{x}} \in -\mathcal{L}\mathbf{x}\,. \tag{7}$$

Remark 7.2.3. A general C^0-nonlinear Hilbert sheaf will not admit a sheaf Laplacian. Moreover, different C^0-nonlinear Hilbert sheaves will require the use of a different generalized gradient. A few examples will be illustrative.

Example 7.2.4. Let $\mathcal{F}$ be a C^0-nonlinear Hilbert sheaf on $\mathcal{G}$ such that all restriction maps are continuously Fréchet differentiable on an open domain of definition. Call such a sheaf **continuously differentiable**. The potential function $V_{\mathcal{F}}(\mathbf{x})$ is continuously Fréchet differentiable as well, and we may take the usual gradient ∇ as our generalized gradient. The corresponding sheaf Laplacian is given by $\mathcal{L}\mathbf{x} := \nabla\frac{1}{2}\|\delta\mathbf{x}\|^2$, and the corresponding heat flow is gradient descent with respect to $V_{\mathcal{F}}$.

Example 7.2.5. As a special case of the previous example, a bounded network Hilbert sheaf $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{\Bbbk}$ can be viewed as a continuously differentiable C^0-nonlinear Hilbert sheaf. In this case, the heat flow reduces to the usual sheaf Laplacian $\mathcal{L}\mathbf{x} = -\delta^*\delta\mathbf{x}$.

Example 7.2.6. Suppose $\mathcal{F}$ is a C^0-nonlinear Hilbert sheaf such that the potential function $V_{\mathcal{F}}(\mathbf{x})$ is proper, convex, and lower semicontinuous. The convex subdifferential

$$\partial V_{\mathcal{F}}(\mathbf{x}) = \{g \in C^0(\mathcal{G};\mathcal{F}) \;:\; V_{\mathcal{F}}(\mathbf{y}) \geqslant V_{\mathcal{F}}(\mathbf{x}) + \langle g, \mathbf{y}-\mathbf{x}\rangle \quad \forall \mathbf{y} \in C^0(\mathcal{G};\mathcal{F})$$

allows us to define the sheaf Laplacian $\mathcal{L}\mathbf{x} := \partial V_{\mathcal{F}}(\mathbf{x})$. The Brézis-Komura theorem will guarantee a unique heat flow.

Example 7.2.7. As a special case of the previous example, the potential function $V_{\mathcal{F}}$ associated to an unbounded network Hilbert sheaf $\mathcal{F} : \mathcal{G} \to \mathbf{Hilb}_{0,\Bbbk}$ is proper, convex, and lower semicontinuous. Moreover, the convex subdifferential $\partial V_{\mathcal{F}}$ recovers the usual sheaf Laplacian $\mathcal{L} = \delta^*\delta$.

Example 7.2.8. Suppose $\mathcal{F}$ is a C^0-nonlinear Hilbert sheaf such that the potential function $V_{\mathcal{F}}(\mathbf{x})$ is locally Lipschitz on an open domain, but not necessarily convex. In this case, we may use the **Clarke generalized gradient** [24, Definition 1.1]:

$$\partial V_{\mathcal{F}}(x) = \overline{\operatorname{cvx}\left\{\lim_{n\to\infty} \nabla V_{\mathcal{F}}(x+h_n) \;:\; h_n \to 0\right\}},$$

where $\operatorname{cvx} S$ denotes the convex hull of points in S. This approach again recovers the usual sheaf Laplacian when applied to a bounded network Hilbert sheaf.

7.2.1 *Local adjoints*

In addition to recovering the usual sheaf Laplacian for network Hilbert sheaves, there is a geometric justification for the Laplacian $\mathcal{L} := \partial V_{\mathcal{F}}$ to be viewed as a nonlinear generalization of the ordinary sheaf Laplacian. In the forthcoming analysis, we assume that our coboundary operator δ is locally Lipschitz continuous, and hence almost everywhere differentiable on a dense G_δ set.

Let X and Y be Hilbert spaces. A continuous function $f : X \to Y$ has a linear adjoint $f^* : Y \to X$ such that $\langle fx, y\rangle_Y = \langle x, f^*y\rangle_X$ for all $x \in X$ and $y \in Y$ if and only if f is a bounded linear function. Consequently, to define a sheaf Laplacian $\mathcal{L} = \delta^*\delta$, the coboundary map δ must be linear. In this section we seek to define a variant of the adjoint for nonlinear maps that allow a second approach to nonlinear Laplacian dynamics on Hilbert spaces.

Definition 7.2.9. Let X, Y be Hilbert spaces, and let $f : D \to Y$ be a continuous map defined on an open subset $D \subseteq X$. A **local linear adjoint** (or simply a **local adjoint**) for f on D is a map $f^* : D \times Y \to X$ such that the following conditions hold.

(i) $f^*_x := f^*(x, -) : Y \to X$ is a bounded linear map for each $x \in D$.

(ii) For all $x \in D$ and $y \in Y$, $\left|\langle f(x+h) - f(x), y\rangle_Y - \langle h, f^*_x(y)\rangle_X\right| = o(\|h\|_X)$.

Proposition 7.2.10. *A map* $f : D \to Y$ *admits at most one local adjoint.*

Proof. We may repeat the argument for the uniqueness of the Fréchet derivative. Fix an $x \in D$ and $y \in Y$. For any $h \in X$, we may compute:

$$\begin{aligned}|\langle h, f_x^* y - g_x^* y\rangle| &\leqslant |\langle h, f_x^* y\rangle - \langle f(x+h) - f(x), y\rangle| + |\langle h, g_x^* y\rangle - \langle f(x+h) - f(x), y\rangle| \\ &= o(\|h\|)\,.\end{aligned}$$

Let $v := f_x^* y - g_x^* y$, and suppose that $v \neq 0$. Then $\frac{|\langle \epsilon v, v\rangle|}{\epsilon} = \|v\|^2$ is not $o(|\epsilon|)$ as $\epsilon \to 0$, so $v = 0$. □

This proof suggests an interpretation of the local adjoint. For small deviations away from x, we have a near equality $\langle f(x+h) - f(x), y\rangle \approx \langle h, f_x^*(y)\rangle$. Moreover, f_x^* is the best approximation, as it is the only $o(\|h\|)$ approximation.

Example 7.2.11. Suppose $A : X \to Y$ is a bounded linear map. Then A has a local adjoint A^* on all of Y, which is given by the usual linear adjoint.

Example 7.2.12. Suppose $f : X \to Y$ is a continuous affine map, given by $f(x) = Ax + b$, where $A : X \to Y$ is a bounded linear operator, and $b \in Y$ is a fixed vector. Then f has a local adjoint f^* all of Y, which is again given by the linear adjoint A^*.

Both of the previous examples may be viewed as special cases of the following proposition.

Proposition 7.2.13. *Suppose* $f : X \to Y$ *is Fréchet differentiable on* D*. Then* f *has a unique local adjoint* f^* *on all of* V *given by* $f_x^* = (D_x f)^*$*.*

Proof. We simply check by the Cauchy-Schwartz inequality that

$$|\langle f(x+h) - f(x), y\rangle - \langle h, (D_x f)^* y\rangle| \leqslant \|f(x+h) - f(x) - D_x f(h)\|\|y\| = o(\|h\|).$$

□

This proposition justifies that the C^0-nonlinear sheaf Laplacian of a C^0-nonlinear Hilbert sheaf is a natural generalization of a Hilbert sheaf. At each point $\mathbf{x} \in C^0(\mathcal{G};\mathcal{F})$ where δ is differentiable, the generalized gradient $\partial V_{\mathcal{F}}(\mathbf{x})$ will agree with the usual gradient

$$\nabla V_{\mathcal{F}}(\mathbf{x}) = \delta_{\mathbf{x}}^* \delta \mathbf{x}\,.$$

When δ is locally Lipschitz, and hence differentiable on a dense G_δ-set (and almost everywhere differentiable with respect to Lebesgue measure when $C^0(\mathcal{G};\mathcal{F})$ is finite dimensional), the dynamics of $\dot{\mathbf{x}} = -\mathcal{L}\mathbf{x}$ look like the linear case locally.

Remark 7.2.14. The converse to Proposition 7.2.13 is not true. In particular, the limit convergence required for the existence of a local adjoint is best thought of as occurring in the weak topology, while the convergence required for differentiation is in the strong topology. Example 7.2.15 provides a schema for functions which admit a local adjoints but are not differentiable. Consequently, one could define a different notion of a nonlinear Laplacian via $\mathcal{L}\mathbf{x} := \delta_{\mathbf{x}}^{*}\delta\mathbf{x}$, whenever $\delta_{\mathbf{x}}$ admits a local adjoint. However, this definition would be too restrictive, as it could not cover the unbounded linear case, and convergence results would be given in the weak topology.

Example 7.2.15. Let $X = Y = \ell^2(\mathbb{N})$, and let $w : X\backslash\{0\} \to Y$ be a continuous function with the following properties.

1. $w(x)$ is a unit vector for all $x \in X\backslash\{0\}$.
2. For all $n \in \mathbb{N}$, there is an $\epsilon = \epsilon(n) > 0$ such that $w(x)$ is supported on basis elements $\{e_j : j > n\}$ whenever $\|x\| < \epsilon$.

Now consider the continuous function $f : X \to Y$ given by

$$f(x) = \begin{cases} \|x\|w(x) & \text{if } x \neq 0 \\ 0 & \text{else.} \end{cases}$$

At $x = 0$, f has a local adjoint given by 0. However, f is not Fréchet differentiable (or even Gateaux differentiable) at 0.

7.2.2 *Generalities on nonlinear heat flow*

A generic C^0-nonlinear Hilbert sheaf—even when admitting a Laplacian—does not have enough structure to meaningfully study dynamics, such as proving global existence of solutions to the heat flow or finding long-time asymptotic behavior. This difficulty is compounded by the ambiguity of the nonspecific choice of a generalized gradient. We now discuss the potential function $V_{\mathcal{F}}$ for different choices of generalized gradients.

Notation 7.2.16. Let ∂^C, ∂^*, and ∇ denote the Clarke gradient, convex subdifferential, and classical gradient respectively. Meanwhile, let ∂ denote a nonspecific (generalized) gradient. In future sections, these superscripts will be omitted when the choice of generalized gradient is clear from context.

Let $\mathcal{F}$ be a C^0-nonlinear Hilbert sheaf with Laplacian $\mathcal{L} = \partial V_{\mathcal{F}}$. A point $\mathbf{x}$ is a **generalized critical point** of $V_{\mathcal{F}}$ if $0 \in \partial V_{\mathcal{F}}(\mathbf{x})$. A generalized critical point $\mathbf{x}$ has zero energy ($V_{\mathcal{F}}(\mathbf{x}) = 0$) if and only if $\mathbf{x}$ is a global section of $\mathcal{F}$ ($\delta\mathbf{x} = 0$). We say that a generalized critical point $\mathbf{x}$ is a **generalized saddle point** if $\mathbf{x}$ is not a local minimum nor maximum of $V_{\mathcal{F}}$.

For the purpose of "consensus seeking", we will be interested in when a heat flow $\mathbf{x}_t$ converges to a global section. It is straightforward to construct examples—even smooth ones—where a C_0-nonlinear heat flow $\mathbf{x}_t$ does not converge to a global section. Indeed, global sections need not exist.

7.2.2.1 *Smooth gradients*

Let $\mathcal{F}$ be a C^0-nonlinear network sheaf with smooth restriction maps. The associated potential function $V_{\mathcal{F}} : C^0(\mathcal{G};\mathcal{F}) \to \mathbb{R}$ will be a smooth function as well; the natural choice of generalized gradient is the usual Fréchet gradient ∇. Call such a network sheaf, equipped with the choice of ∇ for the generalized gradient, a **smooth C^0-nonlinear network sheaf.** The Picard-Lindelöf theorem immediately ensures local existence of heat flows $\mathbf{x}_t$ on $C^0(\mathcal{G};\mathcal{F})$ which satisfy Equation (7).

Proposition 7.2.17. *Let $\mathcal{F}$ be a smooth C^0-nonlinear network sheaf. For every initial value $\mathbf{x}_0$, there is a heat flow $\mathbf{x}_t$ satisfying Equation (7).*

Remark 7.2.18. This local existence will hold more generally if each restriction map is twice continuously Fréchet differentiable. Under this condition, δ, and hence $V_{\mathcal{F}}$, will both be twice continuously differentiable, and the gradient $\nabla V_{\mathcal{F}}$ will itself be locally Lipschitz.

As the solution to a gradient descent, each heat flow $\mathbf{x}_t$ comes equipped with a Lyapunov function. In particular, since $\frac{d}{dt}V_{\mathcal{F}}(\mathbf{x}_t) = -\|\nabla V_{\mathcal{F}}(\mathbf{x}_t)\| \leqslant 0$, the potential function $V_{\mathcal{F}}$ is itself a Lyapunov function for $\mathbf{x}_t$. The asymptotic behavior of heat flow can thus be understood by LaSalle invariance [51]; in particular, every precompact heat flow $\mathbf{x}_t$ converges to a generalized critical point of $V_{\mathcal{F}}$.

Remark 7.2.19. When either $C^0(\mathcal{G};\mathcal{F})$ or $C^1(\mathcal{G};\mathcal{F})$ are finite dimensional, all bounded trajectories are precompact. This finite-dimensionality condition will hold exactly when all vertex stalks or all edge stalks of $\mathcal{F}$ are finite dimensional.

We now turn to understanding these generalized critical points. The C^0-nonlinear Laplacian $\mathcal{L}$ may be computed as $\mathcal{L}\mathbf{x} = \nabla V_{\mathcal{F}}(\mathbf{x}) = (D_{\mathbf{x}}\delta)^*\delta\mathbf{x}$. A cochain $\mathbf{x}$ such that $\delta\mathbf{x} \in \ker\big((D_{\mathbf{x}}\delta)^*\big)$, but $\delta\mathbf{x} \neq 0$ exactly corresponds to a generalized critical point which is not a global section. The derivative $D_{\mathbf{x}}\mathcal{L}$ is, in turn, given by

$$D_{\mathbf{x}}\mathcal{L}(h) = (D_{\mathbf{x}}\delta)^*(D_{\mathbf{x}}\delta)(h) + (D^2_{\mathbf{x}}\delta(h,\cdot))^*\delta(\mathbf{x})\,.$$

At a global section $\mathbf{x}$, $D_{\mathbf{x}}\mathcal{L}(h)$ simplifies to the first component $D_{\mathbf{x}}\mathcal{L}(h) = (D_{\mathbf{x}}\delta)^*(D_{\mathbf{x}}\delta)(h)$. Therefore the derivative $D_{\mathbf{x}}\mathcal{L} : C^0(\mathcal{G};\mathcal{F}) \to C^0(\mathcal{G};\mathcal{F})$ is a positive operator at each global section, and has a non-negative spectrum. In general, this spectrum will fail to be strictly positive; the global section will be point in a **consensus manifold** of points where $V_{\mathcal{F}} = 0$. If 0 is isolated in the spectrum $\sigma(D_{\mathbf{x}}\mathcal{L})$ at a global section $\mathbf{x}$, one may apply the center manifold theorem [104, Theorem 7.1] to study local asymptotic stability of flows.

7.2.2.2 *Convex subdifferentials*

Let $\mathcal{F}$ be a C^0-nonlinear network sheaf such that the associated potential function $V_{\mathcal{F}}$ is lower-semicontinuous and convex. For such a potential function, the convex subdifferential is a suitable choice of generalized gradient. For the potential function $V_{\mathcal{F}}$ to be convex is an extremely strong condition, and the corresponding dynamics are straightforward to analyze.

Proposition 7.2.20. *If $V_{\mathcal{F}}$ is convex, the heat flow $\dot{\mathbf{x}} \in \partial^* V_{\mathcal{F}}(\mathbf{x})$ has a unique global strong solution for all initial values $\mathbf{x}_0 \in \mathrm{Dom}(\delta)$, and a unique global mild solution for all $\mathbf{x}_0 \in \overline{\mathrm{Dom}(\delta)}$. Moreover, $\mathbf{x}_t$ converges to a global minimum of V_F in the weak topology as $t \to \infty$.*

Proof. If $V_{\mathcal{F}}(\mathbf{x}) = \frac{1}{2}\|\delta(\mathbf{x})\|^2$ is convex, then since δ has closed graph, $V_{\mathcal{F}}$ is a proper, convex, lower semicontinuous functional. The Brézis-Komura theorem [74] guarantees a unique heat flow that obeys the dynamics $\dot{\mathbf{x}} \in -\mathcal{L}\mathbf{x} = -\partial^* V_{\mathcal{F}}(\mathbf{x})$ for all initial values $\mathbf{x}_0 \in \overline{\mathrm{Dom}(\delta)}$, as a mild solution. The solution is classical when $\mathbf{x}_0 \in \mathrm{Dom}(\delta)$. By Opial's lemma [97, Theorem 6.3], there are no critical points which are not global minima of $V_{\mathcal{F}}$, and a heat flow $\mathbf{x}_t$ will weakly converge to such a global minimum $\mathbf{x}_*$. □

Remark 7.2.21. Under additional assumptions such as $V_{\mathcal{F}}$ being even [16] or strong convexity [97], $\mathbf{x}_t$ will converge strongly to a global minimum. This also constitutes an alternative proof of the existence and convergence of heat flows for linear Hilbert sheaves.

7.2.2.3 *Clarke gradients*

When the potential function $V_{\mathcal{F}}$ is a globally defined locally Lipschitz continuous, the Clarke gradient ∂^C is an appropriate choice for a generalized gradient. Indeed, for a locally Lipschitz potential function V on a generic Banach space X, the Clarke gradient $\partial^C V(x)$ is non-empty for all $x \in X$ [70]. However, care must be taken with the solution concept for the heat flow $\dot{\mathbf{x}} \in -\partial^C V_{\mathcal{F}}(\mathbf{x})$, as the right hand side is both multivalued and discontinuous. We will primarily consider Filippov solutions [39]. For a friendly introduction to non-smooth and discontinuous differential equations, see the following paper of Cortés [32]. On a Hilbert space X, let $\mathrm{FinSub}(X)$ denote the collection of all finite dimensional linear subspaces of X, and let μ_U denote Lebesgue measure on $U \in \mathrm{FinSub}(X)$.

Definition 7.2.22. Let X be a Hilbert space, and let $f : X \to X$ be a locally Lipschitz function. The **Filippov set-valued map** F[f](x) is defined by

$$F[f](x) := \bigcap_{\epsilon>0} \bigcap_{U \in \mathrm{FinSub}(X)} \bigcap_{\substack{S \subseteq U \\ \mu_U(S)=0}} \overline{\mathrm{cvx}\{f\big((x+V) \cap B_\epsilon(x) \backslash (x+S)\big)\}}\,,$$

for each $x \in X$, where $B_\epsilon(x)$ is the open ball of radius ϵ centered at x. When X is finite dimensional, this definition reduces to

$$F[f](x) := \bigcap_{\epsilon>0} \bigcap_{\substack{S \subseteq X \\ \mu(S)=0}} \overline{\mathrm{cvx}\{f\big(B_\epsilon(x) \backslash S\big)\}}\,.$$

A **Filippov solution** to the differential inclusion $\dot{\mathbf{x}} \in F[f](x)$ is an absolutely continuous curve $x(t)$ for $t \in [0, T]$ such that $\dot{x}(t) \in F[f](x(t))$ for almost every $t \in [0, T]$.

Remark 7.2.23. We make a few remarks about Filippov solutions.

1. The complication in the infinite dimensional definition comes from the fact that there is no canonical nullset structure on an infinite dimensional Hilbert space; we make the standard choice of using the canonical nullsets of all finite dimensional slices via Lebesgue measure.

2. When $f : X \to X$ is the (densely-defined) gradient of a locally Lipschitz potential function $V : X \to \mathbb{R}$, the Filippov function $F[f](x)$ is exactly the Clarke generalized gradient $\partial^C V(x)$.

3. Finally, the multivalued nature of the Filippov function makes the *uniqueness* of Filippov solutions more subtle than existence.

Proposition 7.2.24. *Let $\mathcal{F}$ be a C^0-nonlinear Hilbert sheaf on a network $\mathcal{G}$. If δ is locally Lipschitz and defined on an open set, the heat flow $\dot{\mathbf{x}} \in -\partial^C V_{\mathcal{F}}(\mathbf{x})$ has a local Filippov solution around all $\mathbf{x}_0 \in \mathrm{Dom}(\delta)$.*

Proof. Since δ is locally Lipschitz, then $V_{\mathcal{F}}(\mathbf{x}) = \frac{1}{2}\|\delta(\mathbf{x})\|^2$ is locally Lipschitz as well. It follows that the negative Clarke gradient $-\partial^C V(\mathbf{x})$ is upper semicontinuous [24]. Locally existence of a Filippov solution with initial point $\mathbf{x}_0 \in \mathrm{Dom}(\delta)$ up to the boundary of a closed ball of positive radius inside $C^0(\mathcal{G}; \mathcal{F})$ immediately follows [39, Theorem 7.2]. □

Remark 7.2.25. Under a variety of additional hypotheses, such as bounded sublevel sets of $V_{\mathcal{F}}$, local existence of Filippov solutions may be strengthened to global existence.

The potential function $V_{\mathcal{F}}$ acts as a global Lyapunov function for the heat flow $\mathbf{x}_t$, granting access to a variety of tools for analysis. See [32] for an overview.

7.2.3 *Riemannian consensus*

The structure of C^0-nonlinear Hilbert sheaves can be generalized beyond $\mathbf{Hilb}_{0,\Bbbk}$-valued vertex and edge stalks. When defining the C^0-heat flow via smooth gradient descent, we implicitly identify each vertex stalk with its tangent space at each point. By making this identification explicit, one may work with Riemannian manifolds for stalks.

Definition 7.2.26. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite network. A **Riemannian network sheaf** consists of the following data.

- For each vertex $v \in \mathcal{V}$, a smooth Riemannian manifold $\mathcal{F}(v) := M_v$ with metric g_v.
- For each edge $e \in \mathcal{E}$, a smooth geodesically complete Riemannian manifold $\mathcal{F}(e) := M_e$ with metric g_e.
- For each covering morphism $f : v \to e$, a smooth map $\mathcal{F}_f : M_v \to M_e$.

From the data of a Riemannian network sheaf $\mathcal{F}$, one may construct a coboundary operator as follows. Let $C^0(\mathcal{G};\mathcal{F}) := \prod_{v\in\mathcal{V}} M_v$ and $C^1(\mathcal{G};\mathcal{F}) := \prod_{e\in\mathcal{E}} M_e$ denote the product manifolds of vertex and edge stalks equipped with their product metrics. We abbreviate these spaces to C^0 and C^1 for notational clarity. One may define the coboundary map $\delta : C^0 \to C^1 \times C^1$ as follows. Letting $(\delta\mathbf{x})_e \in M_e \times M_e$ denote the image of δ in the *pair of components* corresponding to the edge e in $C^1 \times C^1$, we define δ by

$$(\delta\mathbf{x})_e := \left(\mathcal{F}_{s(e)}\mathbf{x}_{s(e)}, \mathcal{F}_{t(e)}\mathbf{x}_{t(e)}\right),$$

where $s(e)$ and $t(e)$ denote the source and target of the oriented edge e, and $\mathcal{F}_{s(e)}, \mathcal{F}_{t(e)}$ denote the corresponding covering morphisms in $\mathcal{G}$.

Remark 7.2.27. Since the edge stalks have no intrinsic notion of subtraction, we record both components separately.

Let π_1 and π_2 denote the projections onto the first and second component of $C^1 \times C^1$. The coboundary operator defines a potential function

$$\begin{aligned} V_{\mathcal{F}}(\mathbf{x}) &:= \frac{1}{2} d^2_{C^1}\left(\pi_1\delta\mathbf{x}, \pi_2\delta\mathbf{x}\right) \\ &= \frac{1}{2}\sum_{e\in\mathcal{E}} d^2_e\left(\mathcal{F}_{s(e)}\mathbf{x}_{s(e)}, \mathcal{F}_{t(e)}\mathbf{x}_{t(e)}\right). \end{aligned}$$

Definition 7.2.28. Let $\mathcal{F}$ be a Riemannian network sheaf. A zero-cochain $\mathbf{x} \in C^0(\mathcal{G};\mathcal{F})$ is a **global section** of $\mathcal{F}$ if the potential function $V_{\mathcal{F}}(\mathbf{x}) = 0$. That is, a global section is exactly a choice of point $\mathbf{x}_v$ in each vertex-manifold M_v such that for each edge e, $\mathcal{F}_{s(e)}\mathbf{x}_{s(e)} = \mathcal{F}_{t(e)}\mathbf{x}_{t(e)}$.

Remark 7.2.29. A global section encodes a solution to a consensus problem. Borrowing the language of opinion dynamics à la [55], one may envision each vertex manifold M_v as a space of parameters for some agent, which are then expressed through the restriction maps out of M_v. A global section $\mathbf{x} \in C^0(\mathcal{G};\mathcal{F})$ exactly corresponds to a choice of parameter for each agent that satisfy *expressed consensus* in each edge manifold M_e. Such a problem of finding agreement in a networked collection of points in a manifold is known as the **Riemannian consensus problem** [22, 76, 119, 120]. Distributed algorithms for solving the Riemannian consensus problem has been extensively studied for the constant network sheaf, where there is a fixed manifold M such that $M_\sigma = M$ for all $\sigma \in \mathcal{V} \cup \mathcal{E}$, and all restriction maps are the identity.

Remark 7.2.30. Not every Riemannian network sheaf will admit a global section; indeed, it is possible for two for two restriction maps into the same edge to share no points in their image.

Example 7.2.31. We turn to an example from information geometry [5]. For real parameters $\mu \in \mathbb{R}$ and $\sigma > 0$, let $N(\mu, \sigma)$ denote the normal distribution with mean μ and standard deviation σ. Let $\mathcal{N}_1 := \{N(\mu, \sigma) \ : \ (\mu, \sigma) \in \mathbb{R} \times \mathbb{R}_{>0}\}$ denote the collection of all univariate normal distributions. We endow $\mathcal{N}_1$ with the structure of a Riemannian manifold as follows. For a choice of parameters (μ, σ) let $p(x; \mu, \sigma)$ denote the corresponding Gaussian density function and $\ell(x; \mu, \sigma) := \log p(x; \mu, \sigma)$ its logarithm. The **Fisher information metric** on $\mathcal{N}_1$ may be computed as

$$g = \begin{bmatrix} \mathbb{E}[\partial_\mu \ell \partial_\mu \ell] & \mathbb{E}[\partial_\mu \ell \partial_\sigma \ell] \\ \mathbb{E}[\partial_\sigma \ell \partial_\mu \ell] & \mathbb{E}[\partial_\sigma \ell \partial_\sigma \ell] \end{bmatrix} = \begin{bmatrix} \frac{1}{\sigma^2} & 0 \\ 0 & \frac{2}{\sigma^2} \end{bmatrix}.$$

The Fisher information metric endows $\mathcal{N}_1$ with the structure of a two-dimensional Riemannian manifold.

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite network, and consider the constant Riemannian Hilbert sheaf $\underline{\mathcal{N}_1}$ consisting of the following data.

- A manifold stalk $M_\sigma := \mathcal{N}_1$, equipped with the Fisher information metric, for each $\sigma \in \mathcal{V} \cup \mathcal{E}$.
- The identity map $\mathcal{F}_f := \mathrm{id} : \mathcal{N}_1 \to \mathcal{N}_1$ for each covering map $f : v \to e$.

A global section of this Riemannian network sheaf exactly corresponds to an identical choice of a normal distribution $N(\mu, \sigma) \in M_v$ for each vertex.

Recall that when $f : M \to \mathbb{R}$ is a smooth function on a Riemannian manifold (M, g), one may define the **gradient** of f at $x \in M$ as the unique tangent vector $\nabla_x f(x) \in T_x M$ such that for all $v \in T_x M$ and smooth curves $\gamma : [0, 1] \to M$ with $\gamma(0) = x$ and $\dot{\gamma}(0) = v$, there is an equality:

$$g_x(\nabla_x f(x), v) = \frac{d}{dt} f(\gamma(t))\big|_{t=0}.$$

When the manifold M is geodesically complete, there is a globally defined distance function $d : M \times M \to \mathbb{R}$ for which $d(x, y)$ measures the length of the shortest geodesic connecting x and y. On the product manifold $M \times M$ equipped with the product metric $g \oplus g$ one may define a function $c(x, y) := \frac{1}{2} d^2(x, y)$. For every pair x, y that are connected by a unique length-minimizing geodesic, the function c has a gradient given by

$$\nabla_{(x,y)} c(x, y) = (-\operatorname{Log}_x y, -\operatorname{Log}_y x) \in T_x M \oplus T_y M$$

where $\operatorname{Log}_p q$ is the **Riemannian logarithm**; for $p \neq q$ in M, $\operatorname{Log}_p q := \dot{\gamma}_{pq}(0) d(p, q)$ where γ_{pq} is the unique minimal unit-speed geodesic joining p to q. We adopt the convention that when $p = q$, the Riemannian logarithm $\operatorname{Log}_p q = \operatorname{Log}_p p = 0$. When y is in the cut-locus of x, there may be multiple distinct length-minimizing geodesics connecting x and y. For such a pair, $c(x, y)$ fails to have a well-defined gradient, but admits a Clarke generalized gradient

$$\partial^C_{(x,y)} c(x, y) = \overline{\operatorname{cvx}\{(-\operatorname{Log}^\gamma_x y, -\operatorname{Log}^\gamma_y x) \ : \ \gamma \text{ is a length-minimizing geodesic } x \rightsquigarrow y\}}$$

where $\operatorname{Log}^\gamma$ denotes the Riemannian logarithm with respect to a choice of length-minimizing geodesic.

When $\mathbf{x}_{s(e)}$ and $\mathbf{x}_{t(e)}$ are sufficiently close together for each edge $e \in \mathcal{E}$, the potential function $V_{\mathcal{F}}$ has a well-defined gradient at $\mathbf{x}$ given on each vertex $v \in \mathcal{V}$ by

$$\begin{aligned}
(\nabla_{\mathbf{x}} V_{\mathcal{F}}(\mathbf{x}))_v = &\sum_{\substack{e \in \mathcal{E} \\ s(e) = v}} \left(D_{\mathbf{x}_v} \mathcal{F}_{s(e)}\right)^* \left(-\operatorname{Log}_{\mathcal{F}_{s(e)} \mathbf{x}_v} \left(\mathcal{F}_{t(e)} \mathbf{x}_{t(e)}\right)\right) \\
&+ \sum_{\substack{e \in \mathcal{E} \\ t(e) = v}} \left(D_{\mathbf{x}_v} \mathcal{F}_{t(e)}\right)^* \left(-\operatorname{Log}_{\mathcal{F}_{t(e)} \mathbf{x}_v} \left(\mathcal{F}_{s(e)} \mathbf{x}_{s(e)}\right)\right).
\end{aligned}$$

The gradient $\nabla_{\mathbf{x}} V_{\mathcal{F}}(\mathbf{x})$ defines a vector field on C^0 on the set of zero-cochains $\mathbf{x} \in C^0$ such that there is a unique length-minimizing geodesic between $\mathcal{F}_{s(e)} \mathbf{x}_{s(e)}$ and $\mathcal{F}_{t(e)} \mathbf{x}_{t(e)}$ for all edges $e \in \mathcal{E}$.

Definition 7.2.32. A **Hadamard manifold** is a Riemannian manifold (M, g) that is complete, simply connected, and has non-positive sectional curvature at every point.

Remark 7.2.33. Every Hadamard manifold is a finite dimensional **Hadamard space**—a nonlinear generalization of a Hilbert space. Specifically, every Hadamard space is complete metric space, such that for every pair of points x, y, there is a point m, called the **midpoint** of x and y, such that for all z:

$$d(z, m)^2 + \frac{d(x, y)^2}{4} \leqslant \frac{d(z, x)^2 + d(z, y)^2}{2}.$$

Every Hilbert space with its norm-induced distance function is a Hadamard space with midpoint given by $m = \frac{x+y}{2}$.

Proposition 7.2.34. *Let $\mathcal{F}$ be a Riemannian network manifold on a network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. If M_e is a Hadamard manifold for all $e \in \mathcal{E}$, then the gradient vector field $\nabla_{\mathbf{x}} V_{\mathcal{F}}(\mathbf{x})$ is globally defined.*

Proof. In a Hadamard manifold M, every pair of points $x, y \in M$ is joined by a unique length-minimizing geodesic. It follows that the cut-locus of every point $x \in M$ is empty. The the finite product of Hadamard manifolds, equipped with the product metric, is a Hadamard manifold. Therefore the space of one-cochains C^1 is a Hadamard manifold, and the gradient $\nabla_{\mathbf{x}} V_{\mathcal{F}}(\mathbf{x})$ is well-defined for all $\mathbf{x} \in C^0$. □

When every edge manifold is a Hadamard space, one may consider the gradient descent on C^0

$$\dot{\mathbf{x}} = -\nabla_{\mathbf{x}} V_{\mathcal{F}}(\mathbf{x}), \tag{8}$$

for any choice of initial cochain $\mathbf{x}_0 \in C^0$. Under suitable conditions on the vertex manifolds M_v and the restriction maps $\mathcal{F}_f : M_v \to M_e$, every initial cochain has a globally defined gradient flow $\mathbf{x}_t : \mathbb{R}_{\geqslant 0} \to C^0$ which converges to a global section.

Theorem 7.2.35. *Let $\mathcal{F}$ be a Riemannian network sheaf on a network $\mathcal{G}$ which satisfies the following conditions.*

(i) Every vertex and edge manifold M_σ is a Hadamard manifold.

(ii) The set of global sections $\Gamma(\mathcal{F})$ is non-empty.

(iii) Every restriction map is a totally-geodesic isometry.

For every zero-cochain $\mathbf{x}_0 \in C^0$, there is a globally defined negative gradient flow $\mathbf{x}_t$ satisfying Equation (8), initialized at $\mathbf{x}_0$, which converges to a global section.

Proof. The potential function $V_{\mathcal{F}}$ is a continuous function on C^0. We now check that $V_{\mathcal{F}}$ is geodesically convex. Let $\gamma(t)$ be a geodesic in C^0, with $0 \leqslant t \leqslant 1$. Each component $\gamma_v(t)$ is a geodesic in M_v. Since all restriction maps are geodesically complete, for each edge e, the paths $\eta_{s(e)}(t) := \mathcal{F}_{s(e)}(\gamma_{s(e)}(t))$ and $\eta_{t(e)}(t) := \mathcal{F}_{t(e)}(\gamma_{t(e)}(t))$ are geodesics in M_e. The metric of a Hadamard space is jointly convex along geodesics [9, Section 1.2], giving the inequality

$$d_e\big(\eta_{s(e)}(t), \eta_{t(e)}(t)\big) \leqslant (1-t)d_e\big(\eta_{s(e)}(0), \eta_{t(e)}(0)\big) + td_e\big(\eta_{s(e)}(1), \eta_{t(e)}(1)\big)$$

for all $t \in [0,1]$, where d_e is the distance function on M_e. It follows that the map

$$t \mapsto \frac{1}{2} d_e\left(\eta_{s(e)}(t), \eta_{t(e)}(t)\right)^2$$

is convex in t. Summing over all edges $e \in \mathcal{E}$ yields that $V_{\mathcal{F}}$ is geodesically convex.

Since $V_{\mathcal{F}}$ is a continuous and geodesically convex function on a Hadamard manifold C^0, V has globally defined negative gradient flows $\mathbf{x}_t$ satisfying Equation (8). Moreover, since $V_{\mathcal{F}}$ has a global sections (and thus obtains its minimum), this gradient descent converges to a global section as $t \to \infty$ [9, Theorem 5.1.16]. □

Remark 7.2.36. Many of these conditions can be weakened in practice. Moreover, such results can be extended, with care, to the infinite dimensional setting via Hadamard space theory.

Example 7.2.37. We return to Example 7.2.31. Recall that the constant Riemannian network sheaf $\underline{\mathcal{N}_1}$ on a network $\mathcal{G}$ has all vertex and vertex and edge stalks given by the univariate Gaussian statistical manifold $\mathcal{N}_1$, equipped with the Fisher information metric. All restriction maps are given by the identity function.

The univariate Gaussian statistical manifold $\mathcal{N}_1$ is isometric, up to a constant scaling factor, to the two-dimensional hyperbolic plane [5]. The hyperbolic plane is a Hadamard manifold. Therefore by Theorem 7.2.35, for any choice of an initial zero-cochain $\mathbf{x}_0 \in C^0$, the negative gradient flow $\mathbf{x}_t$ of the potential function $V_{\underline{\mathcal{N}_1}}$ converges to a global section $\mathbf{x}_\infty$ of $\underline{\mathcal{N}_1}$.

7.3 AFFINE SHEAVES

As a straightforward example of a class of C^0-nonlinear Hilbert sheaves, we may consider the class of Hilbert sheaves with affine maps.

Definition 7.3.1. An **affine network sheaf** on a finite graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a C^0-nonlinear network Hilbert sheaf whose restriction maps are densely defined affine. That is, each restriction map $\mathcal{F}_f : \mathcal{F}(v) \to \mathcal{F}(e)$ can be written $\mathcal{F}_f(x) = A_f x + b_f$, where A_f is a densely defined linear operator, and $b_f \in \mathcal{F}(e)$.

Remark 7.3.2. As usual, the most difficult in part in confirming that an assignment of stalks and restriction maps is an affine network sheaf is to check that the coboundary operator $\mathring{\delta}$ is closable. In this case, we observe that $\mathring{\delta}$ is closable if and only if for every edge e with incoming covering maps f, g, we have that $\begin{bmatrix} -A_f & A_g \end{bmatrix}$ is closable.

7.3.1 *Affine dynamics*

When $\mathcal{F}$ is an affine network sheaf, the potential function $V_{\mathcal{F}}$ can be seen to be convex by a straightforward computation. By Proposition 7.2.20, we get heat flows for all initial $\mathbf{x}_0 \in C^0(\mathcal{G};\mathcal{F})$. Under suitable conditions, we can meaningfully analyze this heat flow. Note that the coboundary operator $\delta : C^0(\mathcal{G};\mathcal{F}) \to C^1(\mathcal{G};\mathcal{F})$ is itself an affine map, and can be written as $\delta(\mathbf{x}) := A_\delta \mathbf{x} + b_\delta$, where A_δ is a closed densely defined linear operator. Say that an affine Hilbert sheaf is **proper** if the following conditions hold.

(i) $b_\delta \in \mathrm{Dom}(A_\delta^*)$.

(ii) $\mathcal{R}(A_\delta) \subseteq \mathrm{Dom}(A_\delta^*)$.

When $\mathcal{F}$ is proper affine, the heat flow can be written as

$$\dot{\mathbf{x}} = -A_\delta^*(A_\delta \mathbf{x} + b_\delta)$$

for all initial conditions $\mathbf{x}_0$. We may also prove the following convergence result.

Proposition 7.3.3. *Let $\mathcal{F}$ be a proper affine network sheaf with coboundary operator $\delta\mathbf{x} = A_\delta \mathbf{x} + b_\delta$. We may characterize the asymptotic behavior of a heat flow $\mathbf{x}_t$ with initial value $\mathbf{x}_0$ as follows.*

(i) If $b_\delta \in \mathcal{R}(A_\delta)$, then $\mathbf{x}_t$ converges to the nearest global section to $\mathbf{x}_0$. That is, $\mathbf{x}_\infty := \lim_{t\to\infty} \mathbf{x}_t$ is the nearest point to $\mathbf{x}_0$ such that $A_\delta \mathbf{x}_\infty = -b_\delta$.

(ii) If $\mathcal{R}(A_\delta)$ has closed range, then $\mathbf{x}_\infty$ is the nearest OLS solution to the inconsistent linear system $A_\delta \mathbf{x} = -b_\delta$.

Proof. If $b_\delta \in \mathcal{R}(A_\delta)$, then we may write $-b_\delta = Ac$ for some $c \in C^0(\mathcal{G};\mathcal{F})$, and the space of solutions $\{\mathbf{x} : A_\delta \mathbf{x} + b_\delta = 0\} = c + \ker(A_\delta)$. Letting $\mathbf{y}_t := \mathbf{x}_t + c$, we see that $\dot{\mathbf{y}}_t = -A_\delta^* A_\delta \mathbf{y}_t$, and evolves according to the C_0-semigroup $\mathbf{y}_t = e^{-tA_\delta^* A_\delta}\mathbf{y}_0$, and converges to $\mathbf{y}_\infty = P_{\ker A_\delta}\mathbf{y}_0$, where $P_{\ker A_\delta}$ is the orthogonal projection onto the kernel of A_δ. It follows that $\mathbf{x}_\infty = \mathbf{y}_\infty - c$ is the orthogonal projection of $\mathbf{x}_0$ onto the solution space $c + \ker(A_\delta)$.

If $\mathcal{R}(A_\delta)$ has closed range, then A_δ admits a bounded, globally defined Moore-Penrose pseudoinverse $A_\delta^\dagger : C^1(\mathcal{G};\mathcal{F}) \to C^0(\mathcal{G};\mathcal{F})$. Set $\mathbf{x}^{(s)} := -A_\delta^\dagger b_\delta$. Note that $A_\delta^* A_\delta A_\delta^\dagger b_\delta = A_\delta^* b_\delta$, so $\mathbf{x}^{(s)}$ is a critical point for the heat dynamics. Thus, the flow $\mathbf{x}_t$ can be written as $\mathbf{x}_t := e^{-A_\delta^* A_\delta}\mathbf{x}_0 + \mathbf{x}^{(s)}$, which converges to $\mathbf{x}_t := P_{\ker A_\delta}\mathbf{x}_0 + \mathbf{x}^{(s)}$. This point may be easily identified with the nearest OLS solution to $A_\delta \mathbf{x} = -b_\delta$ to the initial value $\mathbf{x}_0$. □

Remark 7.3.4. These dynamics are already quite useful from the standpoint of cellular sheaf theory. One dynamics-centric perspective on network sheaves is that global sections encode solutions to networked systems of homogeneous linear equations. Heat dynamics then provide a distributed approach to finding a solution to the networked system. When we replace the linear restriction maps with affine maps, the C^0-heat flow now encodes the best OLS solution to a (potentially inconsistent) *inhomogeneous* system of equations. This approach to nonlinear dynamics increases the expressive power of network sheaves.

7.3.2 *Affine cohomology*

While C^0-nonlinear Hilbert sheaves are largely motivated from a "dynamics first" perspective, affine network sheaves still admit a cohomological interpretation. Using the language of torsors, we may understand the structure of affine sheaves and their global sections in terms of the cohomology of the underlying linear maps. This perspective and interpretation are in line with recent work of Ghrist and Cooperband on network torsors, which they used to study visual paradoxes [46]. The broader connection between torsors, sheaves, and cohomology are well established [47, 124]. For a gentle introduction to torsors, see John Baez' expository piece [11].

For simplicity, we assume we are working with an affine network sheaf $\mathcal{F}$ with only finite dimensional stalks, but this cohomology and interpretation can be extended to the bounded infinite dimensional setting given suitable closed-range assumptions. Recall that an **affine space** is a triple (A, W, μ), where A is a set, W is a vector space, and $\mu : W \times A \to A$ is a free, transitive action of the additive topological group of W on A. It immediately follows that the map $\mu(-, x) : W \to A$ is a bijection for each $x \in A$.

Notation 7.3.5. For every $x \in A$ and $w \in W$, we write the action $w + x := \mu(w, x)$. For each pair $x, y \in A$, we denote the unique $w \in W$ such that $w + x = y$ by $y - x$.

We think of the affine space A as a copy of W "without origin"—that is, as a torsor over a one-point space with respect to the additive topological group $(W, +)$, where one can measure differences but not absolute location.

Let (A_0, W_0, μ_0) and (A_1, W_1, μ_1) be affine spaces. A set map $f : A_0 \to A_1$ is an **affine map** if there is a linear map $L : W_0 \to W_1$ such that $f(x + w) = f(x) + Lw$ for all $x \in A_0$ and $w \in W_0$.

Remark 7.3.6. If the affine spaces (A_j, W_j, μ_j) are merely viewed as torsors with respect to the additive group structure of a vector space, the class of affine maps is a strict subset of the the class of morphisms of torsors. We require the map $L : W_0 \to W_1$ to be linear—not merely a group homomorphism of the additive groups $(W_0, +)$ and $(W_1, +)$. However, treating affine spaces as torsors with respect to the topological group structure enforces linearity of all torsor morphisms.

We now apply the recent notion of inhomogeneous network torsors [46] to affine network sheaves.

Definition 7.3.7 (Affine network torsor)**.** Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite network, and $\mathcal{F} : \mathcal{G} \to \mathbf{FinHilb}_{\Bbbk}$ be a finite dimensional weighted network sheaf. A finite dimensional **affine network torsor** over $\mathcal{G}$ is a cellular sheaf of topological spaces $\mathcal{A} : \mathcal{G} \to \mathbf{Top}$ such that the following conditions hold.

(i) Each stalk $\mathcal{A}(\sigma)$ is equipped with a free transitive continuous group action $\mu : \mathcal{F}(\sigma) \times \mathcal{A}(\sigma) \to \mathcal{A}(\sigma)$ of the additive topological group of $\mathcal{F}(\sigma)$. This gives each stalk the structure of an affine space.

(ii) For each covering morphism $f : v \to e$ in $\mathcal{G}$, the restriction map $\mathcal{A}_f : \mathcal{A}(v) \to \mathcal{A}(e)$ is compatible with the topological group actions in the sense that

$$\mathcal{A}_f(w + x) = \mathcal{F}_f(w) + \mathcal{A}_f(x)$$

for all $x \in \mathcal{A}(v)$ and $w \in \mathcal{F}(v)$. That is, $\mathcal{A}_f$ is equivariant up to the linear map $\mathcal{F}_f$.

We call $\mathcal{F}$ the **linear structure sheaf** of the affine network torsor $\mathcal{A}$. When we wish to specify the underlying linear structure sheaf, we call $\mathcal{A}$ an $\mathcal{F}$-**affine network torsor**.

Remark 7.3.8. A finite dimensional $\mathcal{F}$-affine network torsor consists of essentially the same data as an affine network sheaf (Definition 7.3.1) whose affine restriction maps have underlying linear components given by $\mathcal{F}$. However, the lack of objective origin in an affine network torsor means that multiple affine network sheaves may have the same torsor structure.

Remark 7.3.9. Definition 7.3.7 may be viewed as a specialization of inhomogeneous network torsors [46, Definition 6.1] to the additive groups of vector spaces, subject to the additional constraint that all actions and maps are continuous.

Definition 7.3.10. Let $\mathcal{A}, \mathcal{A}'$ be finite dimensional affine network torsors over the same linear structure sheaf $\mathcal{F} : \mathcal{G} \to \mathbf{FinHilb}_{\Bbbk}$. A **morphism** of $\mathcal{F}$-affine network torsors from $\mathcal{A}$ to $\mathcal{A}'$ is a natural transformation $\phi : \mathcal{A} \Rightarrow \mathcal{A}'$ whose component maps $\phi_\sigma : \mathcal{A}(\sigma) \to \mathcal{A}'(\sigma)$ are $\mathcal{F}(\sigma)$-equivariant.

Remark 7.3.11. Since each stalk-map $\phi_\sigma : \mathcal{A}(\sigma) \to \mathcal{A}'(\sigma)$ must be $\mathcal{F}(\sigma)$-equivariant, the linear portion ϕ_σ must be the identity map, making ϕ_σ a translation. That is, there is some $w \in \mathcal{F}(\sigma)$ such that $\phi_\sigma(-) = w + (-)$. Since a translation is a bijection, all morphisms of $\mathcal{F}$-affine network torsors are isomorphisms.

We may recover a classification result for $\mathcal{F}$-affine network torsors similar to [46, Theorem 6.2].

Theorem 7.3.12. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite network, and $\mathcal{F} : \mathcal{G} \to \mathbf{FinHilb}_{\Bbbk}$ be a finite dimensional weighted network sheaf. There is a canonical bijection*

$$H^1(\mathcal{G}; \mathcal{F}) \longleftrightarrow \{\textit{isomorphism classes of } \mathcal{F}\textit{-affine network torsors}\}.$$

Proof. Let $[\mathbf{y}] \in H^1(\mathcal{G};\mathcal{F})$ be a cohomology class with representative $\mathbf{y} \in C^1(\mathcal{G};\mathcal{F})$. We may build an $\mathcal{F}$-affine network torsor $\mathcal{A}^{\mathbf{y}} : \mathcal{G} \to \mathbf{Top}$ with the following data.

- Each stalk $\mathcal{A}^{\mathbf{y}}(\sigma)$ is the topological space underlying the Hilbert space $\mathcal{F}(\sigma)$.
- The group action on $\mathcal{A}^{\mathbf{y}}(\sigma)$ is exactly addition on $\mathcal{F}(\sigma)$.
- For a covering map $f : v \to e$ in $\mathcal{G}$, the restriction map $\mathcal{A}_f$ is given by

$$\mathcal{A}^{\mathbf{y}}_f(x) = \begin{cases} \mathcal{F}_f x & \text{if } f \text{ has positive orientation} \\ \mathcal{F}_f(x) - \mathbf{y}_e & \text{if } f \text{ has negative orientation}, \end{cases}$$

  where $\mathbf{y}_e$ is the component of $\mathbf{y}$ living in $\mathcal{A}^{\mathbf{y}}(e) = \mathcal{F}(e)$.

$\mathcal{A}^{\mathbf{y}}$ is easily seen to be an $\mathcal{F}$-network torsor. Moreover, when $\mathbf{y}, \mathbf{y}' \in [\mathbf{y}]$ are representatives of the same cohomology class of $H^1(\mathcal{G};\mathcal{F})$, there is an element $\mathbf{b} \in C^0(\mathcal{G};\mathcal{F})$ such that $\mathbf{y}' - \mathbf{y} = \delta\mathbf{b}$, where $\delta : C^0(\mathcal{G};\mathcal{F}) \to C^1(\mathcal{G};\mathcal{F})$ is the linear coboundary map induced by $\mathcal{F}$. From this $\mathbf{b}$, we derive an isomorphism of $\mathcal{F}$-affine network torsors $\phi : \mathcal{A}^{\mathbf{y}} \Rightarrow \mathcal{A}^{\mathbf{y}'}$ with naturality squares

$$\begin{array}{ccc}
\mathcal{A}^{\mathbf{y}}(v) & \xrightarrow{\phi_v : x \mapsto x + \mathbf{b}_v} & \mathcal{A}^{\mathbf{y}'}(v) \\
{\scriptstyle \mathcal{A}^{\mathbf{y}}_f}\downarrow & & \downarrow{\scriptstyle \mathcal{A}^{\mathbf{y}'}_f} \\
\mathcal{A}^{\mathbf{y}}(e) & \xrightarrow{\phi_e : x \mapsto x + \mathcal{F}_f \mathbf{b}_v} & \mathcal{A}^{\mathbf{y}'}(e) \\
{\scriptstyle \mathcal{A}^{\mathbf{y}}_g}\uparrow & & \uparrow{\scriptstyle \mathcal{A}^{\mathbf{y}'}_g} \\
\mathcal{A}^{\mathbf{y}}(u) & \xrightarrow{\phi_u : x \mapsto x + \mathbf{b}_u} & \mathcal{A}^{\mathbf{y}'}(u)
\end{array}$$

where $f : v \to e$ and $g : u \to e$ are assigned positive and negative orientations respectively.

Conversely, let $\mathcal{A}$ be an $\mathcal{F}$-affine network torsor. To construct a cocycle from $\mathcal{A}$, fix an "origin" $\mathbf{b}_v \in \mathcal{A}(v)$ for each vertex v. For an oriented edge $e = (u, v)$ with covering morphisms $g : u \to e$ and $f : v \to e$, take $\mathbf{y}_e := f(\mathbf{b}_v) - g(\mathbf{b}_u)$, yielding a cocycle $\mathbf{y} := (\mathbf{y}_e)_{e \in \mathcal{E}}$. For a different choice of origins $\{\mathbf{b}'_v\}_{v \in \mathcal{V}}$, one may check

$$\begin{aligned} f(\mathbf{b}'_v) - g(\mathbf{b}'_u) &= f\big((\mathbf{b}'_v - \mathbf{b}_v) + \mathbf{b}_v\big) - g\big((\mathbf{b}'_u - \mathbf{b}_u) + \mathbf{b}_u\big) \\ &= \big(f(\mathbf{b}_v) - g(\mathbf{b}_u)\big) + \big(\mathcal{F}_f(\mathbf{b}'_v - \mathbf{b}_v) - \mathcal{F}_g(\mathbf{b}'_u - \mathbf{b}_u)\big) \end{aligned}$$

and conclude that $[\mathbf{y}'] = [\mathbf{y}]$. Thus this process yields a well-defined cohomology class $[\mathbf{y}^{\mathcal{A}}] \in H^1(\mathcal{G};\mathcal{F})$. By a similar argument, when $\mathcal{A}$ and $\mathcal{A}'$ are isomorphic as $\mathcal{F}$-affine network torsors, $[\mathbf{y}^{\mathcal{A}}] = [\mathbf{y}^{\mathcal{A}'}]$.

Finally, by observing that $[\mathbf{y}] \mapsto [\mathcal{A}^{\mathbf{y}}]$ and $[\mathcal{A}] \mapsto [\mathbf{y}^{\mathcal{A}}]$ are inverse operations, we establish the desired natural bijection. □

Remark 7.3.13. The correspondence between isomorphism classes of $\mathcal{F}$-affine network torsors and cohomology classes $H^1(\mathcal{G};\mathcal{F})$ gives a way to discuss "the" affine network sheaf corresponding to an affine network torsor. For an $\mathcal{F}$-affine network torsor $\mathcal{A}$, pick a 1-cochain $\mathbf{b} \in H^1(\mathcal{G};\mathcal{F})$ such that $[\mathcal{A}] = [\mathcal{A}^{\mathbf{b}}]$. Since the stalks of $\mathcal{A}^{\mathbf{b}}$ are exactly the vector spaces stalks of the structure sheaf $\mathcal{F}$ with their origins, we may view $\mathcal{A}^{\mathbf{b}}$ as an affine network sheaf (Definition 7.3.1). Different choices of representative $\mathbf{b}$ yield different sheaves.

Let $\mathbb{A}_{\mathcal{F}}$ denote the collection of isomorphism classes of $\mathcal{F}$-affine network torsors. The bijective correspondence $\mathbb{A}_{\mathcal{F}} \cong H^1(\mathcal{G};\mathcal{F})$ endows $\mathbb{A}_{\mathcal{F}}$ with the structure of a vector space. This vector space structure may be described explicitly at the level of affine network torsors.

Definition 7.3.14. Let $\mathcal{A}$ and $\mathcal{A}'$ be $\mathcal{F}$-affine network torsors. The **sum** $\mathcal{A} \boxplus \mathcal{A}'$ is the $\mathcal{F}$-affine network torsor with the following structure.

- For each $\sigma \in \mathcal{V} \amalg \mathcal{E}$, the stalk over σ is given by

$$(\mathcal{A} \boxplus \mathcal{A}')(\sigma) = \big(\mathcal{A}(\sigma) \times \mathcal{A}'(\sigma)\big)/\sim,$$

  where $(x, x') \sim (x + w, x' + (-w))$ for all $w \in \mathcal{F}(\sigma)$. The action of $\mathcal{F}(\sigma)$ on $(\mathcal{A} \boxplus \mathcal{A}')(\sigma)$ is given by $w + [(x, x')] = [(w + x\,,\, x')]$.

- For each covering morphism $f : v \to e$, the restriction map $(\mathcal{A} \boxplus \mathcal{A}')_f : (\mathcal{A} \boxplus \mathcal{A}')(v) \to (\mathcal{A} \boxplus \mathcal{A}')(e)$ is given by

$$(\mathcal{A} \boxplus \mathcal{A}')_f[(x, x')] = [(\mathcal{A}_f(x), \mathcal{A}'_f(x'))]\,.$$

One may straightforwardly check that the sum $\mathcal{A} \boxplus \mathcal{A}'$ is a well-defined $\mathcal{F}$-affine network torsor. After "remembering the origin" and identifying $\mathcal{A}(\sigma) \cong \mathcal{A}'(\sigma) \cong \mathcal{F}(\sigma)$, the stalk $(\mathcal{A} \boxplus \mathcal{A}')(\sigma)$ is given by the vector space quotient $\mathcal{F}(\sigma) \oplus \mathcal{F}(\sigma)/\ker\left(\begin{bmatrix} I & I \end{bmatrix}\right)$.

Remark 7.3.15. This sum operation may be compared to the Baer sum of two group extensions. Given a pair of abelian groups A and B, a group extension

$$0 \to A \xrightarrow{i} E \xrightarrow{p} B \to 1$$

may naturally be viewed as a network A-torsor [46, Definition 4.2] over the base space B, viewed as an edgeless graph with vertex set B. The stalk over $b \in B$ is given by the fiber $p^{-1}(b) \subseteq E$, with A acting freely and transitively via addition through i. Given two group extensions E_1 and E_2, the Baer sum is given by the $\mathbb{Z}$-module quotient

$$E_1 \boxplus_{\text{Baer}} E_2 := \frac{\{(e_1, e_2) \in E_1 \oplus E_2 \,:\, p_1(e_1) = p_2(e_2)\}}{\{(i_1(a), -i_2(a)) \,:\, a \in A\}},$$

and serves as a concrete representation for the abelian group structure of $\mathrm{Ext}^1(B, A)$, which classifies abelian group extensions. The network A-torsor corresponding to the Baer sum $E_1 \boxplus_{\mathrm{Baer}} E_2$ is exactly the sum $E_1 \boxplus E_2$ as network A-torsors. In the nonabelian case, one could plausibly study a networked Schreier theory through inhomogeneous network torsors.

One may also describe the scaling of an $\mathcal{F}$-affine network torsor.

Definition 7.3.16. Let $\mathcal{A}$ be an $\mathcal{F}$-affine network torsor, and $\lambda \in \mathbb{k}$ a scalar. The **scaled** network torsor $\lambda\mathcal{A}$ has the following structure.

- For each $\sigma \in \mathcal{V} \amalg \mathcal{E}$, the stalk over σ is given by

$$(\lambda\mathcal{A})(\sigma) = \big(\mathcal{A}(\sigma) \times \mathcal{F}(\sigma)\big)/\sim,$$

  where $(x, w) \sim (x + w', w + (-\lambda w'))$ for all $w' \in \mathcal{F}(\sigma)$. The action of $\mathcal{F}(\sigma)$ on $(\lambda\mathcal{A})(\sigma)$ is given by $w' + [(x, w)] = [(x\,,\, w' + w)]$.

- For each covering morphism $f : v \to e$, the restriction map $(\lambda\mathcal{A})_f : (\lambda\mathcal{A})(v) \to (\lambda\mathcal{A})(e)$ is given by

$$(\lambda\mathcal{A})_f[(x, w)] = [(\mathcal{A}_f(x), \mathcal{F}_f(w)]\,.$$

Again, it is straightforward to check that this is an $\mathcal{F}$-affine network torsor.

Corollary 7.3.17. *Define operations on $\mathbb{A}_{\mathcal{F}}$ by $[\mathcal{A}] + [\mathcal{A}'] = [\mathcal{A} \boxplus \mathcal{A}']$ and $\lambda[\mathcal{A}] = [\lambda\mathcal{A}]$. These operations define a vector space structure on $\mathbf{A}_{\mathcal{F}}$, which is isomorphic to the vector space structure on $H^1(\mathcal{G}; \mathcal{F})$.*

Proof. Let $\Phi : \mathbb{A}_{\mathcal{F}} \to H^1(\mathcal{G}; \mathcal{F})$ denote the bijective map $\Phi([\mathcal{A}]) = [\mathbf{y}^{\mathcal{A}}]$. It suffices to prove that Φ is linear with respect to these structures. We first prove additivity of $\Phi([\mathcal{A}] + [\mathcal{A}']) = \Phi([\mathcal{A} \boxplus \mathcal{A}'])$. Fix an origin $\mathbf{b}_\sigma = [(a_\sigma, a'_\sigma)] \in \mathcal{A} \boxplus \mathcal{A}'$. Working on an edge e with incoming covering morphisms $f : v \to e$ and $g : u \to e$, we compute:

$$\begin{aligned}(\mathcal{A} \boxplus \mathcal{A}')_g(\mathbf{b}_u) - (\mathcal{A} \boxplus \mathcal{A}')_f(\mathbf{b}_v) &= [(\mathcal{A}_g a_u, \mathcal{A}'_g a'_u)] - [(\mathcal{A}_f a_v, \mathcal{A}'_f a'_v)] \\ &= [(\mathcal{A}_g a_u, \mathcal{A}'_g a_u)] - \big[\big(\mathcal{A}_f a_v + (\mathcal{A}'_f a_v - A'_g a_u)\,,\, \mathcal{A}'_g a_u\big)\big] \\ &= (\mathcal{A}_g a_u - \mathcal{A}_f a_v) + (\mathcal{A}'_g a_u - \mathcal{A}'_f a_v)\,.\end{aligned}$$

It follows that $\Phi([\mathcal{A} \boxplus \mathcal{A}']) = \Phi([\mathcal{A}]) + \Phi([\mathcal{A}'])$. To prove Φ respect scaling, we proceed by a similar argument. Fix a scalar $\lambda \in \mathbb{k}$ and an origin $\mathbf{b}_\sigma = [(a_\sigma, w_\sigma)]$ for each stalk of $(\lambda\mathcal{A})(\sigma)$. Again working on an edge e with incoming covering morphisms $f : v \to e$ and $g : u \to e$, we compute:

$$(\lambda\mathcal{A})_g(\mathbf{b}_u) - (\lambda\mathcal{A})_f(\mathbf{b}_v) = [(\mathcal{A}_g a_u, \mathcal{F}_g w_u)] - [(\mathcal{A}_f a_v, \mathcal{F}_f w_v)]$$

$$= [(\mathcal{A}_g a_u, \mathcal{F}_g w_u)] - \left[\left(\mathcal{A}_g a_u\,,\, \mathcal{F}_f w_v + \lambda(\mathcal{A}_f a_v - \mathcal{A}_g a_u)\right)\right]$$
$$= (\mathcal{F}_g w_u - \mathcal{F}_f w_v) + \lambda(\mathcal{A}_g a_u - \mathcal{A}_f a_v).$$

Writing $\mathbf{c}$ for the 1-cochain with edge components $\mathbf{c}_e = (\mathcal{F}_g w_u - \mathcal{F}_f w_v)$, we see that $\mathbf{c}$ is a coboundary, from which we conclude that $\Phi(\lambda[\mathcal{A}]) = \Phi([\lambda\mathcal{A})] = \lambda\Phi([\mathcal{A}])$. Thus Φ is linear isomorphism, and these operations give $\mathbb{A}_{\mathcal{F}}$ a vector space structure. □

Using the identification $\mathbb{A}_{\mathcal{F}} \cong H^1(\mathcal{G};\mathcal{F})$, we may better understand the structure of global sections of affine network sheaves.

Definition 7.3.18. Let $\mathcal{A}$ be a finite dimensional $\mathcal{F}$-affine network torsor. A point $\mathbf{x} \in \prod_{v\in\mathcal{V}} \mathcal{A}(v)$ is a **global section** of $\mathcal{A}$ if for all edges $e \in \mathcal{E}$ with incoming covering morphisms $f : u \to e$ and $g : v \to e$, we have $\mathcal{A}_f(\mathbf{x}_u) = \mathcal{A}_g(\mathbf{x}_v)$.

Notation 7.3.19. Given an $\mathcal{F}$-affine network torsor $\mathcal{A}$, we may extend the zero-cochain concept and write $C^0(\mathcal{G};\mathcal{A}) := \prod_{v\in\mathcal{V}} \mathcal{A}(v)$. This is itself an affine space under the component-by-component action of the vector space $C^0(\mathcal{G};\mathcal{F})$. One recovers an **affine coboundary map** $\delta^{\text{aff}} : C^0(\mathcal{G};\mathcal{A}) \to C^1(\mathcal{G};\mathcal{F})$ edgewise via

$$(\delta^{\text{aff}}\mathbf{x})_e = \mathcal{A}_g\mathbf{x}_v - \mathcal{A}_f\mathbf{x}_u\,,$$

where e is viewed as a directed edge from u to v, and $f : u \to e$, $g : v \to e$ are the corresponding covering morphisms. Note that the codomain is a 1-cochain of the *structure sheaf* $\mathcal{F}$, as we are taking differences in each affine edge stalk $\mathcal{A}(e)$.

Lemma 7.3.20. *The affine coboundary map δ^{aff} of an $\mathcal{F}$-affine network torsor $\mathcal{A}$ has linear part given by the linear coboundary of the structure of the structure sheaf $\mathcal{F}$. Moreover, a global section of $\mathcal{A}$ is exactly a point $\mathbf{x} \in C^0(\mathcal{G};\mathcal{A})$ such that $\delta^{aff}(\mathbf{x}) = 0$.*

Proposition 7.3.21. *Let $\mathcal{A}$ be a $\mathcal{F}$-affine network torsor. The following are equivalent.*

(i) $\mathcal{A}$ admits a global section.

(ii) $\Phi([\mathcal{A}]) = [0]$ is the is the trivial cohomology class in $H^1(\mathcal{G};\mathcal{F})$.

(iii) For every affine network sheaf $\mathcal{A}^{\mathbf{b}}$ corresponding to the network torsor $\mathcal{A}$, the inhomogeneous linear system $\delta\mathbf{x} = -\mathbf{b}$ is consistent.

Proof. ((i) $\iff$ (ii)): $\mathbf{x}$ is a global section of $\mathcal{A}$ if and only if for every directed edge $e = (u, v)$ with covering morphisms $f : u \to e$ and $g : v \to e$, we have $\mathcal{A}_g\mathbf{x}_v - \mathcal{A}_f\mathbf{x}_u = 0$. Hence $\mathcal{A}$ admits a global section if and only if $\Phi([\mathcal{A}]) = [0]$.

((ii) $\iff$ (iii)): $\Phi([\mathcal{A}]) = [0]$ if and only if $[\mathcal{A}] = [\mathcal{A}^{\mathbf{b}}]$ whenever $\mathbf{b} \in [0] = \mathcal{R}(\delta)$. Meanwhile, $\mathbf{b} \in \mathcal{R}(\delta)$ if and only if $\delta\mathbf{x} = -\mathbf{b}$ is consistent. □

This result provides a direct interpretation of the cohomology of the structure sheaf $\mathcal{F}$. $H^1(\mathcal{G};\mathcal{F})$ exactly encodes obstructions to the consistency of the inhomogeneous linear system $\delta\mathbf{x} = -\mathbf{b}$, or equivalently, the existence of a global section for an $\mathcal{F}$-affine network torsor. These obstructions may be visualized as translations of the affine hyperplane $\mathcal{R}(\delta^{\text{aff}})$ away from the origin. Every nontrivial cohomology class represents a different translation of the affine hyperplane. The 0-cohomology $H^0(\mathcal{G};\mathcal{F})$ similarly represents directions of non-variation for the affine coboundary δ^{aff}. That is, $H^0(\mathcal{G};\mathcal{F})$ encodes what the space of global sections of $\mathcal{A}$ looks like, given that a global section exists. Equivalently, $H^0(\mathcal{G};\mathcal{F})$ describes the directions of variation for the space of OLS solutions to $\delta^{\text{aff}} = 0$.

7.4 CONTINUOUS PIECEWISE AFFINE HILBERT SHEAVES

We now turn our attention to a class of non-smooth network Hilbert sheaves whose restriction maps are continuous piecewise affine functions. Such sheaves will have a coboundary map which is itself continuous piecewise affine, and a corresponding locally quadratic potential function. Clarke gradient descent with respect this potential function yields (non-unique) globally defined C^0-nonlinear heat flows with well behaved long-term behavior.

7.4.1 *Continuous piecewise affine maps*

Let X be a finite dimensional vector space. A **polyhedron** in X is a subset $P \subseteq X$ which is the intersection of finitely many closed affine halfspaces. That is there are a collection of affine functionals $\ell_j : \mathbb{R}^n \to \mathbb{R}$ and constants $b_j \in \mathbb{R}$ such that

$$P = \left\{x \in X \,:\, \ell_j(x) \leqslant b_j \text{ for all } j = 1, \dots, n\right\}.$$

A polyhedron P is necessarily closed and convex, but may be unbounded. P also has a collection of **faces**, which are sets of the form $P \cap H$, where H is an affine hyperplane of codimension 1, where P is entirely contained in one closed affine halfspace determined by H. Both P itself and $\varnothing$ are faces of P.

Definition 7.4.1. Let X be a finite dimensional vector space. A **polyhedral complex** in X is a finite collection of polyhedra $\mathcal{P}$ that satisfy the following axioms.

(i) **Face closure.** If $P \in \mathcal{P}$ and Q is a face of P, then $Q \in \mathcal{P}$.

(ii) **Intersection condition.** If P and Q are polyhedra in $\mathcal{P}$, the intersection $P \cap Q$ is a (possibly empty) shared face of P and Q.

We call each polyhedron $P \in \mathcal{P}$ a **cell** of $\mathcal{P}$.

Given a finite dimensional vector space X, a **polyhedral decomposition** of X is a polyhedral complex $\mathcal{P}$ such that $\bigcup_{P \in \mathcal{P}} P = X$. To each polyhedral decomposition, there are a collection of **top dimensional cells** $\{T_1, \dots, T_k\}$

Using such polyhedral decompositions, we may define the class of piecewise affine maps.

Definition 7.4.2. Let X and Y be finite dimensional vector spaces. A function $f : X \to Y$ is **continuous piecewise affine (CPWA)** if there is a polyhedral decomposition $\mathcal{P} = \{P_j\}_{j=1}^n$ of X, linear maps $\{A_j\}_{j=1}^n$, and constants $\{b_j\}_{j=1}^n$ such that

$$f\big|_{P_j}(x) = A_j x + b_j$$

for all $x \in P_j$.

Notation 7.4.3. Given a continuous piecewise affine map f, we denote the underlying polyhedral decomposition by $\mathcal{P}_f$, and the affine map on the cell $P \in \mathcal{P}_f$ by $f_P(x) = A_P x + b_P$.

By definition, a continuous piecewise affine map is continuous; when two polyhedra P_i and P_j intersect in a shared face P_k, all three maps $f\big|_{P_i}$, $f\big|_{P_j}$, and $f\big|_{P_k}$ agree on P_k. CPWA maps have the following closure properties.

Lemma 7.4.4. *Let $f, f' : X \to Y$, $g : X \to Z$, and $h : W \to Z$, and $i : Y \to Z$ be CPWA maps, and let $\lambda \in \mathbb{R}$. The maps λf, $i \circ f$, $f + f'$, $\begin{bmatrix} f \\ g \end{bmatrix}$, and $f \oplus h$ are all CPWA maps.*

Proof. Let $\mathcal{P}$ and $\mathcal{Q}$ denote the polyhedral decomposition of X and Y underlying the maps $f : X \to Y$ and $i : Y \to Z$. The composition $i \circ f$ is easily seen to be CPWA over the polyhedral decomposition $f^{-1}(\mathcal{Q})$ whose underlying cells are intersections of the form $P \cap f^{-1}(Q)$ where $P \in \mathcal{P}$ and $Q \in \mathcal{Q}$. The scaling λf is CPWA since a scaling of an affine map is affine. Let $\mathcal{P}$ and $\mathcal{Q}$ be the polyhedral decompositions underlying f and f' respectively. Let $\mathcal{P} \vee \mathcal{Q}$ denote their coarsest common refinement via

$$\mathcal{P} \vee \mathcal{Q} = \{P \cap Q \;:\; P \in \mathcal{P} \text{ and } Q \in \mathcal{Q}\}.$$

The map $f + f'$ is CPWA over the common refinement $\mathcal{P} \vee \mathcal{Q}$. The other maps $\begin{bmatrix} f \\ g \end{bmatrix}$, and $f \oplus h$ can be viewed as special cases of the sum. □

Corollary 7.4.5. *Let $X_1, \dots, X_m, Y_1, \dots Y_n$ be a collection of finite dimensional vector spaces, and $f_{ij} : X_j \to Y_i$ be a CPWA map for each pair of indices i, j. Let $\mathcal{X} = \bigoplus_j X_j$ and $\mathcal{Y} = \bigoplus_i Y_i$, and let $\mathbf{M} : \mathcal{X} \to \mathcal{Y}$ be the finite block map $[f_{ij}]$. The map $\mathbf{M}$ is CPWA.*

Proposition 7.4.6. *Every CPWA map* $f : X \to Y$ *is globally Lipschitz.*

Proof. Let $x, y \in X$, and let $\gamma : [0,1] \to X$ be the straight line path from x to y. Let $0 = t_0 < t_1 < \cdots < t_k = 1$ denote the times at which $\gamma(t)$ switches between maximal cells of the underlying polyhedral decomposition $\mathcal{P}$. Letting $M = \max\{\|A_P\|_{op} : P \in \mathcal{P}\}$, we compute:

$$\begin{aligned}
\|f(x) - f(y)\|_Y &= \left\| \sum_{j=1}^{k} \int_{t_{j-1}}^{t_j} \frac{d}{dt}(f \circ \gamma)(t)\, dt \right\| \\
&\leqslant \sum_{j=1}^{k} \int_{t_{j-1}}^{t_j} \left\| \frac{d}{dt}(f \circ \gamma)(t) \right\| dt \\
&\leqslant \sum_{j=1}^{k} \int_{t_{j-1}}^{t_j} M\|x - y\|\, dt \\
&= M\|x - y\| .
\end{aligned}$$

□

Remark 7.4.7. CPWA maps arise in the context of feed-forward neural networks equipped with the **rectifiable linear unit (ReLU)** map as an activation function. Each layer of the neural network is the composition of an affine map $f : \mathbb{R}^n \to \mathbb{R}^m$ and the conic projection onto the positive orthant $\mathrm{ReLU}(x_1, \ldots, x_m)^T = \big(\max(0, x_1), \ldots, \max(0, x_m)\big)^T$. Since the conic projection ReLU is CPWA, the composition of all layers is itself CPWA by Lemma 7.4.4.

7.4.1.1 *Clarke gradients of CPWA maps*

Let $f : X \to Y$ be CPWA with underlying polyhedral decomposition $\mathcal{P}$. Let $\mathcal{T} \subseteq \mathcal{P}$ denote the collection of top dimensional cells. The maximal cells cover the entirety of X, and intersect in lower-dimensional cells. The **interior** of $\mathcal{P}$ is given by the set

$$\mathring{\mathcal{P}} := \bigcup_{T \in \mathcal{T}} \mathring{T},$$

where $\mathring{T}$ is the non-empty topological interior of the maximal cell T.

The CPWA map $f : X \to Y$ is differentiable on $\mathring{\mathcal{P}}$, with derivative $D_x f = A_T x$ for each top cell $T \in \mathcal{T}$ and $x \in \mathring{T}$. It follows that the potential function $V(x) = \frac{1}{2}\|f(x)\|^2$ has a well-defined gradient on $\mathring{\mathcal{P}}$ given by

$$\nabla V(x) = A_T^*(A_T x + b_T)$$

for all $x \in \mathring{T}$.

In general, f (and hence V) will not be differentiable on the boundaries of top cells; there can be a cusp at the boundaries where top cells meet. However, since f is globally Lipschitz, the potential function V is locally Lipschitz and admits a globally non-empty Clarke gradient

$$\partial V(x) = \operatorname{cvx}\left\{A_T^*(A_T x + b_T) \,:\, T \ni x \text{ is a top cell containing } x\right\}.$$

When x is in the interior of a top cell of $\mathcal{P}$, the Clarke gradient is exactly the gradient $\nabla V(x)$ (modulo a pair of set braces). On the other hand, when x is in the intersection of multiple top cells, the Clarke gradient is the convex hull of the gradients of the affine functions of the top cells containing x.

Example 7.4.8. Consider the ReLU map $\mathrm{ReLU} : \mathbb{R} \to \mathbb{R}$ defined by $\mathrm{ReLU}(x) = \max(x, 0)$. This CPWA map has an associated potential function $V(x) = \frac{1}{2}\mathrm{ReLU}(x)^2$. This potential function is continuously differentiable on $\mathbb{R}$, with gradient $\nabla V(x) = \mathrm{ReLU}(x)$.

The shifted ReLU map $f(x) = \mathrm{ReLU}(x) + 1$ is also a CPWA map with potential function $V_f(x) = \frac{1}{2}(1 + \mathrm{ReLU}(x))^2$. This potential function is not differentiable at 0, but has a well-defined Clarke gradient on $\mathbb{R}$ given by

$$\partial V_f(x) = \begin{cases} 0 & \text{if } x < 0 \\ [0, 1] & \text{if } x = 0 \\ x + 1 & \text{if } x > 0. \end{cases}$$

Example 7.4.9. Let $f : \mathbb{R} \to \mathbb{R}$ be the CPWA map $f(x) = 1 + |x|$. The potential function $V_f(x) = \frac{1}{2}\|f(x)\|^2$ has Clarke gradient

$$\partial V_f(x) = \begin{cases} x - 1 & \text{if } x < 0 \\ [-1, 1] & \text{if } x = 0 \\ x + 1 & \text{if } x > 0. \end{cases}$$

Remark 7.4.10. While the Clarke gradient ∂V_f of the potential function associated to a CPWA function $f : X \to Y$ is generically multi-valued, the set of points $x \in X$ such that $\partial V(x)$ is multivalued has Lebesgue measure zero.

7.4.1.2 *CPWA kernels*

We now investigate the zero-sets of continuous piecewise affine maps. While lacking the requisite categorical properties, we adopt the nomenclature of $\ker(f)$ for the zero set $\{x \in X \,:\, f(x) = 0\}$ for a CPWA map $f : X \to Y$.

Recall that the kernel of an affine map $f(x) = Ax + b$ can be written as $k + \ker(A)$, where k is any choice of a point $k \in \ker(f)$. Note that $\ker(f)$ is empty if and only if $b \notin \mathcal{R}(A)$. We may fully characterize the kernel of a CPWA map by the kernel of its component maps; let $f : X \to Y$ be CPWA with underlying polyhedral decomposition $\mathcal{P}$ and affine maps $f_P = A_P x + b_P$ for each $P \in \mathcal{P}$. Let $\ker_P(f) = \ker(f_P) \cap P$, treating f_P as a globally-defined affine map $f_P : X \to Y$. The kernel of f is exactly

$$\ker(f) = \bigcup_{P \in \mathcal{P}} \ker_P(f).$$

For a linear map $A : X \to Y$ between finite dimensional Hilbert spaces, the kernels $\ker(A)$ and $\ker(A^*A)$ agree. We now investigate the corresponding relationship for CPWA maps. We begin with the following lemma about affine maps.

Lemma 7.4.11. *Let* $f(x) = Ax + b$ *be an affine map from* X *to* Y. *We have an agreement of kernels* $\ker(A^*f) = \ker(f)$ *if and only if* $\ker(f) \neq \emptyset$.

Proof. We always have the inclusion $\ker(f) \subseteq \ker(A^*f)$. Since $\ker(A^*A) = \ker(A)$, if $\ker(f)$ is non-empty, there is a $k \in \ker(f)$. We may compute:

$$\begin{aligned}\ker(f) &= k + \ker(A) \\ &= k + \ker(A^*A) \\ &= \ker(A^*f).\end{aligned}$$

Therefore if $\ker(f)$ is non-empty, then $\ker(A^*f) = \ker(A)$. Conversely, a least-squares solution x_{LS} to the equation $Ax = b$ must satisfy the normal equation $A^*Ax = A^*b$, which forces $-x_{LS} \in \ker(A^*f)$. Therefore $\ker(f) \neq \ker(A^*f)$ when $\ker(f) = \varnothing$. □

7.4.2 *CPWA dynamics*

Let $f : X \to Y$ be a CPWA function, and V the potential function $V(x) = \frac{1}{2}\|f(x)\|^2$. We may define the **CPWA Laplacian** $\mathcal{L}$ as the multivalued Clarke gradient $\mathcal{L} := \partial V$. This CPWA Laplacian may be written explicitly as

$$\mathcal{L}(x) = \operatorname{cvx}\{A_T^* f(x) \,:\, T \ni x \text{ is a top dimensional cell}\}.$$

The Laplacian $\mathcal{L}$ is single-valued on the full-measure subset $\mathring{\mathcal{P}} \subseteq X$. Using the CPWA Laplacian, we may define dynamics on X via $\dot{x} \in -\mathcal{L}x$. Since $\mathcal{L}$ is generically discontinuous and multi-valued, we consider Filippov solutions to this differential inclusion.

Definition 7.4.12. Let $f : X \to Y$ be a CPWA function with CPWA Laplacian $\mathcal{L} : X \to X$. A **CPWA heat flow** of f on $[0, T]$ with initial point $x_0 \in X$ is an absolutely continuous curve $x(t) \in X$ for $t \in [0, T]$ such that $x(0) = x_0$ and $\dot{x}(t) \in -\mathcal{L}(x)$ for all t.

Remark 7.4.13. In the language of control theory, the dynamics $\dot{x} \in -\mathcal{L}(x)$ of CPWA heat flow define a **state-dependent switch affine system**. These systems have recently been extensively studied [50, 68, 69, 71, 111], with applications to AC/DC power conversion [3, 105] and neural networks [41, 116]. Most of the literature approaches state-dependent switch affine systems from the perspective of controllability, and the problem of defining a state-dependent switching rule that has well-behaved dynamics. More importantly, CPWA dynamics differ from general switch affine systems in that the potential function V defines a global Lyapunov function, granting additional control over trajectories.

Proposition 7.4.14. *Let $f : X \to Y$ be a CPWA map with CPWA Laplacian $\mathcal{L}$. The Cauchy problem*

$$\begin{aligned} \dot{x} &\in -\mathcal{L}x \\ x(0) &= x_0 \end{aligned}$$

has a globally defined solution for each initial value x_0.

Proof. Since f is Lipschitz continuous, the potential function $V(x) = \frac{1}{2}\|f(x)\|^2$ is locally Lipschitz. It follows that the negative Clarke gradient $-\partial V(x)$ is an upper semicontinuous map [24]. Local existence of a Filippov solution up to the boundary of the closed ball $B(0, r) \subseteq X$ follows [39, Theorem 7.2]. Since f is locally affine, we may bound the norm $\|\dot{x}\| \leqslant M\|x\| + a$ for a suitable choice of constants M and a. Thus the trajectory cannot escape to infinity in finite time, and all Filippov solutions can be extended to the time-interval $[0, \infty)$. $\square$

Uniqueness, on the other hand, cannot be guaranteed in general. While Filippov trajectories inside top-cells are unique, when a trajectory gets stuck in the boundary of a collection of cells, trajectories can be extended in multiple distinct ways according to the choice of a different value in $-\mathcal{L}$.

Example 7.4.15. Consider the CPWA $f : \mathbb{R} \to \mathbb{R}$ given by $f(x) = 1 - |x|$. For every $T \geqslant 0$, the path

$$x_t = \begin{cases} 0 & \text{if } 0 \leqslant t \leqslant T \\ 1 - e^{T-t} & \text{if } t > T \end{cases}$$

is a Filippov solution to the CPWA heat flow of f. This is essentially CPWA modification of Norton's dome [94].

The asymptotic behavior of a bounded CPWA heat flow is governed by the following lemma.

Lemma 7.4.16. *Let x_t be a Filippov solution to the heat flow of a CPWA map $f : X \to Y$. If x_t is bounded, then $x_\infty := \lim_{t\to\infty} x_t$ exists and is a generalized critical point of $V(x) = \frac{1}{2}\|f(x)\|^2$.*

Proof. This follows directly from the work of Drusvyatskiy, Ioffe, and Lewis on generalized gradient descent. Identify $X \cong \mathbb{R}^n$. Since $V(x)$ is piecewise a polynomial in the underlying variables $x_1, \ldots, x_n$, and each polyhedral region on which V is piecewise-defined can be expressed as the solution set to finitely many linear inequalities, the function $V : \mathbb{R}^n \to \mathbb{R}$ is semialgebraic. Fix a trajectory x_t which is constrained to a compact set $C \subseteq \mathbb{R}^n$. $V(x)$ is Lipschitz on C, so x_t is a so-called *curve of near-maximal slope* [37, Proposition 6.4] and moreover converges to a generalized critical point of V [37, Corollary 6.7]. □

By restricting our attention to class of Filippov solutions to the CPWA heat flow that avoid spending superfluous time on the boundaries of cells, we may prove the boundedness of trajectories.

Definition 7.4.17. Let $f : X \to Y$ be a CPWA function, and let x_t be a CPWA heat flow. x_t is **fast** if $\dot{x}_t \in -\mathcal{L}(x_t)$ has maximal norm for all t where $\dot{x}_t$ is defined.

Example 7.4.18. Consider the CPWA Norton's dome (Example 7.4.15) with CPWA function $f(x) = 1 - |x|$. A solution x_t to the CPWA heat flow of f with initial value $x_0 = 0$ is fast if and only if $x_t \neq 0$ for all $t > 0$.

Proposition 7.4.19. *Let $f : X \to Y$ be a CPWA function, and that x_t is a fast solution to the CPWA heat flow. If x_t is in the boundary of a top dimensional cell T for all $t \in [\tau, \tau + \epsilon)$, then every CPWA heat flow y_t such that $y_t = x_t$ for all $t \leqslant \tau$ slides for all $t \in [\tau, \tau + \epsilon)$.*

Proof. Suppose there is such a Filippov solution y_t. Let $t_0 := \inf\{t : y_t \neq x_t\}$. At time t_0, there must be a vector in $-\mathcal{L}(x_{t_0})$ which points outside of the interior of the polyhedral face M on which x_t is sliding. Since every point in $-\mathcal{L}(x_{t_0})$ projects onto the same point M, the velocity selection at time t_0 must be non-maximal, contradicting the fact that x_t is a fast solution. □

Remark 7.4.20. This argument essentially demonstrates that fast solutions avoid superfluous sliding along the boundary of a cell, commonly called a **sliding mode** for a Filippov solution to a discontinuous differential equation. Whenever a fast solution x_t can exit a sliding mode, it does.

Theorem 7.4.21. *Let $f : X \to Y$ be a CPWA function with potential function $V(x) = \frac{1}{2}\|f(x)\|^2$. Every fast solution to the CPWA heat flow is bounded.*

Proof. Let x_t be a fast solution to the CPWA heat flow. We fix the following notation.

- $\mathcal{T}$ is the finite collection of top dimensional cells underlying f.
- $g_T(x) = A_T x + b_T$ is the restriction of f to $T \in \mathcal{T}$.
- $r_t := \|x_t\|$.
- $s_t = \max\{r_t^{-1}\|A_t x_t\| \,:\, T \in \mathcal{T} \text{ and } x \in T\}$.
- $B := \max_{T\in\mathcal{T}} \|b_T\|$.
- $C_0 := \max_{T\in\mathcal{T}} \|A_T^* b_T\|$.
- σ_* is the smallest positive singular value of a linear map A_T with $T \in \mathcal{T}$.

First, by the reverse triangle inequality, we may bound

$$\|\dot{x}_t\| \geqslant \sigma_* r_t s_t - C_0 \tag{9}$$

for all $t \geqslant 0$. When $x_t \neq 0$, by the Cauchy-Schwarz inequality we may bound

$$\dot{r}_t = -\frac{\langle x_t, \dot{x}_t \rangle}{r_t} \leqslant B s_t. \tag{10}$$

Fix an $\epsilon \in (0, 1]$, and set $R := 4C_0(\sigma_*\epsilon)^{-1}$. For each $j \in \mathbb{N}$, set $\epsilon_j := 2^{-j}\epsilon$ and $R_j := 2^j R$. For each pair $j, k \in \mathbb{N}$, let A_{jk} denote the set

$$A_{jk} := \{t \in \mathbb{R}_{\geqslant 0} \,:\, R_j \leqslant r_t < R_{j+1} \text{ and } \epsilon_{k+1} \leqslant s_t < \epsilon_k\}.$$

The union $\bigcup_{j,k} A_{jk} = \{t \,:\, r(t) \geqslant R \text{ and } s_t > 0\}$. We now work to bound the integral

$$\int_{\{t : r_t \geqslant R\}} s_t \, dt = \sum_{j,k} \int_{A_{jk}} s_t \, dt\,.$$

When $t \in A_{jk}$, we may bound $\sigma_* r_t s_t \geqslant \sigma_* R_j \epsilon_{k+1} = C_0 2^{j+1-k}$. For fixed k, when $j \geqslant k+1$, Equation (9) ensures that $\|\dot{x}_t\| \geqslant 3C_0$. We may derive the inequalities

$$\sum_{j\geqslant k+1} \mu(A_{jk}) 3C_0 \leqslant \int_{\bigcup_{j\geqslant k+1} A_{jk}} -\|\dot{x}_t\|^2 \, dt \leqslant \int_0^\infty -\|\dot{x}_t\|^2 \, dt \leqslant V(x_0)\,,$$

where μ denotes Lebesgue measure. Bound the sum $\sum_{j \geqslant k+1} \mu(A_{jk}) \leqslant V(x_0)(3C_0)^{-2}$ and the integral $\sum_{j \geqslant k+1} \int_{A_{jk}} s_t \, dt \leqslant \epsilon_k V(x_0)(3C_0)^{-2}$. A similar argument yields for fixed j an upper bound $\sum_{k \geqslant j} \int_{A_{jk}} s_t \, dt \leqslant \epsilon_j V(x_0) C_0^{-2}$. Using these upper bounds, we compute:

$$\begin{aligned}\int_{\{t : r_t \geqslant R\}} s_t \, dt &= \sum_{j,k} \int_{A_{jk}} s_t \, dt \\ &\leqslant \sum_{k \geqslant 0} \epsilon_k V(x_0)(3C_0)^{-2} + \sum_{j \geqslant 0} \epsilon_j V(x_0) C_0^{-2} \\ &= \frac{20}{9} \epsilon V(x_0) C_0^{-2}.\end{aligned}$$

We may now prove that x_t is bounded. By Equation (10), we have a bound

$$\sup_{t \geqslant 0} r_t \leqslant R + B \int_{\{t : r_t \geqslant R\}} s_t \, dt \leqslant R + \frac{20}{9} \epsilon V(x_0) C_0^{-2},$$

which is finite. Therefore x_t is bounded. □

Remark 7.4.22. In general, fast CPWA heat flows need not converge to local minima of the potential function V. A fast CPWA heat flow may converge to a generalized saddle point of V. However, after finding a saddle point, one may start a new fast solution from the saddle point that immediately enters the interior of a top dimensional cell by the following lemma.

Lemma 7.4.23. *Let* $f : X \to Y$ *be a CPWA map with potential function* $V(x) = \frac{1}{2}\|f(x)\|^2$. *Let* x_0 *be a generalized critical point of* V *such that every CPWA heat flow* x_t *initialized at* x_0 *is constant.* x_0 *is a local minimum of* V.

Proof. Without loss of generality suppose that $x_0 = 0$. If 0 is contained in the interior of a top-dimensional cell T, then 0 is an OLS solution to $A_T x = -b_T$, and hence a local minimum of V. Instead, suppose M is the polyhedral cell of minimal dimension whose relative interior contains 0. Let T be a top-dimensional cell such that $M \subseteq \partial T$. The constrained gradient $\nabla V\big|_M(0)$ is exactly the orthogonal projection of $A_T^* f(0)$ onto M. This quantity is independent of the choice of top-dimensional cell T. Since there can be no non-constant sliding mode for Clarke gradient descent starting from 0 in M, it must be that $\nabla V\big|_M(0) = 0$. Therefore $A_T^* f(0)$ is orthogonal to M, and points outside of T. Therefore for any vector $\nu \in T$, the directional derivative $\langle A_T^* f(0), \nu \rangle \geqslant 0$. This holds for every top cell, so 0 must be a local minimum of V. □

7.4.3 *Piecewise affine sheaves*

A C^0-nonlinear network sheaf whose restriction maps are continuous piecewise affine maps has a coboundary operator which is itself CPWA.

Definition 7.4.24. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a finite network. A **piecewise affine Hilbert sheaf** $\mathcal{F}$ on $\mathcal{G}$ consists of the following data.

- A finite dimensional Hilbert space $\mathcal{F}(\sigma)$ for each $\sigma \in \mathcal{V} \cup \mathcal{E}$ called the **stalk** over σ;
- For each covering map $f : v \to e$ in $\mathcal{G}$, a choice of a continuous piecewise affine map $\mathcal{F}_f : \mathcal{F}(v) \to \mathcal{F}(e)$.

Remark 7.4.25. We may assume, without loss of generality, that all restriction maps out of a stalk $\mathcal{F}(v)$ are defined with respect to a common polyhedral decomposition $\mathcal{P}_\sigma$.

By Lemma 7.4.4 and Corollary 7.4.5, the coboundary map $\delta : C^0(\mathcal{G}; \mathcal{F}) \to C^1(\mathcal{G}; \mathcal{F})$ is itself CPWA. The C^0-nonlinear heat flow with respect to the sheaf Laplacian exactly encodes CPWA heat flow with respect to the coboundary map δ. Under a fast heat flow selection rule, Theorem 7.4.21 ensures that heat flows converge to generalized critical points.

Remark 7.4.26. Following Remark 7.4.7, piecewise affine Hilbert sheaves could potentially serve as the foundation for a different architecture for sheaf neural networks than that of Hansen and Gebhart [53]. By allowing the rectifiable linear unit (and projections onto polyhedral cones more generally), the activation function may be directly incorporated into the restriction maps themselves. Alternatively, one may consider using a smooth activation function like soft-max and a smooth nonlinear Hilbert sheaf.

BIBLIOGRAPHY

[1] S. Abramsky and B. Coecke. "A categorical semantics of quantum protocols." In: *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004.* 2004, pp. 415–425. DOI: 10.1109/LICS.2004.1319636.

[2] Pablo Acuña. "von Neumann's Theorem Revisited." In: *Foundations of Physics* 51.3 (2021), p. 73. DOI: 10.1007/s10701-021-00474-5. URL: https://doi.org/10.1007/s10701-021-00474-5.

[3] Carolina Albea, Germain Garcia, and Luca Zaccarian. "Hybrid dynamic modeling and control of switched affine systems: Application to DC-DC converters." In: *2015 54th IEEE Conference on Decision and Control (CDC)*. 2015, pp. 2264–2269. DOI: 10.1109/CDC.2015.7402544.

[4] Pavel Alexandrov. "Diskrete Raume." In: *Matematicheskii Sbornik* 2 (1937), pp. 501–519.

[5] Shun-ichi Amari and Hiroshi Nagaoka. *Methods of Information Geometry*. Vol. 191. Translations of Mathematical Monographs. American Mathematical Soc., 2000. ISBN: 978-0821843024.

[6] Yu.M. Arlinskiĭ. "Shorting, parallel addition and form sums of nonnegative selfadjoint linear relations." In: *Linear Algebra and its Applications* 599 (2020), pp. 156–200. ISSN: 0024-3795. DOI: https://doi.org/10.1016/j.laa.2020.04.007. URL: https://www.sciencedirect.com/science/article/pii/S0024379520301890.

[7] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther. "Finite element exterior calculus, homological techniques, and applications." In: *Acta Numerica* 15 (2006), 1–155. DOI: 10.1017/S0962492906210018.

[8] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther. "Finite element exterior calculus: from Hodge theory to numerical stability." In: *Bull. Amer. Math. Soc.* 47 (2010), pp. 281–354.

[9] Miroslav Bacak. *Convex Analysis and Optimization in Hadamard Spaces*. Berlin, München, Boston: De Gruyter, 2014. ISBN: 9783110361629. DOI: doi:10.1515/9783110361629. URL: https://doi.org/10.1515/9783110361629.

[10] J. Baez and M. Stay. "Physics, Topology, Logic and Computation: A Rosetta Stone." In: *New Structures for Physics*. Springer Berlin, 2010, 95–172. ISBN: 9783642128219. DOI: 10.1007/978-3-642-12821-9_2. URL: http://dx.doi.org/10.1007/978-3-642-12821-9_2.

[11] John Baez. *Torsors Made Easy*. 2009. URL: https://math.ucr.edu/home/baez/torsors.html.

[12] A. Björner. “Posets, Regular CW Complexes and Bruhat Order.” In: *European Journal of Combinatorics* 5.1 (1984), pp. 7–16. ISSN: 0195-6698. DOI: https://doi.org/10.1016/S0195-6698(84)80012-8. URL: https://www.sciencedirect.com/science/article/pii/S0195669884800128.

[13] Cristian Bodnar, Francesco Di Giovanni, Benjamin Chamberlain, Pietro Lió, and Michael Bronstein. “Neural Sheaf Diffusion: A Topological Perspective on Heterophily and Oversmoothing in GNNs.” In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., 2022, pp. 18527–18541. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/75c45fca2aa416ada062b26cc4fb7641-Paper-Conference.pdf.

[14] Wietse M. Boon, Daniel F. Holmen, Jan M. Nordbotten, and Jon E. Vatne. “The Hodge-Laplacian on the Čech-de Rham complex governs coupled problems.” In: *Journal of Mathematical Analysis and Applications* 551.2, Part 2 (2025), p. 129692. ISSN: 0022-247X. DOI: https://doi.org/10.1016/j.jmaa.2025.129692.

[15] H. Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Universitext. Springer New York, 2010. ISBN: 9780387709130. URL: https://books.google.com/books?id=GAA2XqOIIGoC.

[16] Ronald E Bruck. “Asymptotic convergence of nonlinear contraction semigroups in Hilbert space.” In: *Journal of Functional Analysis* 18.1 (1975), pp. 15–26. ISSN: 0022-1236. DOI: https://doi.org/10.1016/0022-1236(75)90027-0. URL: https://www.sciencedirect.com/science/article/pii/0022123675900270.

[17] J. Brüning and M. Lesch. “Hilbert complexes.” In: *Journal of Functional Analsyis* 108 (1992), pp. 88–132.

[18] M. S. Burgin. “Categories with involution, and correspondences in γ-categories.” In: *Tr. Mosk. Mat. Obs.* (1970), pp. 161–228.

[19] A. Carboni. “Bicategories of partial maps.” eng. In: *Cahiers de Topologie et Géométrie Différentielle Catégoriques* 28.2 (1987), pp. 111–126. URL: http://eudml.org/doc/91396.

[20] Alex Chandler. *Thin Posets, CW Posets, and Categorification*. 2019. arXiv: 1911.05600 [math.CO]. URL: https://arxiv.org/abs/1911.05600.

[21] Salma Charfi and Hanen Ellouz. “Closedness and self-adjointness criteria for block operator matrices involving fully-subordinate entries.” In: *Filomat* 36 (Jan. 2022), pp. 2729–2737. DOI: 10.2298/FIL2208729C.

[22] Sheng Chen, Lindu Zhao, Weigong Zhang, and Peng Shi. "Consensus on compact Riemannian manifolds." In: *Information Sciences* 268 (2014). New Sensing and Processing Technologies for Hand-based Biometrics Authentication, pp. 220–230. ISSN: 0020-0255. DOI: https://doi.org/10.1016/j.ins.2013.12.006. URL: https://www.sciencedirect.com/science/article/pii/S0020025513008505.

[23] Ralph Chill. "On the Łojasiewicz–Simon gradient inequality." In: *Journal of Functional Analysis* 201.2 (2003), pp. 572–601. ISSN: 0022-1236. DOI: https://doi.org/10.1016/S0022-1236(02)00102-7. URL: https://www.sciencedirect.com/science/article/pii/S0022123602001027.

[24] Frank H. Clarke. "Generalized gradients and applications." In: *Transactions of the American Mathematical Society* 205 (1975), pp. 247–262. URL: https://api.semanticscholar.org/CorpusID:120174258.

[25] A. H. Clifford. "Semigroups Admitting Relative Inverses." In: *Annals of Mathematics* 42.4 (1941), pp. 1037–1049. ISSN: 0003486X, 19398980. URL: http://www.jstor.org/stable/1968781 (visited on 07/26/2025).

[26] Robin Cockett and Stephen Lack. "Restriction categories I: categories of partial maps." In: *Theoretical Computer Science* 270.1 (2002), pp. 223–259. ISSN: 0304-3975. DOI: https://doi.org/10.1016/S0304-3975(00)00382-0. URL: https://www.sciencedirect.com/science/article/pii/S0304397500003820.

[27] Robin Cockett and Stephen Lack. "Restriction categories II: partial map classification." In: *Theoretical Computer Science* 294.1 (2003). Category Theory and Computer Science, pp. 61–102. ISSN: 0304-3975. DOI: https://doi.org/10.1016/S0304-3975(01)00245-6. URL: https://www.sciencedirect.com/science/article/pii/S0304397501002456.

[28] Robin Cockett and Stephen Lack. "Restriction categories III: colimits, partial limits and extensivity." In: *Mathematical Structures in Computer Science* 17.4 (2007), 775–817. DOI: 10.1017/S0960129507006056.

[29] George E. Cooke and Ross L. Pinney. *Homology of Cell Complexes*. Princeton University Press, 1967. ISBN: 9780691623139. URL: http://www.jstor.org/stable/j.ctt183pvgt (visited on 07/07/2025).

[30] Zoe Cooperband. "Cellular Cosheaves, Graphic Statics, and Mechanics." PhD thesis. University of Pennsylvania, 2024.

[31] Zoe Cooperband and Robert Ghrist. "Towards Homological Methods in Graphic Statics." In: *Proceedings of the IASS Annual Symposium 2023*. 2023.

[32] Jorge Cortés. "Discontinuous Dynamical Systems: A tutorial on solutions, nonsmooth analysis, and stability." In: *IEEE Control Systems* 28.3 (June 2008), 36–73. ISSN: 1941-000X. DOI: 10.1109/mcs.2008.919306. URL: http://dx.doi.org/10.1109/MCS.2008.919306.

[33] Justin Curry. "Sheaves, Cosheaves and Applications." PhD thesis. University of Pennsylvania, 2014. arXiv: 1303.3255 [math.AT]. URL: https://arxiv.org/abs/1303.3255.

[34] Ju. L. Daleckii and M. G. Krein. *Stability of Solutions of Differential Equations in Banach Space*. American Mathematical Society, 1974.

[35] Gopal Danaraj and Victor Klee. "Shellings of spheres and polytopes." In: *Duke Mathematical Journal* 41.2 (1974), pp. 443 –451. DOI: 10.1215/S0012-7094-74-04150-7. URL: https://doi.org/10.1215/S0012-7094-74-04150-7.

[36] Daniele A. Di Pietro, Jérôme Droniou, and Silvano Pitassi. "Cohomology of the discrete de Rham complex on domains of general topology." In: *Calcolo* 60.2 (2023), p. 32. DOI: 10.1007/s10092-023-00523-7. URL: https://doi.org/10.1007/s10092-023-00523-7.

[37] D. Drusvyatskiy, A. D. Ioffe, and A. S. Lewis. *Curves of Descent*. 2013. arXiv: 1212.1231 [math.OC]. URL: https://arxiv.org/abs/1212.1231.

[38] Lawrence C. Evans. *Partial differential equations*. Providence, R.I.: American Mathematical Society, 2010. ISBN: 9780821849743 0821849743.

[39] A. F. Filippov. *Differential Equations with Discontinuous Righthand Sides*. Springer Dordrecht, 1988.

[40] Uriya A. First and Tali Kaufman. "Cosystolic Expansion of Sheaves on Posets with Applications to Good 2-Query Locally Testable Codes and Lifted Codes." In: *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*. STOC 2024. New York, NY, USA: Association for Computing Machinery, 2024, 1446–1457. ISBN: 9798400703836. DOI: 10.1145/3618260.3649625. URL: https://doi.org/10.1145/3618260.3649625.

[41] Jakob N. Foerster, Justin Gilmer, Jascha Sohl-Dickstein, Jan Chorowski, and David Sussillo. "Input Switched Affine Networks: An RNN Architecture Designed for Interpretability." In: *Proceedings of the 34th International Conference on Machine Learning*. Ed. by Doina Precup and Yee Whye Teh. Vol. 70. Proceedings of Machine Learning Research. PMLR, Aug. 2017, pp. 1136–1145. URL: https://proceedings.mlr.press/v70/foerster17a.html.

[42] Joel Friedman. *Sheaves on Graphs, Their Homological Invariants, and a Proof of the Hanna Neumann Conjecture*. Vol. 233. Memoirs of the American Mathematical Society. American Mathematical Society, 2014.

[43] Rudolf Fritsch and Renzo Piccinini. *Cellular Structures in Topology*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1990.

[44] P. Ghez, R. Lima, and J. E. Roberts. "W*-categories." In: *Pacific Journal of Mathematics* 120.1 (1985), pp. 79 –109.

[45] Robert Ghrist. *Elementary Applied Topology*. Createspace, Sept. 2014. ISBN: 978-1502880857. URL: https://www2.math.upenn.edu/~ghrist/notes.html.

[46] Robert Ghrist and Zoe Cooperband. *Obstructions to Reality: Torsors and Visual Paradox*. 2025. arXiv: 2507.01226 [math.AT]. URL: https://arxiv.org/abs/2507.01226v1.

[47] Jean Giraud. *Cohomologie non Abelienne*. Grundlehren der mathematischen Wissenschaften. Springer Berlin, 1971.

[48] Juan Giribet, Matthias Langer, Francisco Martínez Pería, Friedrich Philipp, and Carsten Trunk. "Spectral enclosures for a class of block operator matrices." In: *Journal of Functional Analysis* 278.10 (2020), p. 108455. ISSN: 0022-1236. DOI: https://doi.org/10.1016/j.jfa.2019.108455. URL: https://www.sciencedirect.com/science/article/pii/S0022123619304483.

[49] Corina Grosu and Florian-Horia Vasilescu. "The Künneth formula for Hilbert complexes." In: *Integral Equations and Operator Theory* 5.1 (1982), pp. 1–17. DOI: 10.1007/BF01694027. URL: https://doi.org/10.1007/BF01694027.

[50] Mohammad Hajiahmadi, Bart De Schutter, and Hans Hellendoorn. "Design of Stabilizing Switching Laws for Mixed Switched Affine Systems." In: *IEEE Transactions on Automatic Control* 61.6 (2016), pp. 1676–1681. DOI: 10.1109/TAC.2015.2480216.

[51] Jack K. Hale. "Dynamical systems and stability." In: *Journal of Mathematical Analysis and Applications* 26.1 (1969), pp. 39–59. ISSN: 0022-247X. DOI: https://doi.org/10.1016/0022-247X(69)90175-9. URL: https://www.sciencedirect.com/science/article/pii/0022247X69901759.

[52] Jakob Hansen. "Laplacians of Cellular Sheaves: theory and applications." PhD thesis. University of Pennsylvania, 2020.

[53] Jakob Hansen and Thomas Gebhart. "Sheaf neural networks." In: *NeurIPS 2020 Workshop on Topological Data Analysis and Beyond*. 2020.

[54] Jakob Hansen and Robert Ghrist. "Toward a spectral theory of cellular sheaves." In: *Journal of Applied and Computational Topology* 3.4 (Aug. 2019), 315–358. ISSN: 2367-1734. DOI: 10.1007/s41468-019-00038-7. URL: http://dx.doi.org/10.1007/s41468-019-00038-7.

[55] Jakob Hansen and Robert Ghrist. *Opinion Dynamics on Discourse Sheaves*. 2020. arXiv: 2005.12798 [math.DS].

[56] Alain Haraux and Mohamed Ali Jendoubi. *The Convergence Problem for Dissipative Autonomous Systems*. SpringerBriefs in Mathematics. Springer Cham, 2003.

[57] Allen Hatcher. *Algebraic topology*. Cambridge: Cambridge University Press, 2002. ISBN: 0-521-79160-X; 0-521-79540-0.

[58] Jens Hemelaer. *Grothendieck topologies on posets*. 2018. arXiv: 1811.10039 [math.CT]. URL: https://arxiv.org/abs/1811.10039.

[59] Patricia Hersh. *CW posets after the Poincare Conjecture*. 2014. arXiv: 1411.1296 [math.CO]. URL: https://arxiv.org/abs/1411.1296.

[60] Patricia Hersh and Richard Kenyon. "Shellability of face posets of electrical networks and the CW poset property." In: *Advances in Applied Mathematics* 127 (2021), p. 102178. ISSN: 0196-8858. DOI: https://doi.org/10.1016/j.aam.2021.102178. URL: https://www.sciencedirect.com/science/article/pii/S0196885821000166.

[61] Chris Heunen and Bart Jacobs. "Quantum Logic in Dagger Kernel Categories." In: *Electronic Notes in Theoretical Computer Science* 270.2 (2011). Proceedings of the 6th International Workshop on Quantum Physics and Logic (QPL 2009), pp. 79–103. ISSN: 1571-0661. DOI: https://doi.org/10.1016/j.entcs.2011.01.024. URL: https://www.sciencedirect.com/science/article/pii/S1571066111000259.

[62] Chris Heunen and Martti Karvonen. "Limits in dagger categories." In: *Theory and Applications of Categories* 34 (18 2019), pp. 468–513.

[63] Chris Heunen and Andre Kornell. "Axioms for the category of Hilbert spaces." In: *Proceedings of the National Academy of Sciences* 119.9 (2022), e2117024119. DOI: 10.1073/pnas.2117024119. eprint: https://www.pnas.org/doi/pdf/10.1073/pnas.2117024119. URL: https://www.pnas.org/doi/abs/10.1073/pnas.2117024119.

[64] Chris Heunen and Jamie Vicary. *Categories for Quantum Theory: An Introduction*. Oxford: Oxford University Press, Nov. 2019. DOI: 10.1093/oso/9780198739623.001.0001.

[65] Michael Holst and Ari Stern. "Geometric Variational Crimes: Hilbert Complexes, Finite Element Exterior Calculus, and Problems on Hypersurfaces." In: *Foundations of Computational Mathematics* 12.3 (2012), pp. 263–293. DOI: 10.1007/s10208-012-9119-7. URL: https://doi.org/10.1007/s10208-012-9119-7.

[66] Chuan-Shen Hu. *A Brief Note for Sheaf Structures on Posets*. Oct. 2020. DOI: 10.48550/arXiv.2010.09651.

[67] Yu-Wei Huang. *Abstract Orientable Incidence Structure and Algorithms for Finite Bounded Acyclic Categories. I. Incidence Structure*. 2023. arXiv: 2303.04306 [math.CO]. URL: https://arxiv.org/abs/2303.04306.

[68] Raffaele Iervolino, Stephan Trenn, and Francesco Vasca. "Stability of piecewise affine systems through discontinuous piecewise quadratic Lyapunov functions." In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. 2017, pp. 5894–5899. DOI: 10.1109/CDC.2017.8264551.

[69] Raffaele Iervolino, Stephan Trenn, and Francesco Vasca. "Asymptotic Stability of Piecewise Affine Systems With Filippov Solutions via Discontinuous Piecewise Lyapunov Functions." In: *IEEE Transactions on Automatic Control* 66.4 (2021), pp. 1513–1528. DOI: 10.1109/TAC.2020.2996597.

[70] Abderrahim Jourani and Lionel Thibault. "Noncoincidence of Approximate and Limiting Subdifferentials of Integral Functionals." In: *SIAM Journal on Control and Optimization* 49.4 (2011), pp. 1435–1453. DOI: 10.1137/090775506. URL: https://doi.org/10.1137/090775506.

[71] Zohra Kader, Christophe Fiter, Laurentiu Hetel, and Lotfi Belkoura. "Stabilization of switched affine systems with disturbed state-dependent switching laws." In: *International Journal of Robust and Nonlinear Control* 28.2 (Jan. 2018), pp. 582–595. DOI: 10.1002/rnc.3887. URL: https://inria.hal.science/hal-01689247.

[72] Martti Karvonen. "The Way of the Dagger." PhD thesis. University of Edenburgh, 2019.

[73] T. Kato. *Perturbation theory for linear operators*. Classics in Mathematics. Reprint of the 1980 edition. Springer, 1995.

[74] Y. Komura. "Nonlinear Semi-Groups in Hilbert Space." In: *Journal of the Mathematical Society of Japan* 19.4 (1967), pp. 493–507.

[75] D. Kozlov. *Combinatorial Algebraic Topology*. Algorithms and Computation in Mathematics. Berlin: Springer, 2008. ISBN: 978-3-540-71961-8.

[76] Spencer Kraisler, Shahriar Talebi, and Mehran Mesbahi. "Distributed Consensus on Manifolds using the Riemannian Center of Mass." In: *2023 IEEE Conference on Control Technology and Applications (CCTA)*. 2023, pp. 130–135. DOI: 10.1109/CCTA54093.2023.10253227.

[77] M. G. Kreĭn. "The theory of self-adjoint extensions of semibounded Hermitian transformations and its applications." In: *Mat. Sbornik* 20 (1947).

[78] S. H. Kulkarni, M. T. Nair, and G. Ramesh. "Some properties of unbounded operators with closed range." In: *Proceedings Mathematical Sciences* 118.4 (2008), pp. 613–625. DOI: 10.1007/s12044-008-0047-z. URL: https://doi.org/10.1007/s12044-008-0047-z.

[79] Serge Lang. *Fundamentals of Differential Geometry*. Graduate Texts in Mathematics. Springer New York, 1999.

[80] John M. Lee. *Introduction to Smooth Manifolds*. eng. New York: Springer, 2012. ISBN: 978-1-4899-9475-2 978-1-4419-9982-5.

[81] Damien Lejay. *Constructible hypersheaves via exit paths*. 2021. arXiv: 2102.12325 [math.AT]. URL: https://arxiv.org/abs/2102.12325.

[82] Bert Lindenhovius. *Grothendieck topologies on a poset*. 2014. arXiv: 1405.4408 [math.CT]. URL: https://arxiv.org/abs/1405.4408.

[83] Jacob Lurie. "Higher algebra." Unpublished. Available online at https://www.math.ias.edu/~lurie/. Sept. 2017.

[84] Saunders MacLane. *Categories for the Working Mathematician*. Graduate Texts in Mathematics, Vol. 5. New York: Springer-Verlag, 1971, pp. ix+262.

[85] R. MacPherson. *Night seminars at the Institute for Advanced Study*. 2012–2014.

[86] Clint McCrory. "Cone Complexes and PL Transversality." In: *Transactions of the American Mathematical Society* 207 (1975), pp. 269–291. ISSN: 00029947. URL: http://www.jstor.org/stable/1997177 (visited on 07/07/2025).

[87] Matthew Di Meglio. R^**-categories: The Hilbert-space analogue of abelian categories*. 2025. arXiv: 2312.02883 [math.CT]. URL: https://arxiv.org/abs/2312.02883.

[88] Peter W. Michor. *Functors and Categories of Banach Spaces*. Lecture Notes in Mathematics. Springer Berlin, 1978.

[89] Haynes R. Miller. "Leray in Oflag XVIIA: The origins of sheaf theory, sheaf cohomology, and spectral sequences." In: *Jean Leray (1906-1998)*. 2000. URL: https://api.semanticscholar.org/CorpusID:13024093.

[90] Michael Moy, Alan Hylton, Robert Kassouf-Short, Jacob Cleveland, Jihun Hwang, Justin Curry, Mark Ronnenberg, Miguel Lopez, and Oliver Chiriac. "A Proposed Clock Synchronization Method for the Solar System Internet." In: *2024 IEEE Aerospace Conference*. 2024, pp. 1–17. DOI: 10.1109/AERO58975.2024.10521325.

[91] Manfred Möller and Franciszek Hugon Szafraniec. "Adjoints and Formal Adjoints of Matrices of Unbounded Operators." In: *Proceedings of the American Mathematical Society* 136.6 (2008), pp. 2165–2176. ISSN: 00029939, 10886826. URL: http://www.jstor.org/stable/20535398 (visited on 07/21/2025).

[92] Rainer Nagel. "Towards a "Matrix Theory" for Unbounded Operator Matrices." In: *Mathematische Zeitschrift* 201.1 (1989), pp. 57–68. URL: http://eudml.org/doc/174034.

[93] Vidit Nanda. "Discrete Morse theory and localization." In: *Journal of Pure and Applied Algebra* 223.2 (2019), pp. 459–488. ISSN: 0022-4049. DOI: https://doi.org/10.1016/j.jpaa.2018.04.001. URL: https://www.sciencedirect.com/science/article/pii/S0022404918300872.

[94] John D. Norton. "Causation as Folk Science." In: *Philosophers' Imprint* 3.4 (Nov. 2003), pp. 1–22.

[95] Dirk Pauly and Michael Schomburg. "Hilbert complexes with mixed boundary conditions part 1: de Rham complex." In: *Mathematical Methods in the Applied Sciences* 45.5 (Jan. 2022), 2465–2507. ISSN: 1099-1476. DOI: 10.1002/mma.7894. URL: http://dx.doi.org/10.1002/mma.7894.

[96] Dirk Pauly and Michael Schomburg. "Hilbert complexes with mixed boundary conditions—Part 2: Elasticity complex." In: *Mathematical Methods in the Applied Sciences* 45.16 (Aug. 2022), 8971–9005. ISSN: 1099-1476. DOI: 10.1002/mma.8242. URL: http://dx.doi.org/10.1002/mma.8242.

[97] Juan Peypouquet. *Convex Optimization in Normed Spaces*. SpringerBriefs in Optimization. Springer Cham, 2015.

[98] G. B. Preston. "Inverse Semi-Groups." In: *Journal of the London Mathematical Society* s1-29.4 (1954), pp. 396–403. DOI: https://doi.org/10.1112/jlms/s1-29.4.396. eprint: https://londmathsoc.onlinelibrary.wiley.com/doi/pdf/10.1112/jlms/s1-29.4.396. URL: https://londmathsoc.onlinelibrary.wiley.com/doi/abs/10.1112/jlms/s1-29.4.396.

[99] David Quinn. "The incidence algebras of posets and acyclic categories." English. In: *Kyushu Journal of Mathematics* 67.1 (Mar. 2013), pp. 117–127. DOI: 10.2206/kyushujm.67.117.

[100] M. Reed and B. Simon. *Methods of Modern Mathematical Physics. I Functional Analysis*. Second. New York: Academic Press, 1980.

[101] E. Robinson and G. Rosolini. "Categories of partial maps." In: *Information and Computation* 79.2 (1988), pp. 95–130. ISSN: 0890-5401. DOI: https://doi.org/10.1016/0890-5401(88)90034-X. URL: https://www.sciencedirect.com/science/article/pii/089054018890034X.

[102] Michael Robinson. "The Nyquist theorem for cellular sheaves." In: *Proceedings of the 10th International Conference on Sampling Theory and Applications*. 2013.

[103] J. Rosicky and W. Tholen. "Factorization, fibration and torsion." eng. In: *Journal of Homotopy and Related Structures* 2.2 (2007), pp. 295–314. URL: http://eudml.org/doc/230964.

[104] David Ruelle. *Elements of Differentiable Dynamics and Bifurcation Theory*. Academic Press, inc., 1989.

[105] Antonio Russo, Gian Paolo Incremona, Alberto Cavallo, and Patrizio Colaneri. "State Dependent Switching Control of Affine Linear Systems With Dwell Time: Application to Power Converters." In: *2022 American Control Conference (ACC)*. 2022, pp. 3807–3813. DOI: 10.23919/ACC53348.2022.9867761.

[106] Konrad Schmüdgen. "Discrete Spectra of Self-adjoint Operators." In: *Unbounded Self-adjoint Operators on Hilbert Space*. Dordrecht: Springer Netherlands, 2012, pp. 265–280. ISBN: 978-94-007-4753-1. DOI: 10.1007/978-94-007-4753-1_12. URL: https://doi.org/10.1007/978-94-007-4753-1_12.

[107] Jean-Pierre Schneiders. "Quasi-abelian categories and sheaves." eng. In: *Mémoires de la Société Mathématique de France* 76 (1999), pp. III1–VI140. URL: http://eudml.org/doc/94927.

[108] Christian Schnell. *Complex manifolds*. 2024. URL: https://www.math.stonybrook.edu/~cschnell/pdf/notes/complex-manifolds.pdf.

[109] Peter Selinger. "Dagger Compact Closed Categories and Completely Positive Maps: (Extended Abstract)." In: *Electronic Notes in Theoretical Computer Science* 170 (2007). Proceedings of the 3rd International Workshop on Quantum Programming Languages (QPL 2005), pp. 139–163. ISSN: 1571-0661. DOI: https://doi.org/10.1016/j.entcs.2006.12.018. URL: https://www.sciencedirect.com/science/article/pii/S1571066107000606.

[110] Peter Selinger. "Finite dimensional Hilbert spaces are complete for dagger compact closed categories." In: *Logical Methods in Computer Science* Volume 8, Issue 3 (Aug. 2012). ISSN: 1860-5974. DOI: 10.2168/lmcs-8(3:6)2012. URL: http://dx.doi.org/10.2168/LMCS-8(3:6)2012.

[111] Alexandre Seuret, Carolina Albea, and Francisco Gordillo. "Practical Stabilization of Switched Affine Systems: Model and Data-Driven Conditions." In: *IEEE Control Systems Letters* 7 (2023), pp. 1628–1633. DOI: 10.1109/LCSYS.2023.3263432.

[112] Allan D. Shepard. "A Cellular Description of the Derived Category of a Stratified Space." PhD thesis. Brown University, 1980.

[113] Leon Simon. "Asymptotics for a Class of Non-Linear Evolution Equations, with Applications to Geometric Problems." In: *Annals of Mathematics* 118.3 (1983), pp. 525–571. ISSN: 0003486X, 19398980. URL: http://www.jstor.org/stable/2006981 (visited on 07/04/2025).

[114] Ivan Singer. *The Theory of Best Approximation and Functional Analysis*. Society for Industrial and Applied Mathematics, 1974.

[115] Julian Suk, Lorenzo Giusti, Tamir Hemo, Miguel Lopez, Konstantinos Barmpas, and Cristian Bodnar. "Surfing on the neural sheaf." In: *NeurIPS 2022 Workshop on Symmetry and Geometry in Neural Representations*. 2022.

[116] Dieter Teichrib and Moritz Schulze Darup. "Reachability analysis for piecewise affine systems with neural network-based controllers." In: *2024 IEEE 63rd Conference on Decision and Control (CDC)*. IEEE, Dec. 2024, 894–900. DOI: 10.1109/cdc56724.2024.10886414. URL: http://dx.doi.org/10.1109/CDC56724.2024.10886414.

[117] C. Tretter. *Spectral Theory of Block Operator Matrices and Applications*. Imperial College Press, 2008. ISBN: 9781848161122. URL: https://books.google.com/books?id=xKJgDQAAQBAJ.

[118] David Treumann. "Exit paths and constructible stacks." In: *Compositio Mathematica* 145.6 (2009), 1504–1532. DOI: 10.1112/S0010437X09004229.

[119] Roberto Tron, Bijan Afsari, and René Vidal. "Average consensus on Riemannian manifolds with bounded curvature." In: *2011 50th IEEE Conference on Decision and Control and European Control Conference*. 2011, pp. 7855–7862. DOI: 10.1109/CDC.2011.6160965.

[120] Roberto Tron, Bijan Afsari, and René Vidal. "Riemannian Consensus for Manifolds With Bounded Curvature." In: *IEEE Transactions on Automatic Control* 58.4 (2013), pp. 921–934. DOI: 10.1109/TAC.2012.2225533.

[121] Lao Tzu. *Tao Te Ching*. Trans. by Arthur Waley. Wordsworth Editions, 1996.

[122] V. Vagner. "Generalized Groups." In: *Doklady Akademii Nauk* 84 (1952), pp. 1119–1122.

[123] Ezio Vasselli. "Bundles of C*-categories." In: *Journal of Functional Analysis* 247.2 (2007), pp. 351–377. ISSN: 0022-1236. DOI: https://doi.org/10.1016/j.jfa.2007.03.016. URL: https://www.sciencedirect.com/science/article/pii/S0022123607001024.

[124] Torsten Wedhorn. *Manifolds, Sheaves, and Cohomology*. Springer Studium Mathematik. Springer Spektrum Wiesbaden, 2016.

[125] Wikipedia. *Functional analysis — Wikipedia, The Free Encyclopedia*. http://en.wikipedia.org/w/index.php?title=Functional%20analysis&oldid=1247960073. [Online; accessed 17-January-2025]. 2025.

[126] nLab authors. *thin category*. https://ncatlab.org/nlab/show/thin+category. Revision 19. June 2025.