

Discrete microlocal Morse theory[☆]Adam Brown^{*}, Ondřej Draganov*Institute of Science and Technology Austria, Am Campus 1, Klosterneuburg, 3400, Austria*

ARTICLE INFO

Article history:

Received 1 October 2022

Received in revised form 1 June 2025

Available online 12 August 2025

Communicated by C.A. Weibel

MSC:

55N30; 55N31; 55-08

Keywords:

Derived sheaf theory

Discrete Morse theory

Computational topology

Microsupport

ABSTRACT

We establish several results combining discrete Morse theory and microlocal sheaf theory in the setting of finite posets and simplicial complexes. Our primary tool is a computationally tractable description of the bounded derived category of sheaves on a poset with the Alexandrov topology. We prove that each bounded complex of sheaves on a finite poset admits a unique (up to isomorphism of complexes) minimal injective resolution, and we provide algorithms for computing minimal injective resolution of an injective complex, as well as several useful functors between derived categories of sheaves. For the constant sheaf on a simplicial complex, we give asymptotically tight bounds on the complexity of computing the minimal injective resolution using those algorithms. Our main result is a novel definition of the discrete microsupport of a bounded complex of sheaves on a finite poset. We detail several foundational properties of the discrete microsupport, as well as a microlocal generalization of the discrete homological Morse theorem and Morse inequalities.

© 2025 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Contents

1. Introduction	2
2. Background and terminology	5
3. Fixing a basis of a complex of injective sheaves	8
3.1. Poset labeled matrices	8
3.2. Change of bases for labeled matrix representations	11
3.3. Complexes of labeled matrices	12
4. Derived categories	13
4.1. Morphisms, null-homotopy, quasi-isomorphism	13
4.2. Minimal complexes of injective sheaves	15
4.3. Recognizing minimality and minimizing an injective complex	18
4.4. Single sheaf in derived category: computing injective resolutions	20
4.5. Injective resolution of the constant sheaf	23

[☆] This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant no. 788183, from the Wittgenstein Prize, Austrian Science Fund (FWF), grant no. Z 342-N31, and from the DFG Collaborative Research Center TRR 109, 'Discretization in Geometry and Dynamics', Austrian Science Fund (FWF), grant no. I 02979-N35.

^{*} Corresponding author.

E-mail addresses: abrown@ist.ac.at (A. Brown), ondrej.draganov@ist.ac.at (O. Draganov).

5.	Derived functors	28
5.1.	Definitions of derived functors	28
5.1.1.	Pushforward	28
5.1.2.	Pullback	29
5.1.3.	Proper pushforward	29
5.1.4.	Proper pullback	30
5.1.5.	Summary	31
5.2.	Derived functors and hypercohomology	31
5.3.	Recollement for posets	33
5.4.	Algorithms for computing derived functors	33
5.4.1.	Poset mapping cylinder of an order preserving map	34
5.4.2.	Computing pushforward	34
5.4.3.	Computing pullback	35
5.4.4.	Computing proper pushforward	36
5.4.5.	Computing proper pullback	37
6.	Microlocal sheaf theory and discrete Morse theory	37
6.1.	Discrete microsupport	38
6.2.	The discrete microlocal Morse theorem and inequalities	40
6.3.	Reduction to classical discrete Morse theory	44
7.	Examples	44
	CRediT authorship contribution statement	53
	Declaration of competing interest	53
	Appendix A. Background: injective sheaves	53
	A.1. Duality between injective and projective resolutions	53
	A.2. Examples of non-injective sheaves	54
	A.3. Proofs of standard facts about injective sheaves	55
	Appendix B. The correspondence between Definition 4.10 and the bounded derived category of sheaves on a finite poset with the Alexandrov topology	56
	Appendix C. Derived category and injective resolution of a sheaf	58
	C.1. Computing a basis for the space of morphisms in a derived category	58
	C.2. Injective resolution via order complex	58
	C.3. Injective hull computation	61
	C.4. Geometrical meaning of multiplicities in the minimal resolution	62
	References	63

1. Introduction

A motif of twentieth-century mathematics is the investigation of global structures via local inquiry. Here, we build on two instances of this global-local theme.

The first instance is Morse theory. Classically, Morse theory studies the global topological structure of a manifold by (locally) analyzing critical points of a real-valued smooth function on the manifold [39,41]. Morse theory has since been generalized to stratified spaces [23] and applied to many areas of mathematics. Forman applied these ideas to study cell complexes, and in doing so initiated the study of discrete Morse theory [19,20]. Forman's reformulation has since led to many advances in applied and computational topology [4, 18,21,28,40,43,45,51].

The second instance is sheaf theory, the poster child of the global-local motif. Introduced by Leray [38] and modernized by Grothendieck, with numerous contributions from Cartan, Serre, Verdier, Deligne, and many others, sheaf theory is a general framework combining homological algebra and topology (see 'A Short History: Les débuts de la théorie des faisceaux' by Houzel in [33] for a beautiful synopsis of the historical development of sheaf theory). Analogous to the recent expansion of discrete Morse theory, a growing body of research studies sheaves on cell complexes with applications in applied and computational topology [12–16,42,44], although sheaves on finite posets were already studied earlier, e.g., in [17]. However, the utility of sheaf theory is most fully realized when working in the setting of derived categories. Computational aspects of derived cellular sheaf theory are explored in [6,7,9,15,49]. Despite these advances, algorithms for many useful computations within derived cellular sheaf theory remain (to the best of our knowledge) underdeveloped.

At the intersection of Morse theory and derived sheaf theory lies microlocal geometry. The ‘microlocal’ perspective of sheaf theory originates from Sato’s study of singularities within systems of linear differential equations [47,50]. This perspective considers the propagation of various local structures (such as local solutions of differential equations or homological vanishing properties of sheaves) along different directions within (the cotangent bundle of) a manifold. Microlocal geometry (and its relatives) have far-reaching applications throughout modern mathematics, including intersection cohomology and perverse sheaf theory [22,24], enumerative geometry [5], symplectic geometry [46], and representation theory [1] (to name a few).

This paper introduces a discrete analog of two important concepts of microlocal geometry, originally introduced by Kashiwara–Schapira for manifolds. The first is the notion of microsupport, which quantifies the cohomological vanishing properties of sheaves on ‘local half-spaces’ [31,32]. The second is a microlocal generalization of the Morse theorem and inequalities, which exchanges singular homology and Euler characteristics with hypercohomology and the Euler–Poincaré index [22,30,52]. Before concluding this introduction with a summary of these results, we first give a brief sketch of an indispensable prerequisite: a systematic treatment of computational derived sheaf theory.

Computational derived sheaf theory Sheaves use algebra to model relationships between local and global properties of a topological space. When the topological space is a poset with the Alexandrov topology, a *sheaf* (of finite-dimensional vector spaces), F , is defined by associating a finite-dimensional vector space, $F(\sigma)$, to each element, σ , and a linear map, $F(\sigma \leq \tau) : F(\sigma) \rightarrow F(\tau)$, to each relation, $\sigma \leq \tau$ (subject to commutativity requirements, see Definition 2.2). The utility of this definition is also its foil: the high level of generality encompasses many pathologies.

A common strategy for analyzing such a complicated mathematical structure is to approximate or represent it with a collection of simpler, or at least more familiar, objects; the goal is to reframe questions concerning the complex structure as questions about the building blocks that represent it. The first tool of this paper is a particular instance of this phenomenon: *injective resolutions*. An *injective resolution* represents a given sheaf (much like a Fourier series represents a periodic function) with an exact sequence, $0 \rightarrow F \rightarrow I^0 \rightarrow I^1 \rightarrow I^2 \rightarrow \dots$, of *injective sheaves* I^d , which admit many desirable properties (see, for example, Lemma 2.8, Proposition 2.9, and Lemma 2.12). Efficient algorithms for computing these sequences are a first step toward applying well-established and powerful theoretical results from derived sheaf theory to computational topology. In this paper, we aim to present this theory in an explicit and computationally amenable framework.

Injective resolutions are used to study sheaves from the ‘derived’ perspective, i.e. as objects in a derived category (Definition 4.10). These derived categories unify and generalize many variants of (co)homology, such as simplicial cohomology, Borel–Moore homology, intersection cohomology, etc. For example, simplicial cohomology (and level-set persistent cohomology, see [7]) can be computed from an injective resolution of the constant sheaf (see Example 2.3 and Section 7), illustrating that even an injective resolution of the constant sheaf contains subtle topological information. Several recent works point to the potential benefits of applying derived sheaf theory to the study of persistent homology [6,7,9,15,34,35]. We approach this subject from a computational perspective to help bridge gaps between applied topology and derived sheaf theory. With this goal in mind, we aim to limit the mathematical prerequisites of our approach whenever possible (a choice that often comes at the cost of brevity).

Main results Broadly, this paper develops computational methods for the bounded derived category sheaves of finite-dimensional vector spaces on finite posets with the Alexandrov topology. Our main contributions are:

1. We establish the existence and uniqueness of a *minimal* injective resolution of a given complex of sheaves (Theorem 4.9). We formulate an inductive algorithm for computing minimal injective resolutions (Algorithm 3) using elementary methods from linear algebra. For the constant sheaf on a simplicial complex, we give an asymptotically tight bound on the complexity of Algorithm 3 (Proposition 4.26 and Corollary 4.27).
2. We give an explicit description of the set of objects and morphisms for (a skeleton of) the bounded derived category of sheaves on a finite poset in terms of ‘labeled matrices’ (Definition 4.10), and provide algorithms for computing right derived (proper) pushforward and (proper) pullback functors along order preserving maps of posets (Section 5). See Appendix B for the correspondence between Definition 4.10 and the standard definition of a derived category.
3. We give a novel definition of the discrete microsupport for a complex of sheaves in the derived category of sheaves on a finite poset (Definition 6.3), and prove a microlocal generalization of the discrete Morse theorem and inequalities (Theorem 6.7 and 6.8).
4. A Python software package implementing the algorithms developed in this paper is freely available at <https://github.com/OnDraganov/desc> [3]. It demonstrates the utility of our approach—for a simplicial complex with a few thousands of simplices, the computation of the minimal injective resolution of the constant sheaf as well as (proper) pushforwards/pullbacks of said resolution take just several seconds on a regular laptop.

One pedagogical benefit of our approach is that while sheaf theory is an indispensable component of our proofs, the computations we describe make no fundamental use of sheaves or category theory, hopefully making these calculations more accessible to anyone with knowledge of linear algebra and elementary poset combinatorics.

Comparison to prior work Both derived sheaf theory and Morse theory are rich subjects that have been thoroughly studied for several decades. There are many textbooks on sheaf theory [10,29,33] and several publications which study sheaves on finite topological spaces. In [49], Shepard relates sheaves on finite cell complexes (viewed as posets) to the classical setting of constructible sheaves on stratified topological spaces. In [36,37], Ladkani studies the homological properties of finite posets and introduces combinatorial criteria guaranteeing derived equivalences between categories of sheaves. In [15], Curry establishes a connection between sheaf theory and persistent homology. More recently, several publications expand on the work initiated by Curry on applications of derived sheaf theory to persistent homology [6,7,9,34,35]. In a set of lecture notes, Goresky beautifully explains derived and perverse sheaf theory in both the Whitney stratified and cellular settings [25]. There are also several closely related works on discrete Morse theory. Notably, Nanda [43] describes a categorical localization procedure for Morse-theoretic cellular simplification which relates to sheaf propagation and our definition of microsupport (Definition 6.3). Additionally, Sköldbberg gives an algebraic approach to discrete Morse theory by studying reductions of chain complexes that preserve the homotopy type of the complex [51]. This work is reminiscent of the derived category approach we take in Section 4. Motivated by these advances and the potential to develop new techniques for computational topology, we aim to establish preliminary results on computational aspects of derived sheaf theory for finite topological spaces. The contributions of this paper are the first of our knowledge to use discrete Morse theory to study the microsupport of derived sheaves from a computational perspective.

Acknowledgments We thank François Petit for the useful discussions and helpful feedback on an earlier version of the manuscript. We thank Herbert Edelsbrunner and Sebastiano Cultrera di Montesano for teaching a course on discrete Morse theory at the Institute of Science and Technology Austria in the spring of 2022, during which much of this work was completed.

2. Background and terminology

In this paper we study finite-dimensional vector space valued sheaves on finite posets, with particular emphasis on face posets of simplicial complexes. To clarify the objects of study and terminology, we begin by recalling definitions and well-known results, drawing heavily from [15,16,25,36,49]. Throughout the paper we fix a field k . In the examples we often choose $k = \mathbb{Z}_2$, the two-element field.

A *poset* is a partially ordered set. We usually denote it by Π with relation \geq . We only consider finite posets. For elements π, τ in a poset Π , we write $\pi <_1 \tau$ if $\pi \leq \tau$ and there is no other element between π and τ . We use some of the standard terminology from simplicial complexes for general posets: the *star* of an element σ is $\text{St } \sigma = \{\tau \in \Pi \mid \sigma \leq \tau\}$, the *boundary* is $\text{bnd}(\sigma) = \{\tau \in \Pi \mid \tau <_1 \sigma\}$, and the *coboundary* is $\text{cobnd}(\sigma) = \{\tau \in \Pi \mid \sigma <_1 \tau\}$. A chain $\pi_0 < \dots < \pi_n$ has *length* n , and the *height* of a poset is the greatest length over all its chains. This corresponds to dimension if $\Pi = (\Sigma, \subseteq)$ is the face poset of a simplicial complex. An *abstract simplicial complex* Σ on a vertex set V is a system of subsets of V such that $A \subseteq B \in \Sigma$ implies $A \in \Sigma$. The combinatorial aspects of posets are tied with geometrical concepts via Alexandrov topology.

Definition 2.1. A finite poset Π is a T_0 -topological space with the *Alexandrov topology*, where a set is *open* when it is upwards closed:

$$U \subset \Pi \text{ is open if and only if } \tau \in U \text{ implies } \gamma \in U \text{ for each } \gamma \geq \tau.$$

The closure of a set $Z \subseteq \Pi$ is $\text{Cl } Z = \{\pi \in \Pi \mid \pi \leq \sigma \text{ for some } \sigma \in Z\}$.

We study vector space representations of posets from the perspective of sheaf theory. Since we only consider finite dimensional vector spaces, what we call a *sheaf* is equivalently also a quiver representation with relations, a module over a bound quiver algebra, a persistence module or a functor from the poset seen as a category to the category of vector spaces (see [15, Theorem 4.2.10]).

Definition 2.2. A *sheaf* F on a finite poset Π is an assignment of a finite-dimensional k -vector space $F(\pi)$ to each element $\pi \in \Pi$, and an assignment of a linear map

$$F(\tau \leq \gamma) : F(\tau) \rightarrow F(\gamma)$$

to each face relation $(\tau \leq \gamma) \in \Pi$, such that

1. $F(\tau \leq \tau)$ is the identity map, $\text{id}_{F(\tau)}$, and
2. $F(\tau \leq \gamma) \circ F(\sigma \leq \tau) = F(\sigma \leq \gamma)$, for each triple $\sigma \leq \tau \leq \gamma \in \Pi$.

The *support* of F is the set of elements $\pi \in \Pi$ for which $F(\pi) \neq 0$.

Example 2.3. The *constant sheaf*, denoted k_Π , on a poset Π , assigns to each element $\pi \in \Pi$ the one-dimensional vector space, k , and to each relation, $(\pi \leq \tau) \in \Pi$, the identity map id_k .

Definition 2.4. A *natural transformation*, $\eta : F \rightarrow G$, between two sheaves on Π , is a collection of linear maps $\eta(\pi) : F(\pi) \rightarrow G(\pi)$ for each $\pi \in \Pi$, such that

$$G(\tau \leq \gamma) \circ \eta(\tau) = \eta(\gamma) \circ F(\tau \leq \gamma),$$

for each $(\tau \leq \gamma) \in \Pi$. For a natural transformation $\eta : F \rightarrow G$, the kernel, cokernel, image, and coimage are taken point-wise, defining sheaves on Π :

$$(\ker \eta)(\pi) := \ker(\eta(\pi)), \quad (\ker \eta)(\pi \leq \tau) := F(\pi \leq \tau)|_{\ker \eta(\pi)}.$$

Moreover, if $\ker \eta(\pi) = 0$ for each $\pi \in \Pi$, we say that η is *injective*. We write $G/F := \operatorname{coker} \eta$ if η is an injection clear from the context.

As kernels, images and cokernels, other concepts can also be easily defined point-wise on sheaves. Importantly the *direct sum* of sheaves is defined as the direct sum at each $\pi \in \Pi$.

Definition 2.5. A sheaf I is called *injective* if for each injective natural transformation $f : A \hookrightarrow B$, any given natural transformation $\alpha : A \rightarrow I$ can be extended to $\beta : B \rightarrow I$ so that $\alpha = \beta \circ f$.

This condition is always satisfied for sheaves over a single point space (i.e., the assignment of a single vector space to a point): we can extend any linear map on a subspace to the whole space by mapping a complement space to 0. See appendix for examples of non-injective sheaves: Example A.1, Example A.2.

In the case of finite-dimensional sheaves over finite posets, each injective sheaf is created from simple building blocks—as a direct sum of *indecomposable* injective sheaves. This allows us to store them efficiently, as discussed in Section 3. We recall standard definitions and results that can be found in [15] and [49]. Although phrased for sheaves on cell complexes in those references, the generalization to any finite poset is straightforward. We revisit proofs of the standard claims below in Appendix A.3.

Definition 2.6 (cf. [15, Definition 7.1.3], [25, Chapter 23]). For each $\pi \in \Pi$, we define an *indecomposable injective sheaf* $[\pi]$ as

$$[\pi](\sigma) := \begin{cases} k & \text{if } \sigma \leq \pi, \\ 0 & \text{otherwise,} \end{cases} \quad \text{with} \quad [\pi](\sigma \leq \tau) := \begin{cases} \operatorname{id} & \text{if } \sigma \leq \tau \leq \pi, \\ 0 & \text{otherwise.} \end{cases}$$

For $n \in \mathbb{Z}_{\geq 0}$, we denote by $[\pi]^n$, the direct sum $\bigoplus_{j=1}^n [\pi]$. For a vector space V , we denote by $[\pi]^V$, the sheaf $[\pi]^{\dim V}$, with an implicitly fixed isomorphism between $k^{\dim V}$ and V .

Lemma 2.7 (cf. [15, Lemma 7.1.5]). *Indecomposable injective sheaves are injective.*

Lemma 2.8 (cf. [49, Lemma 1.3.1]). *A direct sum of injective sheaves is injective. Additionally, if $I \xrightarrow{\alpha} J$ is an injective natural transformation with I, J injective sheaves, then $J \cong I \oplus \operatorname{coker} \alpha$, and $\operatorname{coker} \alpha$ is an injective sheaf.*

Proposition 2.9 (cf. [15, Lemma 7.1.6], [49, Theorem 1.3.2]). *Each injective sheaf on a finite poset is isomorphic to a direct sum of indecomposable injective sheaves.*

The central objects studied in this paper are (cochain) complexes of sheaves. Unless otherwise stated, all complexes we consider are bounded.

Definition 2.10. A (bounded) *complex of sheaves*, denoted (A^\bullet, μ^\bullet) , is a sequence of sheaves A^d and natural transformations μ^d

$$\cdots \rightarrow A^d \xrightarrow{\mu^d} A^{d+1} \xrightarrow{\mu^{d+1}} A^{d+2} \xrightarrow{\mu^{d+2}} \cdots$$

such that $\mu^{d+1} \circ \mu^d = 0$ for each d , and $A^d = 0$ for $|d|$ sufficiently large. A complex of sheaves is *exact* if $\operatorname{im} \mu^d = \ker \mu^{d+1}$ for each d . A morphism $\alpha^\bullet : A^\bullet \rightarrow B^\bullet$ between complexes of sheaves is a collection of natural transformations $\alpha^d : A^d \rightarrow B^d$ such that the diagrams commute:

$$\begin{array}{ccc} A^d & \xrightarrow{\mu^d} & A^{d+1} \\ \downarrow \alpha^d & & \downarrow \alpha^{d+1} \\ B^d & \xrightarrow{\nu^d} & B^{d+1} \end{array}$$

A complex of sheaves is *injective* if each sheaf in it is injective.

We often abbreviate complex of sheaves just to *complex*. Sometimes we abuse the notation and write $A^\bullet = (A^\bullet, \mu^\bullet)$ when the symbols for natural transformations are clear from context or not used.

The tools discussed in this paper are developed to study (co)homology. One version of it, defined for a complex, is the cohomology sheaf in degree d .

Definition 2.11. For a complex of sheaves (A^\bullet, μ^\bullet) on Π , define, for each $d \in \mathbb{Z}$, the *cohomology sheaf* $H^d(A^\bullet)$ by

$$H^d(A^\bullet)(\sigma) := \ker \mu^d(\sigma) / \operatorname{im} \mu^{d-1}(\sigma).$$

Because μ^\bullet is a morphism of complexes, the linear map $A^d(\sigma \leq \tau)$ maps $\ker \mu^d(\sigma)$ to $\ker \mu^d(\tau)$, and $\operatorname{im} \mu^{d-1}(\sigma)$ to $\operatorname{im} \mu^{d-1}(\tau)$. Therefore, $A^d(\sigma \leq \tau)$ induce maps

$$H^d(A^\bullet)(\sigma \leq \tau) : H^d(A^\bullet)(\sigma) \rightarrow H^d(A^\bullet)(\tau).$$

These linear maps give $H^d(A^\bullet)$ the structure of a sheaf on Π .

We finish the section by describing the structure of natural transformations between injective sheaves—this simple observation is crucial for the algorithmic treatment of injective complexes. There are two ingredients: a natural transformation between two direct sums uniquely decomposes to natural transformations between pairs of summands, and a natural transformation between powers of indecomposable injective sheaves is fully described by a single linear map, which can only be non-zero when the support of the codomain is contained in the support of the domain.

Lemma 2.12. Let $\eta : I \rightarrow J$ be a natural transformation between two injective sheaves on a poset Π , and let

$$I \cong \bigoplus_{\pi \in \Pi} [\pi]^{p_\pi} \quad \text{and} \quad J \cong \bigoplus_{\sigma \in \Pi} [\sigma]^{s_\sigma}$$

be the decompositions of I and J into indecomposable injective sheaves with multiplicities p_π, s_σ . Then η can be uniquely described by a collection of linear maps $f_{\pi\sigma} : k^{p_\pi} \rightarrow k^{s_\sigma}$, one for each pair $\pi \leq \sigma$ in Π . On the other hand, each such collection of linear maps defines a natural transformation. In other words,

$$\operatorname{Hom}(I, J) \cong \bigoplus_{\sigma \leq \pi} \operatorname{Hom}(k^{p_\pi}, k^{s_\sigma}) \cong \bigoplus_{\sigma \leq \pi} k^{p_\pi s_\sigma},$$

where $\operatorname{Hom}(I, J)$ denotes the set of natural transformations from I to J and $\operatorname{Hom}(k^{p_\pi}, k^{s_\sigma})$ denotes the set of linear transformations from k^{p_π} to k^{s_σ} .

Proof. Denote by proj_A and incl_A the projection onto and inclusion of a direct summand A , respectively. Then the map $\eta : I \rightarrow J$ decomposes into a sum of maps between the powers of indecomposable injective sheaves:

$$\eta = \sum_{\pi, \sigma \in \Pi} \eta_{\pi\sigma}, \quad \eta_{\pi\sigma} = \operatorname{proj}_{[\sigma]^{s_\sigma}} \circ \eta \circ \operatorname{incl}_{[\pi]^{p_\pi}}.$$

Consider $\tau \in \Pi$. If $\tau \not\leq \sigma$, then $\eta_{\pi\sigma}(\tau) = 0$. Otherwise, $\eta_{\pi\sigma}(\tau) = [\sigma]^{s_\sigma}(\tau \leq \sigma) \circ \eta_{\pi\sigma}(\tau) = \eta_{\pi\sigma}(\sigma) \circ [\pi]^{p_\pi}(\tau \leq \sigma)$, which is equal to $\eta_{\pi\sigma}(\sigma)$ if $\sigma \leq \pi$, and 0 otherwise. This shows that the natural transformation $\eta_{\pi\sigma}$ is determined by the single linear map $f_{\pi\sigma} := \eta_{\pi\sigma}(\sigma)$. Moreover, this map is necessarily 0 whenever $\sigma \not\leq \pi$, and it can be any linear map otherwise. \square

Example 2.13. Consider the three element “meet” poset, Π , given by two relations— $\sigma < \pi$ and $\tau < \pi$ —and two injective sheaves on Π : $I \cong [\pi]^2 \oplus [\sigma]$ and $J \cong [\pi] \oplus [\sigma] \oplus [\tau]^2$. Then any natural transformation $\eta : I \rightarrow J$ is described by four linear maps: $f_{\pi\pi} : k^2 \rightarrow k$, $f_{\pi\sigma} : k^2 \rightarrow k$, $f_{\pi\tau} : k^2 \rightarrow k^2$, and $f_{\sigma\sigma} : k \rightarrow k$. Note that we do not consider, e.g., $f_{\sigma\pi}$, because $\pi \not\leq \sigma$, and so the only natural transformation $[\sigma] \rightarrow [\pi]$ is 0.

3. Fixing a basis of a complex of injective sheaves

The central concept for our algorithmic treatment of injective complexes is a *labeled matrix* representing a natural transformation between injective sheaves. The same way we can represent a linear map by a matrix once we fix bases of the vector spaces, we can represent natural transformations between injective sheaves by matrices, once we fix specific bases. By definition, a natural transformation between sheaves, $\eta : I \rightarrow J$, is described by many matrices: one $\eta(\pi)$ for each $\pi \in \Pi$. By Lemma 2.12, much less information is necessary in case of injective sheaves. In this section we organize this information in a single matrix and show that multiplying the defined matrices corresponds to composition of natural transformations. First, we need to make sure to fix well behaved systems of bases for a given injective sheaf.

Definition 3.1. Let $I = \bigoplus_{i=1}^m [\pi_i]$ be a fixed decomposition of an injective sheaf on a poset Π , with possible repetitions of π_i . Let U_π be a basis of the vector space $I(\pi)$ for each $\pi \in \Pi$. We call the system of bases $(U_\pi)_{\pi \in \Pi}$ a *decomposition-compatible basis* of I , if each map $I(\sigma \leq \pi)$ with respect to bases U_σ, U_π is a projection to a subset of coordinates, respecting the order of the decomposition. As all systems of bases we consider are decomposition-compatible, we also say just *basis* of I .

Example 3.2. If $\sigma < \pi$ and $I = [\pi] \oplus [\sigma] \oplus [\sigma]$, then the matrix representing $I(\sigma \leq \pi)$ with respect to any decomposition-compatible basis is $(1 \ 0 \ 0)$. Note that this still leaves some freedom. For example, if we multiply the basis vector for the “[π]-coordinate” by the same non-zero constant at each $\tau \leq \pi$, we get a formally different system of bases, but the matrices stay the same.

3.1. Poset labeled matrices

We describe how to represent a natural transformation between two injective sheaves with fixed decomposition-compatible bases as a single matrix with labeled columns and rows. The idea is to simply record the information about the individual maps f_{ij} from Lemma 2.12 in a single matrix. Recall that such a map $f_{\pi\sigma} : [\pi]^{p_\pi} \rightarrow [\sigma]^{s_\sigma}$ can be non-zero only if $\sigma \leq \pi$.

Definition 3.3. For a poset Π , a Π -labeled matrix is a matrix M together with a labeling of each column and row by elements of Π , such that the entry on an intersection of a column labeled by π and a row labeled by σ is non-zero only if $\sigma \leq \pi$.

Columns and rows are labeled independently, the labels can be repeated, and not all labels need to be used. If Π is clear from the context, we also say just *labeled matrix*.

For $\pi, \sigma \in \Pi$, we denote by $M[\sigma, \pi]$ the submatrix of M obtained by deleting all rows not labeled by σ and all columns not labeled by π . The same way for $S, P \subseteq \Pi$ we denote by $M[S, P]$ the submatrix obtained by deleting rows and columns labeled by simplices not in S and P , respectively.

We also allow a labeled matrix that has no rows—corresponding to a morphism into a trivial sheaf—or no columns—corresponding to a morphism from a trivial sheaf. Both are special cases of zero matrices.

Note that the difference between a 0×0 and a $0 \times n$ matrix is more than just a formality—their kernels are 0- and n -dimensional, respectively. This will be important when we discuss exactness later.

Example 3.4. Let Π be a three element poset with $\sigma, \tau < \pi$ as in Example 2.13. Then

$$M = \begin{array}{c} \pi \\ \sigma \\ \tau \\ \tau \end{array} \begin{pmatrix} \begin{array}{cc} \pi & \pi \\ a_1 & a_2 \end{array} & 0 \\ \begin{array}{cc} a_3 & a_4 \end{array} & \begin{array}{c} \sigma \\ a_9 \end{array} \\ \begin{array}{cc} a_5 & a_6 \end{array} & 0 \\ \begin{array}{cc} a_7 & a_8 \end{array} & 0 \end{pmatrix}$$

is a Π -labeled matrix. Each of the outlined blocks is a matrix representing one of the four linear maps from Example 2.13, e.g., $M[\sigma, \pi] = (a_3 \ a_4)$ describes $f_{\pi\sigma}$. Note that the zeros are all forced, as non-zero values are only allowed for entries with ‘row-label \leq column-label’.

Example 3.5. The labeled matrices are, in a way, “lower-triangular”: if Π is linearly ordered, $\pi_n > \dots > \pi_1$, then a labeled matrix with both columns and rows labels exactly $(\pi_n, \pi_{n-1}, \dots, \pi_1)$ is a lower-triangular matrix.

Example 3.6. The elements of an *incidence algebra* of a poset Π over a field k are special cases of labeled matrices. By definition, such an element is a function, f , giving each interval $[\sigma, \pi]$ in Π a value in k . This is equivalently described by a labeled matrix, M , with both columns and rows labeled by each element of Π once—the value $f([\sigma, \pi])$ is then the single entry of $M[\sigma, \pi]$. The operations on the incidence algebra agree with the operations we define on labeled matrices. This relation shows that the incidence algebra with field coefficients can be viewed as the space of endomorphisms of the injective sheaf $\bigoplus_{\pi \in \Pi} [\pi]$.

See also Example 3.10 below. For even more examples see matrices in Section 7 labeled by simplices ordered by inclusion, e.g., $13 < 123$, with entries in the field $k = \mathbb{Z}_2$.

Definition 3.7. A Π -labeled matrix M represents a natural transformation $\eta : I \rightarrow J$ between two injective sheaves if the following conditions hold:

- $I \cong \bigoplus_{i=1}^m [\pi_i]$, where π_1, \dots, π_m are the labels of columns of M (in that order),
- $J \cong \bigoplus_{j=1}^n [\sigma_j]$, where $\sigma_1, \dots, \sigma_n$ are the labels of rows of M (in that order),
- For every $\sigma \leq \pi$ from Π , the matrix $M[\sigma, \pi]$ represents the map $f_{\sigma\pi} = \eta_{\sigma\pi}(\sigma)$ described in Lemma 2.12 with respect to fixed decomposition-compatible bases of I and J .

Once a decomposition-compatible basis is chosen for I and J , there is a one-to-one correspondence between natural transformations $\eta : I \rightarrow J$, and labeled matrices with the labeling from the decompositions of I and J . This follows from Lemma 2.12.

The information can be condensed in this way, because the matrices $\eta(\tau)$ overlap. If M represents η , then the matrix representing a linear map $\eta(\tau)$ is a submatrix of M . Recall that $\text{St } \tau$ is the set of elements greater or equal to τ .

Lemma 3.8. If a labeled matrix M represents a natural transformation η , and $\tau \in \Pi$, the submatrix $M[\text{St } \tau, \text{St } \tau]$ is the matrix of the linear map $\eta(\tau)$ with respect to the fixed decomposition-compatible bases.

Proof. As in Lemma 2.12, we have $\eta = \bigoplus_{\sigma \leq \pi} \eta_{\pi\sigma}$. If p_π is the multiplicity of $[\pi]$ in the decomposition of I , then

$$\eta(\tau) = \bigoplus_{\sigma \leq \pi} \eta_{\pi\sigma}(\tau) = \bigoplus_{\tau \leq \sigma \leq \pi} \eta_{\pi\sigma}(\sigma) \circ [\pi]^{p_\pi}(\tau \leq \sigma).$$

By definition, $M[\sigma, \pi]$ is the matrix of $\eta_{\pi\sigma}(\sigma)$ with respect to fixed decomposition-compatible bases for I and J . Moreover, restriction maps $[\pi]^{p_\pi}[\tau \leq \sigma]$ are projections to coordinates consistent with the decomposition. Therefore, $M[\text{St } \tau, \text{St } \tau]$ is the matrix of $\eta(\tau)$ with respect to those fixed bases. \square

Example 3.9. The matrix M from the previous Example 3.4 represents some natural transformation $\eta : I \rightarrow J$. The linear map $\eta(\tau)$ with respect to the fixed bases of I and J is

$$M[\text{St } \tau, \text{St } \tau] = M[\{\tau, \pi\}, \{\tau, \pi\}] = \begin{pmatrix} a_1 & a_2 \\ a_5 & a_6 \\ a_7 & a_8 \end{pmatrix}.$$

Example 3.10. A labeled matrix M representing a natural transformation η on a poset Π , a face poset of a simplicial complex on vertices $\{1, 2, 3, 4, 5, 6\}$ ordered by inclusion—for example $134 < 1234$. The meaning of η is revealed later in Example 4.23. Here we only show how matrices of linear maps $\eta(\sigma)$ overlap in a labeled matrix M .

$$\eta \begin{pmatrix} & 1234 & 1235 & 1245 & 1345 & 16 & 17 \\ 123 & \boxed{1} & \boxed{-1} & \boxed{0} & \boxed{0} & 0 & 0 \\ 124 & \boxed{1} & 0 & \boxed{-1} & \boxed{0} & 0 & 0 \\ 125 & 0 & \boxed{1} & \boxed{-1} & \boxed{0} & 0 & 0 \\ 134 & \boxed{1} & 0 & \boxed{0} & \boxed{-1} & 0 & 0 \\ 135 & 0 & \boxed{1} & 0 & \boxed{-1} & 0 & 0 \\ 145 & \boxed{0} & 0 & \boxed{1} & \boxed{-1} & 0 & 0 \\ 1 & 0 & 0 & 0 & -1 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & -1 & 1 \end{pmatrix} \quad \begin{array}{l} \boxed{\eta(12) = M[\text{St}(12), \text{St}(12)]} \\ \boxed{\eta(13) = M[\text{St}(13), \text{St}(13)]} \\ \boxed{\eta(14) = M[\text{St}(14), \text{St}(14)]} \\ \boxed{\eta(15) = M[\text{St}(15), \text{St}(15)]} \end{array}$$

Multiplying labeled matrices N and M makes sense whenever the column labels of N exactly match the row labels of M . In this case, $N \cdot M$ is the product of the unlabeled matrices with columns labeled as columns of M , and rows labeled as rows of N ; the operation corresponds to the composition of the corresponding natural transformations.

Lemma 3.11. Let I, J, K be injective sheaves, each with a fixed basis. If M represents $\eta : I \rightarrow J$ and N represents $\lambda : J \rightarrow K$ with respect to the fixed bases, then $N \cdot M$ represents $\lambda \circ \eta : I \rightarrow K$.

Proof. Note first that NM is a valid labeled matrix, i.e., it respects the non-zero elements condition: if $0 \neq NM[\sigma, \pi] = N[\sigma, \Pi] \cdot M[\Pi, \pi]$, then there is some τ such that $N[\sigma, \tau] \neq 0$ and $M[\tau, \pi] \neq 0$, which implies $\sigma \leq \tau \leq \pi$.

For any $\tau \in \Pi$, note that $N[\text{St } \tau, \Pi \setminus \text{St } \tau] = 0$ by the non-zero elements condition for N . Then $(NM)[\text{St } \tau, \text{St } \tau] = N[\text{St } \tau, \Pi] \cdot M[\Pi, \text{St } \tau] = N[\text{St } \tau, \text{St } \tau] \cdot M[\text{St } \tau, \text{St } \tau]$, which implies that $N \cdot M$ represents $\lambda \circ \eta$ with respect to the fixed bases on I and K . \square

Example 3.12. Consider M from Example 3.4, and N a labeled matrix representing a natural transformation $\lambda : J \rightarrow K \cong [\pi] \oplus [\sigma]$. The multiplication $N \cdot M$ is as follows:

$$\begin{matrix} & \pi & \sigma & \tau & \tau \\ \pi & \begin{pmatrix} b_1 & 0 & 0 & 0 \\ b_2 & b_3 & 0 & 0 \end{pmatrix} & \cdot & \begin{matrix} \pi & \pi & \sigma \\ \sigma & \begin{pmatrix} a_1 & a_2 & 0 \\ a_3 & a_4 & a_9 \\ \tau & \begin{pmatrix} a_5 & a_6 & 0 \\ a_7 & a_8 & 0 \end{pmatrix} \end{matrix} \end{matrix} = \begin{matrix} \pi & \pi & \sigma \\ \sigma & \begin{pmatrix} b_1 a_1 & b_1 a_2 & 0 \\ b_2 a_1 + b_3 a_3 & b_2 a_2 + b_3 a_4 & b_3 a_9 \end{pmatrix} \end{matrix} \end{matrix}$$

3.2. Change of bases for labeled matrix representations

Just as for matrices between vector spaces, changing the basis of domain or codomain of a natural transformation $\eta : I \rightarrow J$ corresponds to multiplication of the representing labeled matrix, M , by another labeled matrix from the right or the left side, respectively. Consider U, U' two different bases of I . Both can be decomposition-compatible with the same decomposition or with different decompositions—that is, the order of summands can differ. Either way, we can represent the identity $id : I \rightarrow I$ by a labeled matrix, R , with respect to those two bases U and U' . Then R is a change of basis labeled matrix. That is, if M and M' are labeled matrix representations of η with respect to U and U' , respectively, then $M' \cdot R = M$. Analogously, if we change bases in the codomain, J , it corresponds to matrix multiplication from the left side.

The above observation implies that choosing a different basis for the domain or codomain corresponds to column or row operations in the labeled matrix representations, respectively. Since R is itself a labeled matrix, $R[\sigma, \pi]$ is non-zero only if $\sigma \leq \pi$. This restricts the set of operations that are allowed.

Definition 3.13. For a Π -labeled matrix, the *allowed elementary column operations* are

- adding a multiple of a column labeled by σ to a column labeled by π for $\sigma \leq \pi$,
- multiplying a column by a non-zero constant,
- swapping two columns,

and the *allowed elementary row operations* are

- adding a multiple of a row labeled by π to a row labeled by σ for $\sigma \leq \pi$,
- multiplying a row by a non-zero constant,
- swapping two rows.

An allowed elementary row (column) operation corresponds to multiplication from the left (right) side by an *elementary labeled matrix*, which is a special case of an elementary matrix in the standard sense.

Example 3.14. Let Π be a poset with $\sigma < \pi$. Consider $[\pi] \xrightarrow{\eta^0} [\pi] \oplus [\sigma] \xrightarrow{\eta^1} [\sigma]$ with η^i represented by M^i . Change the basis for the middle sheaf so that the identity is represented by R . Then the new matrices $\hat{M}^0 = R \cdot M^0$, $\hat{M}^1 = M^1 \cdot R^{-1}$ represent η^0, η^1 with respect to the new basis, respectively. We performed one allowed elementary operation on each M^i .

$$\begin{matrix} R & \begin{matrix} \pi & \sigma \\ \pi & \begin{pmatrix} 1 & 0 \\ -1 & 1 \end{pmatrix} \end{matrix} & R^{-1} & \begin{matrix} \pi & \sigma \\ \pi & \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \end{matrix} & M^0 & \begin{matrix} \pi \\ \pi & \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{matrix} & M^1 & \begin{matrix} \pi & \sigma \\ \sigma & \begin{pmatrix} 1 & -1 \end{pmatrix} \end{matrix} & \hat{M}^0 & \begin{matrix} \pi \\ \pi & \begin{pmatrix} 1 \\ 0 \end{pmatrix} \end{matrix} & \hat{M}^1 & \begin{matrix} \pi & \sigma \\ \sigma & \begin{pmatrix} 0 & -1 \end{pmatrix} \end{matrix} \end{matrix}$$

Any change of basis matrix, R , is a composition of allowed elementary row operations or column operations. Indeed, we can reduce any such R to the identity matrix using allowed row or column operations. With either option, we first reduce the square “diagonal” blocks $R[\pi, \pi]$ to identity matrices—this is always

possible since we impose no restrictions on operations on those blocks, and $R[\pi, \pi]$ represents an isomorphism of vector spaces. Then we can clear the rest of the matrix, since $R[\sigma, \pi]$ can be non-zero only if $\sigma \leq \pi$, in which case we can clear this block by adding rows labeled by π or by adding columns labeled by σ . In particular, this also implies that the change of bases labeled matrices are invertible, and the inverse is, again, a labeled matrix.

3.3. Complexes of labeled matrices

One labeled matrix represents a single natural transformation between injective sheaves. A series of labeled matrices represents an injective complex.

Definition 3.15. A sequence $M^\bullet = (\dots, M^d, M^{d+1}, \dots)$ of labeled matrices is a *complex of labeled matrices* if

- columns of M^{d+1} have the same labeling as the rows of M^d (including order) for all d ,
- $M^{d+1} \cdot M^d = 0$ for all d .

The sequences we consider are finite, but formally extend indefinitely by empty matrices. Recall that an empty labeled matrix can be a 0×0 matrix, but also an $n \times 0$ matrix, representing a natural transformation from a trivial to a non-trivial sheaf, or a $0 \times n$ matrix.

Note that $M^{d+1} \cdot M^d = 0$ iff $M^{d+1}[\text{St } \tau, \text{St } \tau] \cdot M^d[\text{St } \tau, \text{St } \tau] = 0$ for all $\tau \in \Pi$. This is because any row in either matrix is labeled by some τ , and is non-zero only at columns labeled by $\text{St } \tau$.

Example 3.16. A complex of labeled matrices on $\Delta_3^{(2)}$, the face poset of the 2-skeleton of a tetrahedron 1234 (for example $14 \leq 124$, as $\{1, 4\} \subseteq \{1, 2, 4\}$):

$$\begin{array}{c} M^{-1} \\ \begin{pmatrix} 234 \\ 134 \\ 124 \\ 123 \\ 12 \end{pmatrix} \end{array} \begin{array}{c} M^0 \\ \begin{pmatrix} 234 & 134 & 124 & 123 \\ -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \\ 14 & 0 & -1 & 1 & 0 \\ 13 & 0 & -1 & 0 & 1 \\ 12 & 0 & 0 & -1 & 1 \end{pmatrix} \end{array} \begin{array}{c} M^1 \\ \begin{pmatrix} 34 & 24 & 23 & 14 & 13 & 12 \\ 1 & -1 & 0 & 1 & 0 & 0 \\ 1 & 0 & -1 & 0 & 1 & 0 \\ 0 & 1 & -1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & -1 & 1 \end{pmatrix} \end{array} M^2 \begin{pmatrix} 4 & 3 & 2 & 1 \end{pmatrix}$$

Definition 3.17. An injective complex $(I^\bullet, \eta^\bullet)$ is *represented* by a complex of labeled matrices M^\bullet , if the following hold:

- I^d is isomorphic to the sum of indecomposables given by the labeling of the columns of η^d for every d ,
- M^{d-1} and M^d are representations of the natural transformations η^{d-1} and η^d with respect to the same decomposition-compatible basis of I^d , for every d .

Note that if 0 is the smallest degree with a non-zero entry, I^0 , then we formally require M^\bullet to start with a $n \times 0$ matrix M^{-1} as shown in Example 3.16. In most other examples we do not show the empty matrices at the start or the end—their labeling is determined by the columns of the first or rows of the last non-empty matrix, respectively.

Lemma 3.18. Changing the basis of I^d in a complex corresponds to performing column operations on M^d , and at the same time corresponding inverse row operations on M^{d-1} . If R is the labeled matrix representing the change of basis of I^d , then the new matrices with codomain and domain I^d are $R \cdot M^{d-1}$ and $M^d \cdot R^{-1}$,

respectively. For example, if we add the i -th column to the j -th in M^d , we need to subtract the j -th row from the i -th in M^{d-1} .

Remark. To simplify notation, in the rest of the paper, we use the same symbol both for a natural transformation and for a labeled matrix representation of it. We always assume that a particular decomposition-compatible basis of each injective sheaf is fixed. When we say “ $(I^\bullet, \eta^\bullet)$ is a complex represented by labeled matrices”, we mean that matrices denoted by η^d form a complex of labeled matrices that represents the complex of injective sheaves $(I^\bullet, \eta^\bullet)$.

4. Derived categories

We often use different notions and degrees of homology separately when studying geometrical objects, losing important relationships between the different choices. Derived categories are a tool to take a step back while still abstracting to “just homological information”. The formal definition is technically involved, but it can be understood in simpler terms for our setting of finite dimensional bounded cochain complexes over finite posets—in this section we define essentially the skeleton of the derived category. The objects are minimal injective complexes, and the morphisms are a quotient of morphisms of complexes.

We recall definitions necessary to describe the quotient of morphisms in 4.1. In 4.2 we show why we can restrict our attention to minimal injective complexes and give the definition of (a skeleton of) the derived category. In 4.3 we give an explicit algorithmic way to go from any injective complex to the minimal one. Finally, in 4.4 we recall how to think of a single sheaf as a member in the derived category of minimal injective complexes, and describe a concrete algorithm to compute the minimal injective resolution of a sheaf. See also Appendix C.1 for explicit description of the vector space of morphisms in the derived category with a system of linear equations.

4.1. Morphisms, null-homotopy, quasi-isomorphism

Definition 4.1. A morphism $\alpha^\bullet : F^\bullet \rightarrow G^\bullet$ between two complexes, $(F^\bullet, \eta^\bullet)$ and $(G^\bullet, \delta^\bullet)$, is *null-homotopic* if there exist morphisms of sheaves $h^d : F^d \rightarrow G^{d-1}$:

$$\begin{array}{ccccccc}
 \cdots & \xrightarrow{\eta^{d-2}} & F^{d-1} & \xrightarrow{\eta^{d-1}} & F^d & \xrightarrow{\eta^d} & F^{d+1} & \xrightarrow{\eta^{d+1}} & \cdots \\
 & \swarrow h^{d-1} & \downarrow \alpha^{d-1} & \swarrow h^d & \downarrow \alpha^d & \swarrow h^{d+1} & \downarrow \alpha^{d+1} & \swarrow h^{d+2} & \\
 \cdots & \xrightarrow{\delta^{d-2}} & G^{d-1} & \xrightarrow{\delta^{d-1}} & G^d & \xrightarrow{\delta^d} & G^{d+1} & \xrightarrow{\delta^{d+1}} & \cdots
 \end{array}$$

such that $\alpha^d = \delta^{d-1}h^d + h^{d+1}\eta^d$, for each d .

To get a better intuition about null-homotopic maps, we state an alternate definition. We need to first introduce the mapping cone of a morphism of complexes, which will also be useful in later sections.

Definition 4.2. The *mapping cone*, of a morphism $\alpha^\bullet : F^\bullet \rightarrow G^\bullet$ between two complexes, $(F^\bullet, \eta^\bullet)$ and $(G^\bullet, \delta^\bullet)$, is a complex $C(\alpha^\bullet) = (C^\bullet, \gamma^\bullet)$ given by

$$C^d := F^{d+1} \oplus G^d, \quad \gamma^d = \begin{pmatrix} -\eta^{d+1} & 0 \\ \alpha^{d+1} & \delta^d \end{pmatrix}.$$

$$\begin{array}{ccccccc}
\cdots & \longrightarrow & F^{d-1} & \xrightarrow{\eta^{d-1}} & F^d & \xrightarrow{\eta^d} & F^{d+1} & \xrightarrow{\eta^{d+1}} & F^{d+2} & \longrightarrow & \cdots \\
& & \downarrow \alpha^{d-1} & & \downarrow \alpha^d & & \downarrow \alpha^{d+1} & & \downarrow \alpha^{d+2} & & \\
\cdots & \longrightarrow & G^{d-1} & \xrightarrow{\delta^{d-1}} & G^d & \xrightarrow{\delta^d} & G^{d+1} & \xrightarrow{\delta^{d+1}} & G^{d+2} & \longrightarrow & \cdots
\end{array}$$

$\underbrace{\quad\quad\quad}_{C^{d-1}} \quad \underbrace{\quad\quad\quad}_{C^d} \quad \underbrace{\quad\quad\quad}_{C^{d+1}}$

The following is a useful fact which we need in later sections.

Lemma 4.3 ([33, Def. 1.5.7 and the following discussion]). *A morphism $\alpha^\bullet : F^\bullet \rightarrow G^\bullet$ between two complexes is a quasi-isomorphism if and only if its mapping cone, $C(\alpha^\bullet)$, is exact.*

The mapping cone construction provides a condition equivalent to Definition 4.1. The proof is straightforward if we write $\tilde{\alpha}^\bullet$ in the statement as $(\alpha^\bullet, h^\bullet)$.

Proposition 4.4. *A morphism $\alpha^\bullet : F^\bullet \rightarrow G^\bullet$ between two complexes, $(F^\bullet, \eta^\bullet)$ and $(G^\bullet, \delta^\bullet)$, is null-homotopic if and only if α^\bullet factors through the mapping cone of the identity map on the complex $(G^{\bullet-1}, -\delta^{\bullet-1})$ (with indices shifted by -1):*

$$\begin{array}{ccccc}
& & \alpha^\bullet & & \\
& \nearrow & & \searrow & \\
F^\bullet & \xrightarrow{\tilde{\alpha}^\bullet} & C(\text{id}_{(G^{\bullet-1}, -\delta^{\bullet-1})}) & \xrightarrow{\text{proj}} & G^\bullet
\end{array}$$

Note that, by Lemma 4.3, the mapping cone of an identity map is an exact complex. Hence, the above proposition immediately implies that all null-homotopic maps induce trivial morphisms on cohomology sheaves. This gives some intuition for the following definition, which, on a high level, states that null-homotopic maps should be treated as zero maps for the purposes of homological algebra. Recall the notion of cohomology sheaf from Definition 2.11.

Definition 4.5. A morphism $\alpha^\bullet : F^\bullet \rightarrow G^\bullet$ of complexes of sheaves is a *quasi-isomorphism* if the induced map

$$H^d(\alpha^\bullet) : H^d(F^\bullet) \rightarrow H^d(G^\bullet)$$

between cohomology sheaves is an isomorphism for each $n \in \mathbb{Z}$.

Definition 4.6. An *injective resolution* of a complex of sheaves F^\bullet is a quasi-isomorphism

$$\alpha^\bullet : F^\bullet \rightarrow I^\bullet$$

from F^\bullet to a complex I^\bullet of injective sheaves.

Each complex of sheaves can be identified with a complex of injective sheaves up to quasi-isomorphism—see, e.g., [49, Thm 1.3.4]. In the next section we show that it can be identified with a minimal injective complex, unique up to isomorphism of complexes, and in the following section we give an explicit algorithm for minimizing an injective complex.

4.2. Minimal complexes of injective sheaves

We will now define minimal complexes of injective sheaves, and show that every complex of injective sheaves has a minimal complex of injective sheaves quasi-isomorphic to it, and this minimal complex is unique up to isomorphism of complexes.

Definition 4.7. A complex of injective sheaves I^\bullet is *minimal* if it minimizes the number of indecomposables among all quasi-isomorphic complexes. That is, if for every complex of injective sheaves J^\bullet which is quasi-isomorphic to I^\bullet and for every d , the number of indecomposable summands of I^d is less than or equal to the number of indecomposable summands of J^d .

To formulate one of the equivalent conditions of minimality, we define maximal vectors. In the module perspective, they are exactly the elements of the socle.

Definition 4.8. A vector $s \in F(\pi)$ is *maximal* if $F(\pi \leq \tau)(s) = 0$ for all $\tau > \pi$. We denote by $M_F(\pi)$ the subspace of maximal vectors in $F(\pi)$, i.e.,

$$M_F(\pi) := \bigcap_{\pi < \sigma} \ker F(\pi \leq \sigma).$$

Note that it is sufficient to take only the intersection of $\ker F(\pi \leq \sigma)$ for all $\pi <_1 \sigma$.

Theorem 4.9. Let $(I^\bullet, \eta^\bullet)$ be a complex¹ of injective sheaves on Π . The following are equivalent:

1. I^\bullet is minimal.
2. For any quasi-isomorphic complex $(J^\bullet, \lambda^\bullet)$ of injective sheaves, there exists a morphism of complexes $r^\bullet : J^\bullet \rightarrow I^\bullet$ such that r^d is surjective for each d .
3. For any quasi-isomorphic complex $(J^\bullet, \lambda^\bullet)$ of injective sheaves, there exists a morphism of complexes $q^\bullet : I^\bullet \rightarrow J^\bullet$ such that q^d is injective for each d .
4. For each d , $\pi \in \Pi$, and maximal vector $s \in I^d(\pi)$, $\eta^d(\pi)(s) = 0$.

Moreover, for a given complex of sheaves there exists a unique (up to isomorphism of complexes) injective resolution by a minimal complex of injective sheaves.

Proof. The implications $2 \Rightarrow 1$ and $3 \Rightarrow 1$ are immediate as both imply that $\dim I^d(\pi) \leq \dim J^d(\pi)$ for every complex J^\bullet quasi-isomorphic to I^\bullet , and every $d \in \mathbb{Z}$, $\pi \in \Pi$.

We show implication $1 \Rightarrow 4$, proof of which also yields an argument for existence of minimal injective complexes, and then we show $4 \Rightarrow 3$ and $4 \Rightarrow 2$, which will also yield the uniqueness.

$1 \Rightarrow 4$: We show that if 4 does not hold, then 1 does not hold either. If s is a maximal vector in $I^d(\pi)$ which does not vanish when we apply η^d , we can define $\hat{I}^d := I^d/[\pi]_s$, where $[\pi]_s$ is the indecomposable injective such that $s \in [\pi]_s(\pi)$, and similarly $\hat{I}^{d+1} := I^{d+1}/[\pi]_{\eta^d(\pi)(s)}$. By Lemma 2.8, both are injective sheaves. Denoting by q^d the quotient maps and setting $r = \eta^d(\pi)(s)$, we get:

¹ Recall that the complexes considered in this paper are bounded.

$$\begin{array}{ccccccc}
& & 0 & & 0 & & \\
& & \downarrow & & \downarrow & & \\
& 0 & \longrightarrow & [\pi]_s & \xrightarrow{\eta^d|_{[\pi]_s}} & [\pi]_r & \longrightarrow 0 \\
& \downarrow & & \downarrow & & \downarrow & \\
\cdots & \longrightarrow & I^{d-1} & \xrightarrow{\eta^{d-1}} & I^d & \xrightarrow{\eta^d} & I^{d+1} & \xrightarrow{\eta^{d+1}} & I^{d+2} & \longrightarrow \cdots \\
& \downarrow q^{d-1}=\text{id} & & \downarrow q^d & & \downarrow q^{d+1} & & \downarrow q^{d+2}=\text{id} & \\
\cdots & \longrightarrow & I^{d-1} & \xrightarrow{\hat{\eta}^{d-1}} & \hat{I}^d & \xrightarrow{\hat{\eta}^d} & \hat{I}^{d+1} & \xrightarrow{\hat{\eta}^{d+1}} & I^{d+2} & \longrightarrow \cdots \\
& \downarrow & & \downarrow & & \downarrow & & \downarrow & \\
& 0 & & 0 & & 0 & & 0 &
\end{array}$$

Since the top row has trivial cohomology, and a short exact sequence of complexes induces a long exact sequence in cohomology, we see that

$$0 \rightarrow H^j(I^\bullet) \xrightarrow{q^j} H^j(\hat{I}^\bullet) \rightarrow 0$$

is exact, that is, q^\bullet induces an isomorphism on cohomology. We have found a complex \hat{I}^\bullet quasi-isomorphic to I^\bullet with fewer indecomposable summands in places d and $d+1$, so I^\bullet is not minimal.

By iterating this process, we can produce, from any bounded complex of injective sheaves J^\bullet , a minimal complex of injective sheaves I^\bullet , which proves the existence of minimal injective complexes.

4 \Rightarrow 2, 3: Since I^\bullet , J^\bullet are quasi-isomorphic, there is a quasi-isomorphism $q^\bullet : I^\bullet \rightarrow J^\bullet$. Because both complexes consist of injective sheaves, q^\bullet is a homotopy equivalence. Therefore, let $q^\bullet : I^\bullet \rightarrow J^\bullet$ and $r^\bullet : J^\bullet \rightarrow I^\bullet$ form a homotopy equivalence. We show that if I^\bullet satisfies the condition 4, then q^d is injective and r^d is surjective for each d .

We inductively prove that $r^d \circ q^d$ is injective and that $r^d \circ q^d(\ker \eta^d) = \ker \eta^d$. As the base case, choose any d such that $I^j = 0$ for each $j \leq d$, where the induction hypothesis holds trivially. We show the induction step.

By assumption, $r^d \circ q^d = \text{id} + h^{d+1} \circ \eta^d + \eta^{d-1} \circ h^d$ for a collection of morphisms

$$h^d : I^d \rightarrow I^{d-1}.$$

We begin by showing that the restriction of $r^d \circ q^d$ to $\ker \eta^d$ is injective.

$$\begin{array}{ccccccc}
& & \overset{h^{d-1}}{\curvearrowright} & & \overset{h^d}{\curvearrowright} & & \overset{h^{d+1}}{\curvearrowright} & & \overset{h^{d+2}}{\curvearrowright} & \\
\cdots & \xrightarrow{\eta^{d-2}} & I^{d-1} & \xrightarrow{\eta^{d-1}} & I^d & \xrightarrow{\eta^d} & I^{d+1} & \xrightarrow{\eta^{d+1}} & \cdots \\
& & \downarrow q^{d-1} \uparrow r^{d-1} & & \downarrow q^d \uparrow r^d & & \downarrow q^{d+1} \uparrow r^{d+1} & & \\
\cdots & \xrightarrow{\lambda^{d-2}} & J^{d-1} & \xrightarrow{\lambda^{d-1}} & J^d & \xrightarrow{\lambda^d} & J^{d+1} & \xrightarrow{\lambda^{d+1}} & \cdots
\end{array}$$

Assume $\pi \in \Pi$, $x \in \ker \eta^d(\pi)$, and $x \in \ker(r^d \circ q^d)(\pi)$. Then

$$0 = (r^d \circ q^d)(\pi)(x) = x + 0 + (\eta^{d-1} \circ h^d)(\pi)(x).$$

That is, $x = (\eta^{d-1} \circ h^d)(\pi)(-x)$. Set $y := h^d(\pi)(-x)$. Because q^\bullet is a morphism of complexes, $\lambda^{d-1} \circ q^{d-1} = q^d \circ \eta^{d-1}$. Because $\eta^{d-1}(\pi)(y) = x$, we have that $(\lambda^{d-1} \circ q^{d-1})(\pi)(y) = q^d(\pi)(x)$. Because $(r^d \circ q^d)(\pi)(x) = 0$, we have that

$$(r^d \circ \lambda^{d-1} \circ q^{d-1})(\pi)(y) = 0.$$

Because r^\bullet is a morphism of complexes, $r^d \circ \lambda^{d-1} = \eta^{d-1} \circ r^{d-1}$. Therefore,

$$(\eta^{d-1} \circ r^{d-1} \circ q^{d-1})(\pi)(y) = 0.$$

We have therefore shown that $(r^{d-1} \circ q^{d-1})(\pi)(y) \in \ker \eta^{d-1}(\pi)$.

By the induction assumption, we have that

$$(r^{d-1} \circ q^{d-1})(\ker \eta^{d-1}) = \ker \eta^{d-1}.$$

Therefore, there exists $z \in \ker \eta^{d-1}(\pi)$ such that $(r^{d-1} \circ q^{d-1})(\pi)(z) = (r^{d-1} \circ q^{d-1})(\pi)(y)$. By the induction assumption that $r^{d-1} \circ q^{d-1}$ is injective, we have that $z = y$. Therefore, $y \in \ker \eta^{d-1}(\pi)$, and $x = \eta^{d-1}(\pi)(y) = 0$. Altogether we showed that the restriction of $r^d \circ q^d$ to $\ker \eta^d$ is injective.

For each non-zero vector $u \in I^d(\tau)$, there exists $(\tau \leq \pi) \in \Pi$ and a non-zero maximal vector $s \in I^d(\pi)$ such that $I^d(\tau \leq \pi)(u) = s$. By the condition 4 for I^\bullet , all maximal vectors in $I^d(\pi)$ are elements of $\ker \eta^d(\pi)$. Because the restriction of $r^d \circ q^d$ to $\ker \eta^d$ is injective, $(r^d \circ q^d)(\pi)(s) \neq 0$. By the commutativity assumed in the definition of a natural transformation (Definition 2.4),

$$I^{d+1}(\tau \leq \pi)((r^d \circ q^d)(\tau)(u)) = (r^d \circ q^d)(\pi)(I^d(\tau \leq \pi)(u)) = (r^d \circ q^d)(\pi)(s) \neq 0.$$

Therefore, $(r^d \circ q^d)(\tau)(u) \neq 0$. That is, $r^d \circ q^d$ is injective. Moreover, because r^\bullet and q^\bullet are morphisms of complexes, $(r^d \circ q^d)(\ker \eta^d) \subseteq \ker \eta^d$. Because the restriction of $r^d \circ q^d$ to $\ker \eta^d$ is injective, we have the desired equality: $(r^d \circ q^d)(\ker \eta^d) = \ker \eta^d$.

Since $r^d \circ q^d(\pi)$, for $\pi \in \Pi$, is an endomorphism of finite-dimensional vector spaces, it is injective iff it is surjective. Therefore, q^d is injective and r^d is surjective, as claimed.

Finally, if both complexes satisfy condition 4, then by a symmetric argument, also r^d is injective, and q^d is surjective. In particular, q^d is an isomorphism of sheaves, i.e., q^\bullet is an isomorphism of complexes. \square

Definition 4.10. The *bounded derived category of sheaves* on Π , $D^b(\Pi)$, consists of a set of objects

$$\text{Ob} D^b(\Pi) = \{(I^\bullet, \eta^\bullet) \text{ a minimal complex of injective sheaves}\}$$

and a vector space of morphisms between any two objects

$$\text{Hom}_{D^b(\Pi)}(I^\bullet, J^\bullet)$$

which is defined to be the quotient space of the (finite-dimensional) vector space of all morphisms (of complexes) from I^\bullet to J^\bullet modulo the subspace of null-homotopic morphisms.

Remark. It follows from Theorem 4.9 that the above definition of the bounded derived category of sheaves is in fact a skeleton of the category which is traditionally referred to as the bounded derived category of sheaves [10,33]—see Appendix B for more details. We choose this definition in order to streamline and simplify the computational perspective of derived categories considered below. Any concepts we define for the elements of $D^b(\Pi)$, in the sense above, apply for any (bounded) complex of sheaves on Π . An element $I^\bullet \in D^b(\Pi)$ is a chosen representative of a class of quasi-isomorphic complexes—the unique minimal injective member of it.

4.3. Recognizing minimality and minimizing an injective complex

With labeled matrix representations we can easily recognize whether a complex of injective sheaves is minimal, and obtain a minimal complex in case it is not.

Proposition 4.11. *A complex of injective sheaves $(I^\bullet, \eta^\bullet)$ represented by labeled matrices is minimal iff $\eta^d[\pi, \pi]$ is a zero (or empty) matrix for each $\pi \in \Pi$ and degree d .*

Proof. The claim is a reformulation of the condition (4) in Theorem 4.9: the complex is minimal iff all maximal vectors are sent to zero by the complex maps η^d . If we write $I^d = \bigoplus_{\pi \in \Pi} [\pi]^{p_\pi}$, then the maximal vectors in $I^d(\pi)$ are exactly vectors in $[\pi]^{p_\pi}(\pi)$ (as a subspace of $I^d[\pi]$). In other words, written with respect to the bases fixed for the complex of labeled matrices, a vector $s \in I^d(\pi)$ is maximal iff it is zero at the coordinates labeled by $\text{St } \pi \setminus \{\pi\}$. Then $\eta^d(\pi)(s) = \eta^d[\text{St } \pi, \text{St } \pi] \cdot s$ is zero for all maximal $s \in I^d(\pi)$ iff $\eta^d[\pi, \pi] = 0$. \square

The peeling procedure If $(I^\bullet, \eta^\bullet)$ is not minimal, we perform allowed row and column operations so that the non-zero entries in $\eta^d[\pi, \pi]$ are on the main diagonal, and, for each, the corresponding row and column in the whole labeled matrix η^d only contain that one single non-zero element. This row and column then represents the sequence $0 \rightarrow [\pi]_s \rightarrow [\pi]_{\eta^d(\pi)(s)} \rightarrow 0$ from the proof of Theorem 4.9, and can be safely removed from the complex—unlike in the proof of the theorem, here we see that the sequence actually splits off the complex as a direct summand. We repeat this process until $\eta^d[\pi, \pi] = 0$ for each $\pi \in \Pi$ and degree d , which then verifies minimality by Proposition 4.11. The procedure is described in Algorithm 1.

Algorithm 1 Computing minimal complex: Peeling.

Input: η^0, \dots, η^n labeled matrices representing an injective complex $(I^\bullet, \eta^\bullet)$ on Π

Output: $\hat{\eta}^0, \dots, \hat{\eta}^n$ labeled matrices representing the minimal injective complex quasi-isomorphic to $(I^\bullet, \eta^\bullet)$

1: **while** $\exists d \in \{0, \dots, n\} \exists \pi \in \Pi$ such that $\eta^d[\pi, \pi]$ is a non-zero matrix **do**
 2: find labeled matrices L, R such that:

- $(L \cdot \eta^d \cdot R)[\pi, \pi]$ is a diagonal matrix with entries in $\{0, 1\}$
- every π -labeled column of $L \cdot \eta^d \cdot R$ is either zero on π or on $\Pi \setminus \pi$
- every π -labeled row of $L \cdot \eta^d \cdot R$ is either zero on π or on $\Pi \setminus \pi$

3: $\eta^d \leftarrow L \cdot \eta^d \cdot R, \quad \eta^{d-1} \leftarrow R^{-1} \cdot \eta^{d-1}, \quad \eta^{d+1} \leftarrow \eta^{d+1} \cdot L^{-1}$ \triangleright extended Gaussian elimination on $\eta^d[\pi, \pi]$ and further reduction on η^d
 4: **while** $\exists i$ -th row, j -th column in η^d labeled by π such that $\eta^d[i, j] = 1$ **do** \triangleright Lemma 3.18
 5: delete the i -th row and the j -th column from η^d
 6: delete the j -th row from η^{d-1}
 7: delete the i -th column from η^{d+1}
 8: **return** η^0, \dots, η^n

Proposition 4.12. *Given an injective complex $(I^\bullet, \eta^\bullet)$ represented by labeled matrices, Algorithm 1 outputs a labeled matrix representation of a quasi-isomorphic minimal injective complex.*

Proof. We first need to argue that it is possible to find the desired labeled matrices L, R whenever $\eta^d[\pi, \pi] \neq 0$. All elementary operations are allowed for rows and columns labeled by the same element π . Therefore, following standard Gaussian elimination and further reduction, we reduce $\eta^d[\pi, \pi]$ so that it is a diagonal matrix with entries in $\{0, 1\}$. Then we use the 1 on the diagonal to remove any non-zeros in the corresponding row and column in the whole matrix η^d —indeed, there can be a non-zero exactly when removing it by subtracting our fixed column or row, respectively, is allowed (see Definition 3.13). The operations are

performed as a change of basis of the whole complex, as described in Lemma 3.18. The newly produced complex is (quasi-)isomorphic to the original.

For the new matrices η^\bullet , we argue that the rows and columns with non-zeros in $\eta^d[\pi, \pi]$ can be peeled off the complex. Let (i, j) be an index of η^d with entry 1 such that both the i -th row and the j -th column are indexed by π . By the above construction, both have 0 in all entries other than their intersection (i, j) . Then the complex condition $\eta^d \cdot \eta^{d-1} = 0 = \eta^{d+1} \cdot \eta^d$ implies that the corresponding j -th row in η^{d-1} and i -th column in η^{d+1} are both zero. Therefore, we have identified an exact sequence $0 \rightarrow [\pi] \rightarrow [\pi] \rightarrow 0$ that splits from the complex, and removing it produces a quasi-isomorphic complex, as argued in the proof of Theorem 4.9. The deletions from the labeled matrices described in the algorithm correspond to this removal.

Finally, the algorithm clearly terminates, because each iteration of the main loop removes some non-zero number of the finitely many rows and columns of the complex of labeled matrices. The outputted complex of labeled matrices is minimal, since it satisfies the condition of Proposition 4.11. \square

Example 4.13. Let Π be such that $\sigma < \pi < \mu < \lambda$, and consider an injective complex

$$0 \rightarrow [\lambda] \rightarrow [\mu] \oplus [\pi]^2 \rightarrow [\pi]^3 \oplus [\sigma] \rightarrow [\sigma] \rightarrow 0$$

given by labeled matrices η^d . It is not minimal, because $\eta^1[\pi, \pi] \neq 0$. We minimize it using Algorithm 1. First we reduce the highlighted block $\eta^1[\pi, \pi]$ to a matrix with 1 only on the main diagonal and 0 everywhere else. Row operations suffice, and for each we perform the inverse column operation on η^2 . We get matrices ζ^d . The next step is to clear out the rest of the rows and columns that contain a 1 within the $\zeta^1[\pi, \pi]$ block. We use the diagonal elements to clear out the highlighted non-zeros with two row and two column allowed operations. Finally, the resulting matrices δ^d contain two $0 \rightarrow [\pi] \rightarrow [\pi] \rightarrow 0$ sequences that split off. We can delete all the highlighted rows and columns, obtaining matrices γ^d representing a quasi-isomorphic complex

$$0 \rightarrow [\lambda] \rightarrow [\mu] \rightarrow [\pi] \oplus [\sigma] \rightarrow [\sigma] \rightarrow 0,$$

which is already minimal, because all $\delta^d[x, x]$ blocks are zero.

$$\begin{array}{ccc} \eta^0 \begin{pmatrix} \lambda \\ \mu \\ \pi \\ \pi \end{pmatrix} & \eta^1 \begin{pmatrix} \mu & \pi & \pi \\ \pi \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ -1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix} \end{pmatrix} & \eta^2 \begin{pmatrix} \pi & \pi & \pi & \sigma \\ \sigma & (-1 & 1 & 1 & 0) \end{pmatrix} \\ \zeta^0 \begin{pmatrix} \lambda \\ \mu \\ \pi \\ \pi \end{pmatrix} & \zeta^1 \begin{pmatrix} \mu & \pi & \pi \\ \pi \begin{pmatrix} -1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 1 & 1 \end{pmatrix} \end{pmatrix} & \zeta^2 \begin{pmatrix} \pi & \pi & \pi & \sigma \\ \sigma & (0 & 0 & 1 & 0) \end{pmatrix} \\ \delta^0 \begin{pmatrix} \lambda \\ \mu \\ \pi \\ \pi \end{pmatrix} & \delta^1 \begin{pmatrix} \mu & \pi & \pi \\ \pi \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \end{pmatrix} & \delta^2 \begin{pmatrix} \pi & \pi & \pi & \sigma \\ \sigma & (0 & 0 & 1 & 0) \end{pmatrix} \\ \gamma^0 \begin{pmatrix} \lambda \\ \mu \\ \pi \\ \pi \end{pmatrix} & \gamma^1 \begin{pmatrix} \mu \\ \pi \\ \sigma \end{pmatrix} & \gamma^2 \begin{pmatrix} \pi & \sigma \\ 1 & 0 \end{pmatrix} \end{array}$$

As noted above the proposition, this explicit peeling construction also serves as a proof that the minimal injective resolution splits off any injective complex:

Corollary 4.14. Let J^\bullet be any complex of injective sheaves, and I^\bullet its minimal injective resolution. Then there exists an exact complex of injective sheaves E^\bullet , and an isomorphism of complexes $J^\bullet \cong I^\bullet \oplus E^\bullet$. In other words, the short exact sequence $0 \rightarrow I^\bullet \xrightarrow{q} J^\bullet \rightarrow E^\bullet \rightarrow 0$ splits (in the category of complexes of sheaves).

Apart from the peeling, we also provide an alternative computation to minimize a complex in Section 5.4.3—our computation of $\mathcal{I} \operatorname{id}^* I^\bullet$ (defined there) yields a minimal injective complex quasi-isomorphic to I^\bullet . This can be convenient for implementation, as it requires no column operations.

4.4. Single sheaf in derived category: computing injective resolutions

A single sheaf F can be viewed as a complex $\cdots \rightarrow 0 \rightarrow F \rightarrow 0 \rightarrow \cdots$ with F in degree 0. As for any complex, there exists a quasi-isomorphic complex of injective sheaves, I^\bullet . Replacing F by I^\bullet allows us to study it within the framework of derived categories. The mapping cone (Definition 4.2) of any such quasi-isomorphism is called *injective resolution* of F —it is an exact complex by Lemma 4.3.

Definition 4.15. An *injective resolution* of a sheaf F is an exact sequence

$$0 \rightarrow F \xrightarrow{\alpha} I^0 \xrightarrow{\eta^0} I^1 \xrightarrow{\eta^1} I^2 \xrightarrow{\eta^2} \cdots$$

where I^j is an injective sheaf for each j . We denote by I^\bullet the complex

$$\cdots \rightarrow 0 \rightarrow I^0 \xrightarrow{\eta^0} I^1 \xrightarrow{\eta^1} \cdots.$$

Slightly abusing the naming, we later also say that I^\bullet is an injective resolution of F .

Definition 4.16. An inclusion $F \hookrightarrow I^0$ into an injective sheaf is called an *injective hull* of F .

Any such inclusion can be extended to an injective resolution, and we give an explicit algorithm to compute such extension below.

By Theorem 4.9, there exists a unique minimal injective resolution. It is straightforward to bound its length in terms of the height of the poset Π . First, note that the condition (4) of the theorem can be reformulated in terms of images, since injective resolutions are exact.

Proposition 4.17. An injective resolution $0 \rightarrow F \xrightarrow{\alpha} I^0 \xrightarrow{\eta^0} I^1 \xrightarrow{\eta^1} I^2 \xrightarrow{\eta^2} \cdots$ is minimal iff the following holds:

- Each maximal vector $s \in I^d(\pi)$ is in $\operatorname{im} \eta^{d-1}(\pi)$, for each $d > 0$ and $\pi \in \Pi$. Each maximal vector $s \in I^0(\pi)$ is in $\operatorname{im} \alpha(\pi)$ for each $\pi \in \Pi$.

Corollary 4.18. The minimal injective resolution of a sheaf F on a finite poset Π of height h consists of at most $h + 1$ non-zero injective sheaves.

Proof. The length of the longest chain of non-zero vector spaces in I^d is strictly decreasing in d in the minimal injective resolution. This is implied by Proposition 4.17 and Theorem 4.9 (4): If $I^d(\pi) = 0$, then the former implies that there are no maximal vectors in $I^{d+1}(\pi)$. Therefore, if $I^d(\tau) = 0$ for all $\tau \geq \pi$, then also $I^{d+1}(\tau) = 0$ for all $\tau \geq \pi$. Moreover, if $I^d(\pi) \neq 0$ and $I^d(\tau) = 0$ for all $\tau > \pi$, then all vectors in $I^d(\pi)$ are maximal, and by property (4) and the argument above, $I^j(\pi) = 0$ for all $j > d$. \square

We now turn our attention to the computation of the minimal injective resolution. For a generic sheaf, F , one needs to first compute an injective hull, and then continue the sequence to ensure exactness. There are well known algorithms for both steps found in literature on module theory and representation theory. We present our own version of it for two reasons. Firstly, we formulate the algorithms purely in the language of

linear algebra so that they can be understood without prior knowledge of module or sheaf theory. Second, we generalize the main step of the construction to novel settings in Section 5.

We discuss the construction of an injective hull for a generic sheaf in Appendix C.3. We want to point out that it is feasible to construct it using only basic linear algebra, but one of the strengths of our approach is that we can study interesting geometrical objects without ever explicitly working with a generic non-injective sheaf.

Minimal injective resolution from injective hull We describe how to inductively compute the minimal injective resolution, given the minimal injective hull. That is, given the prefix of the resolution represented by labeled matrices $\eta^0, \dots, \eta^{d-1}$, the goal is to construct its continuation η^d . We need to construct η^d such that

$$I^{d-1} \xrightarrow{\eta^{d-1}} I^d \xrightarrow{\eta^d} I^{d+1}$$

is exact at I^d , and moreover this extension of the complex is minimal in the sense that all maximal vectors in I^{d+1} are in the image of η^d . We construct η^d again by induction, this time over $\pi \in \Pi$ in non-increasing order.

The procedure will only need access to images of $\eta^{d-1}(\pi)$ for $\pi \in \Pi$. In particular, it can also be used to construct η^0 if images of $\alpha(\pi)$ are given—this can be done by running Algorithm 6 in the case of a generic sheaf (Appendix C.3), but it can be done more efficiently in the case of the constant sheaf, as we describe later in Section 4.5.

Recall that we denote labeled matrices by the same symbols as the maps they represent; the linear map $\eta^d(\pi)$ is described by the submatrix $\eta^d(\pi) = \eta^d[\text{St } \pi, \text{St } \pi]$

Procedure MAKEEXACT The construction consists of repeated call of a core procedure MAKEEXACT, described as Algorithm 2. On the input, it takes η^{d-1} , η^d and $\pi \in \Pi$, where η^{d-1} , η^d are two labeled matrices representing a complex

$$I^{d-1} \xrightarrow{\eta^{d-1}} I^d \xrightarrow{\eta^d} I^{d+1},$$

On the output, it returns a labeled matrix representing a map $\hat{\eta}^d : I^d \rightarrow \hat{I}^{d+1}$ such that $\text{im } \eta^{d-1}(\pi) = \ker \hat{\eta}^d(\pi)$, and, importantly, $\hat{\eta}^d(\tau) = \eta^d(\tau)$ for all $\tau \not\leq \pi$.

Algorithm 2 MAKEEXACT(η^{d-1}, η^d, π).

Input: η^{d-1}, η^d labeled matrices over a poset Π , s.t. $\eta^d \cdot \eta^{d-1} = 0$, and $\pi \in \Pi$

Output: $\hat{\eta}^d$ a labeled matrix

1: $B \leftarrow$ basis of $(\text{im}(\eta^{d-1}(\pi)))^\perp$

2: **for all** $b \in B$ **do**

3: **if** b is linearly independent from the rows of $\eta^d(\pi)$ **then**

4: add b as a row of η^d , labeled by π

▷ with zeros at $\Pi \setminus \text{St } \pi$ entries

5: **return** η^d

The goal is to make $I^\bullet(\pi)$ exact in degree d . We achieve this by adding new linearly independent rows to $\eta^d(\pi)$ from the space

$$(\text{im } \eta^{d-1}(\pi))^\perp = \left\{ v \in k^{\dim I^d(\pi)} \mid \forall u \in \text{im } \eta^{d-1}(\pi) : v \cdot u = 0 \right\},$$

until we cannot add any more. To insert the new rows in the full labeled matrix η^d , we fill them by zeros in all columns not labeled by $\text{St } \pi$. Each new row is labeled by π , and represents a new copy of $[\pi]$ in \hat{I}^{d+1} .

Lemma 4.19. Let η^{d-1} , η^d and $\pi \in \Pi$ are as above, and $\hat{\eta}^d := \text{MAKEEXACT}(\eta^{d-1}, \eta^d, \pi)$. Then

- (1) $\hat{\eta}^d$ differs from η^d only by having extra rows labeled by π ,
- (2) $I^{d-1}(\pi) \xrightarrow{\eta^{d-1}(\pi)} I^d(\pi) \xrightarrow{\hat{\eta}^d(\pi)} \hat{I}^{d+1}(\pi)$ is exact at $I^d(\pi)$,
- (3) if $e_i, \dots, e_l \in \hat{I}^{d+1}(\pi)$ are the canonical vectors corresponding to the rows newly added to $\hat{\eta}^d(\pi)$, then $\text{span}(e_i, \dots, e_l) \subseteq \text{im } \hat{\eta}^d(\pi)$.

Proof. The point (1) is clear from the construction; we prove the other two. For (2) we need to show that $\text{im } \eta^{d-1}(\pi) = \ker \hat{\eta}^d(\pi)$. By assumption, $\text{im } \eta^{d-1}(\pi) \subseteq \ker \eta^d(\pi)$. Observe that

$$\begin{aligned} \text{im } \eta^{d-1}(\pi) &= \ker \hat{\eta}^d(\pi) \\ \iff \ker \eta^{d-1}(\pi)^T &= \text{im } \hat{\eta}^d(\pi)^T \\ \iff \text{im } \eta^{d-1}(\pi)^\perp &= \text{span}(\text{rows of } \hat{\eta}^d(\pi)), \end{aligned}$$

where the first equivalence comes from linear duality, and the second is just rewriting both sides of the equality. The same holds for inclusions, with the inclusion direction flipping at the first equivalence—hence, we start with $\text{im } \eta^{d-1}(\pi)^\perp \supseteq \text{span}(\text{rows of } \eta^d(\pi))$. In the algorithm, we add vectors from the left side as new rows in $\hat{\eta}^d(\pi)$, until the equality is achieved, securing the exactness.

To show (3), assume that rows i, \dots, l were added to $\hat{\eta}^d(\pi)$, and let $e_j = (\dots, 0, 1, 0, \dots)$ denote the j -th canonical vector. We show that $\text{span}(e_i, \dots, e_l) \subseteq \text{im } \hat{\eta}^d(\pi)$. Let A_j be the matrix consisting of the first j rows of $\hat{\eta}^d(\pi)$. As we only add new rows when they are linearly independent from the previous, we have $\ker A_j \subsetneq \ker A_{j-1}$ for every $j \in \{i, \dots, l\}$. That is, there exists u_j such that $A_{j-1}u_j = 0$, and $A_j u_j \neq 0$. This is only possible if $A_j u_j = \lambda e_j$ for some $\lambda \neq 0$, which means that $\eta^d(\pi) \cdot u_j = A_l u_j$ is a vector with zeros at positions $1, \dots, j-1$, and a non-zero at the j -th position. Therefore, $\text{span}(e_i, \dots, e_l) = \text{span}(\hat{\eta}^d(\pi) \cdot u_i, \dots, \hat{\eta}^d(\pi) \cdot u_l) \subseteq \text{im } \hat{\eta}^d(\pi)$. \square

Extending the minimal injective resolution from η^{d-1} to η^d The construction of I^{d+1} and η^d is described by Algorithm 3. We start with an empty matrix η^d and repeatedly call MAKEEXACT for all $\pi \in \Pi$ in a non-increasing order. In the next subsection, Example 4.23 goes through the algorithm in detail on a concrete example.

Algorithm 3 A step in the minimal injective resolution.

Input: η^{d-1} a labeled matrix over a poset Π
Output: η^d a labeled matrix
1: $\eta^d \leftarrow$ a labeled matrix with no rows and with columns labeled as rows of η^{d-1}
2: **for** $\pi \in \Pi$ in a non-increasing order **do**
3: $\eta^d \leftarrow \text{MAKEEXACT}(\eta^{d-1}, \eta^d, \pi)$
4: **return** η^d

▷ Algorithm 2

Lemma 4.20. Starting with the minimal injective hull of F , iterative application of Algorithm 3 yields the minimal injective resolution of F .

Proof. Each sheaf is injective by definition. By (2) from Lemma 4.19, we get exactness of $I^\bullet(\pi)$ in degree d right after the call of the procedure MAKEEXACT(η^{d-1}, η^d, π). Since we go through the degrees in the increasing order, and through the poset elements in a non-increasing order, $\eta^{d-1}(\pi)$ and $\eta^d(\pi)$ both stay the same from this point on, and hence $I^\bullet(\pi)$ is still exact in degree d at the end of the algorithm.

To show the minimality, we verify the condition on maximal vectors from Proposition 4.17. Maximal vectors in I^{d+1} over π , $M_{I^{d+1}}(\pi)$, are exactly the new vectors added at step π . That is, if rows i, \dots, l were

added to $\eta^d(\pi)$ at step π , $M_{I^{d+1}}(\pi) = \text{span}(e_i, \dots, e_l)$. By (3) from Lemma 4.19, this space is contained in $\text{im } \eta^d(\pi)$. \square

Remark (*Computing the orthogonal complement*). In Algorithm 2, we purposefully leave out any particular way how to compute the basis of $(\text{im}(\eta^{d-1}(\pi)))^\perp$. One way to compute it is via a standard row reduction algorithm: we start with $U \leftarrow$ identity matrix, $R \leftarrow \eta^{d-1}(\pi)$, and we reduce rows from top to bottom, reducing each by adding the rows above it to push the left-most non-zero as much to the right as possible. Every row operation performed on R is also performed on U , so that $R = U \cdot \eta^{d-1}(\pi)$. We end up with a lower-triangular matrix U such that all its rows corresponding to the zero rows of R form a basis of $(\text{im}(\eta^{d-1}(\pi)))^\perp$.

An immediate advantage of this approach is that we only ever work with rows of $\eta^{d-1}(\pi)$. This means we can represent the matrices η^d in a row-wise sparse representation, e.g., a list of rows, each represented as a “column index \rightarrow value” dictionary. In this representation, when we go from η^d to $\eta^d(\pi)$, we just choose all the rows labeled by $\text{St } \pi$ —there is no need to crop the rows themselves, as all entries not labeled by $\text{St } \pi$ are 0.

4.5. Injective resolution of the constant sheaf

We describe how to use Algorithm 3 to compute the minimal injective resolution of a constant sheaf, give the geometrical meaning of multiplicities of indecomposables in the minimal resolution, showcase the introduced notions with an example, and finally compute the complexity of computing the minimal injective resolution of the constant sheaf.

For the constant sheaf, the construction of the minimal injective hull $k_\Pi \xrightarrow{\alpha} I^0$ is very straightforward: $I^0 = \bigoplus_{j=1}^n [\pi_j]$ where π_1, \dots, π_n are the maximal elements of Π , each exactly once. The map α is given by the diagonal embedding

$$\begin{aligned} \alpha(\sigma) : k_\Pi(\sigma) &\longrightarrow I^0(\sigma) \\ 1 &\mapsto (1, \dots, 1)^T. \end{aligned}$$

Conveniently, we can represent this particular injection α as a labeled matrix: we define η^{-1} as a column of ones with rows labeled by the maximal elements of Π . We label the column by a new “virtual” element greater than all elements of Π . To obtain the minimal injective resolution of k_Π , we now iteratively run Algorithm 3 starting with the input η^{-1} .

Going one step further, we can even find the maximal elements of Π as an iteration of Algorithm 3. Indeed, we start with an (empty) matrix η^{-2} with no columns and one row labeled by the “virtual” element greater than all elements of Π . Running Algorithm 3 on η^{-2} then yields η^{-1} .

To summarize, we can compute the minimal injective resolution of the constant sheaf by iteratively running Algorithm 3 starting with η^{-2} , and then keep matrices η^d only for $d \geq 0$. Note that this is not just a lucky coincidence. This trick is a special case of computing a pullback, which we introduce later in Section 5.4.3. For now, even though we have not defined a pullback yet, we only remark that the constant sheaf k_Π is the pullback f^*k_{pt} for the constant map $f : \Pi \rightarrow \text{pt}$.

Multiplicity of the indecomposable injective sheaves The multiplicity of an indecomposable injective sheaf in the unique minimal injective resolution is a well-defined invariant of F . It is natural to ask what topological information is captured with these multiplicities. Below we answer this question for the constant sheaf on a finite simplicial complex.

Definition 4.21. Let I^\bullet be the minimal injective resolution of F . By $m_F^d(\sigma)$ we denote the multiplicity of $[\sigma]$ in I^d :




Fig. 1. The poset considered in Example 4.23. We omit vertex 1 from the labels.




Fig. 2. The resolution in Example 4.23. The numbers indicate the dimension at each simplex, positioned as in Fig. 1. The pink background indicates the generators. For example, in $I^1(\emptyset)$ we have 6 dimensions coming from generators above \emptyset , and we have 2 more generators at \emptyset .

$$I^d \cong \bigoplus_{\sigma \in \Pi} [\sigma]^{m_F^d(\sigma)}.$$

Equivalently, we can define $m_F^d(\sigma) := \dim M_{I^d}(\sigma)$, where $M_{I^d}(\sigma)$ is as in Definition 4.8.

Theorem 4.22. *Let Σ be a finite simplicial complex, k_Σ the constant sheaf on Σ (viewed as a poset with the face relation), and $H_c^\bullet(|\text{St } \sigma|; k)$ be the singular cohomology with compact support of the geometric realization of $\text{St } \sigma$ (see also Definition 5.12). Then*

$$m_{k_\Sigma}^d(\sigma) = \dim H_c^{d+\dim \sigma}(|\text{St } \sigma|; k).$$

See Appendix C.4 for a proof of the above theorem.

Examples We demonstrate how Algorithm 3 works for constant sheaves with two examples.

Example 4.23. Consider the 3-skeleton of the 4-simplex, with two extra edges attached to vertex 1. We compute the minimal injective resolution of $\Pi := \text{St}(1)$. We describe simplices as lists of vertices, and for brevity omit the vertex 1—e.g., $234 = \{1, 2, 3, 4\}$. See Fig. 1.

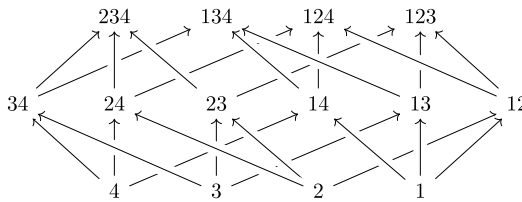
The generators of I^0 are the maximal simplices $(234, 235, 245, 345, 6, 7)$, and $\eta^{-1} : k_\Pi \rightarrow I^0$ is the diagonal embedding for each $\pi \in \Pi$. We construct I^1 and η^0 as in Algorithm 3, with inputs I^0, η^{-1} . Initialize I^1 and η^0 empty, and go through the simplices row-by-row left-to-right as they are in Fig. 1. Starting with 234, the space $I^0(234)$ is 1-dimensional and equal to $\text{im } \eta^{-1}(234)$, so there is nothing to be added, and $\eta^0(234) = 0$. The same happens for all the maximal simplices.

At triangle 23, we have $I^0(23) = k^2$, since two generators are above 23. At the moment, $\eta^0(23)$ is empty, so its kernel is k^2 . We need $\ker \eta^0(23) = \text{im } \eta^{-1}(23) = \text{span}\{(1, 1)\}$. The orthogonal complement of $\text{im } \eta^{-1}(23)$ is generated by the vector $(1, -1)$. We add it as a new row in $\eta^0(23)$. Therefore, we add 23 to I^1 , and add a first row to η^0 ; see Fig. 3. Similarly, we add one row for each other triangle.

$$\eta_0 \begin{pmatrix} 234 & 235 & 245 & 345 & 6 & 7 \\ 23 & \boxed{1} & -1 & 0 & 0 & 0 \\ 24 & \boxed{1} & 0 & -1 & 0 & 0 \\ 25 & \boxed{0} & 1 & -1 & 0 & 0 \\ 34 & \boxed{1} & 0 & 0 & -1 & 0 \\ 35 & \boxed{0} & 1 & 0 & -1 & 0 \\ 45 & \boxed{0} & 0 & 1 & -1 & 0 \\ \emptyset & 0 & 0 & 0 & -1 & 1 \\ \emptyset & 0 & 0 & 0 & 0 & -1 & 1 \end{pmatrix} \quad \eta_1 \begin{pmatrix} 23 & 24 & 25 & 34 & 35 & 45 & \emptyset & \emptyset \\ 2 & \boxed{1} & -1 & 1 & 0 & 0 & 0 & 0 \\ 3 & \boxed{1} & 0 & 0 & -1 & 1 & 0 & 0 \\ 4 & 0 & \boxed{1} & 0 & -1 & 0 & \boxed{1} & 0 \\ 5 & 0 & 0 & \boxed{1} & 0 & -1 & 1 & 0 \end{pmatrix}$$

$$\eta_2 \begin{pmatrix} 2 & 3 & 4 & 5 \\ \emptyset & 1 & -1 & 1 & -1 \end{pmatrix}$$

Fig. 3. Matrices η^0 , η^1 , η^2 in Example 4.23, with highlighted submatrices $\eta^k(2)$ (solid green), $\eta^k(3)$ (dotted blue), $\eta^k(4)$ (rounded corners magenta), $\eta^k(5)$ (dashed red). Recall that $\eta^k(\sigma) = \eta^k[\text{St } \sigma, \text{St } \sigma]$, and note that if $\sigma \not\leq \tau$, then $\eta^k[\sigma, \tau] = 0$. Recall that the labels omit the vertex 1 that is present in all the simplices, e.g., \emptyset is the vertex 1 and 234 is the tetrahedron 1234.



$$\eta^0 \begin{pmatrix} 234 & 134 & 124 & 123 \\ 34 & -1 & 1 & 0 & 0 \\ 24 & -1 & 0 & 1 & 0 \\ 23 & -1 & 0 & 0 & 1 \\ 14 & 0 & -1 & 1 & 0 \\ 13 & 0 & -1 & 0 & 1 \\ 12 & 0 & 0 & -1 & 1 \end{pmatrix} \quad \eta^1 \begin{pmatrix} 34 & 24 & 23 & 14 & 13 & 12 \\ 4 & 1 & -1 & 0 & 1 & 0 & 0 \\ 3 & 1 & 0 & -1 & 0 & 1 & 0 \\ 2 & 0 & 1 & -1 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & -1 & 1 \end{pmatrix}$$

$$I^0 \xrightarrow{\eta^0} I^1 \xrightarrow{\eta^1} I^2$$

$$\begin{pmatrix} 1 & 1 & 1 & 1 \\ 2 & 2 & 2 & 2 & 2 \\ 3 & 3 & 3 & 3 \end{pmatrix} \xrightarrow{\eta^0} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 \\ 3 & 3 & 3 & 3 \end{pmatrix} \xrightarrow{\eta^1} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 \end{pmatrix}$$

Fig. 4. Example 4.24: the minimal injective resolution of the constant sheaf on $\Delta_3^{(2)} \cong S^2$. First is $\Delta_3^{(2)}$ as a poset, second the dimensions in the injective resolution with highlighted generators, as in Fig. 2, and third the matrices describing the natural transformations.

Now for the edges. We have $I^0(2) = k^3$, and $\eta^0(2)$ a 3×3 matrix, highlighted as a green solid rectangle in Fig. 3. We already have $\ker \eta^0(2) = \text{span}\{(1, 1, 1)\} = \text{im } \eta^{-1}(2)$, so we do not add any new generators over 2. The same goes for $\eta^0(3), \eta^0(4), \eta^0(5)$, each of which you can see highlighted in Fig. 3 with a different color and line style.

Finally, we get to the vertex \emptyset , with $\eta^0(\emptyset)$ starting as the part of the matrix in Fig. 3 above the horizontal line. Its rank is 3, and its nullity is 3. We need the kernel to be 1-dimensional, so we need to add two additional rows from $(\text{span}\{(1, 1, 1, 1, 1)\})^\perp$. We also add \emptyset to I^1 twice. This completes the construction of I^1 and η^0 .

The resolution goes on for two more steps: I^2 is generated by $(2, 3, 4, 5)$, I^3 by (\emptyset) . The matrices η^k are in Fig. 3, and the whole resolution is schematically shown in Fig. 2.

Example 4.24. Let $\Sigma := \Delta_3^{(2)}$ be the 2-skeleton of a tetrahedron (whose geometric realization is homeomorphic to the sphere). We give the minimal injective resolution of the constant sheaf k_Σ in Fig. 4.

Complexity analysis We analyze the complexity of finding the minimal injective resolution of the constant sheaf, k_Π , on a poset Π , with n elements and height h , computed by an iterative application of Algorithm 3. That is, we start with $k_\Pi \xrightarrow{\eta^{-1}} I^0$ the minimal injective hull of the constant sheaf as described above, and then iteratively apply Algorithm 3 until $I^d = 0$.

The call of MAKEEXACT consists of finding a basis of $(\text{im}(\eta^{d-1}(\sigma)))^\perp$, and checking for linear independence of rows of $\eta^d(\sigma)$. Both of those operations can be computed in time at most $\mathcal{O}(c^3)$ with c the maximum of the number of rows of $\eta^{d-1}(\sigma)$ and $\eta^d(\sigma)$. In our analysis we ignore the complexity of finding $\text{St } \sigma$ to extract the submatrices from η^d in the first place, since it is less expensive than $\mathcal{O}(c^3)$ when we estimate c by the size of $\text{St } \sigma$.

By Corollary 4.18, the length of the minimal injective resolution is at most $h + 1$. Therefore, we find it in time $\mathcal{O}(h \cdot n \cdot c^3)$, where

$$c = \max_{j, \sigma} \sum_{\pi \in \text{St } \sigma} m_{k_\Pi}^j(\pi)$$

is the maximal number of generators over any star throughout the resolution. This analysis is output-sensitive. To give complexity bounds dependent only on the input, we compare c to the maximal size of a star in Π . How well we can approximate c this way depends on the structure of Π .

Definition 4.25. For $\sigma \in \Pi$ we define

$$m^j(\text{St } \sigma) := \sum_{\pi \in \text{St } \sigma} m_{k_\Pi}^j(\pi)$$

to be the number of generators over $\text{St } \sigma$ in the j -th step of the minimal injective resolution of the constant sheaf on Π (recall Definition 4.21). Furthermore, we define the j -th star complexity of σ as

$$\text{SC}^j(\sigma) := \frac{m^j(\text{St } \sigma)}{\# \text{St } \sigma}.$$

For general posets, $\text{SC}^j(\sigma)$ can be arbitrarily large even when lengths of chains are bounded, because sizes of boundaries and coboundaries can be arbitrarily large. For simplicial complexes, we give an upper bound on $\text{SC}^j(\sigma)$ depending on the dimension.

Proposition 4.26. Let Π be a simplicial complex, $\sigma \in \Pi$ and $d := \dim \text{St } \sigma - \dim \sigma$, where by $\dim \text{St } \sigma$ we mean the maximum dimension of any $\tau \in \text{St } \sigma$. Then

$$\text{SC}^j(\sigma) \leq \binom{d}{j}.$$

This bound is asymptotically tight. If $\Pi = \Delta_n^{(d)}$ is the d -skeleton of the n -simplex, v is a vertex in $\Delta_n^{(d)}$, and j is fixed, then

$$\text{SC}^j(v) \xrightarrow{n \rightarrow \infty} \binom{d}{j}.$$

Proof. We prove the upper bound using Theorem 4.22 and bounding dimensions of homology groups by dimensions of chain groups:

$$\begin{aligned} m^j(\text{St } \sigma) &= \sum_{\tau \in \text{St } \sigma} m^j(\tau) = \sum_{\tau \in \text{St } \sigma} \dim H_c^{j+\dim \tau}(\text{St } \tau) \leq \sum_{\tau \in \text{St } \sigma} \dim C_c^{j+\dim \tau}(\text{St } \tau) \\ &= \sum_{\tau \in \text{St } \sigma} \# \{ \pi \mid \tau <_j \pi \} = \sum_{\pi \in \text{St } \sigma} \# \{ \tau \in \text{St } \sigma \mid \tau <_j \pi \} = \sum_{\pi \in \text{St } \sigma} \# \binom{\pi \setminus \sigma}{j} \leq \# \text{St } \sigma \cdot \binom{d}{j}, \end{aligned}$$

where $d = \dim \text{St } \sigma - \dim \sigma$, and $\tau <_j \pi$ signifies that $\dim \pi - \dim \tau = j$.

Now we analyze $\mathrm{SC}^j(\sigma)$ in the d -skeleton of the n -simplex, $\Delta_n^{(d)}$. We use the fact that $\mathrm{St} \sigma$ in $\Delta_n^{(d)}$ is combinatorially the same as $\Delta_{n'}^{(d')} \cup \{\emptyset\}$, with $n' = n - \dim \sigma - 1$ and $d' = d - \dim \sigma - 1$, using the correspondence $\mathrm{St} \sigma \ni \tau \mapsto \tau \setminus \sigma$. This map induces an isomorphism between the cochain complexes

$$C_c^\bullet(\mathrm{St} \sigma) \cong \tilde{C}^{\bullet - \dim \sigma - 1}(\Delta_{n'}^{(d')}),$$

which, using Theorem 4.22, implies

$$m^j(\sigma) = \dim H_c^{j + \dim \sigma}(\mathrm{St} \sigma) = \dim \tilde{H}^{j-1}(\Delta_{n'}^{(d')}).$$

The (standard simplicial) reduced cohomology $\tilde{H}^i(\Delta_{n'}^{(d')})$ is trivial for all $i \neq d'$, and for $i = d'$, we compute the dimension from the Euler characteristic:

$$\begin{aligned} \dim \tilde{H}^{d'}(\Delta_{n'}^{(d')}) &= (-1)^{d'} \tilde{\chi}(\Delta_{n'}^{(d')}) = (-1)^{d'} \left(1 + \sum_{i=0}^{d'} \binom{n'+1}{i+1} (-1)^i \right) \\ &= (-1)^{d'-1} \left(\sum_{i=0}^{d'+1} \binom{n'+1}{i} (-1)^i \right) = (-1)^{d'-1} \cdot (-1)^{d'+1} \binom{n'}{d'+1} = \binom{n'}{d'+1}. \end{aligned}$$

Therefore,

$$m^j(\sigma) = \begin{cases} \binom{n - \dim \sigma - 1}{d - \dim \sigma} & \text{if } j = d' + 1 = d - \dim \sigma, \\ 0 & \text{otherwise.} \end{cases}$$

Finally, we compute $m^j(\mathrm{St} v)$ for a vertex v :

$$m^j(\mathrm{St} v) = \sum_{\sigma \in \mathrm{St} v} m^j(\sigma) = \sum_{\substack{\sigma \in \mathrm{St} v \\ \dim \sigma = d-j}} \binom{n-d+j-1}{j} = \binom{n}{d-j} \cdot \binom{n-d+j-1}{j}$$

We rearrange this as follows

$$\begin{aligned} m^j(\mathrm{St} v) &= \frac{n!}{(n-d+j)!(d-j)!} \cdot \frac{(n-d+j-1)!}{(n-d-1)!j!} \\ &= \frac{n!}{(n-d)!d!} \cdot \frac{n-d}{n-d+j} \cdot \frac{d!}{(d-j)!j!} = \binom{n}{d} \binom{d}{j} \frac{n-d}{n-d-1}. \end{aligned}$$

Now we can easily compare this with $\#\mathrm{St} v = \sum_{i=0}^d \binom{n}{i}$. When we fix d and j , we get

$$\lim_{n \rightarrow \infty} \mathrm{SC}^j(v) = \lim_{n \rightarrow \infty} \frac{m^j(\mathrm{St} v)}{\#\mathrm{St} v} = \binom{d}{j}. \quad \square$$

Corollary 4.27. *For a fixed dimension h , the Algorithm 3 computes the minimal injective resolution of the constant sheaf on an h -dimensional simplicial complex Σ in time $\mathcal{O}(n \cdot s^3)$, where n is the cardinality of Σ (as an abstract simplicial complex), and s is the cardinality of the largest star in Σ .*

5. Derived functors

A fundamental application of derived sheaf theory is the study of continuous maps. To illustrate these applications, we will recall several functors between categories of sheaves which are induced by continuous maps between topological spaces. In Section 5.1 we define the functors necessary to state the main results of this paper (in Section 6). In Section 5.4 we give detailed algorithms for computing these functors, using the terminology and notation of Section 3.

We define pushforward and pullback for any order preserving map on posets, and proper pushforward and proper pullback for an inclusion of a *locally closed* subset of a poset.

Definition 5.1. A subset Z of a poset Π is *locally closed* if it is an intersection of a closed subset (a down set) and an open set (an upper set). Equivalently, Z is convex: if $a, c \in Z$ and $a \leq b \leq c$, then $b \in Z$.

5.1. Definitions of derived functors

5.1.1. Pushforward

Definition 5.2. Suppose $f : \Pi \rightarrow \Lambda$ is an order preserving map of posets, and F is a sheaf on Π . Then the *pushforward* of F by f is defined by

$$\begin{aligned} f_*F(\lambda) &:= \varprojlim_{\pi \in f^{-1}(\text{St } \lambda)} F(\pi) \\ &= \left\{ v \in \bigoplus_{\pi \in f^{-1}(\text{St } \lambda)} F(\pi) : F(\gamma \leq \tau) \left(\text{proj}_{F(\gamma)}(v) \right) = \text{proj}_{F(\tau)}(v) \text{ for all } (\gamma \leq \tau) \in f^{-1}(\text{St } \lambda) \right\}, \end{aligned}$$

where $\text{proj}_{F(\gamma)} : \bigoplus_{\pi \in f^{-1}(\text{St } \lambda)} F(\pi) \rightarrow F(\gamma)$ denotes projection onto $F(\gamma)$, and the linear maps $f_*F(\kappa \leq \lambda)$ are restrictions of the projections $\bigoplus_{\pi \in f^{-1}(\text{St } \kappa)} F(\pi) \rightarrow \bigoplus_{\pi \in f^{-1}(\text{St } \lambda)} F(\pi)$.

The pushforward is functorial; that is, if $\eta : F \rightarrow G$ is a natural transformation between sheaves F and G on Π , then

$$f_*\eta : f_*F \rightarrow f_*G,$$

is a natural transformation, where $f_*\eta(\lambda)$ is obtained by restricting the domain of the sum of linear maps

$$\sum_{\pi \in f^{-1}(\text{St } \lambda)} \eta(\pi) : \bigoplus_{\pi \in f^{-1}(\text{St } \lambda)} F(\pi) \rightarrow \bigoplus_{\pi \in f^{-1}(\text{St } \lambda)} G(\pi).$$

Notice that if $\eta : F \rightarrow G$ is injective, then $f_*\eta$ is injective. However, the same is not necessarily true for surjectivity. In the language of category theory, f_* is a left exact functor. It is useful to observe that in our finite setting, pushforwards of injective complexes have a straightforward structure.

Lemma 5.3. If $I \cong \bigoplus_{0 \leq j \leq n} [\pi_j]$ is an injective sheaf on Π , then f_*I is an injective sheaf on Λ given by

$$f_*I \cong \bigoplus_{0 \leq j \leq n} [f(\pi_j)].$$

Proof. Because $f_*\left(\bigoplus_j F_j\right) = \bigoplus_j f_*F_j$, it is enough to observe that $f_*[\pi] \cong [f(\pi)]$ for any $\pi \in \Pi$. Indeed, by definition

$$f_*[\pi](\sigma) = \begin{cases} k & \text{if } \pi \in f^{-1}(\text{St } \sigma) \\ 0 & \text{otherwise,} \end{cases}$$

and $\pi \in f^{-1}(\text{St } \sigma)$ if and only if $f(\pi) \in \text{St } \sigma$. \square

5.1.2. Pullback

Definition 5.4. Suppose $f : \Pi \rightarrow \Lambda$ is an order preserving map of posets, and F is a sheaf on Λ . Then the *pullback* of F by f is defined by

$$f^*F(\pi) := F(f(\pi)),$$

with linear maps $f^*F(\pi \leq \tau) = F(f(\pi) \leq f(\tau))$.

As with the pushforward, the pullback is functorial: if $\eta : F \rightarrow G$ is a natural transformation between sheaves F and G on Λ , then η defines a natural transformation

$$f^*\eta : f^*F \rightarrow f^*G,$$

where $f^*\eta(\pi) := \eta(f(\pi))$.

Unlike the pushforward, the pullback preserves both the injectivity and surjectivity of natural transformations, and defines an exact functor. However, (again in contrast to the pushforward) the pullback does not necessarily preserve injectivity of sheaves.

Remark. A classic result in sheaf theory, which continues to hold in the setting of finite posets, is that the functors f^* and f_* form an adjoint pair:

$$\text{Hom}_\Pi(f^*F, G) \cong \text{Hom}_\Lambda(F, f_*G).$$

See [16, Theorem 3.14] for a detailed proof and [49, Theorem 3.4.7] for an analogous theorem of the adjointness of the corresponding derived functors.

Example 5.5. Let $\Pi = (a \leq b, c)$, $\Lambda = (a' \leq b', c' \leq d')$ and $f : \Pi \rightarrow \Lambda$ sends a, b, c to a', b', c' , respectively. Then the constant sheaf on Λ , is injective: $k_\Lambda = [d']$. By the definition, its pullback is $f^*k_\Lambda = k_\Pi$, which is not an injective sheaf on Π .

$$\begin{array}{ccc} & & k \\ & \nearrow & \nwarrow \\ k & & k \\ \nwarrow & & \nearrow \\ & k & \end{array} \xleftarrow{f^*} \begin{array}{ccc} & k & \\ \nearrow & & \nwarrow \\ k & & k \\ \nwarrow & & \nearrow \\ & k & \end{array}$$

In Section 5.4 we describe a pushforward of injective sheaves as a restriction of an injective sheaf over a modified poset.

5.1.3. Proper pushforward

Now we define the proper pushforward. Unlike for the previous functors, here we assume $f := i_Z$ to be an inclusion of a locally closed subset Z .

Definition 5.6. For a locally closed subset, Z , of a poset Π , and a sheaf F on Z , the *proper pushforward* of F by the inclusion $i : Z \hookrightarrow \Pi$ is

$$i_! F(\sigma) := \begin{cases} F(\sigma) & \text{if } \sigma \in Z, \\ 0 & \text{else.} \end{cases}$$

As above, it is straightforward to show that $i_!$ is functorial, and that $i_!(\eta)$ preserves injectivity and surjectivity of η . In other words, $i_!$ is exact.

Unless Z is closed, a proper pushforward, $(i_Z)_! I$, of an injective sheaf is not injective, because there will necessarily exist two poset elements $a \leq b$ with $(i_Z)_! I(a) = 0$ and $(i_Z)_! I(b) \neq 0$. However, it is a subsheaf of an injective sheaf. Indeed, if $I \cong \bigoplus_{0 \leq j \leq n} [\pi_j]$ is an injective sheaf on Z , and $\hat{I} := (i_Z)_* I \cong \bigoplus_{0 \leq j \leq n} [\pi_j]$ is an injective sheaf on Π , then $(i_Z)_! I(\pi) = \hat{I}(\pi)$ for every $\pi \in Z$.

5.1.4. Proper pullback

We define the proper pullback, again only for inclusions of locally closed subsets.

Definition 5.7. For a locally closed subset, Z , of a poset Π , and a sheaf F on Z , the *proper pullback* of F by the inclusion $i : Z \hookrightarrow \Pi$ is defined by

$$i^! F(\zeta) := \left\{ v \in F(\zeta) : F(\zeta \leq \tau)(v) = 0 \text{ for all } \tau \in \text{St}_\Pi \zeta - Z \right\},$$

where $\zeta \in Z$ and $\text{St}_\Pi \zeta$ is the star of ζ in Π .

As above, it is straightforward to show that $i^!$ is functorial, and that $i^!(\eta)$ preserves the injectivity of η but not necessarily surjectivity. As the following lemma shows, computing proper pullbacks of injective sheaves is straightforward.

Lemma 5.8. Let Z be a locally closed subset of Π . If $I \cong \bigoplus_{\pi \in \Pi} [\pi]^{p_\pi}$ is an injective sheaf on Π , then $i_Z^! I$ is an injective sheaf on Z with a decomposition

$$i_Z^! I \cong \bigoplus_{\pi \in Z} [\pi]^{p_\pi}.$$

Moreover, if I^\bullet is a minimal injective complex, then so is $i_Z^! I^\bullet$.

Proof. The functor $i_Z^!$ is additive, so it is enough to consider $i_Z^! [\pi]$ for an indecomposable injective sheaf $[\pi]$.

Assume first that $\pi \notin Z$. Let $\zeta \in Z$. Either $\zeta \leq \pi$, and then $i_Z^! [\pi](\zeta) = 0$, because $[\pi](\zeta \leq \pi) = \text{id}_k$; or $\zeta \not\leq \pi$, and then also $i_Z^! [\pi](\zeta) = 0$, because $[\pi](\zeta) = 0$. Therefore, $i_Z^! [\pi] = 0$ for $\pi \notin Z$.

Now we consider the case when $\pi \in Z$. If $\zeta \in Z$ with $\zeta \not\leq \pi$, then $i_Z^! [\pi](\zeta) = 0$ because $[\pi](\zeta) = 0$. If $\zeta \leq \pi$, then the intersection

$$\{\gamma \in \Pi : \gamma \leq \pi\} \cap \text{St}_\Pi \zeta$$

is a subset of Z (because Z is locally closed). Therefore, if $\tau \in \text{St}_\Pi \zeta - Z$ then $\tau \not\leq \pi$, and hence $[\pi](\zeta \leq \tau) = 0$. It follows that $i_Z^! [\pi](\zeta) = k$. Altogether, we have shown that $i_Z^! ([\pi]_\Pi)$ is isomorphic to $[\pi]_Z$ (we add the subscripts Π and Z to emphasize that $[\pi]_\Pi$ is a sheaf on Π and $[\pi]_Z$ is a sheaf on Z).

The minimality of $i_Z^! I^\bullet$, provided that I^\bullet is minimal, then follows from condition (4) of Theorem 4.9: for a natural transformation $\eta^d : I^d \rightarrow I^{d+1}$ and $\pi \in Z$, the map $i_Z^! \eta^d(\pi)$ is just a restriction of $\eta^d(\pi)$. The maximal vectors in $i_Z^! I(\pi)$ are the same as maximal vectors in $I(\pi)$, so if they are sent to zero by η^d , they are also sent to zero by $i_Z^! \eta^d(\pi)$, verifying the condition (4). \square

Combining this with Lemma 5.3, we see that $(i_Z)_* i_Z^! I$ is the subsheaf of I containing only the indecomposable injective summands $[\pi]$ on Π for $\pi \in Z$.

5.1.5. Summary

For a quick reference, we provide a summary of how the four introduced functors, $(i_Z)_*$, $(i_Z)_!$, $(i_Z)^*$, $(i_Z)^!$, act on injective sheaves in case of an inclusion $i_Z : Z \hookrightarrow \Pi$. The cones in the schema represent indecomposable injectives $[\pi]$.



We see that while $(i_Z)_*$ and $(i_Z)^!$ are built of full indecomposable injectives, this is not the case for the other two— $(i_Z)^*$ contains parts of indecomposable injectives clipped from above, and $(i_Z)_!$ contains parts clipped from below. In Section 5.4 we discuss how to compute injective resolutions of images of injective complexes. While for $(i_Z)_*$ and $(i_Z)^!$ we already have injective complexes, we need to carry out extra computations in case of $(i_Z)^*$ and $(i_Z)_!$, which we describe as Algorithm 4 and Algorithm 5, respectively.

5.2. Derived functors and hypercohomology

We now apply the above operations to complexes of injective sheaves.

Definition 5.9. Let $f : \Sigma \rightarrow \Lambda$ be an order preserving map of posets and $(I^\bullet, \eta^\bullet) \in D^b(\Sigma)$ be a minimal injective complex on Σ . Let φ be one of the functors f_* , f^* , $f_!$, or $f^!$, defined above (we assume f is an inclusion in the proper cases). Define $\mathcal{I}\varphi(I^\bullet) \in D^b(\Lambda)$ to be the minimal injective resolution of the complex

$$\cdots \rightarrow \varphi(I^d) \xrightarrow{\varphi\eta^d} \varphi(I^{d+1}) \xrightarrow{\varphi\eta^{d+1}} \varphi(I^{d+2}) \rightarrow \cdots$$

As with other functors, we usually write $\mathcal{I}\varphi I^\bullet = \mathcal{I}\varphi(I^\bullet)$ for brevity.

Remark. For a functor φ on sheaves, the above defined functor $\mathcal{I}\varphi$ plays the role of the right derived functor $R\varphi$ from the classical theory of derived categories, and we therefore refer to $\mathcal{I}\varphi$ as the *derived functor* of φ .

Remark. While most authors would not use $R\varphi$ for an exact functor φ , we still write $\mathcal{I}\varphi$ to ensure that the result is a minimal injective chain. For example, since f^* is an exact functor it naturally extends to a functor on the derived category, and all of the higher cohomology sheaves vanish for the right derived functor Rf^* . Most authors then use f^* to denote the functor on the category of sheaves as well as on the derived category of sheaves. However, we note that while f^* is exact, it does not preserve the injectivity of sheaves, and therefore does not immediately define a functor on (the chosen skeleton of) the derived category of sheaves as we have defined it. To define such a functor, we must post-compose f^* with a minimal injective resolution. Therefore, we write $\mathcal{I}f^*$ and highlight the algorithmic difference from simply applying f^* to a complex of sheaves. The same remarks apply for $f_!$.

Remark. Secondly, in the most classical settings (such as when studying smooth maps between manifolds), $f^!$ does not exist as a functor on the category of sheaves, but only as a functor on the derived category (this

is indeed one of the motivations for defining the derived category). In these settings, $f^!$ is often introduced as the left adjoint of $Rf_!$. However, for simplicity, when discussing $f^!$ in this paper we restrict our attention to only the poset maps $f : Z \hookrightarrow \Pi$ which are inclusions of locally closed subposets. In this restricted setting, the functors $f^!$ are well-defined as functors on the category of sheaves, and justify the additional notation $Rf^!$ for the corresponding right derived functor on the derived category of sheaves, and our corresponding functor $\mathcal{I}f^!$. It would be interesting and worthwhile to extend our results to more general poset maps, and to give an algorithm for computing $f^!$ in a more general setting when it is classically defined as the left adjoint of $Rf_!$.

Definition 5.10. Let $p : \Sigma \rightarrow \{\star\}$ be the projection map to the single point poset. Let I^\bullet be a complex of injective sheaves on Σ . Define the *hypercohomology* of I^\bullet to be

$$\mathbb{H}^d(I^\bullet) := H^d(\mathcal{I}p_* I^\bullet).$$

Notice that $\mathcal{I}p_*(I^\bullet)$ is simply a complex of vector spaces. As such, the minimality condition forces all maps to be zero maps, which implies that $H^d(\mathcal{I}p_* I^\bullet) = (\mathcal{I}p_* I^\bullet)^d$ when we view the right-hand side as the vector space over the single element.

Corollary 5.11. Assume $I^\bullet \in D^b(\Sigma)$ is a minimal injective complex. Let $m_{I^\bullet}^d(\pi)$ be the multiplicity of the indecomposable summand $[\pi]$ in I^d . If $i_\pi : \{\pi\} \hookrightarrow \Sigma$, then

$$\dim \mathbb{H}^d \mathcal{I}i_\pi^! I^\bullet = m_{I^\bullet}^d(\pi).$$

Proof. This follows directly from Lemma 5.8: we have $(\mathcal{I}i_\pi^! I^\bullet)^d = i_\pi^! I^d \cong [\pi]^{m_{I^\bullet}^d(\pi)}$, and since all vectors on a singleton are maximal, all the complex maps are zero by the minimality and condition (4) of Theorem 4.9. \square

In the setting of simplicial complexes (or, more generally, cellular complexes) the cohomology sheaves of the derived pushforward of a constant sheaf are closely related to singular cohomology groups of level-sets. This allows us to give some intuition behind derived pushforwards with the proposition below. Before formulating it, we recall two classical notions of cohomology for topological spaces (see, e.g., [33, 2.6.8]).

Definition 5.12. Let X be a topological space, $p : \Pi \rightarrow \{\text{pt}\}$ a constant map to a point, and F a sheaf on X with injective resolution I^\bullet . Then

$$\begin{aligned} H^d(X; F) &:= H^d(\mathcal{I}p_* F) = H^d(\mathcal{I}p_* I^\bullet), \\ H_c^d(X; F) &:= H^d(\mathcal{I}p_! F) = H^d(\mathcal{I}p_! I^\bullet). \end{aligned}$$

By $H^d(X; k)$ and $H_c^d(X; k)$ we denote the above with the choice $F = k_X$, the constant sheaf on X . Those definitions yield the standard and compactly supported cohomology (see, e.g., [27]), respectively.

We use these notions mainly for geometric realizations of simplicial complexes and their subsets. Note that we cannot just blindly replace X by a poset—for one, we did not give a definition for $p_!$, and, more importantly, if $\{v\}$ is a vertex and $\{e\}$ an open edge, then they are equal as posets, but $H_c^1(|\{v\}|; k) = 0$ while $H_c^1(|\{e\}|; k) \cong k$. We can, however, replace X by a face poset of a simplicial complex in $H^d(X; F)$.

Proposition 5.13. Suppose $f : \Sigma \rightarrow \Lambda$ is a simplicial map. As sheaves on Λ ,

$$H^d \mathcal{I}f_* k_\Sigma \cong H^d(|f^{-1}(\text{St } -)|; k)$$

where $H^d(|f^{-1}(\text{St } -)|; k)$ is the sheaf defined by associating the simplex λ to the singular cohomology of the geometric realization of $f^{-1}(\text{St } \lambda)$ (with linear maps induced by inclusion).

Proof. We provide a proof at the end of Appendix C.2. \square

5.3. Recollement for posets

Suppose $\Pi = C \sqcup U$ is a partition of the finite poset Π into a closed set C and an open set U (relative to the Alexandrov topology). Let $i_C : C \hookrightarrow \Pi$, $j_U : U \hookrightarrow \Pi$, and p, q, r denote projections to the single element poset $\{\star\}$:

$$\begin{array}{ccccc} C & \xrightarrow{i_C} & \Pi & \xleftarrow{j_U} & U \\ & \searrow p & \downarrow q & \swarrow r & \\ & & \{\star\} & & \end{array}$$

Lemma 5.14. *If $I^\bullet \in D^b(\Pi)$, then there are long exact sequences of hypercohomology groups (recalling Definitions 5.9 and 5.10)*

$$\begin{aligned} \cdots \rightarrow \mathbb{H}^d(\mathcal{I}i_C^! I^\bullet) &\rightarrow \mathbb{H}^d(I^\bullet) \rightarrow \mathbb{H}^d(\mathcal{I}j_U^* I^\bullet) \\ &\rightarrow \mathbb{H}^{d+1}(\mathcal{I}i_C^! I^\bullet) \rightarrow \cdots \end{aligned} \quad (5.1)$$

and

$$\begin{aligned} \cdots \rightarrow \mathbb{H}^d(\mathcal{I}j_U! \mathcal{I}j_U^* I^\bullet) &\rightarrow \mathbb{H}^d(I^\bullet) \rightarrow \mathbb{H}^d(\mathcal{I}i_C^* I^\bullet) \\ &\rightarrow \mathbb{H}^{d+1}(\mathcal{I}j_U! \mathcal{I}j_U^* I^\bullet) \rightarrow \cdots \end{aligned} \quad (5.2)$$

Proof. For any injective sheaf I on Π , there are short exact sequence of sheaves:

$$0 \rightarrow i_{C*} i_C^! I \rightarrow I \rightarrow j_{U*} j_U^* I \rightarrow 0,$$

and

$$0 \rightarrow j_{U!} j_U^* I \rightarrow I \rightarrow i_{C*} i_C^* I \rightarrow 0.$$

These exact sequences define distinguished triangles in $D^b(\Pi)$, which then yield the desired long exact sequence of hypercohomology (see [33, Propositions 1.7.5 and 1.8.8, Remark 2.6.10]). \square

5.4. Algorithms for computing derived functors

In this section we describe how to compute minimal injective resolutions, $\mathcal{I}f_*$, $\mathcal{I}f^*$, $\mathcal{I}f_!$, $\mathcal{I}f^!$, of the introduced functors, f_* , f^* , $f_!$, $f^!$, applied to injective complexes. We first introduce a notion of a poset mapping cylinder, and then describe algorithms to compute each of the four functors. Concrete examples of the computations are presented in Section 7. We assume $f : \Pi \rightarrow \Lambda$ is an order preserving map of posets and $(I^\bullet, \eta^\bullet)$ is a complex of injective sheaves represented by labeled matrices η^\bullet . We will assume each I^\bullet to start in degree 0. If it is not the case, we can shift the degrees in the obvious manner—if n is the smallest degree with non-zero I^n , we consider new matrices $\hat{\eta}^d = \eta^{d+n}$. Note that this means that we start with a zero matrix η^{-1} with rows labeled according to a decomposition of I^0 , and no columns, as discussed after Definition 3.17.

For the discussions we also fix $I = \bigoplus_i [\pi_i]$, $J = \bigoplus_j [\sigma_j]$ and $\eta : I \rightarrow J$ a natural transformation represented by a labeled matrix η .

5.4.1. Poset mapping cylinder of an order preserving map

The strength of our approach is that we can compute all the introduced derived functors using labeled matrices, despite the fact that pullback and proper pushforward do not preserve injectivity. For that purpose, we use poset mapping cylinders, and consider injective complexes on those. The construction is also known as non-Hausdorff mapping cylinder in the study of finite T_0 topological spaces, see, e.g., [2, Definition 2.8.1] or before in [8].

Definition 5.15. Given an order preserving map between two posets, $f : \Pi \rightarrow \Lambda$, its *mapping cylinder*, $\Pi \sqcup_f \Lambda$, is the poset with the underlying set the disjoint union $\Pi \sqcup \Lambda$, and the order defined as follows: $\pi \leq \tau$ in $\Pi \sqcup_f \Lambda$ if and only if one of the following conditions is satisfied

- $\pi, \tau \in \Pi$ and $\pi \leq \tau$ in Π ,
- $\pi, \tau \in \Lambda$ and $\pi \leq \tau$ in Λ , or
- $\pi \in \Pi$, $\tau \in \Lambda$ and $f(\pi) \leq \tau$ in Λ .

We also call it a *poset cylinder* to more clearly distinguish it from the mapping cone of a complex morphism.

In other words, the order relation in $\Pi \sqcup_f \Lambda$ is generated by new relations $\pi \leq f(\pi)$ for every $\pi \in \Pi$ together with the order relations on Π and Λ . Note that every chain in $\Pi \sqcup_f \Lambda$ decomposes into a prefix of elements from Π and suffix of elements from Λ (either possibly empty). The following is an important property of this construction.

Lemma 5.16. For $\pi \in \Pi$, the star of π in $\Pi \sqcup_f \Lambda$ is the union of stars of π in Π and $f(\pi)$ in Λ :

$$\text{St}_{\Pi \sqcup_f \Lambda} \pi = \text{St}_{\Pi} \pi \cup \text{St}_{\Lambda} f(\pi).$$

Proof. The inclusion “ \supseteq ” is immediate since $\pi \leq f(\pi)$ in $\Pi \sqcup_f \Lambda$. For “ \subseteq ”, let $\pi \leq \tau$ in $\Pi \sqcup_f \Lambda$. Then by definition either $\tau \in \Pi$ and $\tau \in \text{St}_{\Pi} \pi$; or $\tau \in \Lambda$ and then $f(\pi) \leq \tau$ in Λ , i.e., $\tau \in \text{St}_{\Lambda} f(\pi)$. \square

In particular, the lemma implies that if η is a labeled matrix representing a natural transformation between two injective sheaves on $\Pi \sqcup_f \Lambda$, and if all its labels are in Λ , then $\eta(\pi) = \eta(f(\pi))$ for all $\pi \in \Pi$.

5.4.2. Computing pushforward

For labeled matrices, computing pushforwards is very simple. By Lemma 5.3, we only need to replace every label π by $f(\pi)$. This yields a labeled matrix representation of an injective complex, but it might not be minimal. Indeed, when f is not injective, $f_*\eta$ might send some maximal vectors to non-zero vectors even if η did not. Therefore, we compute $\mathcal{I}f_*I^\bullet$ by relabeling and then minimizing.

Proposition 5.17. Given an injective complex $(I^\bullet, \eta^\bullet)$ on Π represented by labeled matrices, and an order preserving map $f : \Pi \rightarrow \Lambda$, we obtain $\mathcal{I}f_*I^\bullet$ by performing two steps:

1. change all labels π to $f(\pi)$ in all matrices η^d ,
2. minimize the complex as described in Section 4.3.

We obtain $\mathcal{I}f_*I^\bullet$ regardless of whether I^\bullet is minimal to start with.

5.4.3. Computing pullback

The pullback f^*I of an injective sheaf $I = \bigoplus_j [\lambda_j]$ on Λ might not be injective, but we can describe it as a restriction of an injective sheaf on the poset cylinder $\Pi \sqcup_f \Lambda$ defined above. This injective sheaf is $\hat{I} := \bigoplus_j [\lambda_j]$, where we now use the same symbol $[\lambda_j]$ for two different sheaves: once on Λ , once on $\Pi \sqcup_f \Lambda$. We claim that \hat{I} restricted to Π is f^*I . Indeed, for any $\pi \in \Pi$ we have

$$f^*I(\pi) \stackrel{\text{Def 5.4}}{=} I(f(\pi)) = \bigoplus_{\lambda_j \in \text{St}_\Lambda(f(\pi))}^j [\lambda_j](f(\pi)) \stackrel{\text{Lem 5.16}}{=} \bigoplus_{\lambda_j \in \text{St}_{\Pi \sqcup_f \Lambda}(\pi)}^j [\lambda_j](\pi) = \hat{I}(\pi).$$

The same construction works for the pullback of a natural transformation η . This allows us to describe $f^*\eta$ by the same labeled matrices as η .

Lemma 5.18. *Given an injective complex $(I^\bullet, \eta^\bullet)$ on Λ , the pullback f^*I^\bullet is the restriction of $(\hat{I}^\bullet, \hat{\eta}^\bullet)$ from $\Pi \sqcup_f \Lambda$ to Π . The same complex of labeled matrices represents I^\bullet and \hat{I}^\bullet .*

We describe how to compute a labeled matrix representation of $\mathcal{I}f^*I^\bullet$, given a labeled matrix representation of $(I^\bullet, \eta^\bullet)$. The key insight is a relation between a complex on the poset cylinder and a mapping cone of a complex morphism.

Consider now a complex $(J^\bullet, \delta^\bullet)$ on Π —it can also be viewed as a complex on $\Pi \sqcup_f \Lambda$ by setting $J^\bullet(\lambda) = 0$ for all $\lambda \in \Lambda$ —and a morphism $\hat{\alpha}^\bullet : \hat{I}^\bullet \rightarrow J^\bullet$. We construct a complex $(C^\bullet, \gamma^\bullet)$ on $\Pi \sqcup_f \Lambda$ by setting

$$C^d := \hat{I}^{d+1} \oplus J^d, \quad \gamma^d := \begin{pmatrix} -\hat{\eta}^{d+1} & 0 \\ \hat{\alpha}^{d+1} & \delta^d \end{pmatrix}.$$

Recalling Definition 4.2 together with Lemma 5.18 above gives a meaning to C^\bullet .

Lemma 5.19. *The restriction of $(C^\bullet, \gamma^\bullet)$ from $\Pi \sqcup_f \Lambda$ to Π is the mapping cone of the restriction of $\hat{\alpha}^\bullet$ to Π , which is a morphism $\alpha^\bullet : f^*I^\bullet \rightarrow J^\bullet$.*

This gives us a recipe to compute an injective resolution of f^*I^\bullet . We construct J^\bullet such that C^\bullet restricted to Π is exact. By Lemma 4.3, we then have α^\bullet a quasi-isomorphism, which means that J^\bullet is the resolution we are looking for. The construction itself is analogous to computing injective resolution of a sheaf in Section 4.5. We start with J^\bullet empty, and add indecomposable injective sheaves one by one to ensure exactness at each $\pi \in \Pi$ and for each degree d . The procedure is described as Algorithm 4. We remark that on line 6 we can be adding r instead of $-r$, since complexes given by η^\bullet and $-\eta^\bullet$ are quasi-isomorphic—we switch the sign to be consistent with our choice of signs in the definition of the mapping cone.

Proposition 5.20. *Given an injective complex I^\bullet on Λ and an order preserving map $f : \Pi \rightarrow \Lambda$, Algorithm 4 outputs a complex of labeled matrices representing $\mathcal{I}f^*I^\bullet$. The output is minimal regardless of whether I^\bullet is minimal.*

Proof. We use symbols as defined above the proposition. Since \hat{I}^\bullet only has indecomposables generated by elements in Λ , and all the newly added rows are labeled by Π , the outputted complex of labeled matrices, $\gamma^\bullet[\Pi, \Pi]$, represents the injective complex $(J^\bullet, \delta^\bullet)$ from Lemma 5.19. By its claim and Lemma 4.3, it is quasi-isomorphic to f^*I^\bullet iff C^\bullet is exact for all $\pi \in \Pi$. This is true by construction, which we argue as in Lemma 4.20: by Lemma 4.19, we have exactness right after the application of MAKEEXACT, and due to the induction order this does not change later in the algorithm.

We claim that the resulting complex is minimal (no matter whether I^\bullet was minimal). According to Section 4.3, we need $\delta^d[\pi, \pi]$ to be zero (or empty) for every element $\pi \in \Pi$ and degree d . Since $\delta^d[\pi, \pi] =$

Algorithm 4 Computing $\mathcal{I}f^*I^\bullet$.**Inputs:**

- η^{-1}, \dots, η^n labeled matrices representing a complex of injective sheaves $(I^\bullet, \eta^\bullet)$ on Λ
- $f: \Pi \rightarrow \Lambda$ order preserving map, with Π and Λ disjoint

Output: $\delta^{-1}, \dots, \delta^m$ labeled matrices representing complex $\mathcal{I}f^*(I^\bullet)$ on Π

Notation:

- I^d denotes the tuple of labels of columns of η^d , or equivalently rows of η^{d-1} .
 - C^d denotes the tuple of labels of columns of γ^d , or equivalently rows of γ^{d-1} .
- 1: $\gamma^{-2} \leftarrow$ labeled matrix with no columns, and row labels I^0 ▷ We are building complex $(C^\bullet, \gamma^\bullet)$ on $\Pi \sqcup_f \Lambda$
 - 2: $k \leftarrow -1$
 - 3: **while** $C^d \neq 0$ or $d < n$ **do**
 - 4: $\gamma^d \leftarrow$ labeled matrix with columns labeled by C^d
 - 5: **for each** r row in η^{d+1} **do**
 - 6: add a row $-r$ to γ^d ▷ Including the label; with 0 in columns not labeled by Λ
 - 7: **for** $\pi \in \Pi$ in non-increasing order **do**
 - 8: $\gamma^d \leftarrow \text{MAKEEXACT}(\gamma^{d-1}, \gamma^d, \pi)$ ▷ Algorithm 2
 - 9: $d \leftarrow d + 1$
 - 10: $m \leftarrow \max \{m \mid C^m \neq 0\}$
 - 11: **return** $\gamma^{-1}[\Pi, \Pi], \dots, \gamma^m[\Pi, \Pi]$ ▷ Note that $\gamma^{-2}[\Pi, \Pi]$ has no columns and no rows.

$\gamma^d[\pi, \pi]$, we can prove this for γ^d and $\pi \in \Pi$. All rows labeled by π were added by calling MAKEEXACT. By Lemma 4.19, if we added a row labeled by π in γ^{d-1} as the j -th row of $\gamma^{d-1}(\pi)$, then the canonical vector e_j is in the image of $\gamma^{d-1}(\pi)$. Since we have a complex, this implies $\gamma^d(\pi) \cdot e_j = 0$, which means that the j -th column of $\gamma^d(\pi)$ is 0. This argument is true for all columns labeled by π in γ^d , so we have $\gamma^d[\text{St } \pi, \pi] = 0$, and in particular $\gamma^d[\pi, \pi] = 0$. \square

Alternative minimization algorithm As an alternative to Algorithm 1 for turning a complex of injective sheaves into the minimal one, we can perform the minimization by computing $\mathcal{I}\text{id}^*I^\bullet$. This complex is minimal, because Algorithm 4 always returns a minimal complex, and it is quasi-isomorphic to I^\bullet by the functoriality of (derived) pullback.

5.4.4. Computing proper pushforward

Fix a locally closed subset $Z \xrightarrow{i_Z} \Pi$ and a complex $(I^\bullet, \eta^\bullet)$ of injective sheaves on Z , and recall that $\text{Cl } Z$ denotes the downwards closure of Z in Π . In Algorithm 5 we give a construction of $\mathcal{I}(i_Z)_! I^\bullet$. The construction is very similar to Algorithm 4—we again start with some prefilled labeled matrices, and add new rows to force exactness of the complex on a subset of Π . However, the interpretation of the constructed matrices is different. Whereas the constructed matrices in Algorithm 4 describe a mapping cone, in Algorithm 5 they describe directly the complex of injective sheaves we are looking for.

As noted above in Section 5.1.4, $(i_Z)_! I^\bullet$ is a subcomplex of $(i_Z)_* I^\bullet$, which is the complex represented by the same labeled matrices as I^\bullet , viewed as a complex over Π rather than Z . They differ on $\text{Cl } Z \setminus Z$, where the former is 0 while the latter is non-zero. The idea is to extend $(i_Z)_* I^\bullet$ to J^\bullet such that the inclusion $\alpha^\bullet: (i_Z)_! I^\bullet \hookrightarrow J^\bullet$ becomes a quasi-isomorphism. Consider $(C^\bullet, \gamma^\bullet)$, the mapping cone of α^\bullet , and recall that by Lemma 4.3 we want it to be an exact complex. For $\pi \in Z$ we can leave $J^\bullet(\pi) = (i_Z)_* I^\bullet(\pi)$, and $C^\bullet(\pi)$ is exact, since $\alpha^\bullet(\pi)$ is the identity. Same for $\pi \in \Pi \setminus \text{Cl } Z$ where all considered sheaves are zero. We need to fix the chains for $\pi \in \text{Cl } Z \setminus Z$, which we do by adding indecomposables $[\pi]$ via the MAKEEXACT algorithm. The procedure is described as Algorithm 5.

Proposition 5.21. *Given a locally closed subset $i_Z: Z \hookrightarrow \Pi$, and a minimal injective complex I^\bullet on Z , Algorithm 5 outputs a complex of labeled matrices representing $\mathcal{I}(i_Z)_! I^\bullet$.*

Proof. The complex returned by the algorithm is $(J^\bullet, \delta^\bullet)$. We argued above the proposition that $\alpha^\bullet: (i_Z)_! I^\bullet \hookrightarrow J^\bullet$ is an injection and its mapping cone, $(C^\bullet, \gamma^\bullet)$, is exact for all $\pi \in \Pi \setminus \text{Cl } Z$, as for those π we have $(i_Z)_! I^\bullet(\pi) = J^\bullet(\pi)$. The remaining elements are $\pi \in \text{Cl } Z \setminus Z$, which we loop through on line 7. Let $\pi \in \text{Cl } Z \setminus Z$. Then $(i_Z)_! I^\bullet(\pi)$ is the zero complex, and $\gamma^d(\pi) = \delta^d(\pi)$ for all d . The exactness of

Algorithm 5 Computing $\mathcal{I}(i_Z)_! I^\bullet$.**Inputs:**

- η^{-1}, \dots, η^n labeled matrices representing a complex of injective sheaves $(I^\bullet, \eta^\bullet)$ on Z
- Π a poset, and $Z \subseteq \Pi$ locally closed

Output: $\delta^{-1}, \dots, \delta^m$ labeled matrices representing complex $\mathcal{I}(i_Z)_! I^\bullet$ on Π **Notation:**

- I^d denotes the tuple of labels of columns of η^d , or equivalently rows of η^{d-1} .
- J^d denotes the tuple of labels of columns of δ^d , or equivalently rows of δ^{d-1} .

1: $\delta^{-1} \leftarrow$ labeled matrix with no columns and row labels I^0 2: $d \leftarrow 0$ 3: **while** $J^d \neq 0$ or $d \leq n$ **do**4: $\delta^d \leftarrow$ labeled matrix with columns labeled by J^d 5: **for each** r row in η^d **do**6: add a row r to δ^d ▷ Including the label; with 0 in columns not labeled by Z 7: **for** $\pi \in \text{Cl } Z \setminus Z$ in non-increasing order **do**8: $\delta^d \leftarrow \text{MAKEEXACT}(\delta^{d-1}, \delta^d, \pi)$

▷ Algorithm 2

9: $d \leftarrow d + 1$ 10: $m \leftarrow \max \{m \mid J^m \neq 0\}$ 11: **return** $\delta^{-1}, \dots, \delta^m$

$\delta^\bullet(\pi)$ in degree d is forced by the construction after processing it—see Lemma 4.19—and it does not change afterwards due to the order in which we process the elements.

To show minimality of J^\bullet , we argue that $\delta^d[\pi, \pi] = 0$ for each $\pi \in \Pi$ and degree d . This is trivial for $\pi \in \Pi \setminus \text{Cl } Z$, follows from the minimality of I^\bullet for $\pi \in Z$, and for $\pi \in \text{Cl } Z \setminus Z$ we follow the same argument as for Algorithm 4 in Proposition 5.20. \square

5.4.5. Computing proper pullback

By Lemma 5.8, to compute the proper pullback by an inclusion of a locally closed subset Z , we just need to throw away the indecomposables $[\pi]$ for $\pi \notin Z$.

Proposition 5.22. *Given a locally closed subset $i_Z : Z \hookrightarrow \Pi$, and a minimal injective complex $(I^\bullet, \eta^\bullet)$ on Z represented by labeled matrices, the proper pullback $\mathcal{I}i_Z^! I^\bullet$ is represented by matrices $\eta^d[Z, Z]$.*

An interesting formula Illustrating the utility of our algorithmic treatment of derived categories of sheaves, we make the following observation, which leads to an interesting formula for finite posets. We noted that Algorithm 4 and Algorithm 5 are essentially performing the same procedure. One difference is that in the former we start the construction of the d -th matrix by adding rows from η^{d+1} , while in the latter we add rows from η^d . The second difference is that in Algorithm 4 we take submatrices at the end—this can be interpreted as taking a proper pullback of the complex. Together, this yields an interesting relation between derived functors on posets.

Proposition 5.23. *Let $f : \Pi \rightarrow \Lambda$ be an order preserving map, $p : \Pi \hookrightarrow \Pi \sqcup_f \Lambda$ and $\ell : \Lambda \hookrightarrow \Pi \sqcup_f \Lambda$ be the natural inclusions to the mapping cylinder of f , and let I^\bullet be a complex of injective sheaves on Λ . Then*

$$\mathcal{I}f^* I^\bullet = \mathcal{I}p^! \mathcal{I}\ell_! I^{\bullet+1},$$

where $I^{\bullet+1}$ denotes the complex with $(I^{\bullet+1})^d = I^{d+1}$.

6. Microlocal sheaf theory and discrete Morse theory

In this section, we establish a microlocal generalization of two classical homological results of discrete Morse theory (Theorem 6.7 and Theorem 6.8). These generalizations show that the fibers—preimages of singletons—of an order preserving map between posets which lie in the discrete microsupport of a given sheaf correspond to non-isomorphisms between sheaf hypercohomology groups of adjacent sub-level or super-level

sets. We begin by introducing the discrete microsupport of a complex of sheaves and illustrating some of its basic properties.

6.1. Discrete microsupport

Microsupport was introduced and studied by Sato in the context of systems of linear differential equations and smooth manifolds [33,47,50]. In this setting, the microsupport of a sheaf on a manifold is a conic subset of the cotangent bundle which describes ‘directions’ in which certain homological properties of the sheaf will propagate. This is a natural and powerful invariant to consider, because, among other things, it naturally separates a topological space into strata in which the local properties of a given sheaf are ‘locally constant’. Our goal here is to translate this theory to finite posets, where vectors in the cotangent bundle are replaced by poset relations $\tau < \sigma$. We think of this analogy geometrically by imagining that our poset is the face poset of a triangulation of a smooth manifold. Here the face relation $\sigma < \tau$ corresponds to a family of cotangent vectors, based at a point x in the geometric realization of τ , and locally pointing in the ‘direction’ of the geometric realization of τ . In fact, we hope to make this analogy concrete in future work, by explicitly describing the correspondence between the microsupport of a sheaf on a manifold and the discrete microsupport of that sheaf on the face poset of an appropriate triangulation of the manifold.

Recall that $D^b(\Pi)$ denotes the (skeleton of) derived category whose objects are minimal injective complexes over Π —and that for any complex there exists a unique quasi-isomorphic minimal injective complex. We first define the (non-microlocal) support of a sheaf on a poset, using the tools introduced in Section 5.1.2, 5.1.4 and 5.2—the pullback, proper pullback, and hypercohomology.

Definition 6.1. Let Π be a finite poset, $I^\bullet \in D^b(\Pi)$, and $i_\tau : \{\tau\} \hookrightarrow \Pi$ be the inclusion map. The $*$ -support of $I^\bullet \in D^b(\Pi)$ is defined as

$$\operatorname{supp}^* I^\bullet := \{\tau \in \Pi \mid \mathbb{H}^j \mathcal{I} i_\tau^* I^\bullet \neq 0 \text{ for some } j \in \mathbb{Z}\}.$$

The $!$ -support of $I^\bullet \in D^b(\Pi)$ is defined as

$$\operatorname{supp}^! I^\bullet := \{\tau \in \Pi \mid \mathbb{H}^j \mathcal{I} i_\tau^! I^\bullet \neq 0 \text{ for some } j \in \mathbb{Z}\}.$$

Proposition 6.2. Let $I^\bullet \in D^b(\Pi)$ be a minimal injective complex, and $m_{I^\bullet}^d(\pi)$ be the multiplicity of the indecomposable injective sheaf $[\pi]$ in I^d . Then

$$\pi \notin \operatorname{supp}^! I^\bullet \text{ if and only if } m_{I^\bullet}^d(\pi) = 0 \text{ for all } d \in \mathbb{Z}.$$

Moreover,

$$\operatorname{Cl} \operatorname{supp}^! I^\bullet = \operatorname{Cl} \operatorname{supp}^* I^\bullet,$$

where $\operatorname{Cl} \operatorname{supp}^* I^\bullet$, $\operatorname{Cl} \operatorname{supp}^! I^\bullet$ denote the downwards closure of $\operatorname{supp}^* I^\bullet$, $\operatorname{supp}^! I^\bullet$, respectively, corresponding to closure with respect to the Alexandrov topology (Definition 2.1).

Proof. The first statement is a direct consequence of Corollary 5.11. We begin the proof of the second statement by showing that $\operatorname{supp}^! I^\bullet \subseteq \operatorname{Cl} \operatorname{supp}^* I^\bullet$. Suppose that $m_{I^\bullet}^d(\pi) \neq 0$ for some $d \in \mathbb{Z}$. If $\pi \in \operatorname{supp}^* I^\bullet$, then there is nothing to prove. Assume that $\pi \notin \operatorname{supp}^* I^\bullet$. Because $m_{I^\bullet}^d(\pi) \neq 0$, $I^d(\pi)$ contains a nonzero maximal vector, which we denote by x . Because $\pi \notin \operatorname{supp}^* I^\bullet$, $\mathbb{H}^d(I^\bullet(\pi)) = 0$. Therefore, there exists a vector $v \in I^{d-1}(\pi)$ which maps onto x . Because I^\bullet is assumed to be minimal, v is not maximal. Choose $\tau \geq \pi$ such that $I^{d-1}(\pi \leq \tau)(v)$ is a maximal vector. Then $n_{I^\bullet}^{d-1}(\tau) \neq 0$. If $\tau \in \operatorname{supp}^* I^\bullet$, then $\pi \in \operatorname{Cl} \operatorname{supp}^* I^\bullet$.

Otherwise, we can apply the same argument as above to produce a new simplex $\sigma \geq \tau$ such that $n_{I^\bullet}^{d-2}(\sigma) \neq 0$. Because the complex is bounded (and the poset Π is finite), this process must terminate with an element $\omega \geq \pi$ such that $n_{I^\bullet}^d(\omega) \neq 0$ and $\omega \in \operatorname{supp}^* I^\bullet$, which proves that $\pi \in \operatorname{Cl} \operatorname{supp}^* I^\bullet$.

Second, we prove the inclusion $\operatorname{supp}^* I^\bullet \subseteq \operatorname{Cl} \operatorname{supp}^! I^\bullet$. Suppose $\tau \in \operatorname{supp}^* I^\bullet$. By the definition of i_τ^* , if $\mathbb{H}^d \mathcal{I} i_\tau^* I^\bullet \neq 0$ for some $d \in \mathbb{Z}$, then there exists $\sigma \geq \tau$ such that $m_{I^\bullet}^d(\sigma) \neq 0$. Therefore, $\sigma \in \operatorname{supp}^! I^\bullet$, and $\tau \in \operatorname{Cl} \operatorname{supp}^! I^\bullet$. \square

Remark. The support is traditionally defined to be $\operatorname{Cl} \operatorname{supp}^* I^\bullet$, and is therefore, by definition, closed. As a consequence, there is usually no distinction between the $!$ -support and the $*$ -support. However, for the Morse theoretic results below, we find it useful to keep track of these differences.

We now introduce a novel definition of discrete microsupport, which is a finer invariant than the support(s) defined above. The elements of a microsupport are locally closed subsets of Π rather than its elements.

Definition 6.3. Let Π be a finite poset and $I^\bullet \in D^b(\Pi)$. For a locally closed subset Z of Π , denote the inclusion map by $i_Z : Z \hookrightarrow \Pi$. The *discrete $*$ -microsupport* of I^\bullet is

$$\mu \operatorname{supp}^* I^\bullet := \{Z \text{ locally closed in } \Pi \mid \mathbb{H}^d(\mathcal{I} i_{Z!} \mathcal{I} i_Z^* I^\bullet) \neq 0 \text{ for some } d \in \mathbb{Z}\},$$

and the *discrete $!$ -microsupport* of I^\bullet is

$$\mu \operatorname{supp}^! I^\bullet := \{Z \text{ locally closed in } \Pi \mid \mathbb{H}^d(\mathcal{I} i_Z^! I^\bullet) \neq 0 \text{ for some } d \in \mathbb{Z}\}.$$

By Lemma 5.14, we can equivalently define the discrete microsupport by

$$\begin{aligned} \mu \operatorname{supp}^* I^\bullet &:= \left\{ Z \text{ loc. closed in } \Pi \mid \mathbb{H}^d(\mathcal{I} i_{\operatorname{Cl} Z}^* I^\bullet) \xrightarrow{\sim} \mathbb{H}^d(\mathcal{I} i_{\operatorname{Cl} Z - Z}^* I^\bullet) \text{ for some } d \in \mathbb{Z} \right\} \text{ and} \\ \mu \operatorname{supp}^! I^\bullet &:= \left\{ Z \text{ loc. closed in } \Pi \mid \mathbb{H}^d(\mathcal{I} i_{\operatorname{St} Z}^* I^\bullet) \xrightarrow{\sim} \mathbb{H}^d(\mathcal{I} i_{\operatorname{St} Z - Z}^* I^\bullet) \text{ for some } d \in \mathbb{Z} \right\}. \end{aligned}$$

We use the functors $\mathcal{I} i^*$, $\mathcal{I} i_!$, and $\mathcal{I} i^!$ in the above definition to highlight its computability—indeed, we can easily compute the discrete microsupports defined here with our implementation of the algorithms [3], as showcased in Section 7. For an arbitrary object in the bounded derived category $D^b(k_\Pi)$ as defined, e.g., in [33], one can define the discrete microsupport by replacing $\mathcal{I} i^*$ with i^* , $\mathcal{I} i_!$ with the right derived functor $Ri_!$, and $\mathcal{I} i^!$ with the left adjoint of $Ri_!$.

In a recent work [48], Schapira introduces the propagation-set of a derived sheaf as a means of generalizing the microsupport to presheaves on presites and ind-sheaves. In our setting, the propagation-set admits a simple definition

$$\begin{aligned} \operatorname{Prop} I^\bullet &= \{\sigma \leq \tau \in \Pi \mid \text{for all } \sigma \leq \gamma \leq \tau, \operatorname{St} \sigma - \operatorname{St} \gamma \notin \mu \operatorname{supp}^! I^\bullet \text{ and} \\ &\quad \operatorname{St} \gamma - \operatorname{St} \tau \notin \mu \operatorname{supp}^! I^\bullet\}. \end{aligned}$$

This definition suggests that a more restrictive version of the discrete microsupport might have more desirable theoretical properties (at the expense of increased computational complexity). The more restrictive definition would consider a locally closed set $Z \notin \mu \operatorname{supp}^! I^\bullet$ if and only if $\mathbb{H}^d(\mathcal{I} i_W^! I^\bullet) = 0$ for each locally closed subset W of Z (similarly for $\mu \operatorname{supp}^* I^\bullet$). However, we choose to continue with the less restrictive definition of discrete microsupport, which is sufficient for our results on computational sheaf theory and discrete microlocal Morse theory.

Remark. It is immediate that $\pi \in \operatorname{supp}^! I^\bullet$ if and only if $\{\pi\} \in \mu \operatorname{supp}^! I^\bullet$. Therefore, the discrete $!$ -microsupport generalizes the $!$ -support in an obvious way. However, the discrete $*$ -microsupport is strictly a generalization of the $*$ -support only if Π has the following property: for each $\pi \in \Pi$, if the boundary, $\operatorname{Cl} \pi - \pi$, of π is non-empty, then the hypercohomology of the injective resolution of the constant sheaf on the boundary is non-zero. This happens, for example, whenever Π is a simplicial complex. In general, $\{\pi \in \Pi : \{\pi\} \in \mu \operatorname{supp}^* I^\bullet\} \subset \operatorname{supp}^* I^\bullet$ (for an example, consider the constant sheaf on a totally ordered poset with two elements).

The first result we present ties together $!$ -microsupport with linear maps in cohomology sheaves (recall Definition 2.11).

Theorem 6.4. *If $\operatorname{St} \sigma - \operatorname{St} \tau \notin \mu \operatorname{supp}^! I^\bullet$, for $\sigma \leq \tau \in \Pi$, then the linear maps $H^d(I^\bullet)(\sigma \leq \tau)$ are isomorphisms for each $d \in \mathbb{Z}$.*

Proof. Let $\sigma \leq \tau \in \Pi$. The linear map $H^d(I^\bullet)(\sigma \leq \tau)$ is, by definition,

$$H^d(I^\bullet)(\sigma \leq \tau) : \mathbb{H}^d(i_\sigma^* I^\bullet) \rightarrow \mathbb{H}^d(i_\tau^* I^\bullet),$$

where $i_\gamma : \operatorname{St} \gamma \hookrightarrow \Pi$. Because $\operatorname{St} \tau$ is an open subset of $\operatorname{St} \sigma$ for $\sigma \leq \tau$, we can apply Lemma 5.14 to get a long exact sequence

$$\begin{aligned} \cdots \rightarrow \mathbb{H}^d(\mathcal{I}i_Z^! I^\bullet) &\rightarrow \mathbb{H}^d(i_\sigma^* I^\bullet) \rightarrow \mathbb{H}^d(i_\tau^* I^\bullet) \\ &\rightarrow \mathbb{H}^{d+1}(\mathcal{I}i_Z^! I^\bullet) \rightarrow \cdots \end{aligned}$$

where $Z = \operatorname{St} \sigma - \operatorname{St} \tau$. Because $Z \notin \mu \operatorname{supp}^! I^\bullet$, $\mathbb{H}^d(\mathcal{I}i_Z^! I^\bullet) = 0$ for all $d \in \mathbb{Z}$. Therefore, $H^d(I^\bullet)(\sigma \leq \tau)$ is an isomorphism for each $d \in \mathbb{Z}$. \square

An immediate consequence of the above theorem is the following:

Corollary 6.5. *If $\mu \operatorname{supp}^! I^\bullet$ contains no sets of the form $\operatorname{St} \sigma - \operatorname{St} \tau$, then the cohomology sheaves $H^d(I^\bullet)$ are locally constant.*

This shows that the discrete microsupport can be used to separate our poset into ‘strata’ over which cohomology sheaves are locally constant. We think a natural follow-up to this paper could use this fact to study sheaf theoretic stratification algorithms from the derived perspective, generalizing/building on prior work of Nanda and Brown-Wang [11,43].

6.2. The discrete microlocal Morse theorem and inequalities

Discrete Morse theory identifies topologically relevant parts of discrete objects and lets us skip parts that are “essentially contractible”. A filtration function on a simplicial complex, $f : \Sigma \rightarrow \mathbb{R}$, is called *generalized discrete Morse* if it is order preserving and every non-empty level set, $f^{-1}(r)$, is an interval—a set of the form $[\mu, \tau] = \{\sigma \in \Sigma \mid \mu \leq \sigma \leq \tau\}$; see, e.g., [4,21]. The singletons in the level set decomposition of Σ are called *critical simplices*, and their images *critical values*. The critical values identify changes in cohomology. Namely, comparing growing sublevel sets $f^{-1}((-\infty, s])$ and $f^{-1}((-\infty, r])$, their cohomology (even homotopy type) is the same if the interval $(s, r]$ contains no critical values. Conversely, it always differs if there is exactly one critical value in $(s, r]$. The theorems below generalize the homological claim. Both Σ and \mathbb{R} are replaced by general (finite) posets. Intervals are replaced by locally closed sets, which allows us to

consider any order preserving map f . This relaxation is counterbalanced by a more complicated structure of critical elements.

The results of this section are discrete analogues of work by Kashiwara [30], Schapira–Tose [52], and Witten [53]. In fact, the proofs of Theorem 6.7 and Theorem 6.8 closely follow the main ideas presented in [30] and [52].

Definition 6.6. Suppose that $f : \Lambda \rightarrow \Pi$ is an order preserving map of finite posets. Let $I^\bullet \in D^b(\Lambda)$. An element $\pi \in \Pi$ is

- $(I^\bullet, f, *)$ -critical if $f^{-1}(\pi) \in \mu \operatorname{supp}^* I^\bullet$,
- $(I^\bullet, f, !)$ -critical if $f^{-1}(\pi) \in \mu \operatorname{supp}^! I^\bullet$.

Theorem 6.7 (Discrete Microlocal Morse Theorem). Let $f : \Lambda \rightarrow \Pi$ be an order preserving map of finite posets. For $\pi \in \Pi$, let

$$i_{\leq \pi} : \{\sigma \in \Lambda : f(\sigma) \leq \pi\} \hookrightarrow \Lambda, \quad i_{\geq \pi} : \{\sigma \in \Lambda : f(\sigma) \geq \pi\} \hookrightarrow \Lambda$$

be inclusion maps of sublevel-sets and superlevel-sets of f . For $a \leq b \in \Pi$, let

$$(a, b] = \{z \in \Pi : a < z \leq b\}.$$

For $I^\bullet \in D^b(\Lambda)$, if $(a, b]$ contains no $(I^\bullet, f, !)$ -critical elements, then the morphisms (induced by inclusions)

$$\mathbb{H}^d(\mathcal{I}i_{\leq b}^! I^\bullet) \rightarrow \mathbb{H}^d(\mathcal{I}i_{\leq a}^! I^\bullet) \quad \text{and} \quad (6.1)$$

$$\mathbb{H}^d(\mathcal{I}i_{\geq a}^* I^\bullet) \rightarrow \mathbb{H}^d(\mathcal{I}i_{\geq b}^* I^\bullet) \quad (6.2)$$

are isomorphisms for each $d \in \mathbb{Z}$. If $(a, b]$ contains no $(I^\bullet, f, *)$ -critical elements, then the morphisms

$$\mathbb{H}^d(\mathcal{I}i_{\leq b}^* I^\bullet) \rightarrow \mathbb{H}^d(\mathcal{I}i_{\leq a}^* I^\bullet) \quad (6.3)$$

are isomorphisms for each $d \in \mathbb{Z}$.

Proof. We begin by describing each of the morphisms of the theorem in more detail. By Proposition 5.22, computing $\mathbb{H}^d(\mathcal{I}i_{\leq b}^! I^\bullet)$ is quite simple. First, we consider the complex of labeled submatrices of I^\bullet with labels in the set $\{\sigma \in \Lambda : f(\sigma) \leq b\}$. Then we forget the labels and take the d -cohomology group of the complex of vector spaces these matrices represent. In this basis, morphism (6.1) is the map induced on cohomology groups by the projection map from the complex of injective subsheaves of I^\bullet generated by maximal vectors which are supported in the set $\{\sigma \in \Lambda : f(\sigma) \leq b\}$ to those with maximal vectors supported in the set $\{\sigma \in \Lambda : f(\sigma) \leq a\}$. Morphism (6.2) is similarly induced by projection maps. Morphism (6.3) is slightly more involved. First, let $I^\bullet|_\pi$ be the restriction of I^\bullet to $\{\sigma \in \Lambda : f(\sigma) \leq \pi\}$. (Note that $I^\bullet|_\pi$ is not necessarily a complex of injective sheaves). There is a morphism of complexes of sheaves from $I^\bullet|_b$ to $I^\bullet|_a$ (where both complexes are viewed as complexes of sheaves on Λ by extending sheaves by zero outside of their support) which is the identity map on the set $\{\sigma \in \Lambda : f(\sigma) \leq a\}$ and zero elsewhere. This morphism of complexes of sheaves induces a morphism between the corresponding injective resolutions, which in turn induces morphism (6.3).

We prove that (6.1), (6.2), and (6.3) are isomorphisms by realizing each as a composition isomorphisms. Specifically, there exists $\{a_i\}_{i=0}^k \subset \Pi$ such that $a = a_0 < a_1 < a_2 < \cdots < a_k = b$ and $(a_i, a_{i+1}] = \{a_{i+1}\}$. The morphism (6.1) is equal to the composition of maps of the form

$$\mathbb{H}^d \left(\mathcal{I}i_{\leq a_{i+1}}^! I^\bullet \right) \rightarrow \mathbb{H}^d \left(\mathcal{I}i_{\leq a_i}^! I^\bullet \right),$$

because each is induced by a projection map described in the paragraph above. (A similar argument applies to (6.2) and (6.3)). We proceed by proving that for each pair $a_i < a_{i+1}$, the maps described above are isomorphisms. To make the notation easier to follow, we simply let $a = a_i$ and $b = a_{i+1}$ for the remainder of the proof.

We begin with equation (6.1). Because $f_{\leq a}^{-1}$ is a closed subset of $f_{\leq b}^{-1}$, by Lemma 5.14, and the fact that $\mathcal{I}i_{\leq a}^* \mathcal{I}i_{\leq b}^* I^\bullet = \mathcal{I}i_{\leq a}^* I^\bullet$, we have a long exact sequence of cohomology groups

$$\begin{aligned} \cdots \rightarrow \mathbb{H}^d \left(\mathcal{I}i_{\leq a}^! I^\bullet \right) &\rightarrow \mathbb{H}^d \left(\mathcal{I}i_{\leq b}^! I^\bullet \right) \rightarrow \mathbb{H}^d \left(\mathcal{I}i_{f^{-1}(b)}^! I^\bullet \right) \\ &\rightarrow \mathbb{H}^{d+1} \left(\mathcal{I}i_{\leq a}^! I^\bullet \right) \rightarrow \cdots \end{aligned}$$

By the definition of $(I^\bullet, f, !)$ -critical elements and $\mu \operatorname{supp}^! I^\bullet$, $\mathbb{H}^d \left(\mathcal{I}i_{f^{-1}(b)}^! I^\bullet \right) = 0$. Therefore,

$$\mathbb{H}^d \left(\mathcal{I}i_{\leq a}^! I^\bullet \right) \rightarrow \mathbb{H}^d \left(\mathcal{I}i_{\leq b}^! I^\bullet \right)$$

is an isomorphism for each $d \in \mathbb{Z}$.

We now turn to equation (6.2). Because $f_{\geq a}^{-1}$ is an open subset of Λ ,

$$\mathcal{I}i_{f^{-1}(b)}^! \mathcal{I}i_{\geq a}^* I^\bullet = \mathcal{I}i_{f^{-1}(b)}^! \mathcal{I}i_{\geq a}^! I^\bullet = \mathcal{I}i_{f^{-1}(b)}^! I^\bullet,$$

where $i_{f^{-1}(b)} : f^{-1}(b) \hookrightarrow f_{\geq a}^{-1}$. Moreover, because $f_{\geq b}^{-1}$ is an open subset of $f_{\geq a}^{-1}$, Lemma 5.14 yields

$$\begin{aligned} \cdots \rightarrow \mathbb{H}^d \left(\mathcal{I}i_{f^{-1}(b)}^! I^\bullet \right) &\rightarrow \mathbb{H}^d \left(\mathcal{I}i_{\geq a}^* I^\bullet \right) \rightarrow \mathbb{H}^d \left(\mathcal{I}i_{\geq b}^* I^\bullet \right) \\ &\rightarrow \mathbb{H}^{d+1} \left(\mathcal{I}i_{f^{-1}(b)}^! I^\bullet \right) \rightarrow \cdots \end{aligned}$$

Again, by the definition of $(I^\bullet, f, !)$ -critical elements and $\mu \operatorname{supp}^! I^\bullet$, $\mathbb{H}^d \left(\mathcal{I}i_{f^{-1}(b)}^! I^\bullet \right) = 0$, and equation (6.2) follows.

Finally, we turn to equation (6.3). Because $f^{-1}(b)$ is open in $f_{\leq b}^{-1}$, Lemma 5.14 gives us the long exact sequence

$$\begin{aligned} \cdots \rightarrow \mathbb{H}^d \left(\mathcal{I}i_{f^{-1}(b)}^! \mathcal{I}i_{f^{-1}(b)}^* I^\bullet \right) &\rightarrow \mathbb{H}^d \left(\mathcal{I}i_{\leq b}^* I^\bullet \right) \rightarrow \mathbb{H}^d \left(\mathcal{I}i_{\leq a}^* I^\bullet \right) \\ &\rightarrow \mathbb{H}^{d+1} \left(\mathcal{I}i_{f^{-1}(b)}^! \mathcal{I}i_{f^{-1}(b)}^* I^\bullet \right) \rightarrow \cdots \end{aligned} \quad (6.4)$$

where $i_{f^{-1}(b)} : f^{-1}(b) \hookrightarrow f_{\leq b}^{-1}$. By the definition of $(I^\bullet, f, *)$ -critical elements and $\mu \operatorname{supp}^* I^\bullet$,

$$\mathbb{H}^d \left(\mathcal{I}j_{f^{-1}(b)}^! \mathcal{I}j_{f^{-1}(b)}^* I^\bullet \right) = 0,$$

and equation (6.3) follows. \square

Theorem 6.8 (*Discrete Microlocal Morse Inequality*). *Let $\kappa^! \subset \Pi$ (and $\kappa^* \subset \Pi$) denote the set of $(I^\bullet, f, !)$ -critical elements in Π (and $(I^\bullet, f, *)$ -critical elements, respectively). For each $\ell \in \mathbb{Z}$,*

$$(-1)^\ell \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j(I^\bullet) \leq (-1)^\ell \sum_{\pi \in \kappa^!} \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j \left(\mathcal{I}i_{f^{-1}(\pi)}^! I^\bullet \right), \text{ and} \quad (6.5)$$

$$(-1)^\ell \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j(I^\bullet) \leq (-1)^\ell \sum_{\pi \in \kappa^*} \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j \left(\mathcal{I}i_{f^{-1}(\pi)}^! \mathcal{I}i_{f^{-1}(\pi)}^* I^\bullet \right). \quad (6.6)$$

Moreover, the inequalities become equalities when we sum over all $j \in \mathbb{Z}$:

$$\chi I^\bullet = \sum_{\pi \in \kappa^!} \chi \mathcal{I} i_{f^{-1}(\pi)}^! I^\bullet = \sum_{\pi \in \kappa^*} \chi \mathcal{I} i_{f^{-1}(\pi)}^! \mathcal{I} i_{f^{-1}(\pi)}^* I^\bullet. \quad (6.7)$$

Proof. Suppose

$$\cdots \rightarrow V_C^{d-1} \rightarrow V_A^d \rightarrow V_B^d \rightarrow V_C^d \rightarrow V_A^{d+1} \rightarrow V_B^{d+1} \rightarrow V_C^{d+1} \rightarrow V_A^{d+2} \rightarrow \cdots$$

is a long exact sequence of vector spaces. Because the sequence is exact,

$$\sum_{j \in \mathbb{Z}} (-1)^j (\dim V_A^j - \dim V_B^j + \dim V_C^j) = 0.$$

Moreover, the sequence

$$\cdots \rightarrow V_C^{\ell-1} \rightarrow V_A^\ell \rightarrow V_B^\ell \xrightarrow{\delta} V_C^\ell \rightarrow \operatorname{coker} \delta \rightarrow 0 \rightarrow \cdots$$

is exact, and therefore

$$(-1)^\ell \sum_{j \leq \ell} (-1)^j (\dim V_A^j - \dim V_B^j + \dim V_C^j) = \dim \operatorname{coker} \delta \geq 0.$$

Rearranging the above equations gives

$$\begin{aligned} (-1)^\ell \sum_{j \leq \ell} (-1)^j \dim V_B^j &\leq (-1)^\ell \sum_{j \leq \ell} (-1)^j (\dim V_A^j + \dim V_C^j), \text{ and} \\ \sum_{j \in \mathbb{Z}} (-1)^j \dim V_B^j &= \sum_{j \in \mathbb{Z}} (-1)^j (\dim V_A^j + \dim V_C^j). \end{aligned}$$

Substituting the long exact sequences of Lemma 5.14 to the above inequalities, we have, for each $a \leq b \in \Pi$ such that $(a, b] = \{b\}$,

$$\begin{aligned} (-1)^\ell \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j(\mathcal{I} i_{\leq b}^! I^\bullet) &\leq (-1)^\ell \sum_{j \leq \ell} (-1)^j (\dim \mathbb{H}^j(\mathcal{I} i_{\leq a}^! I^\bullet) + \dim \mathbb{H}^j(\mathcal{I} i_{f^{-1}(b)}^* I^\bullet)), \\ \text{and } \sum_{j \in \mathbb{Z}} (-1)^j \dim \mathbb{H}^j(\mathcal{I} i_{\leq b}^! I^\bullet) &= \sum_{j \in \mathbb{Z}} (-1)^j (\dim \mathbb{H}^j(\mathcal{I} i_{\leq a}^! I^\bullet) + \dim \mathbb{H}^j(\mathcal{I} i_{f^{-1}(b)}^* I^\bullet)). \end{aligned}$$

If $b \in \Pi$ is not $(I^\bullet, f, !)$ -critical, then $\dim \mathbb{H}^j(\mathcal{I} i_{f^{-1}(b)}^* I^\bullet) = 0$, which implies

$$\begin{aligned} (-1)^\ell \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j(\mathcal{I} i_{\leq b}^! I^\bullet) &\leq (-1)^\ell \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j(\mathcal{I} i_{\leq a}^! I^\bullet), \text{ and} \\ \sum_{j \in \mathbb{Z}} (-1)^j \dim \mathbb{H}^j(\mathcal{I} i_{\leq b}^! I^\bullet) &= \sum_{j \in \mathbb{Z}} (-1)^j \dim \mathbb{H}^j(\mathcal{I} i_{\leq a}^! I^\bullet). \end{aligned}$$

By choosing a total order on Π which respects the given partial order, and inductively applying the previous inequalities, we get

$$\begin{aligned} (-1)^\ell \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j(I^\bullet) &\leq \sum_{\pi \in \kappa^!} (-1)^\ell \sum_{j \leq \ell} (-1)^j \dim \mathbb{H}^j(\mathcal{I} i_{f^{-1}(\pi)}^* I^\bullet), \text{ and} \\ \chi I^\bullet &= \sum_{\pi \in \kappa^!} \chi \mathcal{I} i_{f^{-1}(\pi)}^* I^\bullet. \end{aligned}$$

The remaining (in)equalities follow from similar arguments. \square

6.3. Reduction to classical discrete Morse theory

We will conclude this section by illustrating how the prior results generalize classical discrete Morse theory. Let Σ be a finite simplicial complex, and $I_{k\Sigma}^\bullet \in D^b(\Sigma)$ the minimal injective resolution of the constant sheaf on Σ . Let $f : \Sigma \rightarrow \mathbb{R}$ be a generalized discrete Morse function, i.e., an order preserving map such that each level set, $f^{-1}(x)$, is an interval $[\mu, \tau] = \{\sigma \in \Sigma \mid \mu \leq \sigma \leq \tau\}$. For the codomain to be a finite poset, replace \mathbb{R} by the image of f denoted by Π . We have a map $f : \Sigma \rightarrow \Pi$. Let $f_c^{-1} := \{\sigma \in \Sigma : f(\sigma) \leq c\}$ be the sublevel set of f for any $c \in \Pi$.

Proposition 6.9. *The element $c \in \Pi$ is $(I_{k\Sigma}^\bullet, f, *)$ -critical if and only if $f^{-1}(c)$ is a singleton.*

Proof. By assumption $f^{-1}(c) = [\mu, \tau]$. Let $\text{Cl}\tau$ denote the closure of τ in Σ . Because $I_{k\Sigma}^\bullet$ is an injective resolution of the constant sheaf on Σ , the pullback $\mathcal{I}i_{f^{-1}(c)}^* I_{k\Sigma}^\bullet$ is an injective resolution of the constant sheaf on $f^{-1}(c)$ and we have the following isomorphisms

$$\mathbb{H}^d \mathcal{I}(i_{f^{-1}(c)})^* \mathcal{I}i_{f^{-1}(c)}^* I_{k\Sigma}^\bullet \cong H_c^d([\mu, \tau]; k) \cong H_d(\text{Cl}\tau, \text{Cl}\tau - \text{St}\mu; k)$$

between hypercohomology, compactly supported cohomology, and relative homology for each d . Because $\text{Cl}\tau$ is a closed simplex, the hypercohomology groups vanish if and only if $\mu \neq \tau$. \square

The right derived pullback $\mathcal{I}i_c^* I_{k\Sigma}^\bullet$ of the injective resolution of the constant sheaf on Σ by the inclusion map $i_c : f_c^{-1} \hookrightarrow \Sigma$, is an injective resolution of the constant sheaf on f_c^{-1} , and the hypercohomology groups are isomorphic to the usual simplicial cohomology groups:

$$\mathbb{H}^d(\mathcal{I}i_c^* I_{k\Sigma}^\bullet) = H^d(\mathcal{I}p_* \mathcal{I}i_c^* I_{k\Sigma}^\bullet) \cong H^d(f_c^{-1}; k).$$

Therefore, in this context, Theorem 6.7 states that if $\#f^{-1}(b) > 1$ for each element $b \in \Pi$ with $a < b \leq c$, then the natural morphisms between simplicial cohomology groups

$$H^d(f_c^{-1}; k) \rightarrow H^j(f_a^{-1}; k)$$

are isomorphisms for each $d \in \mathbb{Z}$.

7. Examples

In this section we compute simple examples which illustrate the benefits of working with the derived and microlocal perspectives. We will compute three different complexes of sheaves, each obtained via a triangulated (simplicial) map, g , h , and l , from $S^2 \vee S^1$ and D^2 to S^2 . We will then illustrate the discrete microlocal Morse theorem for each of these complexes of sheaves, relative to a discrete Morse function, f , on the triangulation poset Λ of S^2 . All computations in this section are done over $k = \mathbb{Z}_2$.

We consider triangulations of two maps on the wedge sum of a sphere and a circle, with values in a sphere. Both maps send S^2 identically to S^2 . The first map, g , sends the circle to the point where it touches the sphere, and h sends it to the equator. We start by triangulating the spaces as shown in Fig. 5. This gives us two posets, and two order preserving maps, $g, h : \Sigma \rightarrow \Lambda$. Both maps are determined by mapping vertices. Both send $0, 1, 2, 3, 4$ identically, and then g sends $5, 6 \xrightarrow{g} 4$, whereas h sends $5 \xrightarrow{h} 1$ and $6 \xrightarrow{h} 3$. The third map we consider, l , from the closed disk to the sphere, is obtained by identifying all points along




Fig. 5. (a) Triangulations of $S^2 \vee S^1$, the closed disk D^2 , and S^2 . The face relations give rise to the posets Σ , Γ , and Λ , respectively. (b) The poset Σ when both solid and dashed lines are considered, the poset Λ when only solid lines are considered.

the boundary of the closed disk. We triangulate the closed disk according to Fig. 5. Then, the map l is the identity map on the vertices 0, 1, 2, 3, and 4, and $5, 6 \mapsto 4$.

Following the approach of traditional discrete Morse theory, we choose a partition of Λ into intervals, and choose a total order (which we denote by \preceq) on the quotient poset Π by taking the alphabetical order of the labeling in Fig. 6. When we view Π as a totally ordered set, we interpret the corresponding quotient map $f : \Lambda \rightarrow \Pi$ as a discrete Morse function, i.e., an order preserving map from Λ to the real line, and denote by $f_{\preceq X}^{-1}$ the sublevel set $\{\lambda \in \Lambda \mid f(\lambda) \preceq X\}$. (See Fig. 7.)

From now on we work with the posets, and we only refer back to the original spaces when interpreting the results. We denote the elements in the posets by the string of vertices that define the given simplex, e.g., 014 is the triangle with edges 01, 04, 14 and vertices 0, 1, 4. The notation does not distinguish whether we talk about elements in Σ , Γ , or Λ ; in cases where the distinction is not clear from the context, we use subscripts, e.g., $014_\Sigma \in \Sigma$ and $014_\Lambda \in \Lambda$.

We start by computing injective resolutions of the constant sheaves k_Σ and k_Γ , which, in an abuse of notation, we denote again by k_Σ and k_Γ . Then we compute $\mathcal{I}g_*k_\Sigma$, $\mathcal{I}h_*k_\Sigma$, and $\mathcal{I}l_*k_\Gamma$. We give a visualization of several intervals in $\mu \text{supp}^! \mathcal{I}h_*k_\Sigma$ and $\mu \text{supp}^* \mathcal{I}h_*k_\Sigma$. We compute various sets of critical elements of the Morse function f (relative to each of the above complexes of sheaves), and finally compute Betti numbers for hypercohomology groups corresponding to filtrations induced by f .

The injective resolution of k_Σ We compute the minimal injective resolution of k_Σ iteratively using Algorithm 3, starting with a labeled matrix with one column labeled by a virtual element larger than everything

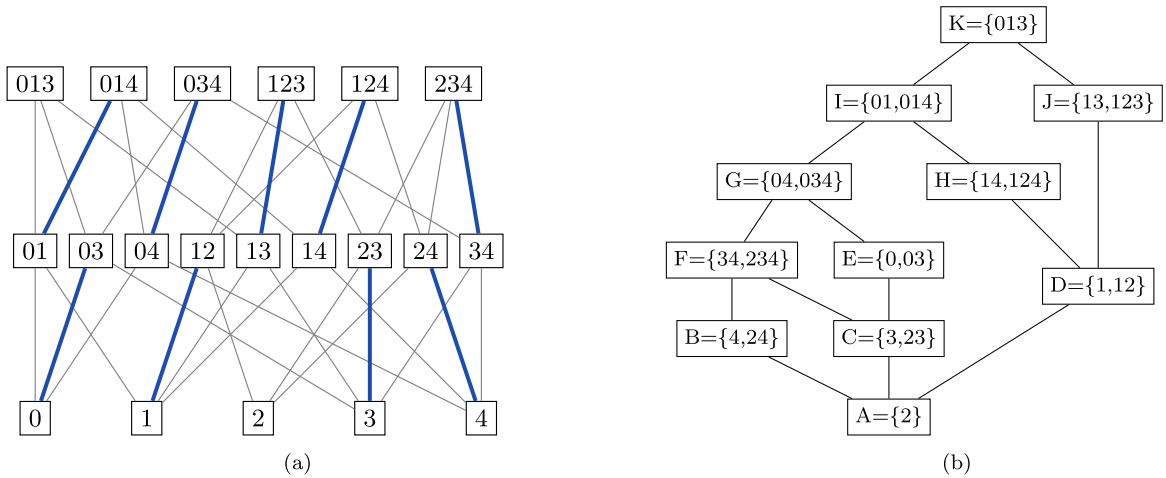


Fig. 6. (a) the Hasse diagram of Λ , with partition by locally closed sets. Each locally closed set has cardinality two, and is illustrated by a blue edge connecting its elements. (b) the Hasse diagram of the corresponding quotient poset Π .




Fig. 7. (a) sublevel set $f_{\leq A}^{-1}$. (b) sublevel set $f_{\leq B}^{-1}$. (c) sublevel set $f_{\leq I}^{-1}$. (d) sublevel set $f_{\leq K}^{-1}$.

in Σ , nine rows labeled by the maximal elements of Σ , and all entries 1 (which will, however, *not* be a part of the complex). Alternatively, we can view the same computation as Algorithm 4 computing the pullback $\mathcal{I}p^*(\cdots \rightarrow 0 \rightarrow k \rightarrow 0 \rightarrow \cdots)$ with respect to the constant map $p: \Sigma \rightarrow \text{pt}$. Either way this yields a minimal complex of injective sheaves

$$\cdots \rightarrow 0 \rightarrow I^0 \xrightarrow{\eta^0} I^1 \xrightarrow{\eta^1} I^2 \rightarrow 0 \rightarrow \cdots$$

with the maps represented by the following two labeled matrices:

$$\eta^0 \begin{pmatrix} & 013 & 014 & 034 & 123 & 124 & 234 & 45 & 46 & 56 \\ 01 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 03 & 1 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 04 & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 12 & \cdot & \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ 13 & 1 & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot \\ 14 & \cdot & 1 & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot \\ 23 & \cdot & \cdot & \cdot & 1 & \cdot & 1 & \cdot & \cdot & \cdot \\ 24 & \cdot & \cdot & \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot \\ 34 & \cdot & \cdot & 1 & \cdot & \cdot & 1 & \cdot & \cdot & \cdot \\ 4 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot \\ 4 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot \\ 5 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 \\ 6 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 \end{pmatrix}$$

$$\eta^1 \begin{pmatrix} & 01 & 03 & 04 & 12 & 13 & 14 & 23 & 24 & 34 & 4 & 4 & 5 & 6 \\ 0 & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & 1 & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 2 & \cdot & \cdot & \cdot & 1 & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 3 & \cdot & 1 & \cdot & \cdot & 1 & \cdot & 1 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot \\ 4 & \cdot & \cdot & 1 & \cdot & \cdot & 1 & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot \end{pmatrix}$$

Notice first that the complex is, indeed, minimal – the only non-empty matrix of the form $\eta^d[\sigma, \sigma]$ is $\eta^1[4, 4] = 0$. We highlighted the submatrices $\eta^0(4) = \eta^0[\text{St } 4, \text{St } 4]$, as the vertex 4 is a particularly interesting part of the example.

*The derived pushforward $\mathcal{I}g_*k_\Sigma$* To compute $\mathcal{I}g_*k_\Sigma$, we only need to relabel the matrices and minimize (see Section 5.4.2). To relabel η^0, η^1 according to g , we need to perform the following changes:

- The columns 45, 46, 56 in η^0 are all relabeled to 4,
- the rows 5, 6 in η^0 are both relabeled to 4,
- the columns 5, 6 in η^1 are both relabeled to 4.

For a moment, we denote the new labeled matrices (which do not necessarily represent a minimal complex) by $\tilde{\eta}^d$. We perform Algorithm 1 on the matrices $\tilde{\eta}^d$ to minimize the complex. The above changes created a new non-empty 4-block within $\tilde{\eta}^0$. We still have $\tilde{\eta}^1[4, 4] = 0$, but $\tilde{\eta}^0[4, 4]$ is now a non-trivial matrix. We reduce $\tilde{\eta}^0[4, 4]$, which is the 4×3 submatrix in the bottom-right corner of $\tilde{\eta}^0$:

$$\begin{pmatrix} 1 & \cdot & \cdot \\ \cdot & 1 & \cdot \\ 1 & \cdot & 1 \\ \cdot & 1 & 1 \end{pmatrix} \rightsquigarrow \begin{pmatrix} \cdot & \cdot & \cdot \\ 1 & \cdot & \cdot \\ \cdot & 1 & \cdot \\ \cdot & \cdot & 1 \end{pmatrix}$$

For the last three rows that are non-zero in the reduced submatrix, we perform further column operations to clear out all remaining places with 1 in the whole matrix. Note that even though it is not necessary, the first of the four rows labeled by 4 is now also zero, because the operations used to reduce the submatrix were performed on the whole matrix. We should now perform corresponding operations also on $\tilde{\eta}^{-1}$ and $\tilde{\eta}^1$. However, $\tilde{\eta}^{-1} = 0$, and all column operations on $\tilde{\eta}^1$ only act on a zero submatrix, so nothing changes. Finally, we remove the last three columns and last three rows in $\tilde{\eta}^0$, and the last three columns in $\tilde{\eta}^1$. The derived pushforward $\mathcal{I}g_*k_\Sigma$ is represented by the resulting labeled matrices.

$$\begin{array}{l} \mathcal{I}g_*\eta^0 \quad \begin{array}{cccccc} & 013 & 014 & 034 & 123 & 124 & 234 \end{array} \\ \begin{array}{l} 01 \\ 03 \\ 04 \\ 12 \\ 13 \\ 14 \\ 23 \\ 24 \\ 34 \\ 4 \end{array} \left(\begin{array}{cccccc} 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & 1 & \cdot & \cdot & \cdot \\ \cdot & \boxed{1} & \boxed{1} & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & 1 & \cdot \\ 1 & \cdot & \cdot & 1 & \cdot & \cdot \\ \cdot & \boxed{1} & \cdot & \cdot & \boxed{1} & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot & 1 & 1 \\ \cdot & \cdot & 1 & \cdot & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{array} \right) \end{array} \quad \begin{array}{l} \mathcal{I}g_*\eta^1 \quad \begin{array}{ccccccccc} & 01 & 03 & 04 & 12 & 13 & 14 & 23 & 24 & 34 & 4 \end{array} \\ \begin{array}{l} 0 \\ 1 \\ 2 \\ 3 \\ 4 \end{array} \left(\begin{array}{ccccccccc} 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & 1 & \cdot & 1 & \cdot & 1 & \cdot & \cdot \\ \cdot & \cdot & \boxed{1} & \cdot & \cdot & \boxed{1} & \cdot & \boxed{1} & \boxed{1} & \cdot & \cdot \end{array} \right) \end{array}$$

*The derived pushforward $\mathcal{I}h_*k_\Sigma$* We compute $\mathcal{I}h_*k_\Sigma$ analogously to above. The relabeling is now as follows:

- The columns 45, 46, and 56 in η^0 are relabeled to 14, 34, and 13, respectively,
- the rows 5 and 6 in η^0 are relabeled to 1 and 3, respectively,

$$\tilde{\lambda}^0 \quad \begin{array}{cccccccc} & 013 & 014 & 034 & 04 & 04 & 123 & 124 & 14 & 234 \\ \begin{array}{l} 01 \\ 03 \\ 04 \\ 04 \\ 04 \\ 12 \\ 13 \\ 14 \\ 14 \\ 23 \\ 24 \\ 34 \end{array} & \left(\begin{array}{cccccccc} 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 & 1 & \cdot & \cdot & \cdot \\ 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 & \cdot \\ \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & 1 \\ \cdot & \cdot & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & 1 \end{array} \right) \end{array}$$

$$\tilde{\lambda}^1 \quad \begin{array}{cccccccccccc} & 01 & 03 & 04 & 04 & 04 & 12 & 13 & 14 & 14 & 23 & 24 & 34 \\ \begin{array}{l} 0 \\ 1 \\ 2 \\ 3 \end{array} & \left(\begin{array}{cccccccccccc} 1 & 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & \cdot & \cdot & \cdot & 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & \cdot & 1 & 1 & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & \cdot & \cdot & 1 & \cdot & \cdot & 1 & \cdot & 1 & \cdot \end{array} \right) \end{array}$$

We highlight the non-trivial 04- and 14-block in blue and red, respectively. Due to their ranks, we see we need to remove 2 columns and 2 rows from the blue 04-block, and 1 column and 1 row from the red 14-block. After the Gaussian elimination and deletion as described in Algorithm 1, we get the following matrices.

$$\mathcal{I}l_*\lambda^0 \quad \begin{array}{ccccccc} & 013 & 014 & 034 & 123 & 124 & 234 \\ \begin{array}{l} 01 \\ 03 \\ 04 \\ 12 \\ 13 \\ 14 \\ 23 \\ 24 \\ 34 \end{array} & \left(\begin{array}{cccccc} 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & 1 & \cdot & \cdot & \cdot \\ \cdot & 1 & 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & 1 & \cdot \\ 1 & \cdot & \cdot & 1 & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & 1 & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & 1 \\ \cdot & \cdot & \cdot & \cdot & 1 & 1 \\ \cdot & \cdot & 1 & \cdot & \cdot & 1 \end{array} \right) \end{array}$$

$$\mathcal{I}l_*\lambda^1 \quad \begin{array}{cccccccc} & 01 & 03 & 04 & 12 & 13 & 14 & 23 & 24 & 34 \\ \begin{array}{l} 0 \\ 1 \\ 2 \\ 3 \end{array} & \left(\begin{array}{cccccccc} 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & \cdot & 1 & 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & 1 & \cdot & \cdot & 1 & 1 & \cdot & \cdot \\ \cdot & 1 & \cdot & \cdot & 1 & \cdot & 1 & \cdot & 1 & \cdot \end{array} \right) \end{array}$$

Now that we computed the derived push-forward of several maps, we have three examples of (complexes of) sheaves on the triangulation of the sphere. We will now use microlocal Morse theory to analyze these examples. First, we choose our analogue of a discrete Morse function: an order preserving map, f (see Fig. 5), from the face-relation poset of the triangulation of the sphere to the real line (or, more precisely, to a totally ordered set). We then compute which fibers $f^{-1}(X)$ lie in the discrete microsupport of each sheaf. For this example we choose f to be an honest discrete Morse function in the sense of [19,20] (see Fig. 6). However, Theorem 6.7 and Theorem 6.8 make no assumptions on f other than the requirement that it is order preserving.

Intervals in the discrete microsupport Giving a full list of all of the locally closed subsets of Λ in the microsupport of $\mathcal{I}g_*k_\Sigma$ (or $\mathcal{I}h_*k_\Sigma$) is cumbersome and mostly unnecessary. In practice, we are mostly interested in the locally closed subsets which arise as fibers of a discrete Morse function (for example, the fibers of f). Below we will compute and visualize the intersection of $\mu \text{supp}^! \mathcal{I}h_*k_\Sigma$ with several intervals of interest. (See Fig. 8.)




Fig. 8. Entries in the tables mark (an incomplete list of) intervals in the microsupport. The rows correspond to three chosen simplices, σ , and the columns correspond to each of the simplices, τ , in Λ . Each entry of the first table is given ✓ if $[\sigma, \tau] \in \mu \text{supp}^! \mathcal{I}h_* k_\Sigma$, ✗ if $[\sigma, \tau] \notin \mu \text{supp}^! \mathcal{I}h_* k_\Sigma$, and are left empty if $\sigma \not\leq \tau$. Similarly, the second table gives the corresponding intervals for $\mu \text{supp}^* \mathcal{I}h_* k_\Sigma$. Below the tables we visualize the intervals contained in the discrete microsupport. In the first row ((a), (b), and (c)) red arrows correspond to intervals $[\sigma, \tau] := \{\gamma \in \Lambda : \sigma \leq \gamma \leq \tau\}$ (where σ is the base of the arrow and τ is the head) contained in $\mu \text{supp}^! \mathcal{I}h_* k_\Sigma$. In the second row ((d), (e), and (f)) red arrows correspond to intervals contained in $\mu \text{supp}^* \mathcal{I}h_* k_\Sigma$. The first column ((a) and (d)) visualize intervals in the star of 0, in the second column ((b) and (e)) the star of 1, and in the third column ((c) and (f)) the star of 4.

Critical elements Here we provide a table of $(\mathcal{I}g_* k_\Sigma, f, *)$, $(\mathcal{I}g_* k_\Sigma, f, !)$, $(\mathcal{I}h_* k_\Sigma, f, *)$, $(\mathcal{I}h_* k_\Sigma, f, !)$, $(\mathcal{I}l_* k_\Gamma, f, *)$, and $(\mathcal{I}l_* k_\Gamma, f, !)$ -critical elements of Π (see Fig. 9). Below we give more details of the four computations for the pair $B = \{4, 24\}$. We fix $j : B \hookrightarrow \Lambda$ to be the inclusion map.

Computing critical elements: $(\mathcal{I}g_ k_\Sigma, f, *)$* Here we check whether $B = \{4, 24\}$ is in $\mu \text{supp}^* \mathcal{I}g_* k_\Sigma$. By definition, this happens iff $\mathbb{H}^d(\mathcal{I}j_* \mathcal{I}j^* \mathcal{I}g_* k_\Sigma) \neq 0$, where $j : B \hookrightarrow \Lambda$.

We first need to compute the pullback $\mathcal{I}j^*(\mathcal{I}g_* k_\Sigma)$ using Algorithm 4 as described in Section 5.4.3. According to the algorithm, we copy $\mathcal{I}g_* \eta^0$ and $\mathcal{I}g_* \eta^1$ as γ^{-1} and γ^0 , viewing them as matrices labeled

by the poset cylinder $\Lambda \sqcup_j B$, and then run the updating step from Algorithm 3 for elements in B . It is, however, enough to only copy the parts above B in $\Lambda \sqcup_j B$, i.e., the submatrices labeled by St 4, which we highlighted in the matrices above. Altogether we get the following matrices γ^d , where one row was added to each of them in the updating step:

$$\gamma^{-1} \begin{array}{c} 014_\Lambda \quad 034_\Lambda \quad 124_\Lambda \quad 234_\Lambda \\ 04_\Lambda \left(\begin{array}{cccc} 1 & 1 & \cdot & \cdot \\ 14_\Lambda & 1 & \cdot & 1 \\ 24_\Lambda & \cdot & \cdot & 1 \\ 34_\Lambda & \cdot & 1 & \cdot \\ 4_\Lambda & \cdot & \cdot & \cdot \\ 24_B & \cdot & \cdot & 1 \end{array} \right) \end{array} \quad \gamma^0 \begin{array}{c} 04_\Lambda \quad 14_\Lambda \quad 24_\Lambda \quad 34_\Lambda \quad 4_\Lambda \quad 24_B \\ 4_\Lambda \left(\begin{array}{ccccc|c} 1 & 1 & 1 & 1 & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 1 & \boxed{\cdot} \end{array} \right) \end{array}$$

We only keep the highlighted submatrix $\gamma^0[B, B]$. It already represents a minimal complex of injective sheaves, and no further minimization is needed. Altogether, we have

$$\mathcal{I}j^*(\mathcal{I}g_*k_\Sigma) = \cdots \rightarrow 0 \rightarrow [24] \xrightarrow{0} [4] \rightarrow 0 \rightarrow \cdots.$$

Next, we need to compute proper pushforward $\mathcal{I}j_!(\mathcal{I}j^*\mathcal{I}g_*k_\Sigma)$. The general construction is described by Algorithm 5 in Section 5.4.4. We first copy the 1×1 matrix from above, and then run the updating step for all elements in $\text{Cl } B \setminus B = \{2\}$. This forces us to add one row in the first matrix, but nothing else in the following. We get

$$\mathcal{I}j_!(\mathcal{I}j^*\mathcal{I}g_*k_\Sigma) = \cdots \rightarrow 0 \rightarrow [24] \xrightarrow{\begin{pmatrix} 0 \\ 1 \end{pmatrix}} [4] \oplus [2] \rightarrow 0 \rightarrow \cdots.$$

Finally,

$$\mathbb{H}^d(\mathcal{I}j_!\mathcal{I}j^*\mathcal{I}g_*k_\Sigma) = \begin{cases} k, & \text{if } d = 1 \\ 0 & \text{otherwise.} \end{cases}$$

The hypercohomology is non-trivial, which means that $B \in \mu \text{supp}^* \mathcal{I}g_*k_\Sigma$.

Computing critical elements: $(\mathcal{I}g_*k_\Sigma, f, !)$ Next, we check whether $B = \{4, 24\}$ is in $\mu \text{supp}^! \mathcal{I}g_*k_\Sigma$, i.e., whether $\mathbb{H}^d(\mathcal{I}j^! \mathcal{I}g_*k_\Sigma) \neq 0$. As discussed in Section 5.4.5, computing the proper pullback is easy – we only need to consider the submatrices $\mathcal{I}g_*\eta^d[B, B]$. We get

$$\mathcal{I}j^! \mathcal{I}g_*k_\Sigma = \cdots \rightarrow 0 \rightarrow 0 \rightarrow [24] \oplus [4] \xrightarrow{\begin{pmatrix} 1 & 0 \end{pmatrix}} [4] \rightarrow 0 \rightarrow \cdots,$$

the hypercohomology of which is

$$\mathbb{H}^d(\mathcal{I}j^! \mathcal{I}g_*k_\Sigma) = \begin{cases} k, & \text{if } d = 1 \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, $B \in \mu \text{supp}^! \mathcal{I}g_*k_\Sigma$.

Critical Elements											
	A	B	C	D	E	F	G	H	I	J	K
$(Ig_*k_\Sigma, f, !)$	✓	✓									✓
$(Ig_*k_\Sigma, f, *)$	✓	✓									✓
$(Ih_*k_\Sigma, f, !)$	✓	✓	✓	✓		✓		✓		✓	✓
$(Ih_*k_\Sigma, f, *)$	✓		✓	✓		✓		✓		✓	✓
$(Il_*k_\Gamma, f, !)$	✓	✓									✓
$(Il_*k_\Gamma, f, *)$	✓	✓									✓

Fig. 9. Entries in the table mark critical elements of the quotient poset Π , for various notions of critical corresponding to row labels.

*Computing critical elements: $(Ih_*k_\Sigma, f, *)$* We follow the same computation as above, but for h . For the pullback, we get matrices

$$\gamma^{-1} \begin{array}{c} 014_\Lambda \ 034_\Lambda \ 124_\Lambda \ 234_\Lambda \ 14_\Lambda \ 34_\Lambda \\ 04_\Lambda \left(\begin{array}{cccccc} 1 & 1 & \cdot & \cdot & \cdot & \cdot \\ 1 & \cdot & 1 & \cdot & \cdot & \cdot \\ \cdot & \cdot & 1 & 1 & \cdot & \cdot \\ \cdot & 1 & \cdot & 1 & \cdot & \cdot \\ 1 & \cdot & \cdot & \cdot & 1 & \cdot \\ 1 & \cdot & \cdot & \cdot & \cdot & 1 \end{array} \right) \\ 24_\Lambda \left(\begin{array}{cccccc} \cdot & \cdot & 1 & \cdot & \cdot & \cdot \end{array} \right) \\ 4_\Lambda \left(\begin{array}{cccccc} \cdot & \cdot & 1 & \cdot & \cdot & \cdot \end{array} \right) \\ 4_\Lambda \left(\begin{array}{cccccc} \cdot & \cdot & 1 & \cdot & \cdot & \cdot \end{array} \right) \\ 24_B \left(\begin{array}{cccccc} \cdot & \cdot & 1 & \cdot & \cdot & \cdot \end{array} \right) \end{array} \quad \gamma^0 \begin{array}{c} 04_\Lambda \ 14_\Lambda \ 24_\Lambda \ 34_\Lambda \ 4_\Lambda \ 4_\Lambda \ 24_B \\ 4_\Lambda \left(\begin{array}{cccccc|c} 1 & 1 & 1 & 1 & \cdot & \cdot & \cdot \end{array} \right) \end{array}$$

The submatrices $\gamma^d[B, B]$ are all empty with one label in degree 0, which yields

$$\mathcal{I}j^*(Ih_*k_\Sigma) = \cdots \rightarrow 0 \rightarrow [24] \rightarrow 0 \rightarrow 0 \rightarrow \cdots.$$

As above, to compute the proper pushforward, we start with the matrices we just computed, and perform the update step for 2. We need to add a row in the matrix going from degree 0, which had one column labeled by 24 and no rows. Hence,

$$\mathcal{I}j_!(\mathcal{I}j^*Ih_*k_\Sigma) = \cdots \rightarrow 0 \rightarrow [24] \xrightarrow{(1)} [2] \rightarrow 0 \rightarrow \cdots.$$

The hypercohomology of this complex is trivial, so $B \notin \mu \operatorname{supp}^* Ih_*k_\Sigma$.

*Computing critical elements: $(Ih_*k_\Sigma, f, !)$* The matrices $Ih_*\eta_0[B, B]$ and $Ih_*\eta_1[B, B]$ give

$$\mathcal{I}j^!Ih_*k_\Sigma = \cdots \rightarrow 0 \rightarrow 0 \rightarrow [24] \oplus [4] \oplus [4] \xrightarrow{\begin{pmatrix} 1 & 0 & 0 \end{pmatrix}} [4] \rightarrow 0 \rightarrow \cdots,$$

with hypercohomology

$$\mathbb{H}^d(\mathcal{I}j^!Ig_*k_\Sigma) = \begin{cases} k^2, & \text{if } d = 1 \\ 0 & \text{otherwise,} \end{cases}$$

and hence $B \in \mu \operatorname{supp}^! Ih_*k_\Sigma$.

Rank of hypercohomology groups in degree 0, 1, and 2											
$\dim \mathbb{H}^d(-)$	A	B	C	D	E	F	G	H	I	J	K
$\mathcal{I}i_Z^! \mathcal{I}g_* k_\Sigma$	0,0,1	0,1,0									1,0,0
$\mathcal{I}i_{\leq x}^! \mathcal{I}g_* k_\Sigma$	0,0,1	0,1,1	0,1,1	0,1,1	0,1,1	0,1,1	0,1,1	0,1,1	0,1,1	0,1,1	1,1,1
$\mathcal{I}i_{\geq x}^* \mathcal{I}g_* k_\Sigma$	1,1,1	1,1,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0
$\mathcal{I}i_Z^! \mathcal{I}i_Z^* \mathcal{I}g_* k_\Sigma$	1,0,0	0,1,0									0,0,1
$\mathcal{I}i_{\leq x}^* \mathcal{I}g_* k_\Sigma$	1,0,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,1
$\mathcal{I}i_Z^! \mathcal{I}h_* k_\Sigma$	0,0,1	0,2,0	0,1,0	0,1,0		1,0,0		1,0,0		1,0,0	1,0,0
$\mathcal{I}i_{\leq x}^! \mathcal{I}h_* k_\Sigma$	0,0,1	0,2,1	0,3,1	0,4,1	0,4,1	0,3,1	0,3,1	0,2,1	0,2,1	0,1,1	1,1,1
$\mathcal{I}i_{\geq x}^* \mathcal{I}h_* k_\Sigma$	1,1,1	1,1,0	2,0,0	3,0,0	4,0,0	4,0,0	3,0,0	3,0,0	2,0,0	2,0,0	1,0,0
$\mathcal{I}i_Z^! \mathcal{I}i_Z^* \mathcal{I}h_* k_\Sigma$	1,0,0		1,0,0	1,0,0		0,1,0		0,1,0		0,1,0	0,0,1
$\mathcal{I}i_{\leq x}^* \mathcal{I}h_* k_\Sigma$	1,0,0	1,0,0	2,0,0	3,0,0	3,0,0	2,0,0	2,0,0	1,0,0	1,0,0	1,1,0	1,1,1
$\mathcal{I}i_Z^! \mathcal{I}l_* k_\Gamma$	0,0,1	0,1,0									1,0,0
$\mathcal{I}i_{\leq x}^! \mathcal{I}l_* k_\Gamma$	0,0,1										1,0,0
$\mathcal{I}i_{\geq x}^* \mathcal{I}l_* k_\Gamma$	1,0,0	1,1,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0	1,0,0
$\mathcal{I}i_Z^! \mathcal{I}i_Z^* \mathcal{I}l_* k_\Gamma$	1,0,0	0,1,0									0,0,1
$\mathcal{I}i_{\leq x}^* \mathcal{I}l_* k_\Gamma$	1,0,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,1,0	1,0,0

Fig. 10. Each entry in the table lists the dimension of hypercohomology groups in degree 0, 1, and 2 (empty entries have trivial hypercohomology). Here, $Z = f^{-1}(x)$ and $i_{\leq x}^!$ ($i_{\geq x}^*$) is the inclusion map of sublevel-sets (superlevel-sets) of f (relative to the total order \leq on Π) into Λ , for x an element of Π corresponding to the column labels.

Filtrations and hypercohomology In Fig. 10 we list Betti numbers for hypercohomology groups associated to the filtration of Λ by sublevel sets of f . As Theorem 6.7 implies, the value in column X for the sublevel filtration differs from the value in its left neighbor only if X is a critical element for the corresponding notion of microlocal support as shown in the table above. Similarly, for the superlevel filtration, non-critical columns X have the same value as their right neighbors.

CRedit authorship contribution statement

Adam Brown: Writing – review & editing, Writing – original draft, Methodology, Formal analysis, Conceptualization. **Ondřej Draganov:** Writing – review & editing, Writing – original draft, Methodology, Investigation, Conceptualization.

Declaration of competing interest

The authors have no financial or personal relationships that may be perceived as influencing their work.

Appendix A. Background: injective sheaves

A.1. Duality between injective and projective resolutions

We should comment on a matter of perspective and terminology. Sheaves over finite posets are closely related to several other mathematical objects studied by various research communities. Specifically, there is a great deal of work within the field of commutative algebra on minimal projective and free resolutions of modules over various kinds of algebras. Here, we choose to focus on the perspective and terminology which

most closely aligns with classical sheaf theory in order to preserve intuition from that discipline. Below we briefly comment on the connections between sheaves, cosheaves, injective resolutions, and projective resolutions, with implications for computing the minimal injective hull of a sheaf.

Given a sheaf F on a poset Π , let \hat{F} denote a sheaf on Π^{op} defined by

$$\hat{F}(\pi) := \text{Hom}_k(F(\pi), k),$$

the k -linear functionals on $F(\pi)$, where $\hat{F}(\pi \leq \tau)$ is the linear map from $\text{Hom}_k(F(\pi), k)$ to $\text{Hom}_k(F(\tau), k)$ defined by precomposing linear functionals on $F(\pi)$ with $F(\tau \leq \pi)$. This defines an exact contravariant functor from the category of sheaves on Π to the category of sheaves on Π^{op} . In fact, this functor can be extended to an equivalence (of triangulated categories) between derived categories $D^b(\Pi)$ and $D^b(\Pi^{\text{op}})^{\text{op}}$ [16, 37]. It is easy to show that I is an injective sheaf on Π if and only if \hat{I} is a projective sheaf on Π^{op} . Moreover, this contravariant functor interchanges injective and projective complexes of sheaves. If $I^\bullet \in D^b(\Pi)$, let $\hat{I}^{-\bullet}$ be the corresponding complex of projective sheaves (where the integer indices are given the opposite sign). It is again straightforward to show that I^\bullet is an injective resolution of a sheaf F on Π if and only if $\hat{I}^{-\bullet}$ is a projective resolution of \hat{F} on Π^{op} . Moreover, this functor preserves minimality of the resolutions. This is to say that computing injective resolutions of sheaves on Π is equivalent to computing projective resolutions of sheaves on Π^{op} . We see two applications of this observation. First, we can compute left derived functors by taking a sheaf F and simply computing the injective resolution of \hat{F} on Π^{op} , taking the transpose of each labeled matrix (Definition 3.3) to obtain a projective resolution of F on Π , and applying the desired right exact functor. Secondly, this reformulation highlights close connections between computational derived sheaf theory and commutative algebra. Particularly, theoretical results and implementations for computing minimal projective resolutions of modules over the incidence algebra of a finite poset (see [26], for example) can be exploited for computations in derived sheaf theory, providing alternatives to Algorithm 6 and Algorithm 3 of the present paper.

A.2. Examples of non-injective sheaves

A sheaf I is injective if any natural transformation from A into it can be extended to any super-sheaf $B \leftarrow A$.

$$\begin{array}{ccccc} 0 & \longrightarrow & A & \longrightarrow & B \\ & & \searrow \forall & & \downarrow \exists \\ & & & & I \end{array}$$

Below are two simple examples of sheaves that are not injective.


Example A.1. We describe a sheaf F on a poset with two elements $\sigma \leq \tau$ which is not injective. Fix any vector space $W \neq 0$, and let $F(\sigma) = 0$ and $F(\tau) = W$. This sheaf injects in a sheaf G given by $G(\sigma) = W = G(\tau)$ with $G(\sigma \leq \tau) = \text{id}$. We claim that $F \xrightarrow{\text{id}} F$ can not be extended to $G \rightarrow F$.

$$\begin{array}{ccccc} & & F & & G & & F \\ & & & \text{id} & & & \\ \tau & W & \xrightarrow{\text{id}} & W & \xrightarrow{?} & W \\ & \uparrow 0 & & \uparrow \text{id} & & \uparrow 0 \\ \sigma & 0 & \xrightarrow{0} & W & \xrightarrow{?} & 0 \\ & & & 0 & & \end{array}$$

Indeed, the triangle on the left diagram commutes iff the right diagram commutes. The only way to make the right square commute is to set both horizontal maps to 0. But then the top triangle does not commute.

The same reasoning applies for any sheaf on any poset with a non-zero vector space one step above a zero vector space. Below we demonstrate one other obstruction to injectivity.

Example A.2. We consider a three-element “V” shaped poset, a vector space $W \neq 0$, and two different endomorphisms $f, g : W \rightarrow W$. We define sheaves A, B and F as follows:



We claim that F is not injective. The sheaf A embeds into both B and F with identity maps. We ask whether we can extend the embedding $A \rightarrow F$ to a natural transformation $\beta : B \rightarrow F$, while respecting the embedding $A \rightarrow B$. The only choice we have is the bottom map $\beta_0 : W \rightarrow W$, as the other two must be the identity. However, commutativity requires $f = \beta_0 = g$, which is impossible to satisfy, since $f \neq g$.

A.3. Proofs of standard facts about injective sheaves

Lemma 2.7 (cf. [15, Lemma 7.1.5]). *Indecomposable injective sheaves are injective.*

Proof. We show that, for a fixed poset Π and $\pi \in \Pi$, $I = [\pi]$ satisfies Definition 2.5. Given an inclusion $A \xhookrightarrow{f} B$ and a natural transformation $\alpha : A \rightarrow I$, we need to find an extension $\beta : B \rightarrow I$. For the linear map $A(\pi) \xrightarrow{f(\pi)} B(\pi)$, there is a projection $A(\pi) \xleftarrow{g} B(\pi)$ such that $gf(\pi) = \text{id}_{A(\pi)}$. We define

$$\beta(\sigma) := \begin{cases} \alpha(\pi) \circ g \circ B(\sigma \leq \pi) & \text{if } \sigma \leq \pi, \\ 0 & \text{otherwise.} \end{cases}$$

For every σ , this satisfies $\beta(\sigma)f(\sigma) = \alpha(\sigma)$, because if $\sigma \leq \pi$, then

$$\beta(\sigma)f(\sigma) = \alpha(\pi)gB(\sigma \leq \pi)f(\sigma) = \alpha(\pi)gf(\pi)A(\sigma \leq \pi) = \alpha(\pi)A(\sigma \leq \pi) = \alpha(\sigma),$$

and otherwise both sides are 0. For the commutativity conditions, consider $\sigma \leq \tau \leq \pi$. Then

$$\beta(\tau)B(\sigma \leq \tau) = \alpha(\pi)gB(\tau \leq \pi)B(\sigma \leq \tau) = \alpha(\pi)gB(\sigma \leq \pi) = \beta(\sigma) = I(\sigma \leq \tau)\beta(\sigma).$$

If $\sigma \leq \tau \not\leq \pi$, then both sides are 0.

Note that this argument works for arbitrary vector spaces,² not only finite-dimensional ones. Indeed, we only need to verify the existence of the map g . Consider a basis $\{b_t\}_{t \in T}$ of $B(\pi)$ such that for some subset $S \subset T$, the subset $\{b_s\}_{s \in S}$ is a basis of $\text{im } f(\pi)$. Then g is a composition of the projection onto $\text{im } f(\pi)$ given by $\sum_{t \in T} \gamma_t b_t \mapsto \sum_{s \in S} \gamma_s b_s$ followed by the inverse of $f(\pi) : A(\pi) \rightarrow \text{im } f(\pi)$. Clearly, $gf(\pi)$ is the identity. The existence of the basis is easily shown with Zorn’s lemma by first obtaining a maximal element, M , of the collection of sets of linearly independent vectors in $\text{im } f(\pi)$, and then a maximal element of the collection of sets of linearly independent vectors in $B(\pi)$ that contain M . \square

² Assuming the Axiom of Choice.

Lemma 2.8 (cf. [49, Lemma 1.3.1]). *A direct sum of injective sheaves is injective. Additionally, if $I \xrightarrow{\alpha} J$ is an injective natural transformation with I, J injective sheaves, then $J \cong I \oplus \operatorname{coker} \alpha$, and $\operatorname{coker} \alpha$ is an injective sheaf.*

Proof. This proof is standard for any abelian category, we include a sketch for completeness. Suppose $I = A \oplus B$, with A, B injective sheaves. Suppose $F \hookrightarrow G$ and $F \rightarrow I$. Then composition with projection gives maps $F \rightarrow A$ and $F \rightarrow B$. By injectivity of A and B , each map extends to $G \rightarrow A$ and $G \rightarrow B$, respectively. The sum of these maps defines an extension $G \rightarrow I$, proving that I is injective. The second claim follows by extending the identity map $I \rightarrow I$ to $J \rightarrow I$ by α and the injectivity of J . Then, the sum of the extension and the quotient map define an isomorphism $J \rightarrow I \oplus \operatorname{coker} \alpha$. The final claim follows by composing a given map $F \rightarrow \operatorname{coker} \alpha$ with the extension by zero map, $\operatorname{coker} \alpha \hookrightarrow J$, to get $F \rightarrow J$. Then, for $F \hookrightarrow G$, we define (by the injectivity of J) an extension $G \rightarrow J$. By post-composing with the projection map, we get the desired extension $G \rightarrow \operatorname{coker} \alpha$. \square

Proposition 2.9 (cf. [15, Lemma 7.1.6], [49, Theorem 1.3.2]). *Every injective sheaf is isomorphic to a direct sum of indecomposable injective sheaves.*

Proof. We adapt the proof of [49, Theorem 1.3.2] to the setting of finite posets on n elements (rather than cell complexes). We fix some linear extension of the partial order, (π_1, \dots, π_n) , and let $\Pi_d = \{\pi_j \mid j \leq d\}$. We will proceed with the proof by working inductively through this filtration of Π . We define support of a sheaf I as

$$\operatorname{supp} I := \{\pi \in \Pi \mid I(\pi) \neq 0\}.$$

Assume that the result holds for injective sheaves supported on Π_{d-1} . Suppose I is an injective sheaf with support contained in Π_d . If $\operatorname{supp} I \subseteq \Pi_{d-1}$, then the inductive assumption implies the result. Therefore, we are left to prove the result for I such that $I(\pi_d) \neq 0$. Set F_{π_d} to be the functor which assigns $I(\pi_d)$ to π_d and the zero vector space to each other poset element (and the zero linear map to each poset relation). Then the identity map induces injective natural transformations

$$F_{\pi_d} \xrightarrow{\alpha} I \quad \text{and} \quad F_{\pi_d} \hookrightarrow \bigoplus_{v \in B} [\pi_d],$$

where B is some basis of $I(\pi_d)$. Because I is injective, we can extend α to a natural transformation $\beta : \bigoplus_{v \in B} [\pi_d] \rightarrow I$. It is injective, because for every $\sigma \leq \pi_d$, the linear map $I(\sigma \leq \pi_d)\beta(\sigma) = \beta(\pi_d) = \alpha(\pi_d)$ is injective. By Lemma 2.8, this implies that

$$I \cong \operatorname{coker} \beta \oplus \bigoplus_{v \in B} [\pi_d],$$

and that $\operatorname{coker} \beta$ is injective. Because $\operatorname{supp} \operatorname{coker} \beta \subseteq \Pi_{d-1}$, the inductive hypothesis completes the proof. \square

Appendix B. The correspondence between Definition 4.10 and the bounded derived category of sheaves on a finite poset with the Alexandrov topology

Here we sketch a correspondence between the bounded derived category of sheaves on a finite poset endowed with the Alexandrov topology and the derived category studied in Section 4. Here we assume the reader is familiar with sheaf theory (for example, [33], [10], [15], and [49]).

Let X be a finite poset endowed with the Alexandrov topology. Let $\operatorname{Mod}(k_X)$ be the category of sheaves of k -vector spaces on X (as defined in [33]), and

$$\mathrm{Mod}_{\mathrm{fin}}(k_X) = \{\mathcal{F} \in \mathrm{Mod}(k_X) \mid \dim_k \mathcal{F}_x < \infty \quad \forall x \in X\},$$

be the subcategory of sheaves with finite-dimensional stalks.

Proposition B.1. $\mathrm{Mod}_{\mathrm{fin}}(k_X)$ is equivalent to the category of functors from X (with morphisms given by poset relations) to the category of finite dimensional k -vector spaces, $\mathrm{Fun}(X, \mathbf{vect}_k)$. Similarly, $\mathrm{Mod}(k_X)$ is equivalent to $\mathrm{Fun}(X, \mathbf{Vect}_k)$.

Proof. See [15, Theorem 4.2.10]. \square

Lemma B.2. $\mathrm{Mod}_{\mathrm{fin}}(k_X)$ is a Serre³ subcategory of $\mathrm{Mod}(k_X)$.

Proof. Assume

$$F_0 \rightarrow F_1 \rightarrow F_2$$

is an exact sequence of sheaves on X , with $F_i \in \mathrm{Mod}_{\mathrm{fin}}(k_X)$ for $i \in \{0, 2\}$. For each $x \in X$ we can apply the pullback i_x^* along the inclusion map $i_x : \{x\} \hookrightarrow X$ to get the exact sequence (recall that i_x^* is an exact functor) of k -vector spaces

$$F_0(x) \rightarrow F_1(x) \rightarrow F_2(x)$$

Because $F_i(x)$ is finite-dimensional for $i \in \{0, 2\}$, $F_1(x)$ is finite-dimensional. \square

Proposition B.3 (Prop 1.7.11 and Remark 1.7.12 in [33]). If C' is a weak Serre subcategory of a category C such that for every monomorphism $f : a \rightarrow b$, $a \in \mathrm{Ob}(C')$, $b \in \mathrm{Ob}(C)$, there exists $g : b \rightarrow c$ with $c \in \mathrm{Ob}(C')$ such that $g \circ f$ is a monomorphism, then $D^b(C') \simeq D_{C'}^b(C)$, where $D_{C'}^b(C)$ denotes the full triangulated subcategory of $D^b(C)$ with cohomology objects in C' .

Let $D^b(k_X)$ be the bounded derived category of sheaves on X (as defined in [33]) and $D_{\mathrm{fin}}^b(k_X)$ be the subcategory whose cohomology objects are in $\mathrm{Mod}_{\mathrm{fin}}(k_X)$.

Theorem B.4. $D_{\mathrm{fin}}^b(k_X)$ is equivalent to $D^b(X)$ (with $D_{\mathrm{fin}}^b(k_X)$ defined above and $D^b(X)$ in Definition 4.10).

Proof. For any monomorphism $f : \mathcal{F} \rightarrow \mathcal{G}$ with $\mathcal{F} \in \mathrm{Mod}_{\mathrm{fin}}(k_X)$, there exists a morphism $g : \mathcal{G} \rightarrow \mathcal{H}$, with $\mathcal{H} \in \mathrm{Mod}_{\mathrm{fin}}(k_X)$, such that $g \circ f$ is a monomorphism. Namely, choose \mathcal{H} to be the injective hull of \mathcal{F} (Definition 4.16). Then, because \mathcal{H} is also an injective object in $\mathrm{Mod}(k_X)$ ([15, Lemma 7.1.5], see also Lemma 2.7), a morphism g exists with the desired properties. Therefore, by [33, Prop 1.7.11] and the previous lemma (Appendix B.2), $D_{\mathrm{fin}}^b(k_X)$ is equivalent to $D^b(\mathrm{Mod}_{\mathrm{fin}}(k_X))$. By the above proposition (Appendix B.1), $D^b(\mathrm{Mod}_{\mathrm{fin}}(k_X))$ is equivalent to $D^b(\mathrm{Fun}(X, \mathbf{vect}_k))$. The result then follows by applying the replacement functor which maps a given complex $\mathcal{C} \in D^b(\mathrm{Fun}(X, \mathbf{vect}_k))$ to its minimal injective resolution in $D^b(X)$. \square

³ In particular, also weak Serre.

Appendix C. Derived category and injective resolution of a sheaf

C.1. Computing a basis for the space of morphisms in a derived category

We describe the spaces of all morphisms and all null-homotopic morphisms between two complexes of injective sheaves as solutions to systems of linear equations. Those systems of equations are smaller compared to the case of general sheaves, because for injective sheaves, we can take advantage of the labeled matrix representations of natural transformations. This said, the size of the systems of linear equations which define these morphisms are impractically large; the goal of this exposition is to give an explicit description of the space of morphisms between two objects in the *derived category of sheaves*, including a way (albeit impractical) to compute their basis.

We fix two complexes of injective sheaves, $(I^\bullet, \eta^\bullet)$, $(J^\bullet, \lambda^\bullet)$, and denote by $m(I^d)$ and $m(J^d)$ the number of indecomposable injective summands in I^d and J^d , respectively. We assume that both complexes are represented as complexes of labeled matrices, and we denote the matrices by the same symbols as the natural transformations they represent, e.g., η^d is a $m(I^{d+1}) \times m(I^d)$ labeled matrix.

We can describe every morphism $\alpha^\bullet : I^\bullet \rightarrow J^\bullet$ as another collection of labeled matrices. The dimensions are fixed: α^d is a $m(J^d) \times m(I^d)$ matrix with columns labeled as columns of η^d , and rows labeled as columns of λ^d . Therefore, we have $\sum_d m(J^d) \cdot m(I^d)$ variables which will be used to define α^\bullet . We observe that the definition for α^\bullet to represent a morphism yields two sets of linear constraints:

- *poset constraints*: $\alpha^d[i, j] = 0$ whenever $\sigma \not\leq \pi$ for σ the label of the i -th row and π the label of the j -th column of α^d ,
- *commutativity constraints*: $\alpha^{d+1} \cdot \eta^d - \lambda^d \cdot \alpha^d = 0$ yields $m(J^{d+1})$ linear constraints for each d .

The system of linear equations for null-homotopic morphisms is the one above together with new variables and new constraints. We have $\sum_d m(I^d) \cdot m(J^{d-1})$ new variables for the matrices h^d representing the natural transformations $I^d \rightarrow J^{d-1}$ as in Definition 4.1. Again, we label each h^d accordingly to fit with the labeling of η^\bullet , λ^\bullet . We get the following two new sets of linear constraints:

- *poset constraints*: $h^d[i, j] = 0$ whenever $\sigma \not\leq \pi$ for σ the label of the i -th row and π the label of the j -th column of h^d ,
- *homotopy constraints*: $\alpha^d - h^{d+1} \cdot \eta^d - \lambda^d \cdot h^{d-1} = 0$ yields $m(J^d)$ linear constraints for each d .

Altogether, we showed the following.

Proposition C.1. *Computing a basis for the space of morphisms in derived category, $\text{Hom}_{D^b(\Pi)}(I^\bullet, J^\bullet)$ (see Definition 4.10), reduces to computing a basis for the space of solutions to the above system of linear equations.*

C.2. Injective resolution via order complex

We present a non-inductive construction of a (not necessarily minimal) injective resolution of a given sheaf F . This section generalizes, from the constant sheaf to general sheaves, Lemma 1.3.17 of [36]. On a practical level, this allows one to compute the cohomology of right derived functors without first computing the full injective resolution (see Section 5).

Definition C.2. The *order complex*, $K(\Pi)$, of a finite poset Π , is the poset of strictly increasing chains $\pi_\bullet = \pi_0 < \pi_1 < \cdots < \pi_d$ in Π . The order complex has the structure of an abstract simplicial complex. Let $K^d(\Pi)$ denote the d -simplices of $K(\Pi)$, i.e. the set of chains $\pi_0 < \pi_1 < \cdots < \pi_d$ of length $d + 1$.

The construction Given a sheaf F on Π , we define (recalling the notation of Definition 2.6)

$$I^d := \bigoplus_{\pi_\bullet \in K^d(\Pi)} [\pi_0]^{F(\pi_d)}.$$

Suppose $\pi_\bullet \in K^d(\Pi)$ and $\pi_\bullet <_1 \tau_\bullet$ (i.e. the chain π_\bullet is obtained from the chain τ_\bullet by removing one element). Then $\pi_0 \geq \tau_0$ and $\pi_d \leq \tau_{d+1}$. Therefore,

$$F(\pi_d \leq \tau_{d+1}) \in \text{Hom}(F(\pi_d), F(\tau_{d+1})) \cong \text{Hom}\left([\pi_0]^{F(\pi_d)}, [\tau_0]^{F(\tau_{d+1})}\right).$$

Definition C.3 ([15, Definition 6.1.9]). A *signed incidence relation* on $K(\Pi)$ is an assignment to each pair of simplices $\sigma_\bullet, \gamma_\bullet \in K(\Pi)$ a number $[\sigma_\bullet : \gamma_\bullet] \in \{-1, 0, 1\}$, such that

1. if $[\sigma_\bullet : \gamma_\bullet] \neq 0$, then $\sigma_\bullet <_1 \gamma_\bullet$, and
2. for each pair of simplices $(\sigma_\bullet, \gamma_\bullet)$,

$$\sum_{\tau_\bullet \in K(\Pi)} [\sigma_\bullet : \tau_\bullet][\tau_\bullet : \gamma_\bullet] = 0.$$

Using this identification, we define the natural transformation $\eta^d : I^d \rightarrow I^{d+1}$ so that on the π_\bullet -summand $[\pi_0]^{F(\pi_d)}$ of I^d ,

$$\eta^d|_{[\pi_0]^{F(\pi_d)}} = \sum_{\pi_\bullet <_1 \tau_\bullet} [\pi_\bullet : \tau_\bullet] F(\pi_d \leq \tau_{d+1}),$$

where $F(\pi_d \leq \tau_{d+1}) \in \text{Hom}([\pi_0]^{F(\pi_d)}, [\tau_0]^{F(\tau_{d+1})})$ is understood to have its codomain as the τ_\bullet -summand of I^{d+1} . Let $\alpha : F \hookrightarrow I^0$ be the natural transformation given by the maps

$$\alpha(\sigma) := \sum_{\sigma \leq \gamma} F(\sigma \leq \gamma) : F(\sigma) \hookrightarrow \bigoplus_{\sigma \leq \gamma} F(\gamma) =: I^0(\sigma).$$

The following lemma is a generalization of [16, Theorem 6.12].

Lemma C.4. Let $t : K(\Pi) \rightarrow \Pi$ denote the poset map which assigns each chain $\pi_0 < \pi_1 < \cdots < \pi_d$ to its terminal element π_d , and $p : \Pi \rightarrow \{pt\}$. Given a sheaf F on a finite poset Π , we have the following isomorphism of sheaf cohomology

$$H^j(\Pi; F) \cong H^j(K(\Pi); t^*F), \text{ for each } j.$$

Proof. When Π is a cell complex, the theorem follows, for example, from [33, Corollary 2.7.7 (iv)]. We include a more general proof for arbitrary finite posets below. Let I^\bullet be an injective resolution of F . Because t^* defines an exact functor from the category of sheaves on Π to the category of sheaves on $K(\Pi)$, we have that

$$0 \rightarrow t^*F \rightarrow t^*I^0 \rightarrow t^*I^1 \rightarrow \cdots \rightarrow t^*I^n \rightarrow 0$$

is an exact sequence. We claim that each t^*I^d is acyclic:

$$H^j(\mathcal{I}(p \circ t)_*(t^*I^d)) = 0, \text{ for } j > 0.$$

Indeed, it is enough to prove that for each indecomposable injective sheaf $[\pi]$ on Π , the sheaf $t^*[\pi]$ is acyclic. Because $t^*[\pi]$ is the constant sheaf on the order complex of the downward closure, $\text{Cl } \pi := \{\tau \in \Pi : \tau \leq \pi\}$, of π in Π , we have

$$H^j(\mathcal{I}(p \circ t)_*(t^*[\pi])) \cong H^j(|K(\text{Cl } \pi)|; k).$$

Notice that $K(\text{Cl } \pi)$, as a simplicial complex, is equal to the cone of the simplicial complex $K(\text{Cl } \pi - \pi)$. Therefore, $|K(\text{Cl } \pi)|$ is contractible, $H^j(|K(\text{Cl } \pi)|; k) = 0$ for $j > 0$, and $t^*[\pi]$ is acyclic. The complex t^*I^\bullet is therefore an acyclic resolution of t^*F , and by standard results of homological algebra (for example, [10, Theorem 4.1]), we have

$$H^j(K(\Pi); t^*F) \cong H^j((p \circ t)_*(t^*I^\bullet)) \cong H^j(\Pi; F). \quad \square$$

Theorem C.5. *The complex $0 \rightarrow F \xrightarrow{\alpha} I^0 \xrightarrow{\eta^0} I^1 \xrightarrow{\eta^1} \dots$ defined above is an injective resolution of F .*

Proof. By construction, each sheaf I^j is injective, and each map η^j (as well as α) is a natural transformation. It remains to show that the sequence is an exact complex. It is enough to show that for each $\pi \in \Pi$, the sequence $0 \rightarrow F(\pi) \xrightarrow{\alpha(\pi)} I^0(\pi) \xrightarrow{\eta^0(\pi)} I^1(\pi) \xrightarrow{\eta^1(\pi)} \dots$ is exact.

We first define a functor t^* from the category of sheaves on Π to the category of sheaves on $K(\Pi)$. To each sheaf F on Π , let t^*F be the sheaf on $K(\Pi)$ defined by associating to each chain $\tau_\bullet = \tau_0 < \dots < \tau_d$ the ‘terminal’ vector space: $t^*F(\tau_\bullet) := F(\tau_d)$ and $t^*F(\tau_\bullet \leq \gamma_\bullet) = F(\tau_d \leq \gamma_j)$ for $\tau_\bullet \in K^d(\Pi)$, $\gamma_\bullet \in K^j(\Pi)$. Because $t^*(\eta)(\tau_\bullet) := \eta(\tau_d) : t^*F(\tau_\bullet) \rightarrow t^*G(\tau_\bullet)$ for any natural transformation $\eta : F \rightarrow G$, it is clear that t^* is an exact functor.

Notice that $0 \rightarrow F(\pi) \xrightarrow{\alpha(\pi)} I^0(\pi) \xrightarrow{\eta^0(\pi)} I^1(\pi) \xrightarrow{\eta^1(\pi)} \dots$ is identical to the compactly supported cochain complex of the sheaf $t^*(F|_{\text{St } \pi})$ on the simplicial complex $K(\text{St } \pi)$ [15, Definition 6.2.1 and Definition 6.2.3]. Therefore, exactness in $I^0(\pi)$ follows from

$$\ker \eta_0(\pi) \cong \Gamma(t^*(F|_{\text{St } \pi})) \cong F(\pi) \cong \text{im } \alpha(\pi),$$

and it remains to prove a vanishing property for the cohomology of $t^*(F|_{\text{St } \pi})$, namely that

$$H^j(K(\text{St } \pi); t^*(F|_{\text{St } \pi})) = 0$$

for $j > 0$. By Lemma C.4, $H^j(K(\text{St } \pi); t^*(F|_{\text{St } \pi})) \cong H^j(\text{St } \pi; F|_{\text{St } \pi})$. Let J^\bullet be an injective resolution of the sheaf $F|_{\text{St } \pi}$ on the poset $\text{St } \pi$. Then $H^j(\text{St } \pi; F|_{\text{St } \pi})$ is, by definition, the j -th cohomology group of the complex of vector spaces $J^\bullet(\pi)$, which, by the exactness of J^\bullet , is zero for $j > 0$. \square

This construction also allows us to prove Proposition 5.13

Proposition 5.13. *Suppose $f : \Sigma \rightarrow \Lambda$ is a simplicial map. As sheaves on Λ ,*

$$H^d \mathcal{I} f_* k_\Sigma \cong H^d(|f^{-1}(\text{St } -)|; k)$$

where $H^d(|f^{-1}(\text{St } -)|; k)$ is the sheaf defined by associating the simplex λ to the singular cohomology of the geometric realization of $f^{-1}(\text{St } \lambda)$ (with linear maps induced by inclusion).

Proof. Let I^\bullet be the injective resolution described above. Then $\mathcal{I}f_* I^\bullet(\lambda)$ is the complex consisting of only the linear combinations of generators for indecomposable sheaves $[\pi] \subset I^\bullet$ such that $\pi \in f^{-1}(\text{St } \lambda)$ and chain maps $\eta^\bullet(f^{-1}(\text{St } \lambda))$. After forgetting the matrix labels representing $\mathcal{I}f_* I^\bullet(\lambda)$, we obtain the simplicial cochain complex of $K(f^{-1}(\text{St } \lambda))$. The cohomology groups of this cochain complex are isomorphic to the singular cohomology of the geometric realization of $f^{-1}(\text{St } \lambda)$:

$$H^d \mathcal{I}f_* k_\Sigma(\lambda) \cong H^d(|f^{-1}(\text{St } \lambda)|; k),$$

and the linear maps $H^d \mathcal{I}f_*(\kappa \leq \lambda)$ are the usual cohomology maps

$$H^d(|K(f^{-1}(\text{St } \kappa))|; k) \rightarrow H^d(|K(f^{-1}(\text{St } \lambda))|; k)$$

induced by inclusion (cf. [29, Chapter II Proposition 5.11]). \square

C.3. Injective hull computation

We describe an explicit construction of the minimal injective hull of a sheaf F on a poset Π . An injective hull of F consists of an injective sheaf I and an inclusion map of F into I . To construct the minimal injective hull of F , we first find M_F , the subsheaf of F with $M_F(\pi)$ the space of maximal vectors in $F(\pi)$ (recall Definition 4.8), and zero maps between the spaces. Recalling the notation described below Definition 2.6, we define

$$I^0 = \bigoplus_{\pi \in \Pi} [\pi]^{M_F(\pi)} \cong \bigoplus_{\pi \in \Pi} [\pi]^{\dim M_F(\pi)},$$

where $[\pi]^{M_F(\pi)}$ is the injective sheaf with $[\pi]^{M_F(\pi)}(\sigma) = M_F(\pi)$ if $\sigma \leq \pi$, and 0 otherwise. We can naturally include $M_F \xrightarrow{\gamma} I^0$, and extend this inclusion to $F \xrightarrow{\alpha} I^0$, using the injectivity of I^0 . We choose the extension $\alpha = \sum_{\pi \in \Pi} \alpha_\pi$, where α_π is given by

$$\alpha_\pi(\sigma) := \begin{cases} \gamma(\pi) \circ \text{proj}_{M_F(\pi)} \circ F(\sigma \leq \pi) & \text{if } \sigma \leq \pi, \\ 0 & \text{otherwise.} \end{cases}$$

Proposition C.6. *The construction above yields the minimal injective hull $F \xrightarrow{\alpha} I^0$.*

Proof. We claim that α is injective. Let $u \in \ker \alpha(\sigma)$. Then

$$0 = \alpha(\sigma)(u) = \sum_{\pi \in \Pi} \alpha_\pi(\sigma)(u) = \sum_{\sigma \leq \pi} \gamma(\pi) \circ \text{proj}_{M_F(\pi)} F(\sigma \leq \pi)(u),$$

which is equivalent to $\text{proj}_{M_F(\pi)} F(\sigma \leq \pi)(u) = 0$ for every $\pi \geq \sigma$, since the images of different α_π have trivial intersections. But this means that $u = 0$, because every non-zero vector is either maximal or maps onto some non-zero maximal vector via the sheaf maps.

An injective hull is minimal iff the minimal injective resolution starts with it. By Proposition 4.17, this minimality is equivalent to the condition that every maximal vector of I^0 is in $\text{im } \alpha$. This is satisfied, as $M_F(\pi)$ are exactly the maximal vectors in $I^0(\pi)$. \square

We give an explicit formulation of an algorithm computing $\alpha(\pi)$ as Algorithm 6. We first fix bases in F . For each $\pi \in \Pi$, we fix a basis $B(\pi) = (v_1, \dots, v_l, w_{l+1}, \dots, w_{l+d})$, with l, k dependent on π , such that $(w_{l+1}, \dots, w_{l+d})$ is a basis of $M_F(\pi)$ —this involves finding a basis of intersection of kernels and its

Algorithm 6 Minimal injective hull.

Input: F with fixed bases as described above, $\pi \in \Pi$
Output: $\alpha(\pi)$ as a $\left(\sum_{\pi \leq \sigma} \dim M_F(\sigma)\right) \times (\dim F(\pi))$ matrix

```

1: procedure INCL( $\sigma, w \in M_F(\sigma)$ )
2:   return inclusion of  $w$  into  $\bigoplus_{\pi < \tau} M_F(\tau)$  ▷ just adding extra zeros
3: for  $v_j \in \{v_1, \dots, v_l\}$  do ▷  $(v_1, \dots, v_l, w_{l+1}, \dots, w_{l+d})$  is the fixed basis of  $F(\pi)$ 
4:    $D \leftarrow$  empty dictionary ▷ keys: elements  $\sigma \in \Pi$ , values: vectors in  $F(\sigma)$ 
5:    $D[\pi] \leftarrow v_j$ 
6:    $u_j \leftarrow 0$  ▷ vector of length  $\sum_{\pi < \tau} \dim M_F(\tau)$ 
7:   for each  $\sigma \geq \pi$  in non-decreasing order do
8:     if  $\sigma \in \text{Keys}(D)$  and  $D(\sigma) \neq 0$  then ▷  $D[\sigma] = F(\pi \leq \sigma)(v_j)$ 
9:        $w \leftarrow D[\sigma]$ 
10:    for each  $\tau >_1 \sigma$  do
11:      if  $\tau \notin \text{Keys}(D)$  then
12:         $D[\tau] \leftarrow F(\sigma \leq \tau)(w)$ 
13:         $u_j \leftarrow u_j + \text{INCL}(\tau, \text{proj}_{M_F(\tau)}(D[\tau]))$ 
14:    clear  $D[\sigma]$  ▷ optional, just to free up memory
15: return a block matrix  $\begin{pmatrix} U & 0 \\ 0 & I \end{pmatrix}$ , where  $U = (u_1 | \dots | u_l)$ ,
    and  $I$  is the identity matrix of order  $d = \dim M_F(\pi)$ 

```

complement. We use the same bases for $I^0(\pi)$. We assume that all maps $F(\pi \leq \sigma)$ are expressed with respect to those bases.

To express $\alpha(\pi)$ with respect to the fixed bases, we need to find the image of each $v_1, \dots, v_l, w_{l+1}, \dots, w_{l+d}$. The maximal vectors w_j are mapped identically to $M_F(\pi) \subseteq I_0(\pi)$. For the other vectors, v_j , we need to find $u_j := \sum_{\pi < \sigma} \text{proj}_{M_F(\sigma)} F(\pi \leq \sigma)(v_j)$. The algorithm does that while avoiding redundant computations. Each σ is added to D at most once. If it is added, then $D[\sigma] = F(\pi \leq \sigma)(v_j)$, and $\text{proj}_{M_F(\pi)} F(\pi \leq \sigma)(v_j)$ is added to u_j . If σ is never added to D , then $F(\pi \leq \sigma)(v_j) = 0$. In the end, u_j contains the desired sum.

C.4. Geometrical meaning of multiplicities in the minimal resolution

Theorem 4.22. *Let Σ be a finite simplicial complex, k_Σ the constant sheaf on Σ (viewed as a poset with the face relation), and $H_c^\bullet(|\text{St } \sigma|; k)$ be the singular cohomology with compact support of the geometric realization of $\text{St } \sigma$. Then*

$$m_{k_\Sigma}^d(\sigma) = \dim H_c^{d+\dim \sigma}(|\text{St } \sigma|; k).$$

Proof. Let I^\bullet be an injective resolution of the constant sheaf on Σ . By Corollary 5.11,

$$m_{k_\Sigma}^d(\sigma) = \dim H^d \mathcal{I}^!_{\sigma} I^\bullet.$$

For $\sigma \in \Sigma$, let

$$C(\sigma) := \sigma^0 * \{\tau \setminus \sigma \mid \tau \in \text{St } \sigma \setminus \{\sigma\}\}$$

be the cone of the link $\{\tau \setminus \sigma \mid \tau \in \text{St } \sigma \setminus \{\sigma\}\}$ (where $\tau \setminus \sigma$ denotes set difference when τ and σ are viewed as sets of vertices in Σ), with cone point given by the vertex σ^0 . Then $\text{St}_{C(\sigma)} \sigma^0 \cong \text{St}_\Sigma \sigma$ as posets. Therefore, $I^\bullet|_{\text{St}_\Sigma \sigma}$, viewed as a complex of sheaves over $\text{St}_{C(\sigma)} \sigma^0$, is an injective resolution of the constant sheaf on $\text{St}_{C(\sigma)} \sigma^0$. Let J^\bullet be the injective resolution of the constant sheaf on the simplicial complex $C(\sigma)$, defined non-inductively via order complex in Appendix C.2. Then $I^\bullet|_{\text{St}_\Sigma \sigma}$ is quasi-isomorphic to $J^\bullet|_{\text{St}_{C(\sigma)} \sigma^0}$. By definition, $i_{\sigma^0}^!(J^\bullet)$ is exactly the complex of vector spaces used to compute compactly supported cohomology of the constant sheaf on $\text{St}_{K(C(\sigma))}(\sigma^0)$, the star of σ^0 in the barycentric subdivision $K(C(\sigma))$ of $C(\sigma)$ (see [15, Definition 6.2.1] and [49]). Because the geometric realization $|\text{St}_{K(C(\sigma))}(\sigma^0)|$ is homeomorphic to $|\text{St}_{C(\sigma)} \sigma^0|$, the cohomology groups of this complex are isomorphic to $H_c^d(|\text{St}_{C(\sigma)}(\sigma^0)|; k)$. Finally, because $|\text{St}_\Sigma \sigma|$ is

homeomorphic to $\mathbb{R}^{\dim \sigma} \times |\mathrm{St}_{C(\sigma)} \sigma^0|$,

$$m_{k\Sigma}^d(\sigma) = \dim H_c^d(|\mathrm{St}_{C(\sigma)}(\sigma^0)|; k) = \dim H_c^{d+\dim \sigma}(|\mathrm{St}_{\Sigma}(\sigma)|; k). \quad \square$$

References

- [1] Jeffrey Adams, Dan Barbasch, David A. Vogan Jr., *The Langlands Classification and Irreducible Characters for Real Reductive Groups*, Progress in Mathematics, vol. 104, Birkhäuser Boston, Inc., Boston, MA, 1992.
- [2] Jonathan Barmak, *Algebraic Topology of Finite Topological Spaces and Applications*, vol. 2032, Springer, 2011, p. 01.
- [3] Adam Brown, Ondřej Draganov, DESC: derived sheaf calculator, <https://github.com/OnDraganov/desc>, 2022.
- [4] Ulrich Bauer, Herbert Edelsbrunner, The Morse theory of Čech and Delaunay complexes, *Trans. Am. Math. Soc.* 369 (1) (2016).
- [5] Kai Behrend, Donaldson-Thomas type invariants via microlocal geometry, *Ann. Math.* (2) 170 (3) (2009) 1307–1338.
- [6] Nicolas Berkouk, Grégory Ginot, A derived isometry theorem for sheaves, *Adv. Math.* 394 (2022) 108033.
- [7] Nicolas Berkouk, Grégory Ginot, Steve Oudot, Level-sets persistence and sheaf theory, arXiv:1907.09759, 2019.
- [8] Jonathan Barmak, Elias Minian, Simple homotopy types and finite spaces, *Adv. Math.* 218 (2008) 87–104.
- [9] Nicolas Berkouk, François Petit, Ephemeral persistence modules and distance comparison, *Algebr. Geom. Topol.* 21 (1) (2021) 247–277.
- [10] Glen E. Bredon, *Sheaf Theory*, second edition, Graduate Texts in Mathematics, vol. 170, Springer-Verlag, New York, 1997.
- [11] Adam Brown, Bei Wang, Sheaf-theoretic stratification learning from geometric and topological perspectives, *Discrete Comput. Geom.* 65 (4) (2021) 1166–1198.
- [12] Justin Curry, Robert Ghrist, Vidit Nanda, Discrete Morse theory for computing cellular sheaf cohomology, *Found. Comput. Math.* 16 (4) (2016) 875–897.
- [13] Justin Curry, Sayan Mukherjee, Katharine Turner, How many directions determine a shape and other sufficiency results for two topological transforms, *Trans. Am. Math. Soc. Ser. B* 9 (32) (2022) 1006–1043.
- [14] Justin Curry, Amit Patel, Classification of constructible cosheaves, *Theory Appl. Categ.* 35 (2020) 27.
- [15] Justin Michael Curry, *Sheaves, cosheaves and applications*, ProQuest LLC, Ann Arbor, MI, 2014, Thesis (Ph.D.)—University of Pennsylvania.
- [16] Justin Michael Curry, Dualities between cellular sheaves and cosheaves, *J. Pure Appl. Algebra* 222 (4) (2018) 966–993.
- [17] René Deheuvels, Homologie des ensembles ordonnés et des espaces topologiques, *Bull. Soc. Math. Fr.* 90 (1962) 261–321.
- [18] Paweł Dłotko, Hubert Wagner, Simplification of complexes for persistent homology computations, *Homol. Homotopy Appl.* 16 (1) (2014) 49–63.
- [19] Robin Forman, Morse theory for cell complexes, *Adv. Math.* 134 (1) (1998) 90–145.
- [20] Robin Forman, A user's guide to discrete Morse theory, *Sémin. Lothar. Comb.* 48 (2002) B48c.
- [21] Ragnar Freij, Equivariant discrete Morse theory, *Discrete Math.* 309 (12) (2009) 3821–3829.
- [22] Mark Goresky, Robert MacPherson, Morse theory and intersection homology theory, in: *Analysis and Topology on Singular Spaces*, II, III, Luminy, 1981, in: *Astérisque*, vol. 101, Soc. Math. France, Paris, 1983, pp. 135–192.
- [23] Mark Goresky, Robert MacPherson, *Stratified Morse Theory*, *Ergebnisse der Mathematik und Ihrer Grenzgebiete* (3) (Results in Mathematics and Related Areas (3)), vol. 14, Springer-Verlag, Berlin, 1988.
- [24] Mark Goresky, Morse theory, stratifications and sheaves, in: *Handbook of Geometry and Topology of Singularities I*, Springer, 2020, pp. 275–319.
- [25] Mark Goresky, *Lecture Notes on Sheaves and Perverse Sheaves*, 2021.
- [26] E.L. Green, Ø. Solberg, D. Zacharia, Minimal projective resolutions, *Trans. Am. Math. Soc.* 353 (7) (2001) 2915–2939.
- [27] Allen Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.
- [28] Shaun Harker, Konstantin Mischaikow, Marian Mrozek, Vidit Nanda, Discrete Morse theoretic algorithms for computing homology of complexes and maps, *Found. Comput. Math.* 14 (1) (2014) 151–184.
- [29] Birger Iversen, *Cohomology of Sheaves*, Universitext, Springer-Verlag, Berlin, 1986.
- [30] Masaki Kashiwara, Index theorem for constructible sheaves, *Astérisque* 130 (193–209) (1985) 4.
- [31] Masaki Kashiwara, Pierre Schapira, Micro-support des faisceaux: application aux modules différentiels, *C. R. Acad. Sci. Paris, Ser. I Math.* 295 (8) (1982) 487–490.
- [32] Masaki Kashiwara, Pierre Schapira, Microlocal study of sheaves, *Astérisque* 128 (1985) 235, Corrections to this article can be found in *Astérisque* No. 130, p. 209.
- [33] Masaki Kashiwara, Pierre Schapira, *Sheaves on Manifolds*, *Grundlehren der Mathematischen Wissenschaften (Fundamental Principles of Mathematical Sciences)*, vol. 292, Springer-Verlag, Berlin, 1994, With a chapter in French by Christian Houzel, Corrected reprint of the 1990 original.
- [34] Masaki Kashiwara, Pierre Schapira, Persistent homology and microlocal sheaf theory, *J. Appl. Comput. Topol.* 2 (1–2) (2018) 83–113.
- [35] Masaki Kashiwara, Pierre Schapira, Piecewise linear sheaves, *Int. Math. Res. Not.* 2021 (15) (2021) 11565–11584.
- [36] Sefi Ladkani, *Homological properties of finite partially ordered sets*, Hebrew University, 2008, Thesis (Ph.D.)—Hebrew University.
- [37] Sefi Ladkani, On derived equivalences of categories of sheaves over finite posets, *J. Pure Appl. Algebra* 212 (2) (2008) 435–451.
- [38] Jean Leray, Sur la forme des espaces topologiques et sur les points fixes des représentations, *J. Math. Pures Appl.* 9 (24) (1945) 95–167.
- [39] John Milnor, *Morse Theory*, *Annals of Mathematics Studies*, vol. 51, Princeton University Press, Princeton, N.J., 1963, Based on lecture notes by M. Spivak and R. Wells.

- [40] Konstantin Mischaikow, Vidit Nanda, Morse theory for filtrations and efficient computation of persistent homology, *Discrete Comput. Geom.* 50 (2) (2013) 330–353.
- [41] Morse Marston, Relations between the critical points of a real function of n independent variables, *Trans. Am. Math. Soc.* 27 (3) (1925) 345–396.
- [42] Robert MacPherson, Amit Patel, Persistent local systems, *Adv. Math.* 386 (2021) 107795.
- [43] Vidit Nanda, Discrete Morse theory and localization, *J. Pure Appl. Algebra* 223 (2) (2019) 459–488.
- [44] Vidit Nanda, Local cohomology and stratification, *Found. Comput. Math.* 20 (2) (2020) 195–222.
- [45] Vidit Nanda, Dai Tamaki, Kohei Tanaka, Discrete Morse theory and classifying spaces, *Adv. Math.* 340 (2018) 723–790.
- [46] David Nadler, Eric Zaslow, Constructible sheaves and the Fukaya category, *J. Am. Math. Soc.* 22 (1) (2009) 233–286.
- [47] Mikio Sato, Regularity of hyperfunction solutions of partial differential equations, *Actes Congr. Int. Nat. Nice* 2 (1970) 785–794.
- [48] Pierre Schapira, Propagation-set for sheaves and microsupport, Unpublished.
- [49] Allen Shepard, A Cellular Description of the Derived Category of a Stratified Space, Brown University, 1985, Thesis (Ph.D.).
- [50] Mikio Sato, Takahiro Kawai, Masaki Kashiwara, Microfunctions and pseudo-differential equations, in: *Hyperfunctions and Pseudo-Differential Equations*, Proc. Conf., Katata, 1971; Dedicated to the Memory of André Martineau, in: *Lecture Notes in Math.*, vol. 287, Springer, Berlin, 1973, pp. 265–529.
- [51] Emil Sköldbberg, Morse theory from an algebraic viewpoint, *Trans. Am. Math. Soc.* 358 (1) (2006) 115–129.
- [52] Pierre Schapira, Nobuyuki Tose, Morse inequalities for \mathbf{R} -constructible sheaves, *Adv. Math.* 93 (1) (1992) 1–8.
- [53] Edward Witten, Supersymmetry and Morse theory, *J. Differ. Geom.* 17 (4) (1983) 661–692 (1982).