Found Comput Math (2016) 16:875–897
DOI 10.1007/s10208-015-9266-8

FOUNDATIONS OF COMPUTATIONAL MATHEMATICS
The Journal of the Society for the Foundations of Computational Mathematics

Discrete Morse Theory for Computing Cellular Sheaf Cohomology

Justin Curry[1] · **Robert Ghrist**[2,3] · **Vidit Nanda**[2]

Received: 31 December 2013 / Revised: 3 April 2015 / Accepted: 20 April 2015 /
Published online: 20 June 2015

Abstract Sheaves and sheaf cohomology are powerful tools in computational topology, greatly generalizing persistent homology. We develop an algorithm for simplifying the computation of cellular sheaf cohomology via (discrete) Morse theoretic techniques. As a consequence, we derive efficient techniques for distributed computation of (ordinary) cohomology of a cell complex.

Keywords Cellular sheaf cohomology · Discrete Morse theory

Mathematics Subject Classification 55-04 · 55N25 · 55N30

1 Introduction

1.1 Computational Topology and Sheaves

It has recently become clear that computation of homology of spaces is of critical importance in several applied contexts. These includebut are not limited to config-

Communicated by Gunnar Carlsson.

✉ Vidit Nanda
vnanda@sas.upenn.edu

Justin Curry
curry@math.duke.edu

Robert Ghrist
ghrist@math.upenn.edu

1 Department of Mathematics, Duke University, Durham, NC, USA

2 Department of Mathematics, University of Pennsylvania, Philadelphia, PA, USA

3 Department of Electrical/Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA

uration spaces in robotics [24,25,27,33], the global qualitative statistics of point cloud data [13,14,23], coverage problems in sensor networks [18,19], circular coordinates for data sets [20] and Conley-type indices for dynamics [4,34,41]. The Euler characteristic—a numerical reduction of homology—is even more ubiquitous, with applications ranging from Gaussian random fields [1,2] to data aggregation problems over networks [5,6] and signal processing [17]. Not coincidentally, development of applications of homological tools has proceeded symbiotically with the development of good algorithms for computational homology [23,34]. Among the best of the latter are methods based on (co)reduction preprocessing [43] and discrete Morse theory [32].

With the parallel success of new applications and fast computations for homology, additional topological structures and techniques are poised to cross the threshold from theory to computation to application. Among the most promising is the theory of **sheaves**. Developed for applications in algebraic topology and matured under a string of breathtaking advances in algebraic geometry, sheaf theory is perhaps best described as a formalization of local-to-global transitions in Mathematics. The margins of this introductory section do not suffice to outline sheaf theory; rather, we present without detailed explanation three principal interpretations of a sheaf $\mathcal{F}$ over a topological space $\mathcal{X}$ taking values in $\mathbf{R}$-modules over some ring $\mathbf{R}$:

1. A sheaf can be thought of as a *data structure* tethered to a space—an assignment to open sets $\mathcal{V} \subset \mathcal{U}$ of $\mathcal{X}$ a homomorphism $\mathcal{F}(\mathcal{U}) \to \mathcal{F}(\mathcal{V})$ between $\mathbf{R}$-modules—the algebraic "data" over the subsets—in a manner that respects composition and gluing (see Sect. 2). Unlike in the case of a bundle, the data sitting atop subsets of $\mathcal{X}$ can change dramatically from place to place.
2. A sheaf can be thought of as a *topological space* in and of itself, together with a projection map $\pi : \mathcal{F} \to \mathcal{X}$ to the base space $\mathcal{X}$. This *étale space* topologizes the data structure and motivates examining its topological features, such as (co)homology.
3. A sheaf can be thought of as a *coefficient system*, assigning to locations in $\mathcal{X}$ the spatially varying $\mathbf{R}$-module coefficients to be used for computing cohomology. This representation of the space within the algebraic category of $\mathbf{R}$-modules provides enough structure to compute cohomology with location-dependent coefficients.

It is these multiple interpretations that portend the ubiquity of sheaves within applied topology. Though sheaves have long been recognized as useful data structures within certain branches of Computer Science (e.g., [31]), sheaf cohomology has a number of emergent applications. These include:

1. **Signal processing:** Sheaf cohomology recovers and extends the classical Nyquist–Shannon sampling theorem [46]; *viz.*, reconstruction from a sample is possible if and only if the appropriate cohomology of an associated ambiguity sheaf vanishes.
2. **Data aggregation:** Data aggregation over a domain can be performed via Euler integrals, an alternating reduction of the cohomology of an associated constructible sheaf over the domain [17].
3. **Network coding:** Various problems in network coding (maximum throughput, merging of networks, rerouting information flow around a failed subnetwork) have interpretations as ranks of cohomologies of a sheaf over the network [29].

4. **Optimization:** The classical max-flow-min-cut theorem has a sheaf-theoretic analog which phrases flow values and cut values as semimodule images of sheaf homology and cohomology, respectively [30,36].
5. **Complexity:** A recent parallel to the Blum–Shub–Smale theory of complexity [9] has emerged for constructible sheaves [7].

These early examples of applications vary greatly in terms of the types of coefficients used (ranging from $\mathbb{Z}$ to $\mathbb{R}$-vector spaces to general commutative monoids) and the types of base spaces. In most applications, however, the relevant sheaves are of a particular discrete form. Topological spaces become computationally tractable substances through a discretization process: This most often takes the form of a simplicial or cell (or CW) complex. A similar modulation exists for sheaves—a sheaf is called *constructible* with respect to a given stratification of the base space if the data assigned to each stratum is locally constant. We will work in the category of **cellular sheaves**, which are constructible with respect to a fixed regular CW stratification of the base space [28].

Motivated by these applications, we establish algorithms for the computation of sheaf cohomology. Our philosophy, inherited from other work on computational homology [32,34,43] is that of reduction of the input structure to a smaller equivalent structure. We do so by means of **discrete Morse theory**, retooling the machinery to work for sheaves.

1.2 Related and Supporting Work

1.2.1 Sheaves

A fair portion of the existing work on computational topology is naturally cast in the language of sheaves, providing novel paths for generalization. For example, Euler integration—integration with respect to Euler characteristic as a valuation—is sheaf-theoretic in nature and in origin, as per [47,48]. It is in fact the decategorification of the cohomology of sheaves associated to constructible functions ([17] gives an exposition of this). More familiar to the reader will be persistent homology [13,22,56], which also has a sheaf-theoretic formulation as follows. The formal dual of a sheaf is a **cosheaf**; in the cellular category, these are quite useful [21,28] and possess a homology theory [16]. The persistent homology of a filtration is the homology of a (certain) cosheaf over a cell complex homeomorphic to an interval [16]. Recent work on *well groups* associated to persistent homology has been expounded in terms of sheaves [39].

1.2.2 Discrete Morse Theory

Discrete Morse theory [15,26] usually begins with the structure of a partial matching on the cells of a CW complex. The unmatched cells serve the same role as critical points do in smooth Morse theory, while the matched cells furnish gradient-like trajectories between them. A Morse cochain complex may be constructed from this data: its cochain groups are freely generated by the critical cells and the boundary operators

may be derived from gradient paths. The fundamental result is that the Morse cochain complex so obtained is homologically equivalent to the original CW complex.

This basic idea has since been vastly generalized and adapted to purely algebraic situations [8,35,51] with only the slightest vestige of its topological origins. One can impose a partial matching directly on the basis elements of a cochain complex and apply discrete Morse theory as usual. This approach has proved useful in the past when simplifying computation of homology groups of abstract cell complexes [32] and the persistent homology groups of their filtrations [42].

1.3 Problem Statement and Results

Our problem centers on the computation of cellular sheaf cohomology. The initial inputs are a cellular sheaf $\mathcal{F}$ over a CW complex $\mathcal{X}$ taking values in free $\mathbf{R}$-modules for some fixed ring $\mathbf{R}$. This input is reprocessed into a cochain complex $F = (C^\bullet, d^\bullet)$ of free $\mathbf{R}$-modules parameterized by a graded poset $(X, \leq)$ [see Definition 2.3]. Our main algorithm, `Scythe` [see Sect. 4], constructs a $\mathcal{F}$-compatible acyclic matching Σ on $(X, \leq)$ and suitably modifies the coboundary operators $d^\bullet$ in order to cut the original cochain complex down to its critical core while preserving its cohomology. The resulting smaller Morse cochain complex $F^\Sigma = (C^\bullet_\Sigma, d^\bullet_\Sigma)$ is parametrized by the poset of critical elements of Σ.

Let $\prec$ denote the covering relation in our graded poset $(X, \leq)$ and define for each $x \in X$ the set of immediate successors $x_+ = \{y \in X \mid x \prec y\}$. The following parameters measure different aspects of the complexity of F:

1. let n be the cardinality $|X|$ of the poset X,
2. let p equal $\max_{x \in X} \{|x_+|\}$,
3. assume that the maximum rank of $F(x)$ as an $\mathbf{R}$-module is $d < \infty$ for $x \in X$,
4. assume that the matching Σ produced by `Scythe` has m_k critical elements of dimension k and define $\tilde{m} = \sum_k m_k^2$, and
5. define ω to be the matrix multiplication exponent[1] over $\mathbf{R}$.

Note that the first three numbers are input parameters, the fourth is an output parameter and the fifth is purely a property of the underlying coefficient ring $\mathbf{R}$. Our main result is as follows.

Theorem *Let F be a cochain complex of free $\mathbf{R}$-modules over a graded poset $(X, \leq)$ and let n, p, d, m and ω be the associated parameters defined above. Then, the time complexity of constructing the Morse complex F^Σ via* `Scythe` *is $\mathrm{O}(np\tilde{m}d^\omega)$, and the space complexity is $\mathrm{O}(n^2pd^2)$.*

Section 2 contains background material on cellular sheaf theory and the fundamentals of discrete Morse theory. In Sect. 3, we provide explicit chain maps that induce isomorphisms on cohomology between the original and reduced complexes. Section 4 contains a description of the algorithm `Scythe`, a verification of its correctness and also a detailed complexity analysis which proves our main theorem above. Finally,

[1] That is, the complexity of composing two $d \times d$ matrices with $\mathbf{R}$-entries is $\mathrm{O}(d^\omega)$.

in Sect. 5, we develop distributed protocols for calculating traditional cohomology groups of a given space by recasting the computations in appropriate sheaf-theoretic frameworks.

2 Background

In this section, we survey preliminary material pertaining to cellular sheaves [16,50, 53] and a purely algebraic version of discrete Morse theory [8,35,51]. Throughout this paper, $\mathbf{R}$ denotes a fixed coefficient ring with identity $1_{\mathbf{R}}$, while $\mathbb{N}$ and $\mathbb{Z}$ denote the natural numbers and integers, respectively.

2.1 Cellular Sheaves and Their Cohomology

Let $\mathcal{X}$ be a finite regular CW complex consisting of cells and their attaching maps [44,52]. For each $n \in \mathbb{N}$, the subcollection of n-dimensional cells will be written $\mathcal{X}^n$. Given cells σ and τ of $\mathcal{X}$, we write $\sigma \leq \tau$ to indicate the face relation in $\mathcal{X}$. Finally, for each pair of cells σ and τ in $\mathcal{X}$, the quantity $[\sigma : \tau] \in \mathbb{Z}$ is defined to equal

- $+1$ if $\sigma \leq \tau$, $\dim \sigma = \dim \tau - 1$, and the local orientations of their attaching maps agree;
- -1 if $\sigma \leq \tau$, $\dim \sigma = \dim \tau - 1$, and the local orientations disagree; and
- 0 otherwise.

It follows from the usual boundary operator axiom that the following relation must hold across each pair of cells $\sigma, \tau \in \mathcal{X}$:

$$\sum_{\sigma \leq \lambda \leq \tau} [\sigma : \lambda][\lambda : \tau] = 0. \tag{1}$$

Definition 2.1 A *cellular sheaf* $\mathcal{F}$ over $\mathcal{X}$ assigns to each cell σ of $\mathcal{X}$ an $\mathbf{R}$-module $\mathcal{F}(\sigma)$ and to each face relation $\sigma \leq \tau$ an $\mathbf{R}$-linear *restriction map* $\mathcal{F}_{\sigma\tau} : \mathcal{F}(\sigma) \to \mathcal{F}(\tau)$ subject to the following compatibility condition: whenever $\sigma \leq \lambda \leq \tau$ in $\mathcal{X}$, we have $\mathcal{F}_{\lambda\tau} \circ \mathcal{F}_{\sigma\lambda} = \mathcal{F}_{\sigma\tau}$.

Simple examples of sheaves include the following:

1. The *constant sheaf*, $\mathbf{R}_{\mathcal{X}}$, assigns the coefficient ring $\mathbf{R}$ to each cell of $\mathcal{X}$ and the identity restriction map $1_{\mathbf{R}} : \mathbf{R} \to \mathbf{R}$ to each face relation.
2. The *skyscraper sheaf* over a single cell σ of $\mathcal{X}$ is a sheaf, $\mathbf{R}_\sigma$, that evaluates to $\mathbf{R}$ on σ and is zero elsewhere, with all restriction maps being zero.
3. An analog of the skyscraper sheaf over a subcomplex $\mathcal{A} \subset \mathcal{X}$ evaluates to $\mathbf{R}$ on all cells of $\mathcal{A}$ and zero elsewhere. The restriction maps are zero except for the identity map from a cell in $\mathcal{A}$ to a face. This sheaf is best described as the *pushforward* $\iota_* \mathbf{R}_{\mathcal{A}}$ of the constant sheaf on $\mathcal{A}$ induced by the inclusion map $\iota : \mathcal{A} \hookrightarrow \mathcal{X}$. This is *not* the same as the sum of skyscraper sheaves over the cells of $\mathcal{A}$, since the restriction maps are not all zero.

Given any cellular sheaf $\mathcal{F}$ on $\mathcal{X}$, we define the n-th *cochain group over* $\mathcal{F}$ to be the direct sum of the $\mathbf{R}$-modules assigned by $\mathcal{F}$ to the n-dimensional cells. That is,

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

The n-th *coboundary operator* $\delta^n : C^n(\mathcal{X};\mathcal{F}) \to C^{n+1}(\mathcal{X};\mathcal{F})$ is completely determined by the following block action. Given $\sigma \in \mathcal{X}^n$ and $\tau \in \mathcal{X}^{n+1}$, the component of δ^n from $\mathcal{F}(\sigma)$ to $\mathcal{F}(\tau)$ precisely equals $[\sigma : \tau]\mathcal{F}_{\sigma\tau}$ and so we obtain a sequence of $\mathbf{R}$-modules

$$0 \to C^0(\mathcal{X};\mathcal{F}) \xrightarrow{\delta^0} C^1(\mathcal{X};\mathcal{F}) \xrightarrow{\delta^1} C^2(\mathcal{X};\mathcal{F}) \xrightarrow{\delta^2} \cdots$$

It follows from a routine calculation involving (1) and the compatibility condition of Definition 2.1 that $\delta^n \circ \delta^{n-1} = 0$ for all $n \in \mathbb{N}$ and hence that $(C^\bullet(\mathcal{X};\mathcal{F}), \delta^\bullet)$ is a cochain complex.

Definition 2.2 Let $\mathcal{F}$ be a cellular sheaf on $\mathcal{X}$. The *cohomology of* $\mathcal{X}$ *with* $\mathcal{F}$ *coefficients* is defined to be the cohomology of the cochain complex $(C^\bullet(\mathcal{X};\mathcal{F}), \delta^\bullet)$. More precisely,

$$H^n(\mathcal{X};\mathcal{F}) = \frac{\ker \delta^n}{\operatorname{img} \delta^{n-1}}.$$

The reader may interpret $H^\bullet(\mathcal{X};\mathcal{F})$ as the cohomology of the data $\mathcal{F}$ over $\mathcal{X}$. The simple examples of sheaves listed above have the following cohomologies:

1. The constant sheaf $\mathbf{R}_\mathcal{X}$ on $\mathcal{X}$ has cohomology $H^\bullet(\mathcal{X};\mathbf{R}_\mathcal{X}) \cong H^\bullet(\mathcal{X};\mathbf{R})$ equal to ordinary cohomology in $\mathbf{R}$ coefficients.
2. The skyscraper sheaf $\mathbf{R}_\sigma$ on $\mathcal{X}$ has cohomology $H^k(\mathcal{X};\mathbf{R}_\sigma) \cong \mathbf{R}$ when $k = \dim\sigma$ and zero otherwise, illustrating that a sheaf can have trivial cohomology even if the underlying base space is noncontractible.
3. The pushforward sheaf $\iota_*\mathbf{R}_\mathcal{A}$ has cohomology $H^\bullet(\mathcal{X};\iota_*\mathbf{R}_\mathcal{A}) \cong H^\bullet(\mathcal{A};\mathbf{R})$, illustrating that a sheaf can have complicated cohomology even if the underlying base space is contractible.

Of course, more intricate examples abound and are the impetus for an effective algorithm for computation.

2.2 Morse Theory for Parametrized Cochain Complexes

Forman's [26] work on Morse theory for CW complexes has been extended to a purely algebraic framework by Batzies and Welker [8], Kozlov [35], and (in greatest generality) by Sköldberg [51]. The central idea is to exploit invertible restriction maps in order to produce a smaller cochain complex with isomorphic cohomology. In order to establish notation compatible with an algorithmic treatment, we provide a brief overview of the main results here.

Recall that given two elements x, y in a poset $(X, \leq)$, we say that y *covers* x whenever $x < y$ and $\{z \in X \mid x < z < y\} = \varnothing$. We denote this covering relation by $x \prec y$ and call a poset $(X, \leq)$ *graded* if it admits a partition $X = \bigcup_{n \in N} X_n$ into subsets indexed by a dimension so that if $x \prec y$ then $\dim y = \dim x + 1$. All graded posets in sight are assumed to be finite.[2]

Definition 2.3 A *parametrization* F of a cochain complex $(C^\bullet, d^\bullet)$ of $\mathbf{R}$-modules over a graded poset $(X, \leq)$ assigns to each $x \in X$ an $\mathbf{R}$-module $F(x)$ and to each covering relation $x \prec y$ a linear map $F_{xy} : F(x) \to F(y)$ so that for all dimensions $n \in \mathbb{N}$,

1. $C^n = \bigoplus_{x \in X_n} F(x)$, and
2. the block of $d^n : C^n \to C^{n+1}$ from $F(x)$ to $F(y)$ is precisely F_{xy}.

By convention, we require $F_{xy} = 0$ whenever $x \not\prec y$.

Cochain complexes parametrized over posets are the basic objects on which discrete Morse theory operates. Before introducing the details, we remark that the cells of a finite regular CW complex $\mathcal{X}$ comprise a graded poset over which the cochain complex $(C^\bullet(\mathcal{X}; \mathcal{F}), \delta^\bullet)$ associated to any sheaf $\mathcal{F}$ is naturally parametrized. Throughout the remainder of this section, we fix a parametrization F of a cochain complex $(C^\bullet, d^\bullet)$ over a graded poset $(X, \leq)$.

The following definition goes back to the work of Chari [15]

Definition 2.4 A *partial matching* on $(X, \leq)$ is a subset $\Sigma \subset X \times X$ of pairs subject to the following axioms:

1. **dimension:** if $(x, y) \in \Sigma$ then $x \prec y$, and
2. **partition:** if $(x, y) \in \Sigma$ then neither x nor y belong to any other pair in Σ.

Moreover, Σ is called **acyclic** if the transitive closure of the relation $\lhd$ defined on pairs in Σ by

$$(x, y) \lhd (x', y') \text{ if and only if } x \prec y',$$

generates a partial order.

We call an acyclic matching Σ on $(X, \leq)$ **compatible** with the parametrization F if for each pair $(x, y) \in \Sigma$ the associated linear map $F_{xy} : F(x) \to F(y)$ is invertible. Let Σ be such a compatible acyclic matching on $(X, \leq)$ and denote by M the *critical* unpaired elements:

$$M = \{m \in X \mid (m, z) \text{ and } (z, m) \text{ are not in } \Sigma \text{ for any } z \in X\}.$$

A *gradient path* γ of Σ is a strictly $\lhd$-increasing sequence $(x_j, y_j)_1^J \subset \Sigma$ arranged as follows:

$$\gamma = y_1 \succ x_1 \prec y_2 \succ x_2 \prec \cdots \prec y_J \succ x_J,$$

[2] When striving for greater generality, one replaces this requirement by the following local finiteness hypothesis on the covering relation: each $x \in X$ can have only finitely many y so that $y \prec x$ or $x \prec y$.

and its *coindex* $F_\gamma : F(y_1) \to F(x_J)$ is the linear map given by

$$F_\gamma = \left(-F^{-1}_{x_J y_J}\right) \circ F_{x_{J-1} y_J} \circ \cdots \circ F_{x_1 y_2} \circ \left(-F^{-1}_{x_1 y_1}\right). \tag{2}$$

For each gradient path $\gamma = (x_j, y_j)_1^J$, we write $\mathsf{s}_\gamma = y_1$ and $\mathsf{t}_\gamma = x_J$ to indicate the source (first) and target (last) elements. Given critical elements $m, m' \in M$, the path γ is said to flow from m to m' whenever the covering relations $m \prec \mathsf{s}(\gamma)$ and $\mathsf{t}(\gamma) \prec m'$ both hold, and a new linear map $F^\Sigma_{mm'} : F(m) \to F(m')$ may be defined by:

$$F^\Sigma_{mm'} = F_{mm'} + \sum_\gamma F_{\mathsf{t}_\gamma m'} \circ F_\gamma \circ F_{m \mathsf{s}_\gamma}, \tag{3}$$

where the sum is taken over all gradient paths γ of Σ flowing from m to m'. If we write $m <_\Sigma m'$ whenever at least one such path exists, then it follows easily from the acyclicity of Σ that the transitive closure of $<_\Sigma$ furnishes a partial order on M which is graded by dimension.

Definition 2.5 The *Morse data* associated to Σ consists of the poset $(M, \leq_\Sigma)$ of critical elements along with a sequence of $\mathbf{R}$-modules

$$0 \to C^0_\Sigma \xrightarrow{d^0_\Sigma} C^1_\Sigma \xrightarrow{d^1_\Sigma} C^2_\Sigma \xrightarrow{d^2_\Sigma} \cdots$$

where $C^n_\Sigma = \bigoplus_{m \in M_n} F(m)$ and the block of $d^n_\Sigma : C^n_\Sigma \to C^{n+1}_\Sigma$ from $F(m)$ to $F(m')$ is $F^\Sigma_{mm'}$.

The following theorem is (dual to) the main result of algebraic Morse theory.

Theorem 2.6 (Sköldberg, [51]) *Let F parametrize a cochain complex $(C^\bullet, d^\bullet)$ of $\mathbf{R}$-modules over a graded poset $(X, \leq)$ and let Σ be a compatible acyclic matching. Then, the Morse data $(C^\bullet_\Sigma, d^\bullet_\Sigma)$ is a cochain complex parametrized over $(M, \leq_\Sigma)$ by F^Σ. Moreover, there are $\mathbf{R}$-module isomorphisms*

$$H^n(C, d) \cong H^n(C_\Sigma, d_\Sigma),$$

on cohomology for each dimension $n \in \mathbb{N}$.

In the next section, we provide a new proof of Theorem 2.6 by constructing explicit cochain equivalences. This proof leads to a recipe for simplifying cohomology computation for an arbitrary cellular sheaf $\mathcal{F}$ given the existence of efficient techniques for constructing compatible matchings and the Morse data. One imposes an acyclic matching Σ on the graded poset of cells in the underlying regular CW complex $\mathcal{X}$ so that for each $(\sigma, \tau) \in \Sigma$ the restriction map $\mathcal{F}_{\sigma\tau}$ is invertible. If the set M of critical cells is much smaller than X, then one simply computes the cohomology of the smaller cochain complex $(C^\bullet_\Sigma, d^\bullet_\Sigma)$.

3 The Cohomological Morse Equivalence

Let F be a parametrization for a cochain complex $(C^\bullet, d^\bullet)$ over a graded poset $(X, \leq)$ and assume that Σ is a compatible acyclic matching on X. We prove Theorem 2.6 via an inductive argument by removing one Σ-pair at a time from X. By suitably updating the parametrization near the removed pair at each step, it is possible to preserve the cohomology until one converges to the Morse parametrization F^Σ over the poset $(M, \leq_\Sigma)$ of critical elements.

3.1 The Reduction Step

The central idea of reducing a cell pair from a CW complex while preserving its homotopy type (and hence, its cohomology) goes back to the work of Whitehead on combinatorial homotopy [55]. Here, we present a suitable version of this reduction step adapted for cellular sheaves and efficient algorithms.

Fix $(x^\star, y^\star) \in \Sigma$ and define $X^\star = X \setminus \{x^\star, y^\star\}$. A graded partial order $\leq_\star$ may be defined on $X^\star$ via the following covering relation: given any cells w and z in $X^\star$, we have $w \prec_\star z$ if either $w \prec z$ in X or if $w \prec y^\star \succ x^\star \prec z$ in X. One obtains a new parametrization $F^\star$ over the reduced poset $(X^\star, \leq_\star)$ as follows: $F^\star(w) = F(w)$ for all $w \in X^\star$, and for each covering relation $w \prec_\star z$, we have the linear map $F^\star_{wz} : F(w) \to F(z)$ given by

$$F^\star_{wz} = F_{wz} - F_{x^\star z} \circ F^{-1}_{x^\star y^\star} \circ F_{wy^\star}. \tag{4}$$

A routine calculation shows that $F^\star$ parametrizes a cochain complex which we denote by $(C^\bullet_\star, d^\bullet_\star)$; moreover, Σ restricts to an acyclic matching $\Sigma^\star$ on $(X^\star, \leq_\star)$.

Proposition 3.1 *Given the restricted acyclic matching $\Sigma^\star$ defined above,*

1. *$\Sigma^\star$ is compatible with the reduced parametrization $F^\star$, and*
2. *the Morse data associated to $\Sigma^\star$ is identical to that of Σ.*

Proof In fact, for any $(x, y) \in \Sigma^\star$ there is an equality $F_{xy} = F^\star_{xy}$ by (4)—otherwise, we violate the acyclicity of Σ as follows. By (4), the nonzeroness of $F^\star_{xy} - F_{xy}$ implies that $F_{xy^\star}$ and $F_{x^\star y}$ do not vanish, which leads to the contradiction $(x, y) \lhd (x^\star, y^\star) \lhd (x, y)$. To prove the second assertion, first note that the critical elements of Σ and $\Sigma^\star$ are identical; and since $F(m) = F^\star(m)$ for each critical m, one obtains an equality of cochain groups $C^n_\Sigma = C^n_{\Sigma^\star}$ for each dimension $n \in N$. Thus, we turn our attention to the linear maps d_Σ and $d_{\Sigma^\star}$. Given any gradient path $\gamma^\star$ of $\Sigma^\star$, say

$$\gamma^\star = y_1 \succ_\star x_1 \prec_\star \cdots \prec_\star y_J \succ_\star x_J,$$

it follows by acyclicity of Σ that there is at most one index $j \in \{1, \ldots, J-1\}$ for which we may have $(x_j, y_j) \lhd (x^\star, y^\star) \lhd (x_{j+1}, y_{j+1})$. Returning to our path $\gamma^\star$, we therefore conclude that there are only two possibilities. Either there is no index j at

which the removed pair (x, y) might fit, in which case $\gamma^\star$ is also a path of Σ with $F_{\gamma^\star} = F^\star_{\gamma^\star}$ from (2). Alternately, there is a single such index j, in which case Σ may have as its paths both $\gamma^\star$ and the unique augmented path γ given by introducing the removed pair $(x^\star, y^\star)$ in the appropriate spot:

$$\gamma = y_1 \succ x_1 \prec \cdots \prec y_j \succ x_j \prec y^\star \succ x^\star \prec y_{j+1} \succ \cdots \prec y_J \succ x_J.$$

It follows from a quick calculation that $F^\star_{\gamma^\star} = F_{\gamma^\star} + F_\gamma$. In both cases, the sum of coindices over all paths (and hence each block of $d^\bullet_\Sigma$) is preserved. Using this information in (3) and Definition 2.5 concludes the argument. □

As a consequence of this proposition, the Morse complex $(C^\bullet_\Sigma, d^\bullet_\Sigma)$ remains invariant under the reduction step. It remains to show that cohomology is preserved when passing from F to the reduced parametrization $F^\star$.

3.2 Cochain Equivalences

For each $n \in N$, define the linear map $\psi^n : C^n \to C^n_\star$ by the following block action. For $w \in X_n$ and $z \in X^\star_n$, the block $\psi_{wz} : F(w) \to F(z)$ is given by:

$$\psi_{wz} = \begin{cases} -F_{x^\star z} \circ F^{-1}_{x^\star y^\star} & w = y^\star, \\ \mathrm{id}_{F(w)} & w = z, \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

Lemma 3.2 *$\psi^\bullet : C^\bullet \to C^\bullet_\star$ is a cochain map. That is, $\psi^{n+1} \circ d^n = d^n_\star \circ \psi^n$ for each $n \in \mathbb{N}$.*

Proof Given $w \in X_n$ and $z \in X^\star_{n+1}$, we show that the blocks of $\psi^{n+1} \circ d^n$ and $d^n_\star \circ \psi^n$ from $F(w)$ to $F^\star(z) = F(z)$ are identical. More precisely, we wish to establish the following:

$$\sum_{w' \in X_{n+1}} \psi_{w'z} \circ F_{ww'} = \sum_{z' \in X^\star_n} F^\star_{z'z} \circ \psi_{wz'}.$$

By (5), we note that the left side is nonzero only for $w' = z$ or for $w' = y^\star$. Combining these contributions, the left side evaluates to $F_{wz} + \psi_{y^\star z} \circ F_{wy^\star}$ which equals $F^\star_{wz}$. Similarly, the right side of the identity above also reduces to $F^\star_{wz}$ immediately at least when $w \neq y^\star$, so it now suffices to show that this right side equals $F^\star_{y^\star z}$ whenever $w = y^\star$. In this case, we calculate

$$\sum_{z' \in X^\star_n} F^\star_{z'z} \circ \psi_{y^\star z'} = -\sum_{z' \in X^\star_n} F^\star_{z'z} \circ F_{x^\star z'} \circ F^{-1}_{x^\star y^\star}$$

Expanding $F^{\star}_{z'z}$ via (4) and distributing terms gives

$$-\sum_{z'\in X^{\star}_n} F_{z'z}\circ F_{x^{\star}z'}\circ F^{-1}_{x^{\star}y^{\star}}+\sum_{z'\in X^{\star}_n} F_{x^{\star}z}\circ F^{-1}_{x^{\star}y^{\star}}\circ F_{z'y^{\star}}\circ F_{x^{\star}z'}\circ F^{-1}_{x^{\star}y^{\star}}.$$

The second sum is zero: since $x^{\star}\prec y^{\star}$ by Definition 2.4, there is no $z'\in X^{\star}$ satisfying $x^{\star}\prec z'\prec y^{\star}$ and so the summand is always trivial. Finally, one can use the fact that $d^{\bullet}$ is a coboundary operator—in particular, that $\sum_{z'\in X_n} F_{z''z}\circ F_{x^{\star}z'}=0$—to show that the first sum equals $F^{\star}_{y^{\star}z}$ as desired. □

We now require a cochain map in the other direction. To this end, define $\phi^n: C^n_{\star}\to C^n$ by the following block action $\phi_{zw}: F(z)\to F(w)$ for each $z\in X^{\star}_n$ and $w\in X_n$:

$$\phi_{zw}=\begin{cases}-F^{-1}_{x^{\star}y^{\star}}\circ F_{zy^{\star}} & w=x^{\star},\\ \mathrm{id}_{F(w)} & w=z,\\ 0 & \text{otherwise.}\end{cases}\tag{6}$$

Lemma 3.3 *$\phi^{\bullet}: C^{\bullet}_{\star}\to C^{\bullet}$ is a cochain map. That is, $\phi^{n+1}\circ d^n_{\star}=d^n\circ\phi^n$ for each $n\in\mathbb{N}$.*

Proof The argument proceeds very similarly to the one in the proof of Lemma 3.2. Given $z\in X^{\star}_n$ and $w\in X_{n+1}$, we establish a block equivalence by showing that the following identity holds:

$$\sum_{z'\in X^{\star}_{n+1}}\phi_{z'w}\circ F^{\star}_{zz'}=\sum_{w'\in X_n} F_{w'w}\circ\phi_{zw'}.$$

By (6), we note that the right side is nontrivial only when $w'=x^{\star}$ or when $w'=z$, and hence, it reduces to $F_{zw}+F_{x^{\star}z}\circ\phi_{zx^{\star}}$, which equals $F^{\star}_{zw}$. The left side also evaluates to the same quantity whenever it is nontrivial provided that $w\neq x^{\star}$. On the other hand, if $w=x^{\star}$, then the left side becomes

$$\sum_{z'\in X^{\star}_{n+1}}\phi_{z'x^{\star}}\circ F^{\star}_{zz'}=-\sum_{z'\in X^{\star}_{n+1}} F^{-1}_{x^{\star}y^{\star}}\circ F_{z'y^{\star}}\circ F^{\star}_{zz'}.$$

Expanding $F^{\star}_{zz'}$ via (4) and distributing terms yields

$$-\sum_{z'\in X^{\star}_{n+1}} F^{-1}_{x^{\star}y^{\star}}\circ F_{z'y^{\star}}\circ F_{zz'}+\sum_{z'\in X^{\star}_{n+1}} F^{-1}_{x^{\star}y^{\star}}\circ F_{z'y^{\star}}\circ F_{x^{\star}z'}\circ F^{-1}_{x^{\star}y^{\star}}\circ F_{zy^{\star}}.$$

The second sum above is always zero, since $(x^{\star},y^{\star})\in\Sigma$ implies $x^{\star}\prec y^{\star}$, and hence, there is no $z'\in X^{\star}$ with $x^{\star}\prec z'\prec y^{\star}$. Finally, the first sum reduces to $F^{\star}_{x^{\star}z}$ since $d^{\bullet}$ is a coboundary operator, and hence, $\sum_{z'\in X_{n+1}} F_{z'y^{\star}}\circ F_{zz'}=0$.

It is easy to verify that $\psi^n \circ \phi^n$ is the identity map on $C^n_\star$ for each $n \in \mathbb{N}$, so in order to conclude that $\psi^\bullet$ and $\phi^\bullet$ are cochain equivalences it suffices to construct a cochain homotopy $\Theta^n : C^n \to C^{n-1}_\star$ between $\phi^n \circ \psi^n$ and the identity on C^n. The following result completes our proof of Theorem 2.6.

Lemma 3.4 *The linear maps $\Theta^n : C^n \to C^{n-1}$ defined by the block action*

$$\Theta_{ww'} = \begin{cases} F^{-1}_{x^\star y^\star} & w' = x^\star \text{ and } w = y^\star, \\ 0 & \text{otherwise,} \end{cases}$$

constitute a cochain homotopy between $\phi^n \circ \psi^n$ and the identity on C^n for each dimension $n \in \mathbb{N}$.

Proof By definition, it suffices to show $\Theta^{n+1} \circ d^n + d^{n-1} \circ \Theta^n = \mathrm{id}_{C^n} - \phi^n \circ \psi^n$. By (5) and (6), we note that $\phi \circ \psi$ has the following block action $F(w) \to F(\overline{w})$ for $w, \overline{w} \in X_n$:

$$(\phi \circ \psi)_{w\overline{w}} = \begin{cases} -F^{-1}_{x^\star y^\star} \circ F_{wy^\star} & \overline{w} = x^\star, \\ -F_{x^\star \overline{w}} \circ F^{-1}_{x^\star y^\star} & w = y^\star, \\ \mathrm{id}_{F(w)} & w = \overline{w} \in X^\star, \\ 0 & \text{otherwise.} \end{cases}$$

A simple calculation confirms that $\mathrm{id}_{C^n} - \Theta^{n+1} \circ d^n - d^{n-1} \circ \Theta^n$ has precisely the same block action and concludes the proof. □

4 Algorithms

In this section, we describe our algorithm `Scythe` which constructs an acyclic matching Σ on $(X, \leq)$ and iteratively implements the reduction step of Sect. 3 in order to reduce a poset-parametrized cochain complex down to the Morse parametrization. Before turning to the details, we recall our main result. Let F be a parametrization of a cochain complex $(C^\bullet, d^\bullet)$ of free $\mathbf{R}$-modules over a graded poset $(X, \leq)$ whose covering relation is denoted by $\prec$ as usual. For each $x \in X$, we define $x_+ = \{y \in X \mid x \prec y\}$ and similarly $x_- = \{y \in X \mid x \succ y\}$. Assume that the acyclic matching imposed by `Scythe` is called Σ. Recall from Sect. 1 the parameters $n = |X|$, $p = \max_{x \in X} \{|x_+|\}$, $d = \max_x \mathrm{rank}\,(F(x))$, $\tilde{m} = \sum_k m_k^2$, and ω.

Note that n, p and d are input parameters. The net critical elements cardinality $\tilde{m}$ is an output parameter, and the multiplication exponent ω is purely a property of the underlying coefficient ring $\mathbf{R}$. The remainder of this section is dedicated to proving our main theorem.

Theorem 4.1 *Let F parametrize a cochain complex of $\mathbf{R}$-modules over a graded poset $(X, \leq)$ and let n, p, d, $\tilde{m}$ and ω be the parameters defined above. Then, the time complexity of constructing the Morse parametrization F^Σ via* `Scythe` *is $\mathrm{O}(np\tilde{m}d^\omega)$, and the space complexity is $\mathrm{O}(n^2 pd^2)$.*

Algorithm: `Scythe`
In: A parametrization F of a cochain complex over a graded poset $(X, \leq)$
Out: Transforms F to the Morse parametrization F^{Σ}, where Σ is an F-compatible acyclic matching on $(X, \leq)$.

```
01  define a queue Que of X-elements
02  while X has noncritical elements
03      select a minimal noncritical element c of X
04      mark c as critical
05      set Que = ∅
06      enqueue c into Que
07      while Que is nonempty
08          dequeue y from Que
09          if y_− has exactly one noncritical x with F_xy invertible
10              enqueue x_+ \ {y} into Que
11              ReducePair(x, y)
12          end if
13          enqueue y_+ into Que
14      end while
15  end while
```

4.1 Description and Verification

The central idea behind our algorithm is derived from iterated breadth-first search[3] and has been exploited on several occasions in similar but less general computational contexts [42,43]. A minimal element $c \in X$ is chosen arbitrarily and declared critical, and elements $y \in c_+$ are scoured for possible pairings. Such a y comprises a viable candidate for pairing if there is a unique uncritical element $x \in y_-$ so that F_{xy} is invertible. As each such pair is found, the reduction step of Sect. 3.1 is applied and both the poset X as well as the parametrization F are locally modified near the reduced pair (x, y) by the subroutine `ReducePair`. The removal of these pairs creates the possibility of new viable candidates for pairings, and we keep track of them using a queue data structure.

Given a pair $x^\star \prec y^\star$ of elements in X with $F_{x^\star y^\star}$ invertible, `ReducePair` performs the reduction step from Sect. 3.1. The key step of this subroutine is Line 04 which corresponds to updating F-values as described in (4). Minor modifications to `ReducePair` along with a few additional data structures would also allow us to catalog and store the cochain equivalences ψ and ϕ as described in Sect. 3.2.

Proposition 4.2 *The collection of those* $(x, y) \in X \times X$ *for which* `ReducePair`(x, y) *is invoked in Line* 12 *of* `Scythe` *constitutes an* F*-compatible acyclic matching* Σ *on* $(X, \leq)$.

Proof The compatibility of the pairing with the parametrization F is enforced in Line 10 of `Scythe` where we check for the invertibility of F_{xy}. The partial matching axioms of Definition 2.4 are easily seen to be satisfied, so we focus here on proving that Σ is acyclic. Returning again to Line 10, note that we only make a pairing (x, y) when x is the last remaining uncritical element in y_-. Now, any pair (x', y') for which

[3] In principle, any method for constructing acyclic partial matchings on graded posets will suffice, provided that it ensures sheaf compatibility by only matching cell pairs whose restriction maps are invertible.

Algorithm: `ReducePair`
In: A pair $(x^\star, y^\star) \in X \times X$ with $x^\star \prec y^\star$ and $F_{x^\star y^\star}$ invertible
Out: Modifies F according to the reduction step

```
01  for each z ∈ x⋆+ \ {y⋆}
02      for each w ∈ y⋆− \ {x⋆}
03          set w ≺ z
04          replace F_wz by F_wz − F_x⋆z ∘ F^{-1}_{x⋆y⋆} ∘ F_wy⋆
05      end for
06  end for
07  remove x⋆ and y⋆ from X
```

$(x', y') \lhd (x, y)$ must by definition satisfy $x' \prec y$, or equivalently, $x' \in y_-$. Since any such x' is manifestly uncritical, it must already have been removed from X along with its paired element y' before the current pair (x, y) was removed. Thus, the order of pair removal is monotonic with respect to $\lhd$, and so, the collection of removed pairs generates an acyclic matching on X. □

It follows immediately from the preceding proposition and the machinery developed in Sect. 3 that the input parametrization F is modified in-place to the Morse parametrization F^Σ: the input poset $(X, \leq)$ is reduced to the critical poset $(M, \leq_\Sigma)$ and the coboundary operator is suitably updated one pair at a time.

4.2 Complexity Analysis

Before performing a thorough analysis of `Scythe` in terms of the complexity parameters introduced in the previous section, we briefly describe some simplifying assumptions. First, the Queue data structure must be managed so that the inner **while** loop spanning Lines 02 through 14 actually terminates. Whenever an element of X is added to the Queue, it is flagged so that it may not be enqueued again in that iteration of the inner **while** loop. But each time the Queue is reinitialized in Line 05, all these flags are cleared. Moreover, we ensure that aside from the critical cell c chosen in Line 03 of `Scythe`, no other critical cells are enqueued. Finally, for the purposes of analyzing complexity, we make the simplifying assumption that we only enqueue those elements of X whose dimension exceeds $\dim c$ by 1. Although this restriction is unnecessary (and indeed, detrimental to performance) in practice, it greatly simplifies the complexity analysis.

Note that the time complexity of calling `ReducePair` with input (x, y) where $\dim x = k$ is $O(pm_k d^\omega)$ as follows. The cardinality of $x_+ \setminus \{y\}$ is at most p by assumption, and since the set $y_- \setminus \{x\}$ only has critical elements by Line 09 of `Scythe`, its cardinality does not exceed m_k. For each pair w and z of elements from these sets, the matrix algebra of Line 04 incurs a further cost of $O(d^\omega)$: the cost of matrix addition is dominated by the costs of inversion and multiplication, which are $O(d^\omega)$ by assumption. Since the inner **while** loop runs at most n times, its total time complexity is $O(npm_k d^\omega)$ where $k = \dim c$ by virtue of our restricted queuing strategy. Finally, since the outer **while** loop executes precisely once per k-dimensional critical element, the total complexity of `Scythe` evaluates to $O(np\tilde{m}d^\omega)$ as claimed in Theorem 4.1.

Remark 4.3 The following observations have practical significance when simplifying computation of cellular sheaf cohomology via discrete Morse theory.

1. Since the output Morse parametrization F^Σ generated by `Scythe` is again a cochain complex of **R**-modules parametrized by a poset, it is possible to *iterate* the simplification scheme. In particular, one may impose an F^Σ-compatible acyclic matching on the critical poset $(M, \leq_\Sigma)$ and so forth, until the Morse parametrization stabilizes. This stabilization is caused by the eventual depletion of cell pairs which may be compatibly matched. In particular, if there are no invertible subblocks in the matrix representation of the Morse coboundary operator, then no further cell pairs may be matched.
2. There is an obvious *dual* algorithm, `CoScythe`, which processes $(X, \leq)$ from the top-down. In particular, a maximal element $c \in X$ may be initially chosen as critical and one may then search for pairings in the set c_- of elements covered by c.

Turning to issues of memory, we recall that F is transformed in-place to F^Σ. Therefore, the only additional overhead is the Que structure. The cost of storing F itself is $\mathrm{O}(npd^2)$: there are n elements in the underlying poset X; for each $x \in X$, there are at most p elements $y \in X$ satisfying $x \prec y$, and for each of these, we must store at most a $d \times d$ matrix F_{xy}. Moreover, since Que itself may get as large as n for each element of X, our worst-case space complexity evaluates to $\mathrm{O}(n^2 pd^2)$.

5 Applications to Distributed Cohomology Computation

The ability to efficiently compute cellular sheaf cohomology will have implications in those emerging applications [signal processing, sampling, tracking, network coding, optimization, etc.] described in Sect. 1. Given the focus of this paper (on computational cohomology), we do not detail such applications. Instead, we demonstrate an application of sheaf cohomology to the more ubiquitous problem of computing ordinary cohomology over a field. Passing from this to a richer coefficient system can and does facilitate a tremendous simplification of the underlying topological space without loss of cohomology.

There are at least two classical examples of this principle in action: the Čech approach and the Leray approach. We describe these classical computational methods below and then present a sheaf-theoretic unification. This has the effect of giving a unified interpretation of persistent [13,56], zig-zag cohomology [14] and the Mayer–Vietoris blowup [49]—all important recent tools in computational topology.

Remark 5.1 Throughout the remainder of this section, we assume:

1. all topological spaces are compact, Hausdorff and locally contractible;
2. all covers consist of finitely many open subsets; and
3. the coefficient ring **R** is a field.

5.1 The Čech Approach

The following classical approach [3] provides a convenient and ubiquitous combinatorial model for representing unwieldy topological spaces.

Definition 5.2 Let $\mathcal{U}$ be a cover of a topological space X. Its *nerve* $N_{\mathcal{U}}$ is the abstract simplicial complex whose n-dimensional simplices are collections $\sigma = (U_0, \ldots, U_n)$ of cover elements with nonempty *support* $\bigcap_0^n U_j$.

Following the usual conflation of an abstract simplicial complex with its cell complex ("geometric") realization, one has the following (simplified version) of the classical theorem of Leray:

Theorem 5.3 (Nerve Theorem [10,37]) *Given a topological space X and a cover $\mathcal{U}$, if the support $U_\sigma \subset X$ of each $\sigma \in N_{\mathcal{U}}$ is acyclic (i.e., the reduced cohomology $\tilde{H}^\bullet(U_\sigma; \mathbf{R}) = 0$ vanishes), then $H^\bullet(N_{\mathcal{U}}; \mathbf{R}) \cong H^\bullet(X; \mathbf{R})$.*

Typically, the cost of guaranteeing acyclicity of supports is that one has to refine substantially the cover $\mathcal{U}$ and hence greatly increase the number of simplices in $N_{\mathcal{U}}$. The following notion is naturally motivated by the desire to compute cohomology with coarser covers and hence fewer simplices.

Definition 5.4 The *Čech cellular sheaves* $\mathcal{C}^n$ associated to the cover $\mathcal{U}$ of a space X are defined on the nerve $N_{\mathcal{U}}$ by the following data. Each $\sigma \in N_{\mathcal{U}}$ is assigned the $\mathbf{R}$-module $\mathcal{C}^n(\sigma) = H^n(U_\sigma; \mathbf{R})$, and each face relation $\sigma \subset \tau$ is assigned the linear map $\mathcal{C}^n_{\sigma\tau} : H^n(U_\sigma; \mathbf{R}) \to H^n(U_\tau; \mathbf{R})$ arising from the inclusion of supports $U_\tau \hookrightarrow U_\sigma$.

If all simplex supports are acyclic, then $\mathcal{C}^0$ reduces to the constant sheaf on $N_{\mathcal{U}}$ and all other $\mathcal{C}^n$s are trivial; in the absence of acyclicity assumptions, the following result yields a simple correction.

Proposition 5.5 *Let X be a topological space and $\mathcal{U}$ a cover whose nerve $N_{\mathcal{U}}$ is at most one-dimensional. Then, for each $n \in \mathbb{N}$,*

$$H^n(X; \mathbf{R}) \cong H^0(N_{\mathcal{U}}; \mathcal{C}^n) \oplus H^1(N_{\mathcal{U}}; \mathcal{C}^{n-1}). \tag{7}$$

We defer the proof to the next section where a more general result is established, but remark here that similar results have been obtained before [12,14] in the context of zig-zag persistent homology. The central difference between these results and ours is that the existing results depend on the direct sum decomposition of zig-zag persistence modules into indecomposable modules (or barcodes). On the other hand, our result makes the recognition that these modules are conceived as sheaves over a linear nerve and moreover that the cohomology of these sheaves can be quickly computed using discrete Morse theory.

5.2 The Leray Approach

One can try to compute the cohomology of X with $\mathbf{R}$ coefficients from a sufficiently nice map $f : X \to Y$ into some simpler space Y. If the image of f comes equipped

with a cover $\mathcal{V}$ having nerve $N_{\mathcal{V}}$, one can try to pull back the associated Čech sheaf on $N_{\mathcal{V}}$ along f to yield local information about X.

Definition 5.6 The *Leray cellular sheaves* $\mathcal{L}^n$ associated with a map $f : X \to Y$ and a cover $\mathcal{V}$ of $f(X) \subset Y$ are defined over the nerve $N_{\mathcal{V}}$ as follows. Each simplex $\sigma \in N_{\mathcal{V}}$ is assigned the cohomology of the preimage of its support, i.e., $\mathcal{L}^n(\sigma) = H^n(f^{-1}(V_\sigma); \mathbf{R})$; furthermore, each face relation $\sigma \subset \tau$ is assigned the map induced on cohomology by the inclusion $f^{-1}(V_\tau) \hookrightarrow f^{-1}(V_\sigma)$.

In the special case where $X = Y$ and f is the identity map, the Leray sheaves clearly coincide with the Čech sheaves associated to the cover $\mathcal{V}$ of X. Thus, the following result generalizes Proposition 5.5.

Theorem 5.7 *Let $f : X \to Y$ be continuous. Assume a cover $\mathcal{V}$ of the image $f(X) \subset Y$ whose nerve $N_{\mathcal{V}}$ is at most one-dimensional. Then, for each $n \in \mathbb{N}$,*

$$H^n(X; \mathbf{R}) \cong H^0(N_{\mathcal{V}}; \mathcal{L}^n) \oplus H^1(N_{\mathcal{V}}; \mathcal{L}^{n-1}). \tag{8}$$

Proof The theorem is a simple consequence of the Leray spectral sequence which packages the cohomology of X into a coefficient system over the space Y from a map $f : X \to Y$ [40]. The restriction to a one-dimensional nerve forces the spectral sequence to collapse on the second page and hence yield the desired isomorphisms. More precisely, for each open $V \subset f(X)$, let $C^n(V; \mathbf{R})$ denote the $\mathbf{R}$-module freely generated by the set of all cochains defined on V. Clearly, if $V \subset U$, then there is a surjection $C^n(U; \mathbf{R}) \to C^n(V; \mathbf{R})$ defined by restriction of cochains. The sheaf $\mathcal{F}$ associated with this presheaf of singular cochains is consequently *flabby* (see [45, p. 97]).

Consider the following double complex of $\mathbf{R}$-modules:

$$\begin{array}{ccccccc}
\vdots & & \vdots & & \vdots & & \vdots \\
\uparrow & & \uparrow & & \uparrow & & \\
C^2(X) & \longrightarrow & \bigoplus_{\dim\sigma=0} \mathcal{F}^2(f^{-1}(V_\sigma)) & \longrightarrow & \bigoplus_{\dim\tau=1} \mathcal{F}^2(f^{-1}(V_\tau)) & \longrightarrow & 0 \\
\uparrow & & \uparrow & & \uparrow & & \\
C^1(X) & \longrightarrow & \bigoplus_{\dim\sigma=0} \mathcal{F}^1(f^{-1}(V_\sigma)) & \longrightarrow & \bigoplus_{\dim\tau=1} \mathcal{F}^1(f^{-1}(V_\tau)) & \longrightarrow & 0 \\
\uparrow & & \uparrow & & \uparrow & & \\
C^0(X) & \longrightarrow & \bigoplus_{\dim\sigma=0} \mathcal{F}^0(f^{-1}(V_\sigma)) & \longrightarrow & \bigoplus_{\dim\tau=1} \mathcal{F}^0(f^{-1}(V_\tau)) & \longrightarrow & 0
\end{array}$$

It follows from standard results [11, Thm II.5.5, Thm III.4.13] that the rows are exact. By the acyclic assembly lemma [54], the spectral sequence converges to the cohomology of the leftmost column, i.e., $H^\bullet(X; \mathbf{R})$. If one takes cohomology in the vertical direction, one obtains the defined cochain groups associated to the Leray cellular sheaves $\mathcal{L}^n$:

$$\begin{array}{ccccc}
\vdots & & \vdots & & \vdots \\
\bigoplus_{\dim\sigma=0} H^2(f^{-1}(V_\sigma)) & \longrightarrow & \bigoplus_{\dim\tau=1} H^2(f^{-1}(V_\tau)) & \longrightarrow & 0 \\
\bigoplus_{\dim\sigma=0} H^1(f^{-1}(V_\sigma)) & \longrightarrow & \bigoplus_{\dim\tau=1} H^1(f^{-1}(V_\tau)) & \longrightarrow & 0 \\
\bigoplus_{\dim\sigma=0} H^0(f^{-1}(V_\sigma)) & \longrightarrow & \bigoplus_{\dim\tau=1} H^0(f^{-1}(V_\tau)) & \longrightarrow & 0
\end{array}$$

Taking cohomology horizontally corresponds precisely to computing separately (in parallel, if one wishes) the cohomology of the Leray sheaves $\mathcal{L}^n$ over $N_{\mathcal{V}}$, thus producing the final stable page of the spectral sequence.

$$\begin{array}{ccccc}
\vdots & & \vdots & & \vdots \\
H^0(N_{\mathcal{V}};\mathcal{L}^2) & & H^1(N_{\mathcal{V}};\mathcal{L}^2) & \nearrow & 0 \\
H^0(N_{\mathcal{V}};\mathcal{L}^1) & & H^1(N_{\mathcal{V}};\mathcal{L}^1) & \nearrow & 0 \\
H^0(N_{\mathcal{V}};\mathcal{L}^0) & & H^1(N_{\mathcal{V}};\mathcal{L}^0) & & 0
\end{array}$$

Over a general ring $\mathbf{R}$, these terms prescribe a filtration of the cohomology, giving rise to extension problems; however, over a field one can read off the cohomology directly. □

Note that the proof indicates precisely where we require the one-dimensional nerve restriction. Without this assumption in place, the second page of the spectral sequence may not be stable and the conclusion of the theorem need not hold.

5.3 An Example

A fairly natural situation where computing cohomology via Theorem 5.7 is advantageous over the obvious alternatives arises when dealing with **Reeb graphs**. Consider a topological space X equipped with a function $f : X \to \mathbb{R}$ and recall that the Reeb graph of the pair (X, f) is a quotient of X by the equivalence relation which identifies two points whenever they lie in the same connected component of $f^{-1}(c)$ for some $c \in \mathbb{R}$.

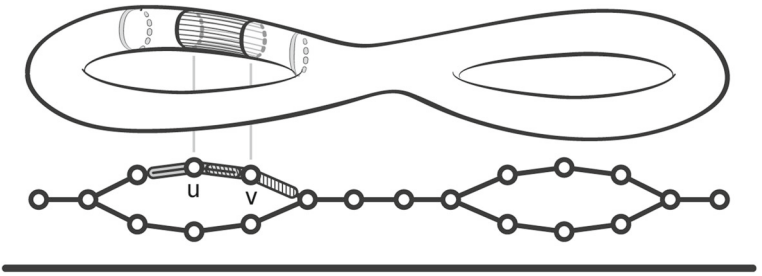

Fig. 1 A genus-2 surface X hovers over (a subdivision of) its Reeb graph Γ associated with downward projection to the real line $\mathbb{R}$. The fibers (lying in X) over nodes u and v of Γ are highlighted, and their intersection comprises the fiber over the edge uv

Let X be a finite CW complex, and consider a continuous function $f : X \to \mathbb{R}$. Given the Reeb graph Γ of (X, f)—for instance, the one illustrated in Fig. 1—one can immediately transform the problem of computing $H^\bullet(X; \mathbf{R})$ to that of computing $H^\bullet(\Gamma; \mathcal{L})$, where $\mathcal{L}$ is the Leray cellular sheaf on a suitable subdivision of Γ associated with the canonical projection $P : X \to \Gamma$. In particular, Theorem 5.7 asserts an isomorphism

$$H^n(X; \mathbf{R}) \cong H^0(\Gamma; \mathcal{L}^n) \oplus H^1(\Gamma; \mathcal{L}^{n-1}),$$

and in cases where P distributes the cells of X almost evenly over those of Γ, it is computationally prudent to evaluate the right side in order to determine the left. In order to estimate the advantage, we employ the following complexity parameters:

1. N is the number of cells in X,
2. d is the dimension of (the maximal cells in) X,
3. g is the number of cells (vertices and edges) in Γ, and
4. $K \leq N$ bounds the number of cells in $P^{-1}(v) \subset X$ across vertices $v \in \Gamma$.

In addition to the usual cost of computing $H^\bullet(\Gamma; \mathcal{L})$, one must also take into account the burden incurred when extracting the data which determines $\mathcal{L}$, i.e., the stalks and restriction maps. To this end, note that the cost of computing a stalk $\mathcal{L}^\bullet(v) = H^\bullet(P^{-1}(v); \mathbf{R})$ over a vertex v of Γ is $O(K^3)$ via Smith diagonalization of a matrix no larger than $K \times K$ in size. Similarly, each stalk $\mathcal{L}(e)$ over an edge e and each restriction map $\mathcal{L}(v) \to \mathcal{L}(e)$ may be evaluated in $O(K^3)$ time since all matrices involved have their sizes bounded above by $K \times K$. More importantly, these local stalk and restriction map computations may be performed in parallel (there are twice as many restriction maps to compute as there are edges in Γ), and hence, the total cost of computing all the $\mathcal{L}$ sheaf data is no more than $O(K^3)$.

Turning now to the computation of sheaf cohomology $H^\bullet(\Gamma; \mathcal{L})$, we note that the relevant cochain complex

$$0 \to \bigoplus_{\dim v=0} \mathcal{L}(v) \xrightarrow{\delta} \bigoplus_{\dim e=1} \mathcal{L}(e) \to 0 \to 0 \to \cdots$$

contains only two interesting cochain groups (parametrized by the vertices and edges of Γ respectively) and a single (potentially) nontrivial coboundary map δ between them. Here, the matrix representation of δ consists of at most $g \times g$ blocks arising from restriction maps over incidence relations of cells in Γ. But each such restriction map furnishes a block no larger than $d \times d$ in size—after all, the domain and codomain of the restrictions are cohomologies of subcomplexes of X, and X itself has dimension d. Thus, the matrix representation of δ has size bounded above by $gd \times gd$. Even in the complete absence of Morse theoretic simplification, one may therefore evaluate $H^\bullet(\Gamma; \mathcal{L})$ in $O(g^3 d^3)$ time. Adding the $O(K^3)$ cost of computing $\mathcal{L}$ data, we confront a combined complexity of $O(K^3 + g^3 d^3)$ for building the Leray sheaf of $P : X \to \Gamma$ in parallel and evaluating its cohomology.

Thus, the sheaf-cohomological method of computing $H^\bullet(X, \mathbf{R})$ is much faster than the traditional methods whenever one has $K^3 + g^3 d^3 \ll N^3$. In particular, this inequality holds when two mild conditions are satisfied by P:

1. P distributes the cells of X evenly across those of Γ, so $N \approx Kg$, and
2. the fibers of P have small cohomology relative to their size, i.e., $d \ll K$.

With these assumptions in place, it is straightforward to estimate the ratio r of worst-case complexity when using the Leray sheaf of P to that of directly computing $H^\bullet(X; \mathbf{R})$. Clearly, we have

$$r = \frac{K^3 + g^3 d^3}{N^3} = \frac{K^3}{N^3} + \frac{g^3 d^3}{N^3}.$$

Using $N \approx Kg$ twice, we have

$$r \approx \frac{1}{g^3} + \frac{d^3}{K^3}.$$

Since g may be increased by subdivision and since $d \ll K$ by assumption on the fibers, $r \ll 1$ and the sheaf-theoretic approach enjoys a significant speedup.

5.4 A Unifying Perspective

There is a more sophisticated version of the nerve described originally by Segal [49] which is homotopically faithful to the underlying space independent of the particulars of the cover. This notion has been used in recent applications [57] and parallelizations for homology computation [38].

Definition 5.8 Let X be a topological space equipped with a cover $\mathcal{U}$ with nerve $N_\mathcal{U}$. The *Mayer Vietoris blowup* $M_\mathcal{U}$ associated to $\mathcal{U}$ is a subset of the product $X \times N_\mathcal{U}$ defined as follows. The pair (x, s) lies in $M_\mathcal{U}$ if and only if there is some simplex $\sigma \in N_\mathcal{U}$ for which $x \in U_\sigma$ and $s \in \sigma$.

Being a subset of the product, $M_{\mathcal{U}}$ is equipped with natural surjective projection maps

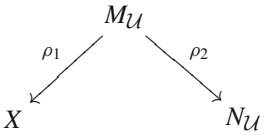

The map ρ_1 has contractible fibers: for any $x \in X$, we have $\rho_1^{-1}(x) = \{x\} \times \sigma_x$ where σ_x is the unique simplex of maximal dimension whose support contains x. Thus, the Mayer–Vietoris blowup is homotopy equivalent to X via ρ_1 in full generality. On the other hand, it is easy to see that the map ρ_2 fails to have contractible fibers precisely when the simplex supports are not contractible. In fact, given $s \in N_{\mathcal{U}}$, the fiber $\rho_2^{-1}(s)$ has the homotopy type of the support of σ_s, which is the unique simplex of maximal dimension whose realization contains s. Since cohomology is a homotopy invariant, this leads to the following observation which unifies the Čech and Leray approaches.

Proposition 5.9 *The Leray cellular sheaves $\mathcal{L}^n$ associated to the map $\rho_2 : M_{\mathcal{U}} \to N_{\mathcal{U}}$, where $N_{\mathcal{U}}$ is covered by (small neighborhoods of the topological) simplices $\{\sigma\}_{\sigma \in N_{\mathcal{U}}}$, are isomorphic to the Čech cellular sheaves $\mathcal{C}^n$ associated to the cover $\mathcal{U}$.*

Remark 5.10 We conclude with the following remarks.

1. The commonality between the Čech and Leray approaches comes as no surprise to anyone sufficiently familiar with spectral sequences (and would have surprised neither Čech nor Leray).
2. Both strategies are examples of *distributed* cohomology computation because in order to determine the sheaf $\mathcal{C}^n$ or $\mathcal{L}^n$, one only needs to compute cohomology locally: of a nontrivial intersection of covering sets in the former case or of a small neighborhood of the fiber $f^{-1}(y)$ in the latter case. In principle, one can assign each local computation to a different processor, compute the appropriate sheaf cohomology over a decidedly nicer space (either $N_{\mathcal{U}}$ or Y depending on the circumstances) and aggregate this information to compute the desired cohomology of X.
3. By taking the appropriate linear duals and working with cosheaves [16], all of our results transform to computations of homology rather than cohomology.

Acknowledgments This work was supported in part by federal contracts FA9550-12-1-0416, FA9550-09-1-0643, and HQ0034-12-C-0027.

References

1. R. Adler, *The Geometry of Random Fields*, (Wiley, 1981 and reprinted by SIAM, 2010).
2. R. Adler and J.E. Taylor, *Random Fields and Geometry* (Springer, 2009).
3. P. Alexandroff. Über den allgemeinen Dimensionsbegriff und seine Beziehungen zur elementaren geometrischen Anschauung. *Math. Ann.*, **98**, 617–635 (1928).

4. Z. Arai, W. Kalies, H. Kokubu, K. Mischaikow, H. Oka, and Pl. Pilarczyk, A Database Schema for the Analysis of Global Dynamics of Multiparameter Systems, *SIAM J. Appl. Dyn. Syst.*, **8**(3), 757–789 (2009).
5. Y. Baryshnikov and R. Ghrist, Target enumeration via Euler characteristic integrals, *SIAM J. Appl. Math.*, **70**(3), 825–844 (2009).
6. Y. Baryshnikov and R. Ghrist, Euler integration over definable functions, *Proc. Natl. Acad. Sci. USA***107**(21), 9525–9530 (2010).
7. S. Basu, A complexity theory of constructible functions and sheaves, *Found. Comput. Math.*, **15**(1), 199–279 (2015).
8. E. Batzies and V. Welker. Discrete Morse theory for cellular resolutions. *J. Reine Angew. Math.*, **543**:147–168 (2002).
9. L. Blum, M. Shub, and S. Smale, On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines, *Bull. Amer. Math. Soc. (N.S.)***21**(1), 1–46 (1989).
10. K. Borsuk. On the imbedding of systems of compacta in simplicial complexes. *Fund. Math.***35**, 217–234 (1948).
11. G. E. Bredon. *Sheaf Theory*, (Springer, 1997)
12. D. Burghelea and T. K. Dey. Topological persistence for circle-valued maps. *Discrete and Computational Geometry*, **50**(1):1–30 (2011).
13. G. Carlsson. Topology and data. *Bull. Amer. Math. Soc. (N.S.)*, **46**(2), 255–308 (2009).
14. G. Carlsson, V. de Silva, and D. Morozov. Zigzag persistent homology and real-valued functions. *Proc. Ann. Sympos. Comp. Geom.*, 247–256 (2009).
15. M. K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Math.*, **217**(1-3) 101–113 (2000).
16. J. Curry. Sheaves, cosheaves and applications. arXiv:1303.3255[math.AT](2013).
17. J. Curry, R. Ghrist, and M. Robinson. Euler calculus and its applications to signals and sensing. *Proc. Sympos. Appl. Math.***70**, 75–145 (2012).
18. V. de Silva and R. Ghrist. Coordinate-free coverage in sensor networks with controlled boundaries via homology, *Intl. J. Robotics Research***25**, 1205–1222 (2006).
19. V. de Silva and R. Ghrist. Coverage in sensor networks via persistent homology, *Algebr. Geom. Topol.*, **7**, 339–358 (2007).
20. V. de Silva, D. Morozov, and M. Vejdemo-Johansson. Persistent Cohomology and Circular Coordinates. *Discrete Comput. Geom.*, **45**(4), 737–759 (2011).
21. V. de Silva, E. Munch, and A. Patel. Categorified Reeb graphs. arXiv:1501.04147[cs.CG], (2015).
22. H. Edelsbrunner, D. Letscher, and A. Zomorodian. Topological persistence and simplification. *Discrete Comput. Geom.*, **28**(4):511–533 (2002).
23. H. Edelsbrunner, J. Harer, *Computational Topology. An Introduction*, (American Mathematical Society, 2010).
24. M. Farber. Topological Complexity of Motion Planning. *Discrete Comput. Geom.***29**, 211–221 (2003).
25. M. Farber. Collision free motion planning on graphs. In *Algorithmic Foundations of Robotics IV*, (M. Erdmann, D. Hsu, M. Overmars, A. F. van der Stappen eds.), Springer, 2005, pp. 123–138.
26. R. Forman. Morse theory for cell complexes. *Advances in Mathematics*, **134**(1), 90–145 (1998).
27. R. Ghrist, Configuration spaces, braids, and robotics. *Lecture Note Series, Inst. Math. Sci., NUS*, vol. 19, World Scientific, 263–304 (2010).
28. R. Ghrist, *Elementary Applied Topology*, (Createspace, 2014).
29. R. Ghrist and Y. Hiraoka. Sheaves for network coding. In *Proc. NOLTA: Nonlinear Theory and Applications*, 266–269, (2011).
30. R. Ghrist and S. Krishnan. A topological max-flow-min-cut theorem. In *Proc. Global Sig. Inf. Proc.*, 815–818 (2013).
31. J. A. Goguen. Sheaf semantics for concurrent interacting objects. *Mathematical Structures in Computer Science*, **2**(2) 159–191, (1992).
32. S. Harker, K. Mischaikow, M. Mrozek, and V. Nanda. Discrete Morse theoretic algorithms for computing homology of complexes and maps. *Found. Comput.l Math.***14**(1), 151–184 (2014).
33. G. Haynes, F. Cohen, and D. Koditschek. Gait Transitions for Quasi-Static Hexapedal Locomotion on Level Ground. in *International Symposium of Robotics Research*, Springer, 2011, pp 105–121.
34. T. Kaczynski, K. Mischaikow, and M. Mrozek. *Computational Homology* (Springer-Verlag, 2004).
35. D. Kozlov. Discrete Morse theory for free chain complexes. *C. R. Math.*, **340**(12), 867–872 (2005).

36. S. Krishnan. Flow-cut duality for sheaves on graphs. arXiv:1409.6712[math.AT], (2014).
37. J. Leray. Sur la forme des espaces topologiques et sur les points fixes des représentations. *J. Math. Pures Appl.*, **24**(9), 95–167 (1945).
38. R. H. Lewis and A. Zomorodian. Multicore homology. arXiv:1407.2275 [cs.CG], (2014).
39. R. MacPherson and A. Patel. Private communication, 2013.
40. J. McCleary. *A User's Guide to Spectral Sequences*, (Cambridge University Press, 2001).
41. K. Mischaikow and M. Mrozek, Conley Index Theory. In *Handbook of Dynamical Systems II: Towards Applications*, (B. Fiedler, ed.) North-Holland, 2002, pp 393–460.
42. K. Mischaikow and V. Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete Comput. Geom.*, **50**(2), 330–353 (2013).
43. M. Mrozek and B. Batko. The coreduction homology algorithm. *Discrete Comput. Geom.*, **41**(1), 96–118 (2009).
44. J. Munkres. *Elements of Algebraic Topology*. (Benjamin/Cummings, 1984).
45. S. Ramanan. *Global Calculus*. (American Mathematical Society, 2005).
46. M. Robinson. The Nyquist theorem for cellular sheaves. *Proc. Sampling Theory and Applications*, 2013, pp 293–296.
47. P. Schapira. Operations on constructible functions. *J. Pure Appl. Algebra*, **72**(1), 83–93 (1991).
48. P. Schapira. Tomography of constructible functions. In *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes*, Springer, 1995, pp. 427–435.
49. G. Segal. Classifying spaces and spectral sequences. *Inst. Hautes Études Sci. Publ. Math.*, **34**, 105–112, (1968).
50. A. Shepard. *A Cellular Description of the Derived Category of a Stratified Space*. Brown University PhD Thesis, 1985.
51. E. Sköldberg. Morse theory from an algebraic viewpoint. *Trans. Amer. Math. Soc.*, **358**(1), 115–129 (2006).
52. E. H. Spanier. *Algebraic Topology*. (McGraw-Hill, 1966).
53. M. Vybornov. Sheaves on triangulated spaces and Koszul duality. arXiv:math/9910150[math.AT], 2000.
54. C. A. Weibel. *An Introduction to Homological Algebra*, (Cambridge University Press, 1995).
55. J. H. C. Whitehead. Combinatorial homotopy I. *Trans. Amer. Math. Soc.*, **55**(5), 453–496 (1949).
56. A. Zomorodian and G. Carlsson. Computing persistent homology. *Discrete Comput. Geom.*, **33**(2), 249–274 (2005).
57. A. Zomorodian and G. Carlsson. Localized homology. *Comput. Geom.*, **41**(3), 126–148, (2008).