

Cellular Bisheaves and Higher Persistence

Vidit Nanda and Amit Patel

Introduction

Our goal in this paper is to describe the combinatorial analogue of a remarkable object which quantifies the algebraic-topological stability of fibers of certain sufficiently tame functions $f : \mathbb{X} \rightarrow \mathbb{M}$. Here $\mathbb{X}$ a cell complex, $\mathbb{M}$ is an oriented manifold, and the fiber of f over some $U \subset \mathbb{M}$ is $f^{-1}(U) \subset \mathbb{X}$. Assuming that $\mathbb{M}$ has a metric $[\cdot, \cdot]$ compatible with its topology, we seek efficiently computable answers to the following question:

if $g : \mathbb{X} \rightarrow \mathbb{M}$ is a function satisfying $\sup_{x \in \mathbb{X}} [f(x), g(x)] < \delta$ (for some small $\delta > 0$), then how can one relate the homology of the fibers of f to the homology of the corresponding fibers of g ? In particular, which homology classes in $f^{-1}(U)$ are guaranteed to persist in $g^{-1}(U)$ for each $U \subset \mathbb{M}$?

Readers who consider this an intrinsically compelling quest are encouraged to begin their journey at Sec 1. The remainder of this introduction is aimed at those who prefer to view such problems through the lens of *persistent homology*, which provides a thoroughly satisfying answer to our question whenever the target manifold $\mathbb{M}$ is $\mathbb{R}$ (i.e., the real line equipped with its standard metric).

Persistent Homology. To each continuous map $f : \mathbb{X} \rightarrow \mathbb{R}$, one can associate an $\mathbb{R}$ -indexed family of (graded) vector spaces via the assignment of sublevelset singular homology with coefficients in some field k :

$$s \mapsto V_f(s) = H_\bullet(\{x \in \mathbb{X} \mid f(x) \leq s\}; k).$$

Inclusion of sublevelsets induces linear maps $V_f(s) \rightarrow V_f(t)$ whenever $s \leq t$. Such a collection of vector spaces and their maps (equivalently, a functor from the poset of real numbers to the category of k -vector spaces), is called a *persistence module*.

Three miracles conspire to render the study of persistence modules tractable; we describe them here for the most familiar and motivating example of modules V_f which arise from the sublevelset homology with k coefficients of a sufficiently tame function $f : \mathbb{X} \rightarrow \mathbb{R}$ (e.g., a piecewise linear map from a finite CW complex, or a Morse function on a compact Riemannian manifold).

- (1) **Structure:** there is a complete combinatorial invariant $\text{Bar}(f)$ of V_f called its *barcode* — here completeness means that two persistence modules are isomorphic if and only if their barcodes coincide.

- (2) **Computability:** if $\mathbb{X}$ is a finite cell complex filtered by locally constant f -values on cells, then $\text{Bar}(f)$ is computable via standard linear algebra (namely, boundary matrix diagonalization).
- (3) **Stability:** there is a natural metric on the set of barcodes under which the assignment $f \mapsto \text{Bar}(f)$ constitutes a 1-Lipschitz map from the tame functions in $L^\infty(\mathbb{X})$ to barcode-space.

Each barcode is a list of subintervals $[b, d)$ of $\mathbb{R} \cup \{\pm\infty\}$, and the dimension of the k -vector space $H_\bullet(f^{-1}(U); k)$ equals the number of intervals in $\text{Bar}(f)$ which contain $U \subset \mathbb{R}$. Stability helps solve our fiber problem as follows. If $g : \mathbb{X} \rightarrow \mathbb{R}$ satisfies $\sup_{\mathbb{X}} |f(x) - g(x)| < \delta$, then only those intervals $[b, d)$ in $\text{Bar}(f)$ which satisfy $U \subset [b + \delta, d - \delta)$ represent homology classes in $f^{-1}(U)$ that are guaranteed to survive in $g^{-1}(U)$.

The structure theorem for tame persistence modules admits a nice representation-theoretic description. The barcode of such a module is (equivalent to) the list of indecomposable summands obtained by discretizing it into a quiver representation whose underlying graph is a type A Dynkin diagram:




By a celebrated result of Gabriel [10], a connected quiver is of *finite type* (meaning that its indecomposable finite-dimensional representations fall into finitely many isomorphism classes) if and only if the underlying graph is a Dynkin diagram of type A, D or E. Moreover, the fact that barcodes can be easily computed via matrix diagonalization follows directly from the observation that if our quiver has all its arrows pointing in the same direction, then its k -vector space valued representations constitute finitely generated modules over the polynomial ring $k[t]$ (which happens to be an Euclidean domain). Thus, our first two miracles appear, at least on the surface, to exist for purely algebraic reasons.

In sharp contrast, reading the arguments which first established the stability theorem for tame persistence modules [6] is a viscerally geometric experience. The most powerful approach to date [5] lifts the barcode-distance to an *interleaving distance* on the space of all persistence modules, and verifies that the resulting metric space admits geodesics. Consequently, if two persistence modules have interleaving distance δ , then for each $t \in [0, \delta]$ there is an intermediate module at distance t from one and $(\delta - t)$ from the other. This *interpolation lemma* has frequently appeared in various guises at the heart of stability and related arguments — see [6, Sec 3.3], [5, Ch 3.4] or [3]. Thus, stability compels us to retreat to the geometric perspective where persistence modules are $\mathbb{R}$ -indexed k -vector spaces, and does not fit naturally within discrete frameworks arising from representations of quivers or modules over polynomial rings.

Higher Persistence. We find it difficult, as card-carrying mathematicians, to witness these three miracle maps and not immediately try to recreate them for maps with more exotic targets (than $\mathbb{R}$). Certainly, we are far from alone — considerable efforts have been recently expended on discovering barcode-like invariants of *higher persistence modules* which arise naturally as sublevelset homologies of tame maps to

$\mathbb{R}^n$ for larger n . The case $n = 2$ is already hard, and contains all the essential difficulties encountered when confronting arbitrary $n > 1$. In our view, the primary source of these troubles is the inherent tension between the algebraic nature of structure and computability on one hand and the inextricable link of stability with geometry on the other.

In any event, the story of higher persistence modules often begins with a grid of two (or more) dimensions whose rows and columns are type-A quivers:



One requires every square in sight to commute, so this becomes a *quiver with relations*, and as such it falls outside the purview of Gabriel's theorem. Even so, it is not of finite type — when all vertical edges point up and all horizontal edges point right, one can view each representation as a module over the polynomial ring $k[s, t]$ in two commuting variables over a field. Since this ring is manifestly not a principal ideal domain, finitely generated $k[s, t]$ -modules are not guaranteed to decompose into pieces sourced from some finite set¹. This lack of higher barcodes was already quantified and lamented in the earliest work on higher persistence [4]; its authors also proffered (as an incomplete but computable alternative) the discrete *rank invariant*, which has been extensively studied thereafter [17, 15].

A plethora of other creative and intricate invariants have been subsequently proposed. The interleaving distance was extended to higher persistence modules in [16], so it is possible to check the stability of these invariants even when the underlying module does not arise from the fiberwise homology of a given function. Here are four examples which we hope will convey the remarkable breadth of ideas that have recently gone into the search for new invariants:

- (1) The authors of [24] have obtained invariants for tame $\mathbb{Q}^n$ -indexed persistence modules by providing a set of algebraic axioms for *noise systems*, which induce new interleaving-type distances between modules. The result of these investigations is a *feature-counting* invariant, which is stable in a certain sense, but NP-hard to compute in general.
- (2) In [11], Hilbert series and prime ideals associated to $k[t_1, \dots, t_n]$ -modules are used to define a natural stratification of $\mathbb{Z}^n$ -indexed persistence modules. This makes it possible to identify, for instance, homology classes that live infinitely along one or more of the coordinate $t_\bullet$ axes. We are not aware of any stability or efficient computability properties enjoyed by this stratification.

¹But we note that the subquiver consisting of two adjacent commuting squares is of finite type; its barcodes were painstakingly described using Auslander-Reiten theory in [13].

- (3) The methods of microlocal sheaf theory are used in [14] to study maps taking values in $\mathbb{R}^n$ (or more generally, k^n for a field k), with sublevelsets being defined relative to a closed convex cone $\gamma \subset \mathbb{R}^n$. For piecewise-linear maps, the resulting invariants are collections of convex polytopes that intersect γ in prescribed ways. Although these polytopes satisfy a derived-categorical stability criterion, they are not easily computable to the best of our knowledge.

This Paper. Here we construct a new invariant for those persistence modules which arise from the fiberwise homology of cellular maps $f : \mathbb{X} \rightarrow \mathbb{M}$ whose target space $\mathbb{M}$ is an oriented finite-dimensional manifold with a fixed regular CW decomposition. Our starting point and guiding principle is the theory of *bisheaves*, which was recently introduced in [20] precisely to tackle the homological stability problem for fibers of constructible maps to oriented manifolds.

Before releasing yet another fish into a teeming pond, we should explain why we expect it to thrive.

- (1) The target $\mathbb{M}$ need not be a vector space (like $\mathbb{R}^n$) or an ordered lattice contained therein (like $\mathbb{Z}^n$). This is not generality for its own sake: for instance, circular coordinates [9] have already played an important role in the theory and applications of persistent homology, and any space $\mathbb{X}$ endowed with two circular coordinates constitutes a torus-valued map $\mathbb{X} \rightarrow \mathbb{T}$. None of the existing invariants would, as far as we can tell, extend naturally to such maps.
- (2) Our invariant is discrete and efficiently computable when $\mathbb{X}$ and $\mathbb{M}$ are finite complexes, requiring nothing more serious than the standard linear algebraic operations which diagonalize boundary matrices of (co)chain complexes to extract (co)homology groups. Moreover, the computability of our invariant is significantly enhanced in the presence of distributed processing architectures.
- (3) Considerations of stability are the *raison d'être* for bisheaves [20]. Thus not only is our invariant stable, but it also provides principled lower bounds on (Betti numbers of) the stable homology classes in every fiber $f^{-1}(U)$.

On the other hand, bisheaves are complicated objects; the amount of abstract machinery required to construct our invariant, even in the relative comfort of the cellular setting, is considerably larger than we would have preferred. Thus, an important secondary goal of our work here is pedagogical — we have eschewed various slick techniques in favor of explicit combinatorial constructions, and indulged ourselves in the occasional geometric interlude. We would like to make this beautiful new theory of bisheaves more accessible to as wide an audience as possible, and we hope that the expository choices made here towards that end will assist the novice without agitating the expert.

The rest of this paper is organized as follows. In Sec blah we blablah...

1. Cellular Bisheaves

Let $\mathbb{M}$ be a regular CW complex² and let A be an abelian category — typical choices of A are the category $\mathbf{Mod}(R)$ of modules over a commutative unital ring R ,

²This means that the closure of every cell in $\mathbb{M}$ is homeomorphic to a closed disk of the correct dimension — see [21, Ch IX.6] for more.

or more generally, the category $\mathbf{Ch}(R)$ of chain complexes of R -modules. By a **sheaf over** $\mathbb{M}$ (taking values in A) we mean a functor

$$\bar{F} : \mathbf{Fc}(\mathbb{M}) \rightarrow A$$

from the poset of cells in $\mathbb{M}$ ordered by the face relation to the abelian category A . Thus, each cell σ of $\mathbb{M}$ is assigned an A -object called the *stalk* of $\bar{F}$ over σ , while each face relation $\sigma \leq \sigma'$ among cells is assigned a corresponding A -morphism $\bar{F}(\sigma \leq \sigma') : \bar{F}(\sigma) \rightarrow \bar{F}(\sigma')$ in A , called its *restriction map*. These assignments of objects and morphisms are subject to the usual laws of associativity and identity. A morphism $\bar{\alpha} : \bar{F} \rightarrow \bar{G}$ of sheaves over $\mathbb{M}$ is prescribed by a collection of A -morphisms $\{\bar{\alpha}_\sigma : \bar{F}(\sigma) \rightarrow \bar{G}(\sigma)\}$, indexed by cells of $\mathbb{M}$, which are required to commute with restriction maps. Sheaf morphisms are composed stalk-wise. We call $\bar{\alpha}$ an epi/mono/isomorphism if each constituent $\bar{\alpha}_\sigma$ is an epi/mono/isomorphism in A .

One similarly has the dual notion of a **cosheaf under** $\mathbb{M}$, which is a functor

$$\underline{F} : \mathbf{Fc}(\mathbb{M})^{\text{op}} \rightarrow A.$$

This assigns to each cell σ a module $\underline{F}(\sigma)$ called its *costalk*, and to each face relation $\sigma \leq \sigma'$ a contravariant A -morphism $\underline{F}(\sigma \leq \sigma') : \underline{F}(\sigma') \rightarrow \underline{F}(\sigma)$, called the *extension map*. As before, a morphism $\underline{\alpha} : \underline{F} \rightarrow \underline{G}$ of cosheaves under $\mathbb{M}$ is a cell-indexed collection of morphisms $\{\underline{\alpha}_\sigma : \underline{F}(\sigma) \rightarrow \underline{G}(\sigma)\}$ in A which must commute with extension maps. For a detailed introduction to cellular (co)sheaves, see [7].

1.1. Definition. The following algebraic-topological object (compare [20, Def 5.1]) coherently intertwines sheaves with cosheaves.

DEFINITION 1.1. A **bisheaf around** $\mathbb{M}$ (taking values in A) is a triple

$$\bar{F} = (\bar{F}, \underline{F}, F),$$

defined as follows. Here $\bar{F}$ is an A -valued sheaf over $\mathbb{M}$, while $\underline{F}$ is an A -valued cosheaf under $\mathbb{M}$, and

$$F = \{F_\sigma : \bar{F}(\sigma) \rightarrow \underline{F}(\sigma)\}$$

is a collection of A -morphisms indexed by the cells of $\mathbb{M}$ so that the following diagram commutes in A for each face relation $\sigma \leq \sigma'$:

$$\begin{array}{ccc} \bar{F}(\sigma) & \xrightarrow{\bar{F}(\sigma \leq \sigma')} & \bar{F}(\sigma') \\ F_\sigma \downarrow & & \downarrow F_{\sigma'} \\ \underline{F}(\sigma) & \xleftarrow{\underline{F}(\sigma \leq \sigma')} & \underline{F}(\sigma') \end{array}$$

(The right-pointing map is the restriction map of the sheaf $\bar{F}$, while the left-pointing map is the corestriction map of the cosheaf $\underline{F}$).

The reader who urgently craves interesting examples of bisheaves is requested to delay gratification until Sec 2; for now, let us note that bisheaves around $\mathbb{M}$ naturally form a category as follows. A morphism $\bar{F} \rightarrow \bar{G}$ is a pair $\bar{\alpha} = (\bar{\alpha}, \underline{\alpha})$ where $\bar{\alpha} : \bar{F} \rightarrow \bar{G}$

is a morphism of sheaves, $\underline{\alpha} : \underline{G} \rightarrow \underline{F}$ is a morphism of cosheaves, and the following diagram commutes in $\mathcal{A}$ for every cell σ in $\mathbb{M}$:

$$\begin{array}{ccc} \overline{F}(\sigma) & \xrightarrow{\overline{\alpha}_\sigma} & \overline{G}(\sigma) \\ F_\sigma \downarrow & & \downarrow G_\sigma \\ \underline{F}(\sigma) & \xleftarrow{\underline{\alpha}_\sigma} & \underline{G}(\sigma) \end{array}$$

We call $\overline{\alpha}$ an isomorphism if both $\overline{\alpha}$ and $\underline{\alpha}$ are isomorphisms. Compositions of bisheaf morphisms are defined in the obvious manner (one composes the sheaf maps and cosheaf maps independently), so we obtain the desired category $\mathbf{Bsh}(\mathbb{M})$ of bisheaves around $\mathbb{M}$.

REMARK 1.2. Bisheaves may be regarded as ordinary sheaves taking values in the *twisted arrow* category [19, Ex IX.6.3] of $\mathcal{A}$. Unfortunately, the twisted arrow category does not inherit an abelian structure from $\mathcal{A}$, so one can not expect $\mathbf{Bsh}(\mathbb{M})$ to be an abelian category. Fortunately, Def 1.1 has some rather concrete specializations. For instance, if $\mathcal{A}$ is the category of finite dimensional real vector spaces, then an $\mathcal{A}$ -valued bisheaf around $\mathbb{M}$ is completely determined by the assignment one real matrix to every cell of $\mathbb{M}$ plus two real matrices to every face relation. Of course, these matrices must satisfy various commutativity conditions described above.

1.2. Isification and Local Systems. (Maybe do this using comma categories?)

2. Bisheaves from Fibers

Let $\mathbb{X}$ and $\mathbb{M}$ be regular CW complexes, and let $f : \mathbb{X} \rightarrow \mathbb{M}$ be a *combinatorial* cellular map. That is to say, f is specified here entirely in terms of a poset map $\mathbf{Fc}(\mathbb{X}) \rightarrow \mathbf{Fc}(\mathbb{M})$ that sends cells to cells, rather than as a continuous function between the underlying topological spaces; all we require is $\dim f(\tau) \leq \dim \tau$ for every cell τ of $\mathbb{X}$. Also let $\mathbf{Mod}(R)$ and $\mathbf{Ch}(R)$ denote the categories of R -modules and chain complexes of R -modules over some fixed commutative unital ring R .


Our goal here is to explicitly construct a certain $\mathbf{Mod}(R)$ -valued bisheaf

$$\overline{F} = (\overline{F}, \underline{F}, F)$$

around $\mathbb{M}$ arising from the homology of fibers of f , as in [20, Ex 5.3]. Both the constituent sheaf $\overline{F}$ and the cosheaf $\underline{F}$ can be constructed without any further assumptions, but constructing the intertwining maps F requires additional structure to be imposed on the target CW complex $\mathbb{M}$.

EXAMPLE 2.1. Consider, as a running example throughout this section, the cellular map whose domain is the union of a torus and a sphere along an equatorial circle S : and whose codomain is the subcomplex consisting only of the sphere. Our map sends the toral part of its domain radially down to the underlying circle S , and acts as the identity on the remaining spherical part of the domain.

The following (standard) terminology will be used throughout: the *open star* of a cell σ in the CW complex $\mathbb{M}$ is the subposet of $\mathbf{Fc}(\mathbb{M})$ consisting of all cells lying



above σ with respect to the face partial order:

$$\mathbf{st} \sigma = \{\sigma' \in \mathbf{Fc}(\mathbb{M}) \mid \sigma' \geq \sigma\}.$$

The sheaf and cosheaf that we are about to define involve homology of the fibers of f over such open stars; rather than writing $f^{-1}(\mathbf{st} \sigma)$ for each such fiber, we will write

$$f//\sigma = \{\tau \in \mathbf{Fc}(\mathbb{X}) \mid f(\tau) \geq \sigma\}.$$

While fibers are clearly subposets of $\mathbf{Fc}(\mathbb{X})$, in general the cells in a given $f//\sigma$ do not constitute a subcomplex of $\mathbb{X}$. We work throughout with coefficients in a commutative unital ring R , which will remain suppressed from the notation for chain and homology modules.

2.1. The Borel-Moore Homology Sheaf. The first component required by the bisheaf $\overline{F}$ associated to our cellular map $f : \mathbb{X} \rightarrow \mathbb{M}$ is a sheaf $\overline{F}$ over $\mathbb{M}$. This sheaf will have as its stalk over each cell σ the *Borel-Moore homology* of the corresponding fiber $f//\sigma$, i.e.,

$$\overline{F}_\bullet(\sigma) = H_\bullet(\mathbb{X}, \mathbb{X} - f//\sigma).$$

We describe an explicit cellular model of this sheaf here.

DEFINITION 2.2. The complex $\overline{C}_\bullet(\sigma)$ of cellular **Borel-Moore chains** in $f//\sigma$ is the chain complex


$$\cdots \rightarrow \overline{C}_d(\sigma) \rightarrow \overline{C}_{d-1}(\sigma) \rightarrow \cdots \rightarrow \overline{C}_1(\sigma) \rightarrow \overline{C}_0(\sigma)$$

of R -modules defined as follows. The d -th chain module is generated freely by the set of all d -cells in $f//\sigma$:

$$\overline{C}_d(\sigma) = R[\{\tau \in f//\sigma \mid \dim \tau = d\}],$$

while the boundary operator $\overline{C}_d(\sigma) \rightarrow \overline{C}_{d-1}(\sigma)$ is obtained by restricting the standard cellular boundary map of $\mathbb{X}$ to (chains generated by) the cells in $f//\sigma$.

EXAMPLE 2.3. When σ is a 1-cell lying in the equator $S \subset \mathbb{M}$ from Ex 2.1, the only nontrivial chain modules are $\overline{C}_1(\sigma)$ and $\overline{C}_2(\sigma)$, which have ranks 4 and 6 respectively. Illustrated above are the cells which participate in the Borel-Moore chain complex of σ (note that the dotted cells in the boundary are excluded.)



If a face relation $\sigma \leq \sigma'$ holds among cells of $\mathbb{M}$, then there is an inclusion $f//\sigma' \hookrightarrow f//\sigma$ of fibers because $f(\tau) \geq \sigma'$ automatically implies $f(\tau) \geq \sigma$. Thus, in this case we have a *based epimorphism* of chain complexes $\overline{\mathbf{C}}_{\bullet}(\sigma) \twoheadrightarrow \overline{\mathbf{C}}_{\bullet}(\sigma')$ induced by the following simple action on each cell τ of $f//\sigma$:

$$\tau \mapsto \begin{cases} \tau & \text{if } \tau \text{ is in } f//\sigma' \\ 0 & \text{otherwise.} \end{cases}$$

The *sheaf of Borel-Moore chains* associated to $f : \mathbb{X} \rightarrow \mathbb{M}$ is the functor $\overline{\mathbf{C}}_{\bullet} : \mathbf{Fc}(\mathbb{M}) \rightarrow \mathbf{Ch}(R)$ whose stalk over each cell σ is the chain complex $\overline{\mathbf{C}}_{\bullet}(\sigma)$ from Def 2.2 and whose restriction map over each face relation $\sigma \leq \sigma'$ is the based epimorphism $\overline{\mathbf{C}}_{\bullet}(\sigma) \twoheadrightarrow \overline{\mathbf{C}}_{\bullet}(\sigma')$ defined above.

DEFINITION 2.4. The **Borel-Moore homology sheaf** $\overline{\mathbf{F}}_{\bullet} : \mathbf{Fc}(\mathbb{M}) \rightarrow \mathbf{Mod}(R)$ associated to the cellular map $f : \mathbb{X} \rightarrow \mathbb{M}$ is the homology of $\overline{\mathbf{C}}_{\bullet}$. In other words, it is given by the composite

$$\mathbf{Fc}(\mathbb{M}) \xrightarrow{\overline{\mathbf{C}}_{\bullet}} \mathbf{Ch}(R) \xrightarrow{H_{\bullet}} \mathbf{Mod}(R)$$

of $\overline{\mathbf{C}}_{\bullet}$ with the usual homology functor defined on chain complexes of R -modules.

Thus, for each dimension $d \geq 0$, the stalk $\overline{\mathbf{F}}_d(\sigma)$ over a cell σ of $\mathbb{M}$ is the d -th homology of the chain complex from Def 2.2, while the restriction map $\overline{\mathbf{F}}(\sigma \leq \sigma')$ associated to a face relation $\sigma \leq \sigma'$ is induced on homology by the based epimorphism of chains. Here are the Betti numbers (i.e., ranks of homology modules) of $\overline{\mathbf{F}}$ -stalks over simplices in the codomain of the cellular map from Ex 2.1:

$$\sigma \mapsto \begin{cases} (0, 0, 2, 0, 0, \dots) & \sigma \in \mathbb{S} \\ (0, 0, 1, 0, 0, \dots) & \sigma \notin \mathbb{S}. \end{cases}$$

We conclude our description of the Borel-Moore sheaf by proving that it is appropriately named.

PROPOSITION 2.5. For each dimension $d \geq 0$ and simplex $\sigma \in \mathbb{M}$, there exists an isomorphism in $\mathbf{Mod}(R)$

$$\overline{\mathbf{F}}_d(\sigma) \simeq H_d(\mathbb{X}, \mathbb{X} - f//\sigma)$$

between the d -th stalk over σ of the cellular Borel-Moore homology sheaf and the d -th relative homology module of the pair $(\mathbb{X}, \mathbb{X} - f//\sigma)$.

PROOF. By excision, the relative homology of the pair $(\mathbb{X}, \mathbb{X} - f//\sigma)$ is the same as that of $(\overline{f//\sigma}, \partial \overline{f//\sigma})$, where the overline indicates closure in $\mathbb{X}$ while ∂ denotes the cellular boundary (e.g., the subcomplex spanned by dotted cells in the illustration above).

It is easily checked that the chain complex $\overline{\mathcal{C}}_\bullet(\sigma)$ is obtained precisely by quotienting the cellular chain complex associated to $f//\sigma$ by the subcomplex corresponding to its boundary $\partial f//\sigma$. $\square$

2.2. The Singular Homology Cosheaf. The second ingredient required to build the bisheaf $\overline{\mathcal{F}}$ associated to $f : \mathbb{X} \rightarrow \mathbb{M}$ is a cosheaf $\underline{\mathcal{F}}$ under $\mathbb{M}$. To each simplex σ this cosheaf will simply assign the singular homology of its fiber:

$$\underline{\mathcal{F}}_\bullet(\sigma) = H_\bullet(f//\sigma),$$

so conceptually this may be a far more familiar object than the sheaf $\overline{\mathcal{F}}$ from the preceding subsection. However, since fibers $f//\sigma$ may not be CW complexes in their own right, building an explicit cellular chain-level model for $\underline{\mathcal{F}}$ is a slightly more intricate process.

DEFINITION 2.6. The **barycentric subdivision** ΔP of a poset $(P, <)$ is that abstract simplicial complex whose d -simplices are all finite ascending sequences of $(d + 1)$ distinct P -elements of the form

$$\gamma = (p_0 < p_1 < \cdots < p_{d-1} < p_d),$$

with the face relation arising from the inclusion of subsequences. (Note that ΔP is also called the *order complex* of P , particularly in topological combinatorics [1, Sec 3]).

We will call simplices of ΔP the **exit paths** of P in accordance with [25, 8, 22]. If $P = \mathbf{Fc}(\mathbb{X})$, then ΔP is the honest barycentric subdivision of the CW complex $\mathbb{X}$, so there is a homeomorphism between $\Delta \mathbf{Fc}(\mathbb{X})$ and $\mathbb{X}$ — for details, see the proof of [18, Thm III.1.7]. It will also be useful to observe that the barycentric subdivision $\Delta(f//\sigma)$ of a fiber $f//\sigma \subset \mathbf{Fc}(\mathbb{X})$ consists of precisely those exit paths $(\tau_0 < \cdots < \tau_k)$ in $\Delta \mathbf{Fc}(\mathbb{X})$ whose first cell satisfies $f(\tau_0) \geq \sigma$.

DEFINITION 2.7. The complex $\underline{\mathcal{C}}_\bullet(\sigma)$ of cellular **singular chains** in the fiber $f//\sigma$ is the R -module chain complex

$$\cdots \rightarrow \underline{\mathcal{C}}_d(\sigma) \rightarrow \underline{\mathcal{C}}_{d-1}(\sigma) \rightarrow \cdots \rightarrow \underline{\mathcal{C}}_1(\sigma) \rightarrow \underline{\mathcal{C}}_0(\sigma),$$


where each $\underline{\mathcal{C}}_d(\sigma)$ is freely generated by the d -simplices of $\Delta(f//\sigma)$:

$$\underline{\mathcal{C}}_d(\sigma) = R[\{(\tau_0 < \cdots < \tau_d) \in \Delta \mathbf{Fc}(\mathbb{X}) \mid f(\tau_0) \geq \sigma\}],$$

and the boundary operator is obtained by restricting the usual simplicial boundary map of $\Delta \mathbf{Fc}(\mathbb{X})$.

EXAMPLE 2.8. Returning once more to Ex 2.1, let us describe the (subdivided) cells in $f//\sigma$ which contribute to $\underline{\mathcal{C}}_\bullet(\sigma)$ whenever σ is a 1-cell lying in the equatorial circle $\mathbb{S}$. The modules $\underline{\mathcal{C}}_0(\sigma)$ and $\underline{\mathcal{C}}_1(\sigma)$ both have rank 10, while all the other chain modules are trivial. The figure below illustrates the ten vertices and ten edges which generate the nontrivial modules. These assume the shape of a circle which has two tendrils emanating from the basepoint:

If $\sigma \leq \sigma'$ holds in $\mathbb{M}$, then the barycentric subdivision $\Delta(f//\sigma')$ of the fiber over σ' is a simplicial subcomplex of $\Delta(f//\sigma)$ — to see this, note that if some exit path $\gamma = (\tau_0 < \cdots < \tau_k)$ in $\mathbf{Fc}(\mathbb{X})$ satisfies $f(\tau_0) \geq \sigma'$, then we automatically also obtain



$f(\tau_0) \geq \sigma$. Thus, there is a *based monomorphism* $\underline{C}_\bullet(\sigma') \hookrightarrow \underline{C}_\bullet(\sigma)$ of the corresponding singular chain complexes.

The *cosheaf of singular chains* for our cellular map $f : \mathbb{X} \rightarrow \mathbb{M}$ is the functor $\underline{C}_\bullet : \mathbf{Fc}(\mathbb{M})^{\text{op}} \rightarrow \mathbf{Ch}(R)$ that assigns to each cell σ the costalk $\underline{C}_\bullet(\sigma)$ from Def 2.7, and to each face relation $\sigma \leq \sigma'$ the based monomorphism $\underline{C}_\bullet(\sigma') \hookrightarrow \underline{C}_\bullet(\sigma)$ defined above.

DEFINITION 2.9. The **singular homology cosheaf** $\underline{E}_\bullet : \mathbf{Fc}(\mathbb{M}) \rightarrow \mathbf{Mod}(R)$ associated to the cellular map $f : \mathbb{X} \rightarrow \mathbb{M}$ is the homology of $\underline{C}_\bullet$. Namely, it is the composite

$$\mathbf{Fc}(\mathbb{M})^{\text{op}} \xrightarrow{\underline{C}_\bullet} \mathbf{Ch}(R) \xrightarrow{H_\bullet} \mathbf{Mod}(R)$$

of $\underline{C}_\bullet$ with the usual homology functor for chain complexes of R -modules.

As promised, the Betti numbers of $\underline{E}$ -costalks under the cells of the codomain from Ex 2.1 are precisely what one would expect from the singular homology of the corresponding fibers:

$$\sigma \mapsto \begin{cases} (1, 1, 0, 0, 0, \dots) & \sigma \in S \\ (1, 0, 0, 0, 0, \dots) & \sigma \notin S. \end{cases}$$

A geometric reason for this equivalence is already apparent from the illustration in Ex 2.8 — the fiber over σ , which looks like of a cylinder-with-two-fins, deformation retracts horizontally onto our circle-with-two-tendrils. This observation forms the core of the following argument.

PROPOSITION 2.10. *For each dimension $d \geq 0$ and simplex $\sigma \in \mathbb{M}$, there is an R -module isomorphism*

$$\underline{E}_d(\sigma) \simeq H_d(f//\sigma)$$

between the d -th costalk of the singular homology cosheaf under σ and the d -th singular homology module of the fiber $f//\sigma$.

PROOF. Let V_σ denote the poset of simplices in the barycentric subdivision $\Delta\mathbf{Fc}(\mathbb{X})$ which intersects the fiber $f//\sigma$:

$$V_\sigma = \{(\tau_0 < \dots < \tau_k) \in \Delta\mathbf{Fc}(\mathbb{X}) \mid f(\tau_k) \geq \sigma\}.$$

Since the geometric realization of V_σ is homeomorphic to $f//\sigma$, it suffices to show that V_σ is homotopy-equivalent to the simplicial complex $\Delta(f//\sigma)$ whose homology is computed by the chain complex $\underline{C}_\bullet(\sigma)$ from Def 2.7. To this end, associate to each exit path

$$\gamma = (\tau_0 < \dots < \tau_k) \text{ in } V_\sigma$$

the smallest index $i(\gamma)$ in $\{0, \dots, k\}$ for which $f(\tau_i) \geq \sigma$ holds. Let $r_\sigma(\gamma)$ be the exit path obtained by removing all the cells in γ which precede $\tau_{i(\gamma)}$:

$$r_\sigma(\gamma) = (\tau_{i(\gamma)} < \dots < \tau_k).$$

Now, $r_\sigma(\gamma)$ lies in $\Delta(f//\sigma)$ because its first cell satisfies $f(\tau_{i(\gamma)}) \geq \sigma$. It is easily confirmed that $r_\sigma : V_\sigma \rightarrow \Delta(f//\sigma)$ defines an order-preserving map. To finish the argument, we will establish that r_σ has contractible fibers and appeal to the fact that poset maps with contractible fibers induce homotopy equivalence of order complexes [1, Thm 2]. One can verify that the fiber of r_σ over each exit path $\eta = (\rho_0 < \dots < \rho_\ell)$ in $\Delta(f//\sigma)$ is precisely the open star of η in V_σ , and the desired result follows from the fact that open stars are contractible. $\square$

2.3. The Intertwining Maps. In order to build the final piece of the bisheaf $\bar{F}$, we require the target CW complex $\mathbb{M}$ to have Poincaré duality. Such complexes are best described in terms of their orientation cosheaves (see [7, Def 7.1] or [23, Sec 3]), which we now define.

DEFINITION 2.11. The **orientation cosheaf** $\underline{O}^\bullet$ on $\mathbb{M}$ takes values in the category of R -modules, and is defined as follows.

- (1) The stalk $\underline{O}^\bullet(\sigma)$ over a cell σ in $\mathbb{M}$ is the cohomology of a cochain complex:

$$\underline{L}^0(\sigma) \rightarrow \underline{L}^1(\sigma) \rightarrow \dots \rightarrow \underline{L}^k(\sigma) \rightarrow \dots,$$

where $\underline{L}^k(\sigma)$ is the free R -module spanned by all the k -dimensional cells in the open star $\mathbf{st} \sigma \subset \mathbf{Fc}(\mathbb{M})$. The coboundary maps in the complex above are obtained by restricting the usual cellular coboundary operator on $\mathbb{M}$.

- (2) The extension map $\underline{O}^\bullet(\sigma \leq \sigma')$ associated to a face relation $\sigma \leq \sigma'$ in $\mathbf{Fc}(\mathbb{M})$ is simply the map on cohomology induced by the based monomorphism

$$\underline{L}^\bullet(\sigma') \hookrightarrow \underline{L}^\bullet(\sigma)$$

arising from the fact that $\mathbf{st} \sigma'$ is a subposet of $\mathbf{st} \sigma$.

By construction, $\underline{O}^\bullet(\sigma)$ is always (isomorphic to) the compactly-supported cohomology of σ 's open star $\mathbf{st} \sigma$ — indeed, one perspective on $\underline{O}$ is that it is the Borel-Moore homology sheaf from Sec 2.1 for the identity cellular map $\mathbb{M} \rightarrow \mathbb{M}$, except that the boundary morphisms of the underlying chain complexes have been systematically replaced by their adjoints.

DEFINITION 2.12. We call the m -dimensional regular CW complex $\mathbb{M}$ a **Poincaré complex** if two conditions hold:

- (1) its top compactly supported cohomology satisfies $H_c^m(\mathbb{M}; R) = R$, and
- (2) the inclusion of the open star $\mathbf{st} \sigma \subset \mathbb{M}$ of any cell σ in $\mathbb{M}$ induces an isomorphism between $\underline{O}^m(\sigma)$ and $H_c^m(\mathbb{M}; R)$.

Henceforth, we will assume that $\mathbb{M}$ is a **Poincaré complex** of dimension m , noting that the class of such complexes includes all m -dimensional topological manifolds without boundary (and more generally, all m -dimensional boundaryless R -homology manifolds). Armed with Poincaré duality, we will complete our construction of the

bisheaf $\bar{F}$ associated to $f : \mathbb{X} \rightarrow \mathbb{M}$ by describing the desired intertwining maps over a fixed cell $\sigma \in \mathbb{M}$. These will take the form

$$F_\sigma : \bar{F}_{m+\bullet}(\sigma) \rightarrow \underline{F}_\bullet(\sigma)$$

between (an m -shifted version of) the Borel-Moore homology sheaf from Sec 2.1 and the singular homology cosheaf from Sec 2.2.

REMARK 2.13. The following construction is described in [20, Ex 5.3]. Let o be a generator of the top cohomology $H_c^m(\mathbb{M}; R)$, and let $o|_\sigma$ be its image in $\underline{Q}^m(\sigma)$ under an isomorphism guaranteed by Def 2.12. Then the pullback $f^*(o|_\sigma)$ is an element in the m -th compactly-supported cohomology of the fiber $f//\sigma$. The cap product with $f^*(o|_\sigma)$ therefore constitutes a map

$$\bar{F}_{m+\bullet}(\sigma) \xrightarrow{\cap f^*(o|_\sigma)} \bar{F}_\bullet(\sigma).$$

from the Borel-Moore homology to the singular homology of the fiber $f//\sigma$. When f is the identity on $\mathbb{M}$, this map reduces to the standard Poincaré duality isomorphism on $\mathbf{st} \, m$. What follows below is a combinatorial avatar of this basic construction. An advantage of working in the cellular setting is that one obtains an explicit chain-level formula for this cap product map that can be used to perform machine computations.

Since $\mathbb{M}$ is assumed to be an m -dimensional Poincaré complex, each cell σ in $\mathbb{M}$ is the face of some m -dimensional cell, so we can always select a *maximal exit path through* σ in $\mathbf{Fc}(\mathbb{M})$, i.e., a strictly ascending sequence of cells

$$\eta = (\sigma_0 < \sigma_1 < \cdots < \sigma_m) \tag{1}$$

with $\dim \sigma_i = i$, so that the cell in position $\dim \sigma$ is σ itself. Let us assume that such an η has been chosen for each cell σ of $\mathbb{M}$. This selection can be completely arbitrary, and in particular we do not require any compatibility between η and the maximal path chosen for any other cell σ' in $\mathbb{M}$ (even if $\sigma \leq \sigma'$).

Let τ be a generating cell of the chain module $\bar{C}_{m+d}(\sigma)$ from Def 2.2 — this means $\dim \tau = m + d$ and $f(\tau) \geq \sigma$. Define the collection of *maximal exit paths through* τ over η as the following subset of $\Delta \mathbf{Fc}(\mathbb{X})$:

$$P_\eta(\tau) = \{(\tau_0 < \tau_1 < \cdots < \tau_{m+d} = \tau) \mid \dim \tau_i = i \text{ and } f(\tau_j) = \sigma_j\}. \tag{2}$$

(Here i runs from 0 to $m + d$ while j runs from 0 to m). In other words, $P_\eta(\tau)$ is the collection of all maximal exit paths in $\mathbf{Fc}(\mathbb{X})$ whose

- (1) last cell is τ , and
- (2) first m cells are mapped by f to the cells of η .

To each exit path $\gamma = (\tau_0 < \cdots < \tau_{m+d})$ in this set $P_\eta(\tau)$, we assign a *weight* $w_\gamma \in R$, given by the product

$$w_\gamma = - \prod_{i=1}^{m+d} [\tau_i : \tau_{i-1}], \tag{3}$$

where each factor $[\tau_i : \tau_{i-1}]$ is the R -valued degree of the attaching map from the boundary of the cell τ_i to its face τ_{i-1} in $\mathbb{X}$. Since $\mathbb{X}$ is assumed to be a *regular* CW complex, each degree is either 1_R or -1_R , so our weights always take values in $\pm 1_R$.

DEFINITION 2.14. The d -th **intertwining map** $C_\sigma : \overline{\mathbf{C}}_{m+d}(\sigma) \rightarrow \underline{\mathbf{C}}_d(\sigma)$ is defined (as a linear map of R -modules) by the following action on a basis cell τ of its domain:

$$C_\sigma(\tau) = \sum_{\gamma} w_{\gamma} \cdot \gamma_{\geq m}.$$

The sum is taken over the set of all exit paths $\gamma \in P_\eta(\tau)$ as defined in (2). For each such path, $w_{\gamma} \in \{\pm 1_R\}$ is the weight as described in (3), while $\gamma_{\geq m}$ is the subpath obtained by removing the first $(m-1)$ cells.

Let us confirm that $C_\sigma(\tau)$ is indeed a chain in $\underline{\mathbf{C}}_{\bullet}(\sigma)$. Letting $i \in \{0, \dots, m\}$ denote the dimension of σ , it follows from (2) and (1) that each exit path $\gamma \in P_\eta(\tau)$ has as its i -th cell some $\tau_i \in \mathbb{X}$ whose f -image is σ . Therefore, $f(\tau_m) \geq f(\tau_i) = \sigma$ holds, so the first cell τ_m of $\gamma_{\geq m}$ lies in $f//\sigma$, whence $\gamma_{\geq m}$ is a basis element of $\underline{\mathbf{C}}_{\bullet}(\sigma)$ as desired.

PROPOSITION 2.15. *The following properties hold for the maps C_σ from Def 2.14.*

- (1) C_σ prescribes a chain map $\overline{\mathbf{C}}_{m+\bullet}(\sigma) \rightarrow \underline{\mathbf{C}}_{\bullet}(\sigma)$, and hence induces well-defined maps on homology.
- (2) These induced maps $F_\sigma : \overline{\mathbf{F}}_{m+\bullet}(\sigma) \rightarrow \underline{\mathbf{F}}_{\bullet}(\sigma)$ on homology are independent of the choice of maximal path η made in (1).
- (3) This collection $\{F_\sigma\}$ of maps on homology indexed by cells σ of $\mathbb{M}$ satisfies the defining property of a bisheaf. Namely, the diagram

$$\begin{array}{ccc} \overline{\mathbf{F}}_{m+\bullet}(\sigma) & \xrightarrow{\overline{\mathbf{F}}_{m+\bullet}(\sigma \leq \sigma')} & \overline{\mathbf{F}}_{m+\bullet}(\sigma') \\ F_\sigma \downarrow & & \downarrow F_{\sigma'} \\ \underline{\mathbf{F}}_{\bullet}(\sigma) & \xleftarrow{\underline{\mathbf{F}}_{\bullet}(\sigma \leq \sigma')} & \underline{\mathbf{F}}_{\bullet}(\sigma') \end{array}$$

commutes for every face relation $\sigma \leq \sigma'$ in $\mathbb{M}$

PROOF. All three assertions follow quickly once we establish the fact that C_σ factors into two very well-known chain maps

$$\overline{\mathbf{C}}_{m+\bullet}(\sigma) \rightarrow \mathbf{D}_{m+\bullet}(\sigma) \rightarrow \underline{\mathbf{C}}_{\bullet}(\sigma),$$

where the intermediate $\mathbf{D}_{m+\bullet}(\sigma)$ is generated freely by exit paths lying in the poset V_σ which appeared in the proof of Prop 2.10. More precisely, the generators of the chain module $\mathbf{D}_k(\sigma)$ are all exit paths $(\tau_0 < \dots < \tau_k)$ in $\mathbf{Fc}(\mathbb{X})$ whose last cell τ_k lies in $f//\sigma$ (i.e., satisfies $f(\tau_k) \geq \sigma$). Our first map $\overline{\mathbf{C}}_{m+\bullet}(\sigma) \rightarrow \mathbf{D}_{m+\bullet}(\sigma)$ is simply the canonical chain equivalence between the chains of $f//\sigma$ and those of V_σ ; it acts on each $(m+d)$ -dimensional cell τ via:

$$\tau \mapsto \sum_{\gamma} w_{\gamma} \cdot \gamma.$$

The sum here is over all maximal exit paths in V_σ of length $(m+d)$ whose last cell is τ , with the weights w_{γ} being given by the formula in (3). And the second chain map $\mathbf{D}_{m+\bullet}(\sigma) \rightarrow \underline{\mathbf{C}}_{\bullet}(\sigma)$ is obtained via the simplicial cap product with a certain m -dimensional simplicial cochain $f^*(\eta)$ in V_σ , where $\eta = (\sigma_0 < \dots < \sigma_m)$ is the chosen

maximal path from (1). This cochain is defined by the following action on each exit path $\delta = (\tau_0 < \cdots < \tau_m)$ in V_σ :

$$\langle f^*(\eta), \delta \rangle = \begin{cases} 1_R, & f(\tau_i) = \sigma_i \text{ for all } i \in \{0, \dots, m\}, \\ 0_R, & \text{otherwise.} \end{cases}$$

By the cap product formula for simplicial complexes [12, Sec 3.3], we obtain our second chain map $D_{m+\bullet}(\sigma) \rightarrow \underline{C}_\bullet(\sigma)$; it acts on each exit path $\gamma = (\tau_0 < \cdots < \tau_{m+d})$ in V_σ via

$$\gamma \mapsto \langle f^*(\eta), \gamma_{\leq m} \rangle \cdot \gamma_{\geq m}.$$

Here $\gamma_{\leq m}$ is the subpath $(\tau_0 < \cdots < \tau_m)$ and $\gamma_{\geq m}$ is defined similarly. A brief glance at Def 2.14 will confirm that C_σ factors into the two chain maps described above, which immediately proves (1). To see (2), we use the fact that $\mathbb{M}$ is a Poincaré complex: any two choices of η are cohomologous (being top-dimensional simplices) as generators of the top cohomology in the barycentric subdivision of $f//\sigma$. Thus, the η -invariance of F_σ follows from well-definedness of the cap product on (co)homology. And finally, to establish (3) for a face relation $\sigma \leq \sigma'$, we can always choose a maximal exit path η that goes through both σ and σ' , thus making the desired diagram commute directly on the chain level. $\square$

References

- [1] A Björner. Nerves, fibers and homotopy groups. *Journal of Combinatorial Theory A* 102, 88–93 (2003).
- [2] P Bubenik and M Catanzaro. A geometric approach to multi-parameter persistence. *In preparation*, (2018).
- [3] P Bubenik, V de Silva and V Nanda. Higher interpolation and extension for persistence modules. *SIAM Journal on Applied Algebra and Geometry* 1(1), 172–184 (2017).
- [4] G Carlsson and A Zomorodian. The theory of multidimensional persistence. *Discrete and Computational Geometry* 42(1), 71–93 (2009).
- [5] F Chazal, V de Silva, M Glisse and S Oudot. *Structure and stability of persistence modules*, Springer (2016).
- [6] D Cohen-Steiner, H Edelsbrunner and J Harer. Stability of persistence diagrams. *Discrete and Computational Geometry*, 37(1), 103–120 (2007).
- [7] J Curry. Sheaves, cosheaves and applications. *arXiv:1303.3255 [math.AT]*, (2013).
- [8] J Curry and A Patel. Classifying constructible cosheaves. *arXiv:1603.01587 [math.AT]*, (2016).
- [9] V de Silva, D Morozov and M Vejdemo-Johansson. Persistent cohomology and circular coordinates. *Discrete and Computational Geometry*, 45(4), 737–759 (2011).
- [10] P Gabriel. Unzerlegbare Darstellungen I, *Manuscripta Mathematica* 6, 71–103 (1972).
- [11] H Harrington, N Otter, H Schenk and U Tillmann. Stratifying multiparameter persistent homology. *arXiv:1708.07390 [math.AT]* (2017).
- [12] A Hatcher. *Algebraic Topology*, Cambridge University Press (2002).
- [13] Y Hiraoka and E Escobar. Persistence modules on commutative ladders of finite type. *Discrete and Computational Geometry* 55(1), 100–157, (2016).
- [14] M Kashiwara and P Schapira. Persistent homology and microlocal sheaf theory. *Journal of Applied and Computational Topology* doi:10.1007/s41468-018-0019-z (2018).
- [15] C Landi. The rank invariant stability via interleavings. In: Chambers E., Fasy B., Ziegelmeier L. (eds) *Research in Computational Topology*. Association for Women in Mathematics Series, vol 13. Springer, (2018).
- [16] M Lesnick. The theory of the interleaving distance on multidimensional persistence modules *Foundations of Computational Mathematics* 15(3), 613–650 (2015).

- [17] M Lesnick and M Wright. Interactive visualization of 2-D persistence modules *arXiv:1512.00180 [math.AT]* (2015).
- [18] A Lundell and S Weingram *The Topology of CW Complexes*, Van Nostrand Reinhold Co (1969).
- [19] S Mac Lane. *Categories for the Working Mathematician*. Springer Heidelberg (1998).
- [20] R MacPherson and A Patel. Persistent local systems. *arXiv:1805.02539v1 [math.AT]* (2018).
- [21] W Massey. *A Basic Course in Algebraic Topology*, Springer-Verlag (1991).
- [22] V Nanda. Discrete Morse theory and localization. *Journal of Pure and Applied Algebra*, 223(2), 459–488 (2019).
- [23] V Nanda. Local cohomology and stratification. *arXiv:1707.00354v1 [math.AT]* (2017).
- [24] M Scolamiero, W Chachólski, A Lundman, R Ramanujam and S Öberg. Multidimensional persistence and noise. *Foundations of Computational Mathematics* 17(6), 1367–1406 (2017).
- [25] D Treumann. Exit paths and constructible stacks. *Compositio Mathematica* 145, 1504–1532 (2009).