Chapter 9

Sheaves

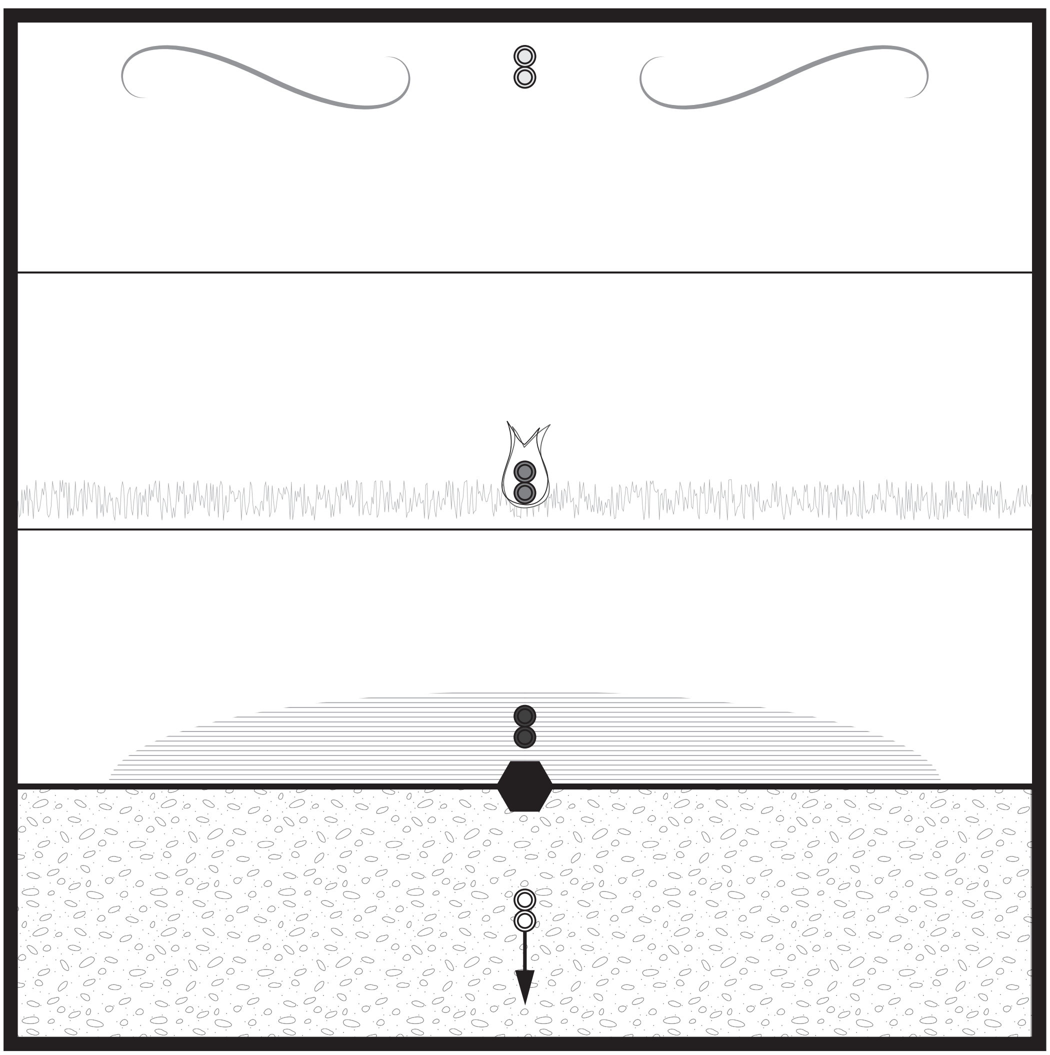

Data has, in most settings, ceased to be a scarce resource; the problem of how to get data has been eclipsed by how to manage its abundance and variety. Topology possesses several tools relevant to the aggregation and fusion of local data. Among the most powerful and flexible of these is the theory of sheaves, a structure for the collation of data parameterized by a space.

9.1 Cellular sheaves

A fiber bundle associates to sufficiently small open sets U in the base space B a product $U \times F$ with a fiber F. Very often, these fibers are vector spaces, modules, or groups, whose algebraic structure is respected by the ensuing topology of the total space E. Though useful and common, bundles (and even fibrations) are too restrictive to handle the general phenomenon of merging different forms of algebraic data over a base space.

Consider the following simple example. A robot locomotes through a system of narrow hallways. Using an omnidirectional laser scanner, it records real-valued sensor data – distance to the wall, say – along hallway directions. In the interior of a hallway, this directional data has rank two: forwards and backwards. At a branch point, where hallways meet, the number of feasible directions jumps instantaneously (in the conceptual limit where the hallways form a planar graph). The rank of data in the robot's *sensorium* changes. This idealized setting attaches to each point of a planar graph a real vector space whose dimension equals the number of directions along which one can move from that point. One can then imagine more complex sensors that store feature-detection, bearings, or other data in algebraic structures that can vary from place-to-place. Furthermore, it is possible to correlate data locally – as the robot moves along a hallway, there is consistent notion of *ahead* and *behind*. This amalgamation of algebraic data along a space is at the heart of the notion of a sheaf.

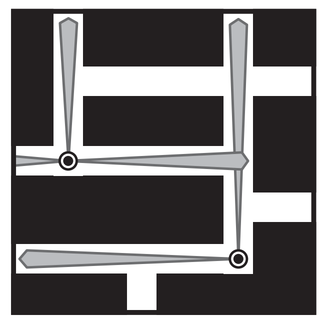

Though sheaf theory is a remarkably intricate language, the following treatment is, in keeping with the spirit of this text, elementary, emphasizing *cellular* sheaves; these possess computational and intuitive advantages reminiscent of cellular homology.

Fix X a regular cell complex with $\trianglelefteq$ denoting the face relation: $\sigma \trianglelefteq \tau$ iff $\sigma \subset \overline{\tau}$. As a model for *data* over X, consider the setting of abelian groups and homomorphisms (or, if preferred, vector spaces and linear transformations). A **cellular sheaf** over X, $\mathcal{F}$, is generated by an assignment to each cell σ of X an abelian group $\mathcal{F}(\sigma)$ and to each face $\sigma \trianglelefteq \tau$ of τ a **restriction**[1] **map** – a homomorphism $\mathcal{F}(\sigma \trianglelefteq \tau)\colon \mathcal{F}(\sigma) \to \mathcal{F}(\tau)$ such that faces of faces satisfy the composition rule:

$$\rho \trianglelefteq \sigma \trianglelefteq \tau \quad \Rightarrow \quad \mathcal{F}(\rho \trianglelefteq \tau) = \mathcal{F}(\sigma \trianglelefteq \tau) \circ \mathcal{F}(\rho \trianglelefteq \sigma). \tag{9.1}$$

[1]Yes, it seems backwards to call this a restriction. The terminology comes from the topological perspective of §9.6, of which the cellular case acts as a nerve.

The *trivial* face $\tau \trianglelefteq \tau$ by default induces the identity isomorphism $\mathcal{F}(\tau \trianglelefteq \tau) = \mathsf{Id}$. This simple definition of a sheaf as a representation of the face structure belies a powerful depth, one that is appreciated only later.

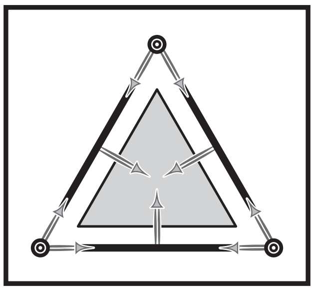

One says that the sheaf is *generated* by its values on individual cells of X: this data $\mathcal{F}(\tau)$ over a cell τ is also called the group of **local sections** of $\mathcal{F}$ over τ: one writes $\mathsf{s}_\tau \in \mathcal{F}(\tau)$ for a local section over τ. Though the sheaf is generated by local sections, there is more to a sheaf than its generating data, just as there is more to a vector space than its basis. The restriction maps of a sheaf encode how local sections are continued into more global objects – sections defined over larger subsets of X. The value of the sheaf $\mathcal{F}$ on all of X is defined to be collections of local sections that *continue* according to the restriction maps on faces:

$$\mathcal{F}(X) := \{(\mathsf{s}_\tau)_{\tau \in X} : \mathsf{s}_\sigma = \mathcal{F}(\rho \trianglelefteq \sigma)(\mathsf{s}_\rho) \ \forall\, \rho \trianglelefteq \sigma\} \subset \prod_\tau \mathcal{F}(\tau)\,. \tag{9.2}$$

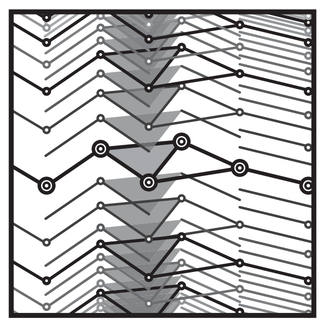

That is, $\mathcal{F}(X)$ consists of all choices of local data over cells which are compatible with respect to restriction maps: the **global sections**. In general, one thinks of $\mathcal{F}$ as a data structure that assigns to any subset $A \subset X$ the corresponding group $\mathcal{F}(A)$ of sections *over* A, where Equation (9.2) is modified to use the smallest collection of cells in X containing A and consistency is enforced on all faces within this A-containing subcomplex. In the case of a single point $x \in X$, one speaks of the **stalk** of $\mathcal{F}$ at x, $\mathcal{F}_x$. In the present setting, this is simply $\mathcal{F}_x = \mathcal{F}(\sigma)$, for σ the unique cell in whose interior x lies.

9.2 Examples of cellular sheaves

The simplest example of a cellular sheaf is the **constant** sheaf $\mathbf{G}_X$ on X, which assigns the group $\mathbf{G}$ to each cell of X and the identity homomorphism to each face relation.

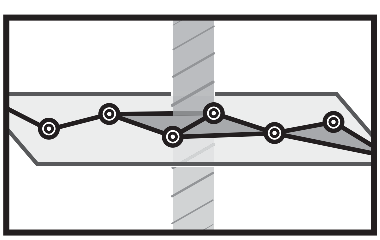

The data remain constant over X: every point has stalk $\mathbf{G}$ and the global sections match local sections, $\mathbf{G}_X(X) \cong \mathbf{G}$. The antipodal example to this is the **skyscraper** sheaf $\mathbf{G}_\sigma$ that assigns the group $\mathbf{G}$ to the cell σ and zero to all other cells. In this case, all the restriction maps are the zero map (except the identity $\sigma \trianglelefteq \sigma$). The stalks vanish except on the cell σ. This sheaf has no nonzero global sections if $\dim \sigma > 0$, but for a vertex v, $\mathbf{G}_v(X) \cong \mathbf{G}$.

Note that the restriction maps are a crucial component of a sheaf. Given a pair of sheaves over X, $\mathcal{F}$ and $\mathcal{G}$, one defines the sum $\mathcal{F} \oplus \mathcal{G}$ to be the sheaf on X whose data on σ is $\mathcal{F}(\sigma) \oplus \mathcal{G}(\sigma)$ and whose restriction maps are likewise direct sums of the form $\mathcal{F}(\sigma \trianglelefteq \tau) \oplus \mathcal{G}(\sigma \trianglelefteq \tau)$. Note that summing up skyscraper sheaves over every cell of X is *not* the same as the constant sheaf,

$$\bigoplus_{\sigma \in X} \mathbf{G}_\sigma \neq \mathbf{G}_X,$$

even though all the stalks agree: the restriction maps of this city full of skyscrapers are all zero. The value of the sum-of-skyscrapers sheaf on all of X is $\oplus_v \mathbf{G}$ – an anarchic assignment of any $\mathbf{G}$-element to each vertex v of X agrees, via the restriction maps, to the null on higher-dimensional cells. On the other hand, $\mathbf{G}_X(X) \cong \mathbf{G}$, since all of the identity restriction maps force a perfect consensus among cells.

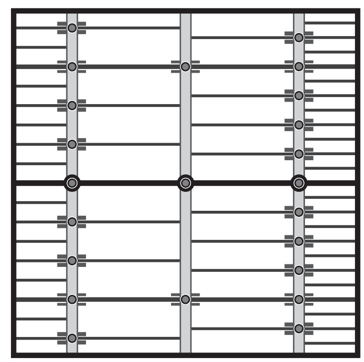

Nontrivial restriction maps lead to interesting situations vis-a-vis local versus global. Consider the case of a sheaf $\mathcal{F}$ whose stalks are all $\mathbb{Z}$. The restriction maps $\mathcal{F}(\sigma \trianglelefteq \tau)$ must therefore be homomorphisms $\mathbb{Z} \to \mathbb{Z}$; *i.e.*, multiplication by some constant. Note that these constants cannot be arbitrary, since by definition, composition must hold, meaning that the multiplication constants must factor according to the restriction maps. Then, assuming that none of these constants equals zero, and that the base complex X is connected, then there are always nonzero global sections – in fact, a $\mathbb{Z}$'s worth. However, it is *not* the case that local sections can always be patched together to give a global section.

Example 9.1 (Recurrence equations) ⊚

Consider the simple linear recurrence $u_{n+1} = A_n u_n$, where $u_n \in \mathbb{R}^k$ is a vector of states and A_n is a k-by-k real matrix. These discrete-time analogues of nonautonomous linear ODE systems are important in everything from population models to audio filtering. Such a system can be thought of in terms of sheaves over the time-line $\mathbb{R}$ with the cell structure on $\mathbb{R}$ (or $\mathbb{R}^+$) having $\mathbb{Z}$ (or $\mathbb{N}$) as vertices. Let $\mathcal{F}$ be the sheaf that assigns to each cell the vector space $\mathbb{R}^k$. One can encode the dynamics of the recurrence relation as follows. Let $\mathcal{F}(\{n\} \trianglelefteq (n, n+1))$ be the update map $u \mapsto A_n u$ and let $\mathcal{F}(\{n+1\} \trianglelefteq (n, n+1))$ be the identity. Then global sections u of $\mathcal{F}$ correspond precisely to solutions to the recurrence equation, since u restricted to $\{n+1\}$ must equal A_n times u restricted to $\{n\}$ in order to be a consistent section. To be a sheaf taking values in $\mathbb{R}$-vector spaces, the dynamics have to be linear, so that the space of solutions is a linear subspace. ⊚

Example 9.2 (Local cohomology) ⊚

An excellent example for building intuition is to be found in local cohomology. Fix X a cell complex and consider the cellular sheaf $\mathcal{F}$ that assigns to the cell σ the (singular) k^{th} **local cohomology** $H^k(X, X - \overline{\sigma})$, where $\overline{\sigma}$ is the closure of σ in X and the particular coefficient ring is left to taste. If $\sigma \trianglelefteq \tau$, then the inclusion $\iota\colon (X, X - \overline{\tau}) \hookrightarrow (X, X - \overline{\sigma})$ induces on H^k the homomorphism $\mathcal{F}(\sigma \trianglelefteq \tau)$, via

$$\mathcal{F}(\sigma \trianglelefteq \tau) = H(\iota)\colon H^k(X, X - \overline{\sigma}) \to H^k(X, X - \overline{\tau}).$$

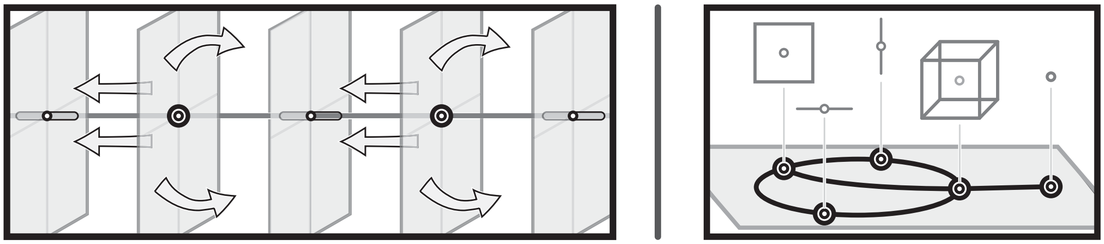

The stalk of this sheaf encodes local H^n. For example, on an n-dimensional manifold, this process yields a constant sheaf (of dimension one), called the **orientation sheaf**, *cf.* Example 4.18. The manifold is orientable if and only if the orientation sheaf has a global section. For a finite graph, the local H^1 sheaf has stalk dimension equal to 1 on edges and equal to $\deg(v) - 1$ on each vertex v. The restriction maps are projections for vertices of degree greater than two. ⊚

Example 9.3 (Fiber homology) ⊚

Consider a cellular map $f\colon X \to Y$ between cell complexes. There is an induced cellular sheaf on Y that encodes f. Given σ a simplex of Y, define the **fiber homology sheaf** $\mathcal{F}(\sigma) = H_k(f^{-1}\overline{\sigma})$. For $\sigma \trianglelefteq \tau$, one has $\overline{\sigma} \subset \overline{\tau}$ and likewise with the f-inverse image in X; hence the restriction map $\mathcal{F}(\sigma \trianglelefteq \tau)$ comes from the induced map $H_k(f^{-1}\overline{\sigma}) \to H_k(f^{-1}\overline{\tau})$. For example, if $h\colon X \to \mathbb{R}$ is a (cellular) height function, then the fiber homology sheaf on $\mathbb{R}$ records the homology of the level sets of h as stalks over vertices of $\mathbb{R}$. ⊚

Example 9.4 (Logic gates) ⊚

Consider a simple XOR gate, with two binary inputs and one binary output given by exclusive conjunction. Topologize this gate as a directed Y-graph Y. Let $\mathcal{F}$ be the sheaf taking values in $\mathbb{F}_2$ vector spaces over Y with stalk dimension one everywhere except at the central vertex, where it equals two.

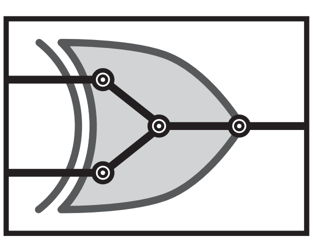

The restriction maps from the central vertex to the three edges are as follows. On the two input edges, the restriction map is projection to the first and second factors, respectively. The restriction map to the output edge is addition: $+\colon \mathbb{F}_2^2 \to \mathbb{F}_2$. This instantiates an exclusive-OR gate – the global sections correspond precisely to the truth table of inputs and outputs. A similar approach does *not* work for an AND gate, since the operation $\mathbb{F}_2^2 \to \mathbb{F}_2$ encoded by AND is no longer a homomorphism; neither is the involutive NOT, nor OR, nor NOR, nor NAND. See [159, 254] for other approaches to sheaf circuitry. ⊚

9.3 Cellular sheaf cohomology

Sheaves are interpretable as an algebraic data structure over a space (cell complex, for the moment). Equally illustrative is the interpretation as a coefficient system for cohomology. A sheaf permits *local* coefficients that can change from cell-to-cell: $\mathcal{F}(\sigma)$ is the coefficient group over the cell σ and the restriction maps $\mathcal{F}(\sigma \trianglelefteq \tau)$ encode the switching required to glue together local cochains into global cocycles.

The definition of cellular sheaf cohomology is straightforward and similar to that of ordinary cellular cohomology. Given $\mathcal{F}$ over a compact cell complex X, let $C^n(X;\mathcal{F})$ be the product of $\mathcal{F}(\sigma)$ over all n-cells σ: note that this aligns with the definition of cellular cohomology $C^n_{\text{cell}}(X;\mathbf{G})$ in the case of the constant sheaf $\mathbf{G}_X$. These cochains are connected by coboundary maps defined as follows:

$$0 \longrightarrow \prod_{\dim\ \sigma=0} \mathcal{F}(\sigma) \xrightarrow{d} \prod_{\dim\ \sigma=1} \mathcal{F}(\sigma) \xrightarrow{d} \prod_{\dim\ \sigma=2} \mathcal{F}(\sigma) \xrightarrow{d} \cdots \tag{9.3}$$

where, by analogy with (4.6), the coboundary map is given as

$$d(\sigma) := \sum_{\sigma \trianglelefteq \tau} [\sigma\colon \tau]\, \mathcal{F}(\sigma \trianglelefteq \tau). \tag{9.4}$$

Note that $d\colon C^n(X;\mathcal{F}) \to C^{n+1}(X;\mathcal{F})$, since $[\sigma\colon \tau] = 0$ unless σ is a codimension-1 face of τ. This gives a cochain complex: in the computation of d^2, the incidence numbers factor from the restriction maps, and the computation from cellular co/homology suffices to yield 0. The resulting **cellular sheaf cohomology** is denoted $H^\bullet(X;\mathcal{F})$. The cohomology of the constant sheaf $\mathbf{G}_X$ on X is, clearly, $H^\bullet_{\text{cell}}(X;\mathbf{G})$. A skyscraper sheaf $\mathbf{G}_\sigma$ over a cell σ has a (perfect) cochain complex with $d = 0$. As such, the cohomology is isomorphic to the cochain complex: $H^\bullet(X;\mathbf{G}_\sigma)$ vanishes except for grading $\dim \sigma$, at which the cohomology has rank equal to one.

One particularly useful set of interpretations for sheaf cohomology uses the terminology of local and global sections. Recall from Example 6.8 the interpretation of ordinary H^0 as connected components of a space; replacing the space with a sheaf over the space yields the following.

Lemma 9.5. *Sheaf cohomology in grading zero classifies global sections:*

$$H^0(X;\mathcal{F}) = \mathcal{F}(X). \tag{9.5}$$

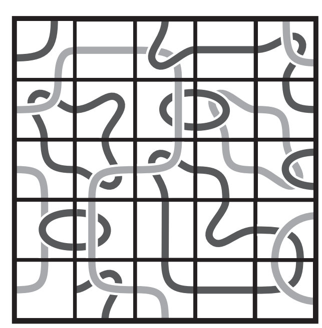

It may seem puzzling to the beginner that $\mathcal{F}(X)$ is determined completely by the data on the vertices and edges of X. Is the data on the higher-dimensional cells ignored? No: the compatibility of data on higher-dimensional cells is built into the definition of a sheaf. In this, it is analogous to a planar tiling by square tiles with colored edges. Consider the plane as a cubical complex. One places square tiles at the vertices and demands that incident tiles have compatible edges. The global solutions (legal tilings of the plane) are classified completely by vertex and edge data: there are no additional constraints where the four corners of the tiles meet (on the 2-cells of the complex).

Example 9.6 (Matrix equations) ⊚

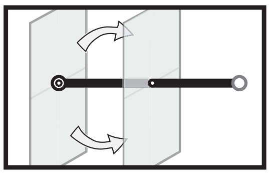

The elements of linear algebra recur throughout topology, including sheaf cohomology. Consider the following sheaf $\mathcal{F}$ over the closed interval with two vertices, a and b, and one edge e. The stalks are given as $\mathcal{F}(a) = \mathbb{R}^m$, $\mathcal{F}(b) = 0$, and $\mathcal{F}(e) = \mathbb{R}^n$. The restriction maps are $\mathcal{F}(b \trianglelefteq e) = 0$ and $\mathcal{F}(a \trianglelefteq e) = A$, where A is a linear transformation. Then, by definition, the sheaf cohomology is $H^0 \cong \ker A$ and $H^1 \cong \operatorname{coker} A$.

Cellular sheaf cohomology taking values in vector spaces is really a characterization of solutions to complex networks of linear equations. If one modifies $\mathcal{F}(b) = \mathbb{R}^p$ with $\mathcal{F}(b \trianglelefteq e) = B$ another linear transformation, then the cochain complex takes the form

$$0 \longrightarrow \mathbb{R}^m \times \mathbb{R}^p \xrightarrow{[A|-B]} \mathbb{R}^n \longrightarrow 0 \longrightarrow \cdots ,$$

where $d = [A|-B] : \mathbb{R}^{m+p} \to \mathbb{R}^n$. The zero$^{\text{th}}$ sheaf cohomology H^0 is precisely the set of solutions to the equation $Ax = By$, for $x \in \mathbb{R}^m$ and $y \in \mathbb{R}^p$. These are the global sections over the closed edge. The first sheaf cohomology measures the degree to which $Ax - By$ does not span $\mathbb{R}^n$. As before, all higher sheaf cohomology groups vanish. ⊚

As sheaves provide a means of doing cohomology, one anticipates that the methods of Chapters 5 *ff.* generalize as well. This is relatively straightforward, viewing a sheaf as a generalized coefficient system. By analogy with Lemma 5.5, consider a short exact sequence of sheaves over X,

$$0 \longrightarrow \mathcal{F} \xrightarrow{i} \mathcal{G} \xrightarrow{j} \mathcal{H} \longrightarrow 0 ,$$

where exactness is enforced cell-by-cell (or stalk-by-stalk). This leads to a connecting homomorphism and long exact sequence on sheaf cohomology:

$$\longrightarrow H^{n-1}(X;\mathcal{H}) \xrightarrow{\delta} H^n(X;\mathcal{F}) \xrightarrow{H(i)} H^n(X;\mathcal{G}) \xrightarrow{H(j)} H^n(X;\mathcal{H}) \longrightarrow . \qquad (9.6)$$

For example, consider a sheaf $\mathcal{F}$ on X and a closed subcomplex $A \subset X$. Let $\mathcal{F}_A$ be the restriction of $\mathcal{F}$ to A (so that stalks are zero off of A and all restriction maps $\mathcal{F}(\sigma \trianglelefteq \tau)$ are set to zero unless $\sigma, \tau \subset A$), and let $\mathcal{F}_{X-A}$ be the complementary restriction of $\mathcal{F}$ to $X-A$ (so that it vanishes on A). Then the following sequence of cellular sheaves on X is exact:

$$0 \longrightarrow \mathcal{F}_{X-A} \xrightarrow{i} \mathcal{F} \xrightarrow{j} \mathcal{F}_A \longrightarrow 0 .$$

Here, j is a projection map that sends stalks of $X-A$ to zero; i is the inclusion of $\mathcal{F}$ restricted to $X-A$, the kernel of j. Equation (9.6) yields an analogue of the long exact sequence of $\mathcal{F}$ over the pair (X, A):

$$\longrightarrow H^{n-1}(A;\mathcal{F}) \xrightarrow{\delta} H^n(X,A;\mathcal{F}) \xrightarrow{H(j)} H^n(X;\mathcal{F}) \xrightarrow{H(i)} H^n(A;\mathcal{F}) \longrightarrow .$$

Here, a convenient notation is adopted: $H^\bullet(A;\mathcal{F}) := H^\bullet(X;\mathcal{F}_A)$ and $H^\bullet(X,A;\mathcal{F}) := H^\bullet(X;\mathcal{F}_{X-A})$. This relative sheaf cohomology is significant. The reader may take the following as an exercise in definitions and the long exact sequence:

Corollary 9.7. *For A a subcomplex of X, $H^0(X,A;\mathcal{F})$ classifies global sections of $\mathcal{F}$ on X which vanish on A.*

This interpretation is in keeping with the notion of cohomology as an obstruction to solving certain problems of extension and restriction.

9.4 Flow sheaves and obstructions

Even the simple setting of a cellular sheaf over a network (a 1-d cell complex) has interesting applications. Consider a flow network, as per Example 6.5, in which a directed acyclic network X from a source node s to a target node t has an assignment of capacity constraints, $\mathsf{cap}(e) \in \mathbb{N}$ for each edge $e \in E(X)$. A flow $\varphi\colon E(X) \to \mathbb{N}$ assigns to edges a flow rate $0 \leq \varphi(e) \leq \mathsf{cap}(e)$ in a conservative manner – at each node $v \in V(X)$,

$$\sum_{e\to v} \varphi(e) = \sum_{v\to e'} \varphi(e') =: \varphi_v, \qquad (9.7)$$

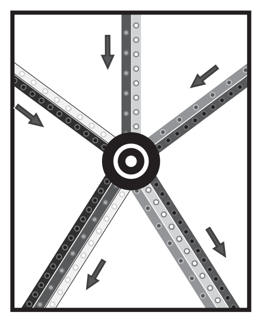

where $e \to v$ denotes an input edge and $v \to e'$ denotes an output edge at v. The **flow value**, $\mathsf{val}(\varphi)$, equals the net flow out of s (or equivalently into t).

It is possible to interpret flows in terms of the cohomology of a **flow sheaf**, following the ideas of Hiraoka (with additional input from Robinson). Fix a capacity-respecting flow φ on X and choose a consistent **routing protocol** at each vertex. Namely, for each vertex v, choose a binary matrix Ψ_v that encodes which portions of the incoming flow are sent to which portions of the outgoing flow at v,

$$\Psi_v\colon \bigoplus_{e\to v} \mathbb{R}^{\mathsf{cap}(e)} \longrightarrow \bigoplus_{v\to e'} \mathbb{R}^{\mathsf{cap}(e')},$$

where $(\Psi_v)_{i,j} = 1$ means that the j^{th} unit of incoming flow is sent to the i^{th} unit of outgoing flow; conservation is imposed by insisting that Ψ_v is a permutation matrix padded with extra zero rows/columns. For purposes of building the appropriate cellular sheaf, subdivide X to $\tilde{X}$ as follows:

1. Each edge of X is split in two with an additional vertex in the middle.
2. Each such vertex v added has trivial routing: $\Psi_v = \mathsf{Id}$.
3. There is a single *feedback edge* from t to s with sufficiently large capacity.

Let $\mathcal{F}$ be the sheaf on $\tilde{X}$ taking values in $\mathbb{R}$-vector spaces whose stalk on each edge e is $\mathcal{F}(e) = \mathbb{R}^{\mathsf{cap}(e)}$. At a vertex v, the stalk is $\mathcal{F}(v) = \mathbb{R}^{\mathsf{cap}(v)}$, where $\mathsf{cap}(v)$ is defined to be the sum of $\mathsf{cap}(e')$ for all incoming edges $e' \to v$. The sheaf restriction maps enforce conservation:

$$\mathcal{F}(v \trianglelefteq e)\colon \bigoplus_{e' \to v} \mathbb{R}^{\mathsf{cap}(e')} \longrightarrow \mathbb{R}^{\mathsf{cap}(e)}$$

where for $e \to v$, the map $\mathcal{F}(v \trianglelefteq e)$ is projection to the $\mathbb{R}^{\mathsf{cap}(e)}$-factor in $\mathcal{F}(v)$, and for $v \to e$, the map $\mathcal{F}(v \trianglelefteq e)$ is projection of the image of Ψ_v to the $\mathbb{R}^{\mathsf{cap}(e)}$-factor. This flow sheaf satisfies a type of Poincaré duality:

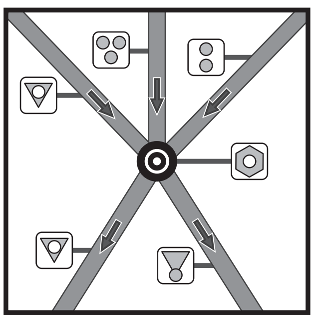

Lemma 9.8. *For a flow-sheaf $\mathcal{F}$ as above, $H^0(\tilde{X};\mathcal{F}) \cong H^1(\tilde{X};\mathcal{F})$ and $\dim H^0(\tilde{X};\mathcal{F})$ equals the net flow value $\mathsf{val}(\varphi)$.*

Proof. The simplest proof of this duality is by counting: $C^0(\tilde{X};\mathcal{F}) \cong C^1(\tilde{X};\mathcal{F})$, since each vertex has stalk isomorphic to the sums of stalks of incoming edges, and this enumeration exhausts all cells of $\tilde{X}$. Passing to cohomology, the isomorphism comes from exactness (*i.e.*, Lemma 1.10): $\dim \ker d = \dim \operatorname{coker} d$. By Lemma 9.5, $\dim H^0(\tilde{X};\mathcal{F})$ equals the number of global sections of $\mathcal{F}$; by conservation, each such section counts a unit of flow that circulates from source to target. ⊙

For a fixed flow φ on $\tilde{X}$ there is a way to view the max-flow-min-cut theorem of Example 6.5 in the language of the long exact sequence of a pair in sheaf cohomology. Let $C \subset E(X)$ be a collection of edges in X, thought of as a putative cut-set. Since C is not a subcomplex of X, consider instead $\tilde{C} \subset \tilde{X}$, the corresponding subcomplex of refined midpoint vertices of C in $\tilde{X}$. The long exact sequence of the pair $(\tilde{X}, \tilde{C})$ in $\mathcal{F}$-cohomology yields:

$$0 \longrightarrow H^0(\tilde{X},\tilde{C}) \xrightarrow{H(j)} H^0(\tilde{X}) \xrightarrow{H(i)} H^0(\tilde{C}) \xrightarrow{\delta} H^1(\tilde{X},\tilde{C}) \xrightarrow{H(j)} H^1(\tilde{X}) \xrightarrow{H(i)} 0\ .$$

Consider the zero$^{\text{th}}$ cohomology $H^0(\tilde{C};\mathcal{F})$ of $\tilde{C}$. By definition, this cohomology has dimension equal to the net capacity of the edge set C:

$$\dim H^0(\tilde{C};\mathcal{F}) = \sum_{e \in C} \mathsf{cap}(e) =: \mathsf{val}(C).$$

By Lemma 9.8, the terms $H^1(\tilde{X};\mathcal{F}) \cong H^0(\tilde{X};\mathcal{F})$ each have dimension $\mathsf{val}(\varphi)$ the flow value. What connects these is the long exact sequence. By Corollary 9.7, the relative cohomology $H^0(\tilde{X},\tilde{C})$ measures how much of the flow does *not* pass through C. Thus, by definition, if C is a *cut*, then $H^0(\tilde{X},\tilde{C}) = 0$, and one can say that this relative cohomology is the *obstruction* to being a cut. For C a cut, this relative H^0 vanishes, and $H(i)$ above is injective by exactness, meaning that for any flow and cut,

$\mathsf{val}(\varphi) \leq \mathsf{val}(C)$: *the flow value is bounded by cut values* – the weak form of the max-flow min-cut theorem [146].

The crucial term in the long exact sequence is the connecting homomorphism δ. Assuming C a cut, then, by exactness, $H^0(\tilde{X}) \cong \ker \delta$ and $H^1(\tilde{X}) \cong \operatorname{coker} \delta$. There is a splitting $H^1(\tilde{X}, \tilde{C}; \mathcal{F}) \cong H^1(\tilde{X}; \mathcal{F}) \oplus \operatorname{im} \delta$. If a nonzero class is in the image of δ, then it corresponds to a local section of C that does not come from any global flow on $\tilde{X}$. Otherwise, it corresponds to a class in H^1 – a flow, by Lemma 9.8. Summarizing this discussion:

1. $\dim H^0(\tilde{X}; \mathcal{F}) = \mathsf{val}(\varphi)$ is the flow value;
2. $\dim H^0(\tilde{C}; \mathcal{F}) = \mathsf{val}(C)$ is the cut value;
3. $H^0(\tilde{X}, \tilde{C}; \mathcal{F})$ is the obstruction to being a cut; and
4. δ is the obstruction to $\mathsf{val}(\varphi) = \mathsf{val}(C)$.

When $\delta \neq 0$, the cut is not minimal or the flow is not maximal. When the two obstructions, δ and $H^0(\tilde{X}, \tilde{C}; \mathcal{F})$ both vanish, then the given flow and cut satisfy $\dim H^0(\tilde{X}; \mathcal{F}) = \dim H^0(\tilde{C}; \mathcal{F})$, or, maxflow-equals-mincut.

This construction of flow sheaves may seem an over-complicated restatement of conservation: net inputs equals net outputs. By adapting sheaves and cohomology to this elementary setting, one may properly generalize.

9.5 Information flows and network coding

The following is an application to **network coding** [146] – a branch of network information flow problems in which algebraic coding can be performed. This is motivated by communications: when sending data over a network, the data is split into packets that are individually routed from source to target(s) and then re-integrated. Packets of data are sent over edges at a particular bit rate, and packets are routed and/or coded at nodes to be broadcast along other edges.

Fix a communications network modeled as a finite directed acyclic graph X from a fixed source (or sender) node s to multiple target (or receiver) nodes t_i in X. The data sent by s lies in a vector space over $\mathbb{F}_q$, a fixed finite field of q elements. Data is transmitted along edges of X and acted upon at nodes via linear transformations and rebroadcast, respecting directedness of X. The fundamental problems of network coding concern data throughput at the targets t_i, given constraints on X, on q, on codings, and on bit-rate capacities of edges.

Example 9.9 (Butterfly network) ⊚

The classical *butterfly network* demonstrates that network coding can improve transmission rates [127, 196]. In this network, there is a single source s and two targets t_1, t_2. Each edge has unit capacity (only one packet in $\mathbb{F}_q$ can be sent per unit time). The goal is to send two packets of data, $a, b \in \mathbb{F}_q$ from the source to both targets as quickly as possible, assuming each edge can carry one bit of data per unit time. If the nodes act as routers – data are redirected along edges unchanged – then it requires at least five units of time to transmit a and b to both targets, since one central node must switch from transmitting a to b. If, however, the central node receives both a

and b and transmits $a+b$, then, in four time steps, t_1 receives a and $a+b$ and t_2 receives b and $a+b$. From this, the original signals can be *decoded* algebraically at the targets. ⊙

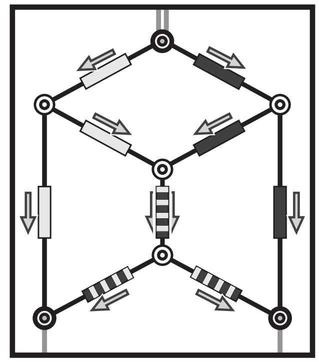

The general situation is similar. Fix X a directed network with source s and multiple targets t_i as above. The network data consists of vectors over a sufficiently large finite field $\mathbb{F} = \mathbb{F}_q$. Each edge e in X has a capacity $\mathsf{cap}(e) \in \mathbb{N}$, representing the bit-rate transmission capacity of e. The network coding is given in terms of $\mathbb{F}$-linear transformations at the nodes of X. To each node $v \in X$ is associated a **local coding map** ,

$$\Psi_v : \bigoplus_{e \to v} \mathbb{F}^{\mathsf{cap}(e)} \longrightarrow \bigoplus_{v \to e'} \mathbb{F}^{\mathsf{cap}(e')},$$

that maps the vector space of data entering v to the vector space of data exiting v. For example, routing protocols are the (zero-padded) permutation matrices of §9.4. The central node v of the butterfly network of Example 9.9 has $\Psi_v = [1|1]$ to encode via addition of signals. The decoding of messages received at t_i is given by a decoding edge from $\mathsf{t}_i \to \mathsf{s}$ with corresponding decoding map Ψ_{t_i} that disentangles coded messages at the targets. Let $\tilde{X}$ be the network obtained from X by subdividing each edge (with inherited capacities and routings as in §9.4) and adding the decoding edges.

The following construction is due to Hiraoka and parallels his flow sheaf construction from §9.4. Given the local coding maps $\{\Psi_v\}$, build a **network coding sheaf** $\mathcal{F}$ on $\tilde{X}$ as follows. On each edge e, $\mathcal{F}(e) = \mathbb{F}^{\mathsf{cap}(e)}$. At a vertex v, the stalk is $\mathcal{F}(v) = \mathbb{F}^{\mathsf{cap}(v)}$, where $\mathsf{cap}(v) = \sum_{e' \to v} \mathsf{cap}(e')$. For $e \to v$, the map $\mathcal{F}(v \trianglelefteq e)$ is projection to the $\mathbb{F}^{\mathsf{cap}(e)}$-factor in $\mathcal{F}(v)$. For $v \to e$, the map $\mathcal{F}(v \trianglelefteq e)$ is projection of the image of Ψ_v to the $\mathbb{F}^{\mathsf{cap}(e)}$-factor. The following results are straightforward generalizations of those of §9.4:

Theorem 9.10 ([146]). *With a network coding sheaf $\mathcal{F}$ for single-source-multi-target networks with decoding edges as above, $H^\bullet(\tilde{X}; \mathcal{F})$ has the following interpretations:*

1. *The net information flow rate equals* $\dim H^0(\tilde{X}; \mathcal{F})$.
2. *The net information flow rate which persists if a subnetwork A becomes inoperative equals* $\dim H^0(\tilde{X}, \tilde{A}; \mathcal{F})$.
3. *The obstruction to extending an information flow on a subnetwork A to all of X respecting the coding and capacity constraints is $\delta \colon H^1(\tilde{X}, \tilde{A}; \mathcal{F}) \to H^1(\tilde{X}; \mathcal{F})$. If $\delta = 0$, all information flows on A are extendable to X.*

9.6 From cellular to topological

Recall from Chapters 4-6 that it is relatively easy to define, visualize, and compute cellular co/homology, whereas it is of limited use in proving theorems such as homotopy invariance; the more powerful methods implicate the less-computable less-intuitive singular theory. A similar dichotomy exists in sheaf theory: all the standard texts [47, 95, 183] present definitions in the setting where the base space is a topological space with no fixed cell structure. Instead of stalks changing from cell-to-cell, they can change from point-to-point.

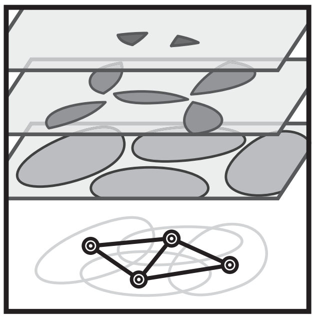

As an interpolative step, consider a space X outfitted with a locally finite cover $\mathcal{U}$ by open sets $\{U_\alpha\}$. A sheaf $\mathcal{F}$ on X *subordinate to the cover* $\mathcal{U}$ is, precisely, a cellular sheaf on the nerve $\mathcal{N}(\mathcal{U})$ of the cover. That is, there is an assignment to each cover element U_α of an abelian group $\mathcal{F}(U_\alpha)$, and for every non-empty intersection of cover elements, *e.g.*, $U_{\alpha\beta\gamma} := U_\alpha \cap U_\beta \cap U_\gamma$, there is the corresponding data $\mathcal{F}(U_{\alpha\beta\gamma})$. In addition, there are restriction homomorphisms according to the pattern of intersections in $\mathcal{U}$. For example, if $U_\alpha \cap U_\beta \neq \varnothing$, then $\mathcal{F}(U_\alpha \trianglelefteq U_{\alpha\beta})\colon \mathcal{F}(U_\alpha) \to \mathcal{F}(U_{\alpha\beta})$. Note that these restriction homomorphisms are indeed induced by restrictions of the cover elements, and the direction of the homomorphism is in the direction of the restriction. In the language of this chapter, one thinks of $\mathcal{F}$ as an assignment of data to $\mathcal{U}$ and the restriction maps as encoding how data changes when one narrows the field of view.

The cohomology of a sheaf subordinate to a cover is often called the **Čech cohomology**, $\check{H}^\bullet(\mathcal{U};\mathcal{F}) := H^\bullet(\mathcal{N}(\mathcal{U});\mathcal{F})$, of the sheaf. This has a natural parallel with the Čech co/homology of Chapter 4, and is stable under refinement of covers. The progression from sheaves on cellular complexes to topological spaces is clear in principle: one takes a *limit* of finer and finer covers. Since one can glue together local sections on cover elements, then, assuming a limit is possible, one anticipates the ability to assign *data* in the form of a group of local sections to *any* open subset of a topological space X. It would be confusing/redundant to call this a *singular sheaf* or even a *topological sheaf*. Perhaps *sheaf over a topology* would be the most descriptive term.

This discussion motivates the following definition. A **presheaf**, $\mathcal{P}$, on a space X, is an assignment to each open set $U \subset X$ of an abelian group $\mathcal{P}(U)$ of local data (or *sections*) along with restriction homomorphisms $\mathcal{P}(U \trianglelefteq V)\colon \mathcal{P}(U) \to \mathcal{P}(V)$ for every $V \subset U$. These restriction homomorphisms must satisfy composition and identity relations: $\mathcal{P}(U \trianglelefteq U) = \mathrm{Id}$ and

$$W \subset V \subset U \quad \Rightarrow \quad \mathcal{P}(U \trianglelefteq W) = \mathcal{P}(V \trianglelefteq W) \circ \mathcal{P}(U \trianglelefteq V),$$

cf. Equation (9.1). It is easy to confuse the directions of the inclusions in switching from celllular sheaves to sheaves on a topology. The reader will rightly suspect that a presheaf is *not quite* a sheaf. There are a number of subtleties associated with passing from a presheaf over a topological space to a sheaf. As a first example, note that stalks – pointwise data assignments – require a limiting process. For $x \in X$ and $\{U_i\}$ a nested sequence of open neighborhoods converging to x (that is, any neighborhood of x contains all U_i for i sufficiently large), one has a sequence of groups of sections and restriction homomorphisms

$$\cdots \longrightarrow \mathcal{P}(U_{i-1}) \longrightarrow \mathcal{P}(U_i) \longrightarrow \mathcal{P}(U_{i+1}) \longrightarrow \cdots \tag{9.8}$$

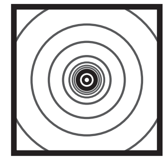

The **stalk** of $\mathcal{P}$ at x, $\mathcal{P}_x$, is the group of equivalence classes of sequences $(\mathsf{s}_i)_i$ with $\mathsf{s}_{i+1} = \mathcal{P}(U_i \trianglelefteq U_{i+1})(\mathsf{s}_i)$ for all i; and where two such sequences are equivalent if they eventually agree: $[(\mathsf{s}_i)] = [(\mathsf{s}'_i)]$ if and only if $\mathsf{s}_i = \mathsf{s}'_i$ for all i large. One shows that $\mathcal{P}_x$ is a group and is independent of the system of neighborhoods chosen to limit to x. Examples of presheaves and their stalks include the following:

1. **Skyscrapers:** Fix a point $p \in X$ and consider the presheaf that sends an open set U to 0 unless $p \in U$, for which it sends U to $\mathbf{G}$, with all restriction homomorphisms being either $\mathbf{Id}$ or 0 depending on the presence of p; this presheaf has stalk 0 everywhere except for $\mathbf{G}$ at p.
2. **Constant functions:** The presheaf that assigns to open sets U the group of constant functions $U \to \mathbf{G}$ is a presheaf with all restrictions being literal restrictions and all stalks isomorphic to $\mathbf{G}$.
3. **Smooth functions:** The presheaf that assigns to open sets U of an n-manfiold M the group $C^k(U;\mathbb{R})$ of k-times-differentiable real-valued functions (for $0 \leq k \leq \infty$) is a presheaf whose restriction homomorphisms are again restrictions; however, the stalks are the *germs*[2] of functions $\mathbb{R}^n \to \mathbb{R}$.
4. **Local homology:** The presheaf that assigns to open sets $U \subset X$ the local k^{th} homology $H_k(X, X-U)$ has restriction homomorphisms induced via functoriality: for $V \subset U$, the induced homomorphism $H_k(X, X-U) \to H_k(X, X-V)$ behaves appropriately. The stalk of this presheaf at $p \in X$ is the local homology $H_k(X, X-p)$, which can vary greatly from point-to-point.
5. **Orientations:** For an n-manifold M, the (homology) **orientation presheaf** $\mathcal{O}$ assigns to U open the local homology $H_n(M, M-U;\mathbb{Z})$. When this is a constant sheaf, the manifold is orientable (*cf.* Examples 4.18 and 9.2).
6. **Fibers:** Given $h\colon X \to Y$, the k^{th} cohomology fiber presheaf over Y assigns to $U \subset Y$ the cohomology $H^k(h^{-1}(U))$ of the inverse image. This presheaf can vary dramatically with change in U. For a sufficiently tame h, the stalk at $y \in Y$ is equal to $H^k(h^{-1}(y))$.

A sheaf is a presheaf that respects *gluings* as well as restrictions. In the cellular setting, this is accomplished by *fiat* as per Equation (9.2). In the topological setting, gluing, as with stalks, requires a limiting process. The definition is this: a presheaf $\mathcal{F}$ is a **sheaf** if and only if for any open U and $\mathcal{U} = \{U_i\}$ a finite open cover of U, and for any local sections $\mathsf{s}_i \in \mathcal{F}(U_i)$, which agree on all pairwise overlaps $U_{ij} = U_i \cap U_j$ via restriction maps, there *exists* a *unique* global section $\mathsf{s}_U \in \mathcal{F}(U)$ which restricts to s_i on each U_i. Note the double criterion of existence and uniqueness of global sections from local. This **gluing axiom** can be succinctly summarized as follows – for each such cover of $U = \cup_i U_i$, the following sequence is exact:

$$0 \longrightarrow \mathcal{F}(U) \xrightarrow{\mathcal{F}(U \trianglelefteq U_k)} \prod_k \mathcal{F}(U_k) \xrightarrow{d} \prod_{i,j} \mathcal{F}(U_{ij}) \ . \tag{9.9}$$

[2]Such germs are generalizations of Taylor series consisting of equivalence classes of functions that locally agree.

Here, the map d is the coboundary map for cellular sheaf cohomology on the nerve $\mathcal{N}(\mathcal{U})$, *cf.* Equation (9.4): the U_i correspond to vertices, the U_{ij} to edges, and an orientation is chosen to determine the correct $\pm$ sign. Exactness means that $\mathcal{F}(U) = \ker d$. Thus, a good way to remember the gluing axiom is in terms of cohomology: to be a sheaf, $\mathcal{F}(U)$ must agree with the cellular sheaf cohomology $H^0(\mathcal{N}(\mathcal{U}); \mathcal{F})$ for any finite open cover $\mathcal{U}$ of U.

Every example of a presheaf in the list above is a sheaf *except* one. The canonical example of a non-sheaf presheaf is that of *constant* functions. This violates the existence property of gluing as follows. Given a disjoint union $U = U_1 \sqcup U_2$ and the cover $\mathcal{U} = \{U_1, U_2\}$, any constant functions on U_1 and U_2 trivially agree on the (empty) overlap. However, these local sections do not extend to a global section – a constant function on U – if the values on U_1 and U_2 differ. In contrast, the presheaf of *locally* constant functions *is* a sheaf: the **constant sheaf**. Bounded real-valued maps are also examples of presheaves that are not sheaves, since maps that are locally bounded may not be globally bounded: *locally* bounded maps do form a sheaf.

The definition of cohomology $H^\bullet(X; \mathcal{F})$ for a sheaf over a topological space X is more implicit than in the cellular setting. One approach takes $\mathcal{U}$ an open cover of X and computes the Cech cohomology $\check{H}^\bullet(\mathcal{U}; \mathcal{F})$ of the sheaf subordinated to the cover. This clearly works for H^0 by dint of the gluing axiom, above. One can show that for sufficiently fine covers over sufficiently well-behaved sheaves, the higher cohomology stabilizes as well, and the limit is well-defined. The details of this and parallel constructions for sheaf cohomology can be found in [157, 183]. Suffice to say that most of the forms of cohomology seen in this text are expressible as a sheaf cohomology – either cellular or topological.

Example 9.11 (Differential equations) ⊙

Many sheaves on manifolds can be generated by partial differential equations. Consider for example the sheaf of holomorphic functions on $\mathbb{C}$ that assigns to $U \subset \mathbb{C}$ open the functions $f : U \to \mathbb{C}$ satisfying $\partial f / \partial \bar{z} = 0$, where $\partial / \partial \bar{z}$ is the linear first-order partial differential operator which, when written out in real/imaginary components yields the more familiar *Cauchy-Riemann* PDEs. The condition of being holomorphic is local in nature, and the restriction of a holomorphic function is clearly holomorphic. Crucially, extension of holomorphic functions that agree on overlaps – *analytic continuation* – also holds, guaranteeing that holomorphic functions form a sheaf. Much of the impetus for sheaf cohomology came from problems of analytic continuation, viewed as producing global sections of this sheaf. Sheaf cohomology provides obstructions to analytic continuation, not just on $\mathbb{C}$, but on complex manifolds locally modelled on $\mathbb{C}^n$.

Other systems of PDEs arise in calculus on manifolds. The sheaf that associates to open sets U of a manifold M the p-forms $\Omega^p(U)$ is well-defined – restriction and extension of forms is sensible. These p-forms are linear objects that satisfy the partial differential equations $d^2 = 0$. The resulting sheaf cohomology is identical to the de Rham cohomology ${}_{dR}H^\bullet(M)$ of §6.9. ⊙

9.7 Operations on sheaves

After working with enough examples of sheaves, certain common features resolve into canonical constructions. These steps are the beginnings of sheaf *theory* as opposed to working with sheaves one-at-a-time. This theory – a means of working with data over spaces in a platform-independent manner – quickly demands the development of tools and language beyond the scope of this chapter. A few basic definitions and examples will have to suffice until the next chapter brings that language to bear in §10.9.

Morphisms: A **sheaf morphism** $\eta\colon \mathcal{F} \to \mathcal{F}'$ between two sheaves on X is a transformation between local sections that respects restrictions. That is, for each cell τ (or open set U) there is a homomorphism $\eta\colon \mathcal{F}(\tau) \to \mathcal{F}'(\tau)$ (resp. $\eta\colon \mathcal{F}(U) \to \mathcal{F}'(U)$) and for each $\sigma \trianglelefteq \tau$ (resp. $U \supset V$) the sheaf restriction maps commute with η:

$$\begin{array}{ccc} \mathcal{F}(\sigma) & \xrightarrow{\mathcal{F}(\sigma\trianglelefteq\tau)} & \mathcal{F}(\tau) \\ \eta\downarrow & & \downarrow\eta \\ \mathcal{F}'(\sigma) & \xrightarrow[\mathcal{F}'(\sigma\trianglelefteq\tau)]{} & \mathcal{F}'(\tau) \end{array} \qquad \begin{array}{ccc} \mathcal{F}(U) & \xrightarrow{\mathcal{F}(U\trianglelefteq V)} & \mathcal{F}(V) \\ \eta\downarrow & & \downarrow\eta \\ \mathcal{F}'(U) & \xrightarrow[\mathcal{F}'(U\trianglelefteq V)]{} & \mathcal{F}'(V) \end{array} \tag{9.10}$$

Sheaf morphisms are ways of transforming or evolving data over the same base space. Other sheaf operations answer the question of how to translate data on one space over to data on another by way of a mapping of base spaces.

Direct image: Assume a map $h\colon X \to Y$ of spaces or a cellular map on cell complexes. For $\mathcal{F}_X$ a sheaf over X, the **direct image** of h is the sheaf $h_*\mathcal{F}_X$ on Y that *pushes* data from X to Y by pulling back space via h^{-1}. In the topological setting, this means $h_*\mathcal{F}_X(U) = \mathcal{F}_X(h^{-1}(U))$, for $U \subset Y$ open. In the cellular setting, this means $h_*\mathcal{F}_X(\sigma) = \mathcal{F}_X(h^{-1}(\sigma))$ for σ a cell of Y. For h injective, one visualizes the direct image as embedding the data over X into Y; for h surjective, one might think of the direct image sheaf as being the data on X, *folded* atop Y.

Inverse image: For $h\colon X \to Y$ as before, one can *pull back* a sheaf from Y to X. The **inverse image** of a sheaf $\mathcal{F}_Y$ on Y is the sheaf on X defined neatly in the cellular setting by $h^*\mathcal{F}_Y(\sigma) = \mathcal{F}_Y(h(\sigma))$ for σ a cell of X. For h cellular, this is well-defined. The difficulty arises in the topological setting. One wants to define $h^*\mathcal{F}_Y(U) = \mathcal{F}_Y(h(U))$ for $U \subset X$ open; however, $h(U)$ may not be open. Therefore, one takes a limit. Let V_i be a system of nested open sets in X with $\cap_i V_i = \overline{h(U)}$. Then the restriction maps give a sequence of abelian groups $\mathcal{F}_Y(V_i) \to \mathcal{F}_Y(V_{i+1})$. In the same manner that stalks over a point were defined in terms of equivalence classes of sequences of group elements, the larger subset $\overline{h(U)}$ has as its inverse image the 'limit' $h^*\mathcal{F}_Y(U) := \lim_{i\to\infty} \mathcal{F}_Y(V_i)$ consisting of equivalence classes of sequences of elements in $\mathcal{F}_Y(V_i)$.

In the case of direct and inverse images, one shows that the definitions lead to obvious restriction maps and the resulting presheaves are in fact sheaves. At first, these operations serve to increase the precision of language in defining and characterizing sheaves. Only later is it clear that these operations are at the heart of sheaf theory.

Example 9.12 (Compact cohomology, redux) ⊚

Recall from §9.3 that for A a subcomplex of X, the computation of cellular sheaf cohomology of A proceeds by *"restricting"* the sheaf $\mathcal{F}$ to A, yielding $\mathcal{F}_A$, then computing $H^\bullet(X; \mathcal{F}_A)$. This *ad hoc* definition can now be specified and improved. Let $\iota\colon A \hookrightarrow X$ be the inclusion. Then one defines the cohomology of $\mathcal{F}$ on A to be $H^\bullet(A; \iota^*\mathcal{F})$. This not only has the benefit of providing a sheaf on A, it easily extends to the setting where A is not a subcomplex of X. Given any collection of (open) cells $A \subset X$, the **sheaf cohomology with compact supports** on A is defined to be $H^\bullet_c(A; \mathcal{F}) := H^\bullet(X; \iota_*\iota^*\mathcal{F})$.

The notation $H^\bullet_c$ denotes, as in §6.4, compact supports. The need for this distinction is easily discerned with, say, a constant sheaf. If $\mathcal{F} = \mathbf{G}_X$ is a constant sheaf, then the cohomology $H^\bullet_c(A; \mathcal{F})$ is isomorphic to the singular compactly supported cohomology $H^\bullet_c(A; \mathbf{G})$ of A as a topological space, so that for an open n-simplex σ in X, $H^\bullet_c(\sigma; \mathbf{G}_X)$ vanishes except in grading n. Of course, if A is a closed subcomplex of X, then $H^\bullet_c(A; \mathcal{F}) \cong H^\bullet(A; \mathcal{F})$. ⊚

One strategy in sheaf theory is to manipulate data structures over spaces by pushing or pulling sheaves in a way that the impact on sheaf cohomology is controlled. In certain instances, simple general conditions guarantee such control.

Theorem 9.13 (Vietoris Mapping Theorem). *Assume $\mathcal{F}$ a sheaf on X and $h\colon X \to Y$ a proper map with fibers satisfying $H^n(h^{-1}(y); \mathcal{F}) = 0$ for all $n > 0$. Then h induces an isomorphism on sheaf cohomology: $H^\bullet(X; \mathcal{F}) \cong H^\bullet(Y; h_*\mathcal{F})$.*

In particular, if the fibers of h are all zero-dimensional, then the acyclicity hypothesis of the theorem is always satisfied. The theorem holds for cellular sheaves also with obvious modifications.

Example 9.14 (Flow sheaves, redux) ⊚

The flow sheaves of §9.4 can be represented as pushforwards of simple component sheaves of two types. To model a single unit of flow that circulates over a network (with feedback edge), consider first S, a circle, subdivided into a directed graph (with a single cycle), and let $\mathcal{F}_S = \mathbb{R}_S$ be the constant sheaf on S. Next, let $I = [0, 1]$ be an interval, subdivided into a directed linear graph with (potentially many) edges, and let $\mathcal{F}_I = \mathbb{R}_{(0,1]}$ denote the constant sheaf on I that is set to $\mathbb{R}$ everywhere except zero at the left endpoint. With a bit of insight, one sees that every flow sheaf is precisely the pushforward of some finite number of copies of $\mathcal{F}_S$ and $\mathcal{F}_I$ over a disjoint union of circles and intervals, with the projection map p gluing together subintervals in an orientation-preserving manner.

This clarifies the computation of flow sheaf cohomology. By Poincaré duality, $\mathcal{F}_S$ has cohomology $H^0 \cong H^1 \cong \mathbb{R}$. It is clear that $H^\bullet(I; \mathcal{F}_I) = 0$, since there are no global sections and the coboundary map d is surjective. The projection map p has discrete fibers; hence, by Theorem 9.13, flow sheaves have $H^0 \cong H^1$ and the dimension of this cohomology is precisely the flow value: *cf.* Lemma 9.8. ⊚

Example 9.15 (Euler characteristic, redux) ⊚

There are other operations that apply to sheaves. Following the perspective that a sheaf is an algebraic enhancement of a topological space, one repeats standard topological constructs. Among the most primal is that of the Euler characteristic. For a cellular sheaf $\mathcal{F}$ on a cell complex X taking values in vector spaces, one defines the Euler characteristic in the obvious way:

$$\chi(\mathcal{F}) := \sum_{\sigma}(-1)^{\dim\,\sigma}\dim\,\mathcal{F}(\sigma) = \sum_{n=0}^{\infty}(-1)^{n}\dim\,H^{n}(X;\mathcal{F}). \tag{9.11}$$

How does the sheaf-theoretic Euler characteristic relate to the familiar χ of previous chapters? Recall that the combinatorial definition of Euler characteristic used in the o-minimal setting of Chapter 3 is *not* the homological definition of Corollary 5.18 in general: the homological formula holds only for compact spaces, on which χ is, like homology, a homotopy invariant. The combinatorial χ has an interpretation in terms of compactly-supported cohomology. For A a locally-compact definable space,

$$\chi(A) = \sum_{n=0}^{\infty}(-1)^{n}\dim\,H_c^{n}(A). \tag{9.12}$$

This can be interpreted at the level of sheaves. For A a definable set represented as a subcollection $\iota\colon A \hookrightarrow X$ of open cells in a definable triangulation of X, the Euler characteristic of A can be expressed as the Euler characteristic of the constant sheaf $\mathbb{R}_A$ on A:

$$\chi(A) = \chi(A;\mathbb{R}_A) = \chi(X;\iota_*\iota^*\mathbb{R}_X) = \sum_{n=0}^{\infty}(-1)^{n}\dim\,H_c^{n}(A;\mathbb{R}_A). \tag{9.13}$$

⊚

9.8 Sampling and reconstruction

Sheaf morphisms play a role in Robinson's recent sheaf-theoretic reinterpretation of the Nyquist-Shannon sampling theorem [255]. The classical result in sampling theory says the following. Consider a signal $f\colon \mathbb{R} \to \mathbb{R}$ whose values are known only on some periodic sampling, say $f(n)$ for $n \in \mathbb{Z}$. Under what conditions can f be reconstructed from the samples? A little Fourier analysis yields the classical result: if f is **bandlimited** – that is, the Fourier transform $\hat{f}\colon \mathbb{R} \to \mathbb{C}$ of f has compact support on the interval $[-\frac{1}{2}, \frac{1}{2}]$ – then reconstruction of $f(t)$ from $f(\mathbb{Z})$ is (constructively) possible. This is useful in knowing, *e.g.*, how to encode and transmit human voice for maximum clarity over the telephone/internet; how to filter signals for transmission; and how to compress images and video. The Nyquist rate – that one must sample at a frequency at least twice the signal bandwidth – is iconic in signal processing, and all manner of attempts to generalize it abound in the literature on signal and image processing.

This venerable barrier to reconstruction has a modern topological interpretation [255]. Let X be a cell complex and $\mathcal{F}$ a cellular sheaf on X taking values in abelian

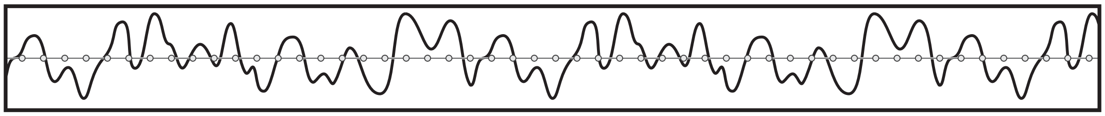

groups. One considers the global sections, $\mathcal{F}(X)$, to be the *signals*: in this section, the desiderata. Sampling will be interpreted not necessarily as a restriction to a subcomplex, but rather as a potentially more complex sheaf morphism, supported on a subcomplex. Define a **sampling** of $\mathcal{F}$ on a (closed) subcomplex $X_0 \subset X$ to be a surjective sheaf morphism $\mathcal{S}\colon \mathcal{F} \to \mathcal{F}_0$ from $\mathcal{F}$ to a **sampling sheaf** $\mathcal{F}_0$ on X, supported on X_0. Surjectivity means that on each cell of X_0, $\mathcal{S}$ is surjective onto that stalk. Such a sampling induces an **ambiguity sheaf**, $\mathcal{A} := \ker \mathcal{S} = \ker(\mathcal{F} \to \mathcal{F}_0)$, on X; the restriction maps of $\mathcal{A}$ are inherited from $\mathcal{F}$. Because of surjectivity of $\mathcal{S}$ and the definition of $\mathcal{A}$ as a kernel, one has the short exact sequence of sheaves on X: $0 \to \mathcal{A} \to \mathcal{F} \xrightarrow{\mathcal{S}} \mathcal{F}_0 \to 0$. This generates as per Equation (9.6) the long exact sequence

$$0 \longrightarrow H^0(X;\mathcal{A}) \longrightarrow H^0(X;\mathcal{F}) \xrightarrow{H(\mathcal{S})} H^0(X_0;\mathcal{F}_0) \xrightarrow{\delta} H^1(X;\mathcal{A}) \longrightarrow \cdots$$

One has the following immediate interpretation. The given sampling is a global section of $\mathcal{F}_0$. The desired signal is a corresponding global section of $\mathcal{F}$. The ability to reconstruct the original signal from the sample is equivalent to having $H(\mathcal{S})$ invertible. By exactness, one observes the following: (1) $H(\mathcal{S})$ is injective iff $H^0(X;\mathcal{A}) = 0$; (2) $H(\mathcal{S})$ is surjective iff $\delta = 0$. Otherwise said, $H^0(X;\mathcal{A})$ is the obstruction to sampling reconstruction, whereas the connecting homomorphism δ indicates the degree of redundancy in the sampling.

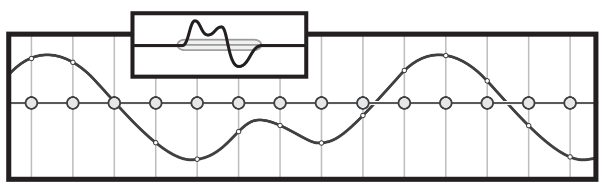

For the example of the classical Nyquist-Shannon theorem: $X = \mathbb{R}$, with the cell structure induced by $X_0 = \mathbb{Z}$; $\mathcal{F}$ is the constant sheaf of continuous $\mathbb{C}$-valued functions supported on a fixed interval $[-B, B]$; $\mathcal{F}_0 = \mathbb{C}_{X_0}$ is the constant $\mathbb{C}$-sheaf over X_0; and the sampling morphism $\mathcal{S}\colon \mathcal{F} \to \mathcal{F}_0$ is the inverse Fourier transform $\mathcal{S}(n, \hat{f}) = \int_{-B}^{B} \hat{f}(\xi)e^{-2\pi i n\xi}d\xi$. One thinks of X as the frequency domain, with a section of $\mathcal{F}$ *not* as the original signal $f(x)$, but as its Fourier transform $\hat{f}(\xi)$. Then, the sampling morphism, being the inverse Fourier transform, yields the values of $f(n)$ for $n \in \mathbb{Z}$ – these are the samples of the original signal $f(x)$.

In this instantiation, the ambiguity sheaf, $\mathcal{A} = \ker \mathcal{S}$, has stalk over $n \in \mathbb{Z}$ equal to the subgroup of all band-limited functions whose n^{th} Fourier coefficient vanishes. The global sections of $\mathcal{A}$, elements of $H^0(X;\mathcal{A})$, therefore consist of bandlimited functions $\hat{f}\colon [-B, B] \to \mathbb{C}$ for which $\int_{-\infty}^{\infty} \hat{f}(\xi)e^{-2\pi i n\xi}d\xi = 0$ for all n. Basic Fourier theory says that for $B \leq \frac{1}{2}$, all the Fourier coefficients of $\hat{f}$ vanish, as must $\hat{f}$. Thus, $H^0(X;\mathcal{A}) = 0$ and signal reconstruction is possible. When $B < \frac{1}{2}$, $\delta \neq 0$, meaning that oversampling has occurred: the $B = \frac{1}{2}$ case is sharp.

This is simply a reinterpretation of Nyquist-Shannon into sheaf-theoretic terms, and neither replaces nor improves the harmonic-analysis proof; as well, using the constant sheaf on $\mathbb{R}$ does not reveal what the long-exact sequence on sheaf cohomology is really capable of. However, the benefits of this more general interpretation are worth noting: the result holds for other kinds of samplings which are not amenable to harmonic analysis, as well as to other base spaces than Euclidean $\mathbb{R}^n$ [255].

Example 9.16 (Cut samples; Flow signals) ⦾
If $\tilde{X}$ is taken to be the subdivision of a directed network X as in §9.4 and $\mathcal{F}$ a network flow sheaf, then a valid sampling morphism is $\mathcal{S}\colon \mathcal{F} \to \mathcal{F}_{\tilde{C}}$, where $\mathcal{F}_{\tilde{C}}$ is the restriction of $\mathcal{F}$ to an edge-cut C, suitably subdivided so that $\tilde{C}$ is a closed subcomplex of $\tilde{X}$. In this case, $H^\bullet(\tilde{X};\mathcal{F}_{\tilde{C}}) = H^\bullet(\tilde{C};\mathcal{F})$. The ambiguity sheaf in this case is $\ker \mathcal{S}$, which yields the relative cohomology $H^\bullet(\tilde{X};\mathcal{A}) = H^\bullet(\tilde{X},\tilde{C};\mathcal{F})$. One recovers the results from §9.4: $H^0(\tilde{X},\tilde{C};\mathcal{F})$ classifies the obstruction to signal reconstruction (*is C a cut?*) and the redundancy of the sampling (*is the cut minimal?*) is measured by δ. The expressiveness of sheaf-theoretic language allows one to think of both the Max-Flow-Min-Cut and the Nyquist-Shannon theorems as expressions of the same principle, wherein the cut is a *sampling* of the flow sheaf *signal*. ⦾

9.9 Euler integration, redux

The Euler integration of Chapter 3, though simple enough to define combinatorially, has its origins in sheaf theory, through independent work of Kashiwara [190] and MacPherson [215] in the 1970s. The interpretation as a calculus was envisaged and promoted by Schapira [269] and Viro [297] independently years later. The subject was independently rediscovered as combinatorics [262, 268] in a more limited form. The sheaf-theoretic version has impacted algebraic geometry [172] and motivic integration [67, 92] independent of its newfound utility in signal/data processing.

Euler calculus is built on a basis of constructible sheaves – these differ slightly from the cellular sheaves introduced thus far in that cellular sheaves begin with a fixed cell structure, but constructible sheaves do not fix the cell structure in advance, but use tameness to imply a cell structure. For simplicity, fix an o-minimal structure as in §3.5 and let X be a definable space. A **constructible sheaf** on X is a cellular sheaf for some definable triangulation of X.

The idea behind the Euler integral is that one converts a constructible function $h\colon X \to \mathbb{Z}$ into a sheaf $\mathcal{F}_h$ whose Euler characteristic *is* the integral of h with respect to $d\chi$. This procedure is simple in the case where $h\colon X \to \mathbb{N}$. On each simplex σ in a tame triangulation of X, define $\mathcal{F}_h$ to be the $\mathbb{R}$-vector space of dimension $h(\sigma)$. All the restriction maps vanish. Put simply, $\mathcal{F}_h$ is a sum of skyscraper sheaves over the simplices that converts h-values to dimensions. Then, by construction,

$$\chi(\mathcal{F}_h) = \chi\left(\bigoplus_\sigma \mathbb{R}^{h(\sigma)}\right) = \sum_\sigma (-1)^{\dim\sigma} h(\sigma) = \int_X h\, d\chi.$$

The only difficulty is in how to adapt the definition to integrands $h\colon X \to \mathbb{Z}$ that can take on both positive and negative values. This is accomplished by switching from individual sheaves $\mathcal{F}_h$ to complexes of sheaves $\mathcal{F}_h^\bullet$ of the form

$$\mathcal{F}_h^\bullet = \cdots \longrightarrow 0 \longrightarrow \mathcal{F}_h^- \longrightarrow \mathcal{F}_h^+ \longrightarrow 0 \longrightarrow \cdots$$

where $\mathcal{F}_h^+$ the sum of skyscraper sheaves for the restriction of h to $\{h > 0\}$ and $\mathcal{F}_h^-$ the sum of skyscraper sheaves for the restriction of $-h$ to $\{h < 0\}$. The grading in the sequence above is chosen so that

$$\chi(\mathcal{F}_h^\bullet) = \chi(\mathcal{F}_h^+) - \chi(\mathcal{F}_h^-) = \int_X h\,d\chi.$$

The sheaf-theoretic perspective provides motivation for some of the otherwise esoteric-seeming results of Chapter 3. The Fubini-type result (Theorem 3.11) is nothing more than the commutativity of χ-on-sheaves with the direct image operation. Likewise, the Fourier-Sato transform, $\mathcal{F}_{\mathcal{S}}$, of Example 3.18, is a particular instance of the general **Fourier-Sato transform** of sheaves [191] in the setting of constructible functions. The student who progresses in sheaf theory may find it helpful to translate as much of the general machinery as possible into specific examples from Euler calculus.

9.10 Cosheaves

Sheaf theory is built for cohomology; one suspects that the homological counterpart should be and be by no means unimportant. This dual is called a **cosheaf** and has its own set of emerging applications.

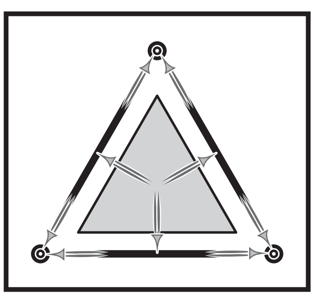

A cosheaf is a data structure over a space that encodes extension rather than restriction. For a cell complex X, a cosheaf $\hat{\mathcal{F}}$ assigns to each cell σ an abelian group $\hat{\mathcal{F}}(\sigma)$ and to face pairs $\sigma \trianglelefteq \tau$ an **extension map** – a homomorphism $\hat{\mathcal{F}}(\sigma \trianglelefteq \tau)$ which sends $\hat{\mathcal{F}}(\tau) \to \hat{\mathcal{F}}(\sigma)$ and composes properly: for $\rho \trianglelefteq \sigma \trianglelefteq \tau$, the extension maps satisfy $\hat{\mathcal{F}}(\rho \trianglelefteq \tau) = \hat{\mathcal{F}}(\rho \trianglelefteq \sigma) \circ \hat{\mathcal{F}}(\sigma \trianglelefteq \tau)$.

One imagines extending data from cells to faces. Note how this flips directions in the definition of a sheaf. Continuing the pattern of appending the *co-* prefix to connote this contravariance, one calls the data $\hat{\mathcal{F}}(\sigma)$ over a cell the **costalk**. From this basis of costalks, one assembles the full cosheaf structure on X. In analogy with Equation (9.2), one defines:

$$\hat{\mathcal{F}}(X) := \bigoplus_\tau \hat{\mathcal{F}}(\tau) \Big/ \sim \quad : \quad s_\rho \sim \hat{\mathcal{F}}(\rho \trianglelefteq \sigma)(s_\sigma). \tag{9.14}$$

This is much harder to parse than the global-sections intuition of a sheaf over X: the cosheaf over X is an equivalence class of data assignments to cells of X, where the equivalence is up to compatible extensions that fall from higher-dimensional to lower-dimensional cells.

As a simple example, consider the skyscraper cosheaf $\hat{\mathbf{G}}_\sigma$ that assigns an abelian group $\mathbf{G}$ to a cell σ and 0 to all other cells. The value of this cosheaf over all of X is zero *unless* σ has no cofaces (*i.e.*, it is locally top-dimensional), in which case the global sections are $\hat{\mathcal{F}}(X) \cong \mathbf{G}$. Recall: global nonzero sections of a skyscraper *sheaf* exist only if the cell is minimal in the face poset (a vertex); for a skyscraper *cosheaf*, this is turned upside down and only maxima in the face poset contribute globally.

Example 9.17 (Sensing and inference) ⊚

The duality between sheaves and cosheaves is expressed in applications to sensing. Consider the problem of integrating sensor data into an inference. Assume that for, say, a finite collection of sensors $x_i \in \mathcal{Q}$, each returns some sensed attributes $S(x_i) \subset \mathcal{S}$ about a fixed object $z \in \mathcal{O}$. The set of nodes defines a simplex X (of dimension $|\mathcal{Q}|-1$) that will act as a base space. Consider the sheaf on X taking values in vector spaces defined as follows. To each vertex x_i is assigned the vector space whose basis is the finite set $S(x_i)$ of attributes: for example, *the target is grey and large*; *the target has a long nose*. The restriction maps act as inclusions to vector spaces obtained by taking the unions over the vertices of the associated simplex. The global sections $\mathcal{F}(X)$ give the full sensorium. *Sensing yields a sheaf.*

Dually, there is a cosheaf defined in terms of what the sensors allow one to conclude. Consider the assignment of a vector space $\hat{\mathcal{F}}(x_i)$ to each vertex $x_i \in \mathcal{Q}$ whose basis is all objects in $\mathcal{O}$ consistent with $S(x_i)$: *rhinoceros, elephant, whale*; *anteater, elephant*. The extension maps are projections to the subspace defined by a subbasis generated by intersection: *elephant*. The global sections of this cosheaf $\hat{\mathcal{F}}(X)$ return constraint satisfactions consistent with sensing. *Inference is a cosheaf.* One suspects this cartoonish example can be greatly generalized. ⊚

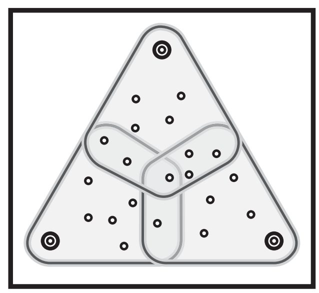

As foreshadowed, cosheaves are built for homology. The obvious cellular cosheaf chain complex is given as,

$$\cdots \xrightarrow{\partial} \bigoplus_{\dim\ \sigma=2} \hat{\mathcal{F}}(\sigma) \xrightarrow{\partial} \bigoplus_{\dim\ \sigma=1} \hat{\mathcal{F}}(\sigma) \xrightarrow{\partial} \bigoplus_{\dim\ \sigma=0} \hat{\mathcal{F}}(\sigma) \longrightarrow 0 \tag{9.15}$$

where the boundary map is:

$$\partial(\tau) = \sum_{\sigma \trianglelefteq \tau} [\sigma : \tau]\, \hat{\mathcal{F}}(\sigma \trianglelefteq \tau). \tag{9.16}$$

Notice that ∂ *decreases* the grading by one and that $\partial^2 = 0$, as a boundary operator must. The resulting **cosheaf homology**, $H_\bullet(X; \hat{\mathcal{F}})$ gives a collation and classification of homological features in the data structure $\hat{\mathcal{F}}$.

If $\mathbf{G}_X$ is the constant cosheaf on X, then the cosheaf homology agrees with the cellular homology with $\mathbf{G}$ coefficients. For a skyscraper cosheaf $\hat{\mathcal{F}}_\sigma$ the cosheaf homology vanishes unless σ is a maximal cell in the face poset (a cell of locally top dimension), in which case $H^\bullet(X; \hat{\mathcal{F}}_\sigma) \cong \mathbf{G}$ in grading $\dim\ \sigma$. This duality between dimensions and sheaf cohomology / cosheaf homology on skyscrapers is reminiscent of the delicacies of duality between homology and compactly supported cohomology from Chapter 6.

Example 9.18 (Precosheaves) ⊚

In the topological setting, a **precosheaf** assigns data to open sets $U \subset X$ and for a subset $V \subset U$ the extension map $\hat{\mathfrak{F}}(U \trianglelefteqslant V)$ acts as $\hat{\mathfrak{F}}(V) \to \hat{\mathfrak{F}}(U)$ with the corresponding composition in the case of $W \subset V \subset U$. To be a cosheaf, a precosheaf must behave well under limits of covers: The analogue of Equation (9.9) for cosheaves is the dual condition, that

$$\bigoplus_{i,j} \hat{\mathfrak{F}}(U_{ij}) \xrightarrow{\partial} \bigoplus_k \hat{\mathfrak{F}}(U_k) \xrightarrow{\hat{\mathfrak{F}}(U \trianglelefteqslant U_k)} \hat{\mathfrak{F}}(U) \longrightarrow 0 \ . \tag{9.17}$$

is exact for any finite open cover $\mathcal{U} = \{U_i\}$ of U. In other words, $\hat{\mathfrak{F}}(U)$ agrees with the cellular cosheaf homology $H_0(\mathcal{N}(\mathcal{U}); \hat{\mathfrak{F}})$ for any cover $\mathcal{U}$. In practice, this agreement *cannot* be assumed and can be challenging to confirm.

Many of the examples of sheaves over a topology come from functions: continuous functions, smooth functions, vector fields, differential forms, etc. The dual cosheaf notions are generated (taking a hint from Poincaré duality) by restricting to functions of compact support. For example, let M be a manifold. Then the precosheaf on M of compactly supported continuous functions $\hat{\mathfrak{F}}(U) = C_c(U, \mathbb{R})$ with the extension map from V to $U \supset V$ being extension-by-zero is a cosheaf. Likewise, the sheaf of differential p-forms Ω^p on M is complemented by the cosheaf of p-currents Ω_p (see §6.11). ⊚

Example 9.19 (Homology fiber cosheaves) ⊚

One simple example of a cosheaf comes from a map $f : X \to Y$ where Y is a locally connected space. The precosheaf that assigns to $U \subset Y$ the homology $H_k(f^{-1}(U))$ is in fact a cosheaf on Y, with extension arising from the induced homomorphism on H_k. The cellular case is analogous (*cf.* Example 9.3). ⊚

9.11 Bézier curves and splines

The problem of patching together local data over a cell complex is commonly encountered by architects and designers working with polyhedral representations of objects. The subject is classical, as exemplified by the simple problem of how to specify a polynomial planar curve γ between two endpoints with global control over the degree and local control at the endpoints. This is precisely the context in which a **Bézier curve** is appropriate.

For example, a planar Bézier curve is specified by the locations of the two endpoints, along with additional **control points**, each of which may be interpreted as a *handle* (or tangent vector at the endpoint) specifying derivative data of the resulting curve at each endpoint. The reader who has used any modern drawing software will have a visceral understanding of the control that these handles give over the resulting smooth curve. Most programs use a cubic Bézier curve in the plane – the image of a cubic polynomial $\vec{p}(t)$ for $0 \leq t \leq 1$. In these programs, the specification of the endpoints and the endpoint handles (tangent vectors) completely determines the curve uniquely.

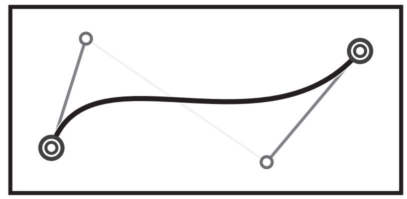

This can be viewed from the perspective of a cosheaf $\hat{\mathcal{S}}$ over the closed interval $I = [0, 1]$. The costalk over the interior $(0, 1)$ is the space of all cubic polynomials from $[0, 1] \to \mathbb{R}^2$, which is isomorphic to $\mathbb{R}^4 \oplus \mathbb{R}^4$ (one cubic polynomial for each of the x and y coordinates). If one sets the costalks at the endpoints 0, 1 to be $\mathbb{R}^2$, the physical locations of the endpoints, then the obvious extension maps to the endpoints are nothing more than evaluation at 0 and 1 respectively. The corresponding cosheaf chain complex is:

$$\cdots \longrightarrow 0 \longrightarrow \mathbb{R}^4 \oplus \mathbb{R}^4 \xrightarrow{\partial} \mathbb{R}^2 \oplus \mathbb{R}^2 \longrightarrow 0.$$

Here, the boundary operator ∂ computes how far the cubic polynomial (edge costalk) 'misses' the specified endpoints (vertex costalks). It is clear that $H_0(\hat{\mathcal{S}}) = 0$, since ∂ is surjective. It is also clear that the space of global solutions is $H_1(\hat{\mathcal{S}}) \cong \mathbb{R}^2 \oplus \mathbb{R}^2$, meaning that there are four degrees of freedom available for a cubic planar Bézier curve with fixed endpoints: these degrees of freedom are captured precisely by the pair of handles, each of which is specified by a (planar) tangent vector. Note the interesting duality: the global solutions with boundary condition are characterized by the top-dimensional homology of the cosheaf, instead of the zero-dimensional cohomology of a sheaf.

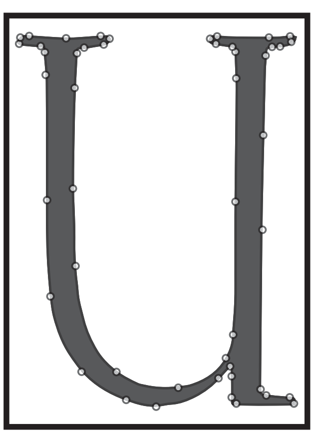

In practice, a compound curve drawn in any vector graphics software is a **spline** obtained by patching together segments of Bézier curves, usually cubic. Such curves are specified by handles at each endpoint of each segment. For a C^1 (differentiable) curve, one aligns the directions of the handles at segment endpoints; corners can also be specified, yielding a C^0 curve. This finite-dimensional representation of complex systems of curves is of foundational importance: all modern font systems[3] are based on such splines.

A cosheaf for a polynomial Bézier surface over a polygonal 2-cell with constraints on boundary points would be more challenging to write out, but would have meaningful homology in degree two. Patched together, such higher-dimensional polynomial splines are crucial in architecture and design, contributing immensely to both aesthetic and technical aspects of what can be built. This prompts the study and classification of splines with various constraints.

The following is a classification of simpler, $\mathbb{R}$-valued splines over simplicial manifolds, based on independent work of Billera [36] and Yuzvinsky [304] reinterpreted in

[3]True type fonts use quadratic Bézier splines; Postscript and Metafont use cubic Bézier splines.

the context of cellular cosheaves. Given $X \subset \mathbb{R}^n$ a simplicial complex realized with convex simplices, consider P_m the ring of $\mathbb{R}$-valued polynomials in n variables which have degree $\leq m$. By S^r_m denote the (vector) space of degree $\leq m$ $\mathbb{R}$-valued splines on X which are of global smoothness class C^r, for fixed $m > r \geq 0$. For example, a dome over a triangulated disc might be specified as an element of S^1_4 if it were built of quartic surface patches which must meet preserving first derivatives.

The following cosheaf captures the spline constraints. Let $\hat{\mathcal{S}}^r_m$ denote the cellular cosheaf on X whose costalk on a simplex σ is the quotient

$$\hat{\mathcal{S}}^r_m(\sigma) = P_m / (P_m \cap I^r(\sigma)) ,$$

where $I^r(\sigma)$ is the subspace of polynomials in n variables which vanish on σ (and thus on the affine space spanned by this convex simplex) but which have complementary degree at least $r+1$. The extension map is given as follows: for $\sigma \trianglelefteq \tau$, then $I^r(\tau) < I^r(\sigma)$ and there is a natural map on the quotient. This satisfies the cosheaf composition law and thus yields a well-defined cosheaf homology.

Theorem 9.20 ([36]). *If X as above is an n-dimensional simplicial manifold-with-boundary, then $S^r_m \cong H_n(X - \partial X; \hat{\mathcal{S}}^r_m)$.*

Note that $\hat{\mathcal{S}}^0_1$ is precisely the setting of continuous piecewise-linear functions. In this case, it can be shown that $\dim S^0_1 = |V|$, the number of vertices of X. One can imagine the use of a long exact sequence to characterize constraints, the use of Euler characteristic to bound dimensions, and the use of exact sequences of cosheaves to decompose S^r_m: all this and more appear (implicitly) in [36].

9.12 Barcodes, redux

The full story of cosheaves and their applications has yet to be written. This chapter closes with an example from Curry [77] related to topological data analysis. Recall the setting of §5.13-5.15, in which a linear sequence of homologies and induced homorphisms forms, via indecomposables, infographics of parameterized homologies: barcodes. For example, consider the *zigzag* sequence of the form

$$\cdots \longleftarrow X_i \longrightarrow X_i \cup X_{i+1} \longleftarrow X_{i+1} \longrightarrow \cdots .$$

Passing to k^{th} homology H_k in field-$\mathbb{F}$ coefficients gives a sequence of $\mathbb{F}$-vector spaces with alternating linear transformations. This k^{th} homology of the sequence is precisely a cosheaf over a base space $\mathbb{R}$, discretized by $\mathbb{Z} \subset \mathbb{R}$. Costalks over edges, $H_k(X_i; \mathbb{F})$, are mapped to costalks over vertices, $H_k(X_i \cup X_{i-1}; \mathbb{F})$ and $H_k(X_i \cup X_{i+1}; \mathbb{F})$. There are no higher-dimensional cells, and so composition is trivially satisfied.

It is not necessary to begin with alternating zigzag sequences. For example, the monotone sequences of §5.13 can be recast as a cosheaf over $\mathbb{R}$ by interweaving backward-pointing identity maps, as done with recurrence relations in Example 9.1. By such means, one can recast the theory of barcodes into the language of cosheaves. *A homology barcode is a cosheaf over $\mathbb{R}$.*

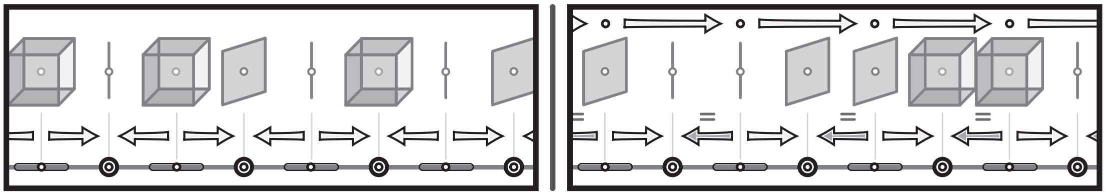

Homology cosheaves with field coefficients are classified completely by the Structure Theorem 5.21 into **cosheaf interval indecomposables** which are the constant cosheaves with costalk $\mathbb{F}$ over some interval. Note, however, that because of the decomposition of $\mathbb{R}$ into vertices and edges, these indecomposable intervals can have endpoints corresponding either to vertices or edges of the cell structure on $\mathbb{R}$. There are exactly four different types of bars that can arise as cosheaf interval indecomposables: a closed interval, an open interval, or (left/right) half-open intervals. Since cosheaves form coefficients for homology, one can sensibly speak of the homology of a cosheaf interval indecomposable. Over a base space of $\mathbb{R}$, the only homology that can be nonvanishing is H_0 and H_1. It is a simple exercise to compute:

Lemma 9.21 ([77]). *Cosheaf interval indecomposables and their homologies are:*

interval	type	H_0	H_1
closed bar	$\hat{\mathfrak{F}} = \bullet --- \bullet$	$H_0(\mathbb{R};\hat{\mathfrak{F}}) \cong \mathbb{F}$	$H_1(\mathbb{R};\hat{\mathfrak{F}}) \cong 0$
open bar	$\hat{\mathfrak{F}} = \circ --- \circ$	$H_0(\mathbb{R};\hat{\mathfrak{F}}) \cong 0$	$H_1(\mathbb{R};\hat{\mathfrak{F}}) \cong \mathbb{F}$
left-open bar	$\hat{\mathfrak{F}} = \circ --- \bullet$	$H_0(\mathbb{R};\hat{\mathfrak{F}}) \cong 0$	$H_1(\mathbb{R};\hat{\mathfrak{F}}) \cong 0$
right-open bar	$\hat{\mathfrak{F}} = \bullet --- \circ$	$H_0(\mathbb{R};\hat{\mathfrak{F}}) \cong 0$	$H_1(\mathbb{R};\hat{\mathfrak{F}}) \cong 0$

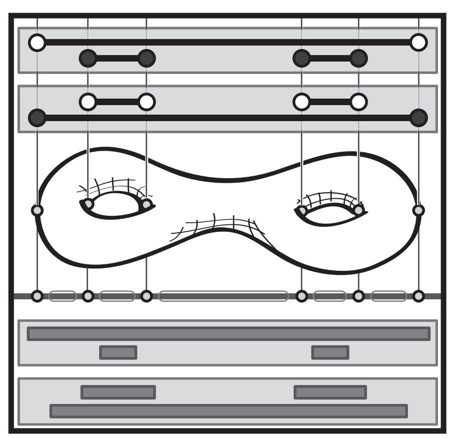

This is neither frivolous nor inconsequential, as evidenced in the following cosheaf interpretation of level-set persistence. Let $h\colon X \to \mathbb{R}$ be a cellular map from a compact cell complex X to the reals, outfitted with a cell structure whose vertex set $\{v_i\}$ includes critical values of h. Let $\hat{\mathcal{B}}_k$ be the k^{th} homology fiber cosheaf of h with coefficients in a field $\mathbb{F}$. Specifically, given the vertices $v_0 < v_1 < \cdots < v_N \subset \mathbb{R}$, define the k^{th} homology cosheaf $\hat{\mathcal{B}}_k$ as the cosheaf of h-level-set preimage homologies whose extension maps are induced on homology by inclusions as in the following diagram:

$$\cdots \longleftarrow H_k(v_{i-1}<h<v_i) \longrightarrow H_k(v_{i-1}<h<v_{i+1}) \longleftarrow H_k(v_i<h<v_{i+1}) \longrightarrow \cdots .$$

The following theorem blends homology, persistence, cosheaves, and Morse theory into a compact package:

Theorem 9.22 ([77], cf. [78]). *For $h\colon X \to \mathbb{R}$ and $\hat{\mathcal{B}}_k$ as above, the homology of X in $\mathbb{F}$-coefficients can be computed in terms of the homology of $\mathbb{R}$ with coefficients in the homology cosheaf: for any n,*

$$H_n(X;\mathbb{F}) \cong H_0(\mathbb{R};\hat{\mathcal{B}}_n) \oplus H_1(\mathbb{R};\hat{\mathcal{B}}_{n-1}).$$

As a simple example, consider the canonical Morse function $h\colon S \to \mathbb{R}$ on a compact genus-2 surface with one minimum, four saddle critical points, and one maximum. The corresponding critical values induce a cell structure on $\mathbb{R}$ and the nonzero homology cosheaves are $\hat{\mathcal{B}}_0$ and $\hat{\mathcal{B}}_1$. These have interval decompositions as follows: $\hat{\mathcal{B}}_0$ has one closed bar and two open bars; $\hat{\mathcal{B}}_1$ has one open bar and two closed bars. Theorem 9.22 allows one to read off $H_\bullet(S;\mathbb{F})$ trivially by counting these intervals. For example, $\dim H_1(S;\mathbb{F})$ equals the number of closed intervals in $\hat{\mathcal{B}}_1$ plus the number of open intervals in $\hat{\mathcal{B}}_0$. The reader will note that this graphical language also expresses Poincaré duality as a neatly visceral symmetry on the level of cosheaf barcodes.

Notes

1. Sheaf theory has had an enormous impact within mathematics, in algebraic geometry, number theory, complex analysis, logic, and more. The agricultural terminology that suffuses this subject is entwined with France and the fascinating history of sheaves.
2. This chapter pretends to survey sheaf theory; in reality, the subject is so broad and deep as to evade compression. Experts in sheaf theory will find this chapter ridiculously elementary. This, however, is how mathematics transitions from pure to applied: slowly and through simple examples. It is the author's hope that the themes of elementary homological algebra, persistence, and sheaves will make their way via applications into the undergraduate linear algebra curriculum, where they belong.
3. The perspective of using cellular sheaves is unorthodox, but hopefully edifying. Cellular sheaves and cosheaves have developed in fits and starts. Their first explicit instantiation (so far as the author can tell) was by Fulton, Goresky, MacPherson, and McCrory in their famous 1977-78 seminar. Relevant references include the theses of Shepard [276], Vybornov [298] and Curry [77].
4. Sheaves, like simplicial complexes, are defined abstractly but have a geometric realization. To any sheaf $\mathcal{F}$ is associated an **étale space**: a topological space (also denoted $\mathcal{F}$) and a projection $\pi\colon \mathcal{F} \to X$ that, like a covering space, is a local homeomorphism with discrete fibers. Sections of the sheaf are precisely sections to the map π. This perspective can be useful; however, as the topology on the étale space is almost always non-Hausdorff, it is very unenlightening from the point of view of visualizing sheaves and their cohomology.
5. Sheaves taking values in vector spaces or abelian groups are convenient for doing cohomology; if one is willing to give up a simply-defined H^k for $k > 0$, then it is possible to work with sheaves taking values in sets, spaces, or other categories (see §10.9). Such non-abelian data types are certainly natural and useful, though not covered in this chapter.
6. Most of the examples of this chapter are intentionally elementary, with networks as base spaces. Even so, it is a harbinger that such simple sheaves have clean applications. A number of applications of sheaves to data have recently emerged, using networks as a base space [146, 150, 201, 254, 255, 256].

7. Several applications of sheaves have been suggested in the computer science literature: see, *e.g.*, [159] and subsequents. Structural aspects seem to have displaced applications as the focus of this line of inquiry. This author believes that sheaf cohomology is the missing ingredient to make sheaf theory applied as opposed to applicable.
8. The full statement and proof of the Lefschetz formula of Goresky-MacPherson in §7.7 uses sheaf theory extensively.
9. The treatment of flow sheaves in §9.4 does not provide an independent proof of the classical Max-Flow/Min-Cut theorem, since the interpretations of the relative sheaf cohomology as obstructions subtly entail the theorem, and a careful proof of the interpretation and computation of obstructions would re-prove the result. There is a wholly sheaf-theoretic proof of the max-flow min-cut theorem by Krishnan [201]: see Example 10.25.
10. Constructible functions on a definable X are locally defined, and thus form a sheaf CF_X. Euler integration, the Fubini Theorem, and more follow directly from canonical sheaf operations [270].
11. The use of cosheaves in splines is not in the literature, but it is an obvious reformulation of existing work. The initial work of Billera [36] presented a chain complex by fiat and argued that the homology was meaningful. Yuzvinsky [304] reformulated the problem as a sheaf over the dual poset to a certain hperplane arrangement associated to the cell complex. Schenck and collaborators [223, 271] applied methods from spectral sequences and more classical commutative algebra to push the theory further. This subject blends softly into algebraic geometry; the treatment in this text is intended to emphasize topological aspects of the problem.
12. The reformulation of persistent homology and barcodes in terms of cosheaves is from the thesis of Curry [77]. Robinson [256] likewise recently interpreted persistent cohomology as a sheaf over $\mathbb{R}$. Theorem 9.22 can be derived from earlier works [57, 52] without using cosheaves: both of these, and the dual sheaf-theoretic version [78] are consequences of the **Leray spectral sequence**.
13. It is to be hoped that the language of sheaves and cosheaves provides the key to understanding and computing multi-dimensional persistence by means of generalized barcodes. This ambitious project will likely require nontrivial contributions from both sheaf theory and representation theory.
14. It is not an exaggeration to say that in sheaf theory, individual sheaves are of secondary importance. The real power in sheaf theory comes from dextrous use of complexes of sheaves and sheaf morphisms. This text has too little room to demonstrate this fully, as well as to unfold the rest of the six canonical operations of which direct and inverse image are two.