

Chapter 4

Discrete Meshes

In this chapter we will talk about taking the smooth concepts we learned about so far and applying them to the discrete setting of meshes. We'll begin with a more formal treatment of what makes a mesh before describing the discrete version of exterior calculus. We will also talk about interpolation of discrete k -forms and conversely the sampling of smooth k -forms for the discrete setting.

4.1 Abstract Simplicial Complexes and Manifolds

When we speak informally of a triangle mesh or a tetrahedral mesh we have a pretty clear picture before our mental eye: an embedded simplicial manifold (2D or 3D) made of piecewise linear (PL) triangles resp. tetrahedra glued together “in the obvious way.” Here we want to fix our notation, make these notions a bit more precise, and articulate them in a way that makes them more general even if most of the time the 2D and 3D setting will be sufficient.

Lets start with the example of a triangle mesh and describe its connectivity (not yet the geometry). It is a complex K made of (triangle) facets, edges, and vertices, $K = \{V, E, F\}$. We'll “name” the vertices with the integers from 1 to n (assuming we have $|V| = n$). Depending on context we may use just an integer i to refer to a vertex or say v_i to emphasize that we are referring to a vertex. For example $V = \{v_i\}_{i=1,\dots,n}$. With this in place edges are pairs of integers $E = \{e_{ij}\}_{i,j \in V}$, where e_{ij} connects vertex v_i (tail) to v_j (head). Triangle facets are triples of integers $F = \{f_{ijk}\}_{i,j,k \in V}$. Note that $f_{ijk} = f_{jki} = f_{kij}$, which has orientation opposite to $f_{jik} = f_{kji} = f_{ikj}$. For tetrahedral meshes $K = \{V, E, F, T\}$ we'll also have tetrahedra which are quadruples of indices $T = \{t_{ijkl}\}_{i,j,k,l \in V}$. There are again two possible orientations depending on the sign of the permutation of the indices. And more generally k -simplices are given as $(k+1)$ -tuples of indices into the vertex set.

Such complexes are called **abstract simplicial complexes** if for every simplex $\sigma \in K$ its subsets are also in K . So in the case of a triangle mesh we require that for every facet $f \in F$ all its edges $e \subset f$ are in E and for every edge its vertices are in V . For tetrahedral meshes one adds the corresponding requirement that for all tetrahedra $t \in T$ their respective triangle facets must be contained in F .

Definition 4.1. A (finite) simplicial complex $\mathcal{K}$ is a set of subsets of the integers $\{1, \dots, n\}$ such that any subset of a given subset is also in the complex. The dimension m of $\mathcal{K}$ is one less than the cardinality of the largest subset in $\mathcal{K}$ (“the highest dimensional simplex”).

To do differential geometry we need a little bit more though. We want not just a 2-dimensional simplicial complex (triangle mesh) or a 3-dimensional simplicial complex (tetrahedral mesh), we also want them to be 2- resp. 3-dimensional orientable manifolds, possibly with boundary. For example, no dangling edges, shark fins, or “double cones” (see Fig. 4.1).




Figure 4.1: Examples of non-manifold simplicial complex configurations we must exclude: dangling edges (left), “shark fins” (middle), and “double cones” (right).

To be able to talk about a continuous surface (or volume) we define the **carrier of a simplex**. For a vertex it is just the vertex itself. For an edge e_{ij} the carrier consists of the formal convex combinations of its vertices, i.e., the set $\{v(\theta) := \theta v_i + (1 - \theta)v_j | \theta \in [0, 1]\}$. For a given point along an edge, θ will be called its barycentric coordinate. Similarly the carrier of a facet f_{ijk} is the set $\{\alpha v_i + \beta v_j + \gamma v_k | \alpha, \beta, \gamma \in [0, 1], \alpha + \beta + \gamma = 1\}$ and any point inside a triangle is described by its barycentric coordinates (α, β, γ) .

With this in place it makes sense to ask whether a given triangle mesh is 2-manifold. For that every point in the mesh must have a disk like neighborhood within the mesh (see Fig. 4.2). This is obviously the case for any point in the strict interior of a triangle facet. What about points in the strict interior of an edge? If the edge has exactly two triangles incident on it, those points too satisfy the 2-manifold property. Finally a vertex has a disk like neighborhood if the set of triangles incident on it is singly connected. If we allow the surface to have a boundary, i.e., if it is a 2-manifold with boundary, the boundary will be characterized by, possibly multiple, closed loops of edges along which all points have only a half-disk as a neighborhood.




Figure 4.2: For a triangle mesh to be 2-manifold each point must have an open neighborhood which is a disk. On the left a point in the open interior of a triangle, in the middle a point on the open interior of an edge, and on the right the neighborhood of a vertex.

As a practical aside, what does this mean for specifying a valid surface triangle mesh for input into a program? A standard way to describe such meshes is some variant of the following format. First a list of vertices with their attributes, say, point positions, followed by a list of integer triples giving the triangles as indices into the list of vertices. Here one must be careful to check whether the indexing is 0 or 1 based. Both conventions exist.

Exercise 4.1. What tests would an input parser need to perform to ascertain that the file describes a valid 2-manifold? What about being orientable?

So far the carrier of the simplicial complex, consisting of the union of the carriers of all its simplices is still abstract. In many settings though we actually have an embedding, typically in $\mathbb{R}^3$. Suppose we have point positions $\{p_i \in \mathbb{R}^3\}_{i \in V}$ associated with all vertices. In that case the carrier of an abstract edge is mapped to the line segment $p_{ij}(\theta) = \theta p_i + (1 - \theta)p_j$, $\theta \in [0, 1]$. Similarly triangle facets become the usual linear triangles given by the point positions of their respective vertices with $p_{ijk}(\alpha, \beta, \gamma) = \alpha p_i + \beta p_j + \gamma p_k$, $\alpha, \beta, \gamma \in [0, 1]$, and $\alpha + \beta + \gamma = 1$. This would be the standard piecewise linear (PL) triangle mesh we usually envision when we think of a triangle mesh in space. In a tetrahedral mesh the carrier of a tetrahedron consists of all points which are convex linear combinations of its four vertex points.

Exercise 4.2. Given a point inside an embedded triangle, how does one find its corresponding barycentric coordinate? What about a line or a tetrahedron?

In order to integrate k -forms our manifold must be orientable. What does this mean? Recall that the defining property of a smooth m -manifold M is the existence of a smooth parameterization (a diffeomorphism) of a neighborhood of any $p \in M$ over a region in $\mathbb{R}^m$. Since $\mathbb{R}^m$ has an orientation the parameterization induces an orientation on the neighborhood of $p \in M$. A manifold is orientable if we can choose parameterizations such that whenever they overlap they agree on the orientation. Suppose we have done this. Then we have a globally consistent orientation on the entire manifold which is induced by the canonical orientation on $\mathbb{R}^m$ around each point. The discrete case is no different and we orient the highest dimensional simplices

using the induced orientation from the parameterization (see Fig. 4.3).

As a practical matter we also need to figure out how to store/access simplices and their associated data. For example, the six permutations of $[ijk]$ all refer to the same simplex (triangle), falling into two groups, “even” and “odd,” depending on the sign of the permutation relative to some canonical order. Consequently the integral of some 2-form α on the triangle $[ijk]$ will yield a value with sign opposite to that same 2-form integrated over the triangle $[jik]$. A simple naming/storage policy to address this uses sorted index order as the canonical name of a given simplex and maps all other names (the permutations) to the canonical one keeping track of the sign. Consider for example the normals associated with triangles in a triangle mesh. (Note that a consistent assignment of normals requires orientability and we assume that everything has in fact been oriented.) Assume that $[ijk]$ represents sorted order for some triangle. The normals would reside in a container of $|T|$ elements accessed by integer triples. Retrieving N_{jik} , e.g., would return $-N_{ijk}$.

4.2 Discrete Exterior Calculus

Now that we have an understanding of what we mean by discrete geometry in the form of an orientable m -manifold (possibly with boundary) simplicial complex, we are ready to transfer the concepts of smooth exterior calculus to this setting. Our treatment will be independent of dimension though the reader is invited to visualize surface meshes (“triangle meshes”) and/or volume meshes (“tet(rahedral) meshes”) as the most common case.




Figure 4.3: The top level simplices are oriented according to an orientation induced by the parameterization. This can be done consistently if and only if the manifold is orientable. On the left a triangulation which is orientable while the Möbius strip on the right does not admit a consistent orientation.

4.2.1 Discrete Differential

Our main task is the appropriate definition of the differential, and with it the concept of a discrete differential k -form, in the simplicial mesh setting. To this end we will employ Stokes' theorem to *define* the differential in terms of the boundary operator

$$\int_M d\omega := \int_{\partial M} \omega.$$

Thinking of integration as a bi-linear pairing of k -domains and k -forms the **differential is the adjoint of the boundary operator**

$$\langle M, d\omega \rangle := \langle \partial M, \omega \rangle.$$

With this definition we do not need to worry about the details of d , rather we need to ensure that taking a boundary and subsequent integration are well defined. This requires less smoothness in the form than if we insisted on the ability to evaluate the form pointwise. Another, tautological, consequence of using Stokes' theorem to *define* the discrete exterior derivative is that Stokes' theorem holds exactly at the discrete level. This in turn means that any fact relying on Stokes' theorem in the smooth setting will have its counterpart at the discrete level.

Remark 4.1. Let's tease apart in a bit more formal manner the notion of d as the adjoint of ∂ , i.e., d as the **co-boundary operator** in our finite dimensional discrete setting. To talk about the domain we define the linear space of k -chains $C_k(K; \mathbb{R})$, i.e., formal linear combinations of k -simplices from the complex K . This is the space on which ∂ acts as a linear operator, $\partial: C_k(K; \mathbb{R}) \rightarrow C_{k-1}(K; \mathbb{R})$. A non-degenerate bi-linear pairing then induces a duality relationship. For us that is the pairing of discrete k -domains with discrete k -forms via integration. Let ω be a k -form and notice that $\langle \cdot, \omega \rangle = \int_{(\cdot)} \omega$ is a linear functional acting on elements of $C_k(K; \mathbb{R})$. This makes discrete k -forms into elements of the dual space $C_k^*(K; \mathbb{R}) =: C^k(K; \mathbb{R})$, so-called **co-chains** and we can talk about the adjoint of ∂ as the map $\partial^*: C_{k-1}^*(K; \mathbb{R}) \rightarrow C_k^*(K; \mathbb{R})$. Using Stokes' theorem to relate d and ∂ , as well as the identification $C_k^*(K; \mathbb{R}) = C^k(K; \mathbb{R})$, we arrive at $d: C^{k-1}(K; \mathbb{R}) \rightarrow C^k(K; \mathbb{R})$. To express these linear operators as matrices we use k -simplices of our mesh as a basis for the space of k -chains $C_k(K; \mathbb{R})$. The pairing via integration then induces a dual basis for the space of k -co-chains $C^k(K; \mathbb{R}) = C_k^*(K; \mathbb{R})$, which are our discrete k -forms. With these bases the discrete differential is the transpose of the incidence matrix, $d_{k-1} = \partial_k^T$ (here the subscript indicates the order resp. dimension of the argument).

Luckily the boundary of a simplicial complex is straightforward and has a simple definition. Let $\sigma^k = [i_0, \dots, i_k]$ denote a k -simplex given by $k+1$ vertices $i_0 \dots i_k \in V$, then

$$\partial \sigma^k := \sum_{j=0}^k (-1)^j [i_0, \dots, \hat{i}_j, \dots, i_k].$$




Figure 4.4: Example of the boundary operator when applied to the tetrahedron $[1234]$. The four boundary triangles are found by dropping one index at a time. The orientation of a boundary facet either agrees (+) or disagrees (−) with that of the tetrahedron.

In other words, the boundary of a k -simplex consists of the $(k-1)$ -simplices one gets by dropping (indicated by the hat notation) one of the indices and adjusting the sign (see Fig. 4.4).

Example 4.1. The boundary of a tetrahedron consists of the 4 triangles one gets when dropping one of the 4 indices making up the original tetrahedron (see Fig. 4.4). The boundary of a triangle results as the 3 edges one gets when dropping one each of the original indices. For an edge we get the two incident vertices and for vertices the boundary is empty.

With this approach, a discrete k -form becomes an assignment of values to k -simplices on the oriented manifold simplicial complex we are working with. Consider for example a tetrahedral simplicial complex $K = \{V, E, F, T\}$ made of the set of vertices (V), edges (E), facets (F), and tetrahedra (T) and the corresponding discrete 0-, 1-, 2-, 3-forms

$$\begin{array}{ll}
 \alpha^0: V \rightarrow \mathbb{R} & [i] \mapsto \alpha_i^0 \\
 \alpha^1: E \rightarrow \mathbb{R} & [ij] \mapsto \alpha_{ij}^1 \\
 \alpha^2: F \rightarrow \mathbb{R} & [ijk] \mapsto \alpha_{ijk}^2 \\
 \alpha^3: T \rightarrow \mathbb{R} & [ijkl] \mapsto \alpha_{ijkl}^3,
 \end{array}$$

and correspondingly for higher dimensional simplicial complexes.

To get a better sense as to what this all means imagine an implementation. Each k -form is a vector of coefficients, each coefficient corresponding to the value associated with a particular k -simplex according to some arbitrary but fixed numbering scheme. This setup turns the discrete d operator into a sequence of sparse matrices which are transposes of incidence matrices. For example, the boundary operator for edges, ∂_1 , has dimension $|V| \times |E|$ with a single -1 and $+1$ in each column denoting the beginning resp. end of a given edge. The differential of discrete 0-forms is then $d_0 = \partial_1^T$. Applying d_0 to a function now amounts to taking the difference of values of the function along each oriented edge. This difference is not an approximation of the differential, but is exact by virtue of us using *integrated forms* and the fundamental theorem of

calculus (Stokes' theorem in 1D). Similarly the boundary operator ∂_2 for facets (triangles) is of dimension $|E| \times |F|$ with each column containing three non-zero entries of ± 1 according to the intrinsic orientation of each edge relative to the orientation of the containing triangle

$$(\partial_2)_{[jk],[ijk]} = +1 \quad (\partial_2)_{[ik],[ijk]} = -1 \quad (\partial_2)_{[ij],[ijk]} = +1,$$

leading to $d_1 = \partial_2^\top$. Here we used simplices to index the rows and columns of the associated matrix. This pattern of construction continues as we go up in dimension. The important part is that the d matrices are sparse and their entries easily derived from incidence relations between the elements of the mesh.




Figure 4.5: A mesh consisting of two tetrahedra together with the enumeration of all its simplices.

As an example consider the tetrahedral mesh in Fig. 4.5. There we used the convention of sorted index order as the canonical representative of a given simplex. Since discrete k -forms are just maps from a simplex to a number we can store them as vectors with the same order as the vectors of simplices. Using this ordering we get the incidence matrices, and hence their transposes, the discrete d operators, as

$$\partial_3 = \begin{pmatrix} -1 & 0 \\ +1 & 0 \\ -1 & 0 \\ +1 & -1 \\ 0 & +1 \\ 0 & -1 \\ 0 & +1 \end{pmatrix} \quad \partial_2 = \begin{pmatrix} +1 & +1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & +1 & 0 & 0 & 0 & 0 \\ 0 & -1 & -1 & 0 & 0 & 0 & 0 \\ +1 & 0 & 0 & +1 & +1 & 0 & 0 \\ 0 & +1 & 0 & -1 & 0 & +1 & 0 \\ 0 & 0 & 0 & 0 & -1 & -1 & 0 \\ 0 & 0 & +1 & +1 & 0 & 0 & +1 \\ 0 & 0 & 0 & 0 & +1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & +1 & +1 \end{pmatrix}$$

$$\partial_1 = \begin{pmatrix} -1 & -1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ +1 & 0 & 0 & -1 & -1 & -1 & 0 & 0 & 0 \\ 0 & +1 & 0 & +1 & 0 & 0 & -1 & -1 & 0 \\ 0 & 0 & +1 & 0 & +1 & 0 & +1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & +1 & 0 & +1 & +1 \end{pmatrix}$$

4.2.2 Discrete Hodge Star

Another discrete operator we need is the Hodge star. Recall that the Hodge star maps a k -form to an $(m - k)$ -form. While the former is evaluated on a k -dimensional domain, the latter is evaluated on an $(m - k)$ -dimensional domain. In the discrete setting we begin by creating this pairing between k -dimensional and $(m - k)$ -dimensional pieces of domain. This is accomplished via the **dual mesh** (see Fig. 4.6). For example, in the case of a tetrahedral mesh the dual mesh associates to each tetrahedron a vertex, to each triangle an edge, to each edge a polygon, and to each vertex a dual volume cell (see Fig. 4.6, center). Note that the dual of a simplicial complex is generally no longer simplicial but instead a *cell complex*, however the incidences are preserved under this duality mapping. For example, two tetrahedra sharing a triangle map to two vertices connected by an edge, which is the edge dual to the shared triangle. More precisely, the incidence matrices (boundary operators) of the dual mesh, ∂^* , are given as $\partial_{m-k+1}^* = \partial_k^\top$ for $k = 1 \dots m$. Consequently, the dual d operators are given by $d_k^* = (\partial_{k+1}^*)^\top = \partial_{m-k}$. In other words, the incidence matrices of the primal mesh are all that's needed to define the discrete differential both on the primal simplicial complex and on the dual cell complex.

So far we have only defined the dual mesh combinatorially and not given it an embedding. The dual mesh is not always required to be embedded. If we need an embedding there are two common choices. One places the dual vertex at the barycenter of its containing highest dimensional simplex, the other places it at the center of the circumsphere of the containing simplex.

Having mapped the domain pieces we must now describe how the values of discrete forms transform under this mapping. There are a number of different possible approaches and we will use something called the **diagonal Hodge star**. Given an embedding of the dual mesh, the diagonal Hodge star ensures that the integral over the respective pieces of domain agree on average

$$\frac{1}{|\sigma|} \int_{\sigma} \alpha = \frac{s}{|\sigma^*|} \int_{\sigma^*} *\alpha,$$

where $|\cdot|$ denotes the intrinsic volume, σ an appropriate simplex, σ^* its dual cell, and $s = \pm 1$ indicates whether the orientation of (σ, σ^*) is positive ($s = +1$) or negative ($s = -1$). The orientation of (σ, σ^*) is the orientation of the basis which results when concatenating a basis for the linear space which contains σ and similarly for σ^* . The intrinsic volume of a vertex is 1, that of an edge its length, that of a facet its area, and that of a tetrahedron its volume (and so




Figure 4.6: The dual mesh associates to every primal k -simplex a dual $(m - k)$ -cell. For example, a primal (black) triangle mesh associates a dual (red) vertex to each triangle, a dual edge to each primal edge, and a dual polygon to every primal vertex (left). In the case of a tetrahedral mesh (middle) each primal tetrahedron is mapped to a dual vertex, each primal triangle to a dual edge, each primal edge to a dual polygon, and finally each primal vertex to a dual 3-cell. If the dual vertices have an embedding, then all other dual cells have corresponding embeddings. For example, the polygon dual to a primal edge arises as the cycle of dual vertices belonging to the primal tetrahedra incident on the primal edge (right).

on). Correspondingly discrete coefficients α_σ , representing the integral, are mapped to discrete coefficients $*\alpha_{*\sigma}$ as

$$*\alpha_{*\sigma} := \frac{s|*\sigma|}{|\sigma|} \alpha_\sigma.$$

This relationship corresponds to a diagonal matrix whose entries are the quotients of intrinsic volumes of k -dimensional simplices resp. their duals

$$(*_k)_{\sigma^*, \sigma} = \frac{s|\sigma^*|}{|\sigma|} \quad (*_k)^{-1}_{\sigma, \sigma^*} = \frac{|\sigma|}{s|\sigma^*|}$$

Note that σ and σ^* are in 1-to-1 correspondence, which implies that a discrete differential form and its Hodge dual can be stored in vectors of the same length.

Example 4.2. Consider $\mathbb{R}^3$ and a velocity 1-form η represented on a tetrahedral mesh. The discrete version of η may live on facets (triangles) as a variable indicating the flux across that facet. Or η may live on the edge dual to a facet and represent the circulation along the edge between the two defining tetrahedra. To first order the two quantities are related by dividing the flux by triangle area and multiplying the result by the length of the dual edge (or vice versa).

Given that we now have both a discrete d and a discrete $*$ we can also build the **discrete co-differential**. Recall its definition for a k -form in the smooth setting

$$\delta = (-1)^k *^{-1} d *.$$

In the discrete setting, for a concrete k this amounts to the product

$$\delta_k = (-1)^k *_{k-1}^{-1} \partial_k *_{k-1},$$

where ∂_k is the incidence matrix (of the primal graph) mapping each k -simplex to its (signed) boundary. In other words the d operator acting on the dual mesh is the transpose of the d operator acting on the primal mesh.

4.2.3 An Example: The Laplacian

To get a little practice with these discrete operators, let's consider the Laplacian for 0-forms on a triangle mesh. Suppose we are given an orientable 2-manifold simplicial complex $K = \{V, E, F\}$ (without boundary) together with an embedding $\{p_i \in \mathbb{R}^3\}_{i \in V}$. Let's place the dual vertices at the circumcenter of their primal triangle circumcircles. This is all we need to define the discrete differential on the primal and dual mesh as well as the diagonal Hodge star operators. In the smooth setting the Poisson equation is written as

$$\Delta u = g$$

for $u, g: M \rightarrow \mathbb{R}$. Expanding it in terms of (still smooth) exterior calculus we have

$$*d*du = g \iff d*du = *g.$$

Note that we moved the leftmost Hodge star in the definition of Δ to the right hand side so as to get a symmetric linear system, which admits much faster solution methods than a general matrix. For the discrete version, let $\{*g_i\}_{i \in V}$ be the vector of integrals of g over the 2-cells dual to each primal vertex, i.e., $*g_i = \int_{[i]^*} *g$ and $\{u_i\}_{i \in V}$ represent the discrete primal 0-form to be solved for. Including all dimension/order subscripts we have to solve the linear algebra problem

$$*_2 g = d_1^* *_1 d_0 u = \partial_1 \frac{s|*e|}{|e|} \partial_1^T u = Lu$$

where $\frac{s|*e|}{|e|}$ is the diagonal matrix of ratios of dual to primal edge lengths and ∂_1 the $|V| \times |E|$ incidence matrix of vertices on edges. Notice that the matrix L is symmetric by construction. We can visualize the individual steps for a given vertex v_i and its neighborhood as shown in Fig. 4.7.

Exercise 4.3. Consider an embedded surface triangle mesh and use circumcentric duals, i.e., place the dual vertex of a triangle at the center of its circumcircle. Show that in that case

$$\frac{s|*e_{ij}|}{|e_{ij}|} = \frac{1}{2}(\cot \alpha_{ij}^k + \cot \alpha_{ji}^l)$$





Figure 4.7: The Laplacian $d_0^* \circ d_0$ for 0-forms on a triangle mesh acting on data $\{u_i\}_{i \in V}$ proceeds by first forming $(du)_{ij}$ on each edge. The subsequent Hodge star moves this data to the corresponding dual edges. Applying another d we get $(d_0^* \circ d_0 u)_{i*}$ on each cell dual to a primal vertex. This is the (signed) sum of the values on the surrounding boundary dual edges. The final result is a 2-form which should match the rhs, $*_2 g$, in a Poisson problem.



for edge e_{ij} shared by facets f_{ijk} and f_{jil} . The angle α_{ij}^k is across from e_{ij} at vertex v_k and correspondingly for α_{ji}^l . Under what circumstances is the sum of the two cotans negative? (Recall the discussion of the orientation of (σ, σ^*) .)

4.3 Whitney Elements

So far we have built the surface (or volume) through piecewise linear interpolation. Now we want to talk about interpolating discrete k -forms on such a 2- or 3-manifold, and more generally on m -manifolds. Whitney elements allow us to do this, bootstrapping PL interpolation of the carrier to PL interpolation of forms on the manifold. For now we'll stick to the case of 2- and 3-manifolds before giving the general construction.

Given values h_i of some function at the vertices $i \in V$ of a triangle mesh, it is an easy matter to interpolate the values in a piecewise linear fashion in the same way we interpolate vertex positions over the triangle

$$h_{ijk}(v) = h_i \phi_i(v) + h_j \phi_j(v) + h_k \phi_k(v) = h_i \alpha + h_j \beta + h_k \gamma$$

for $v = \alpha v_i + \beta v_j + \gamma v_k$ with (α, β, γ) the barycentric coordinates of points in the carrier of f_{ijk} as before. The new functions $\{\phi_i\}_{i \in V}$ appearing here are associated with vertices and sometimes referred to as hat (or tent) functions or barycentric interpolation functions. They are piecewise linear, supported on all triangles incident on a given vertex, and, when paired with vertices (this is a bit formal, but bear with me), take on values

$$\langle v_j, \phi_i \rangle := \int_{v_j} \phi_i := \delta_{ij}$$

Notice that the definition of the ϕ_i extends to arbitrary dimension simplices and is thus general.

Given a discrete 0-form $\{h_i\}_{i \in V}$ we can now construct a PL continuous 0-form

$$h(v) = \sum_{i \in V} h_i \phi_i(v)$$

for any v in the carrier of our simplicial 2-, 3- or more general m -manifold. Incidentally the PL embedding of the simplicial complex follows in the same manner

$$p(v) = \sum_{i \in V} p_i \phi_i(v).$$

Now we can bootstrap this construction into producing PL interpolants for 1-forms. Given an edge e_{ij} define

$$\phi_{ij}(v)(X) = \phi_i(v) d\phi_j(X) - \phi_j(v) d\phi_i(X)$$

for v some point in the carrier of the simplicial manifold and $X \in T_v M$ a tangent vector in the tangent space at v . To avoid clutter we will generally leave the base point and tangent arguments out writing instead

$$\phi_{ij} = \phi_i d\phi_j - \phi_j d\phi_i.$$

Because $d\phi_i$ and $d\phi_j$ are piecewise constant and they are multiplied by the PL hat functions at v_j resp. v_i we see that ϕ_{ij} is supported on the simplices incident on the edge e_{ij} , i.e., the intersection of the supports of ϕ_i resp. $d\phi_i$ and ϕ_j resp. $d\phi_j$ (see Fig. 4.8). For a 2-manifold this would be the two (or one at the boundary) incident triangle facets, while for a tetrahedral mesh the support would encompass all the tetrahedra incident on e_{ij} . Obviously ϕ_{ij} is piecewise linear. As we will see below with a more general computation, the edge bases also satisfy the interpolation property

$$\langle e_{kl}, \phi_{ij} \rangle = \int_{e_{kl}} \phi_{ij} = \delta_{ki} \delta_{lj}$$

for all $e_{kl} \in E$.

We can visualize ϕ_{ij} by drawing the vector field $\phi_{ij}^\sharp$ (see Fig. 4.8). With the edge basis forms we can PL interpolate a discrete 1-form $\{\omega_{ij}\}_{e_{ij} \in E}$ as

$$\omega = \sum_{e_{ij} \in E} \omega_{ij} \phi_{ij}.$$




Figure 4.8: Support of ϕ_i and ϕ_j (left) and the corresponding $(d\phi)^\sharp$ (middle). Taken together with the definition $\phi_{ij} = \phi_i d\phi_j - \phi_j d\phi_i$ results in the piecewise linear, interpolating 1-form basis $\{\phi_{ij}\}_{e_{ij} \in E}$ visualized by its vector proxy $\phi_{ij}^\sharp$ (right).

For 2-forms the interpolant is

$$\phi_{ijk} = 2(\phi_i d\phi_j \wedge d\phi_k - \phi_j d\phi_i \wedge d\phi_k + \phi_k d\phi_i \wedge d\phi_j).$$

On a triangle mesh this is the highest order form and

$$\langle f_{lmn}, \phi_{ijk} \rangle = \int_{f_{lmn}} \phi_{ijk} = \delta_{il} \delta_{jm} \delta_{kn}.$$

For a tetrahedral mesh it is PL on the two (or one at the boundary) tetrahedra incident on f_{ijk} .

Fig. 4.9 (right) shows a visualization of $(*\phi_{ijk})^\sharp$ as a vector field.




Figure 4.9: The edge basis form ϕ_{ij} in a tetrahedral mesh, visualized on a single tetrahedron incident on the given edge (left). The facet 2-form basis ϕ_{ijk} visualized as the vector field $(*\phi_{ijk})^\sharp$ on a single tetrahedron incident on the selected facet (right).

Exercise 4.4. Earlier we said that $f_{ijk} = f_{jki} = f_{kij}$ and the other three permutation have the opposite orientation. Show that the above interpolants are well defined in this sense. And, in a little bit, convince yourself that this is true for the general version of Whitney bases as well.

4.3.1 General form of Whitney Bases

The pattern we have seen in the construction is an instance of a general construction of such interpolating bases. Let $\sigma = [i_0 \dots i_j]$ be a j -dimensional simplex in an m -dimensional simplicial complex. The corresponding interpolant is then given by

$$\phi_\sigma = j! \sum_{l=0}^j (-1)^l \phi_{i_l} d\phi_{i_0} \wedge \dots \wedge \widehat{d\phi_{i_l}} \wedge \dots \wedge d\phi_{i_j}$$

where $\widehat{d\phi_{i_l}}$ denotes that $d\phi_{i_l}$ is dropped from the wedge product.

First we show that

$$\int_\sigma \phi_\sigma = 1$$

for which we use $1 = \sum_{l=0}^j \phi_{i_l}$ and hence $0 = \sum_{l=0}^j d\phi_{i_l}$. Consider the first summand of ϕ_σ

$$(-1)^0 \phi_{i_0} d\phi_{i_1} \wedge \dots \wedge d\phi_{i_j}.$$

All other summands contain $d\phi_{i_0} = -\sum_{l=1}^j d\phi_{i_l}$. Replace $d\phi_{i_0}$ thus and notice that only the $d\phi_{i_l}$ which is otherwise dropped in that summand “survives” (why?)

$$(-1)^l \phi_{i_l} (-1) d\phi_{i_l} \wedge d\phi_{i_1} \wedge \dots \wedge \widehat{d\phi_{i_l}} \wedge \dots \wedge d\phi_{i_j}.$$

Now note that it takes $(l-1)$ interchanges to move $d\phi_{i_l}$ from the front to the “missing” slot. Taken together we then find

$$\int_\sigma \phi_\sigma = j! \int_\sigma \sum_{l=0}^j \phi_{i_l} d\phi_{i_1} \wedge \dots \wedge d\phi_{i_j} = j! \int_\sigma d\phi_{i_1} \wedge \dots \wedge d\phi_{i_j} = \frac{j!}{j!} = 1$$

since the integral just measures the volume of σ . To see this use a basis $\{X_{i_l}\}_{l=1,\dots,j}$ for the tangent space of σ which is dual to the $\{d\phi_{i_l}\}_{l=1,\dots,j}$.

To complete the argument for interpolation suppose $\tilde{\sigma} \neq \sigma$. Then there exists an $n \in \sigma$ which is not contained in $\tilde{\sigma}$. For the corresponding ϕ_n we have $\phi_n(v) = 0$ for $v \in \tilde{\sigma}$. Being thus constant on $\tilde{\sigma}$ we also have $d\phi_n = 0$ on $\tilde{\sigma}$. These two facts together conspire to ensure that each summand of ϕ_σ vanishes on $\tilde{\sigma}$ and thus

$$\int_{\tilde{\sigma}} \phi_\sigma = 0.$$

Exercise 4.5. Show that interpolation after taking the discrete exterior derivative gives the same result as taking the exterior derivative after interpolation. Your argument should only use the fact that Whitney bases are interpolating and not involve any calculations.

4.4 Sampling

The interpolation of discrete data into continuous forms is the counterpart to sampling continuous forms to produce discrete data. For a function, say $h: M \rightarrow \mathbb{R}$, it is pretty clear what this means. Sample its value at vertices to arrive at a discrete function $\{h_i\}_{i \in V}$. This is already fundamental to creating a triangle mesh (simplicial complex) from a smooth surface. Imagine drawing a triangle mesh on a smooth surface. Now build the corresponding simplicial complex $K = \{V, E, F\}$ and sample the smooth surface at all vertices. The PL reconstruction of that discrete mesh will yield the triangle mesh approximation of the original surface. Not surprisingly, if we want this PL approximation to be good we should sample more often (make smaller triangles) where the surface changes rapidly and can afford larger triangles where it is relatively flat to achieve a given threshold of approximation. So far so good. What about higher degree forms? For example we might have an $\mathbb{R}$ -valued 1-form $\omega \in \Omega^1(M; \mathbb{R})$ on the surface. Perhaps it corresponds to a tangent vector field $\omega^\sharp = X \in \Gamma(TM)$ (a section of the tangent bundle, since for every point $p \in M$ we have a vector $X_p \in T_p M$). We can then sample it on edges by computing discrete 1-form coefficients

$$\omega_{ij} = \int_{e_{ij}} \omega = \langle e_{ij}, \omega \rangle.$$

And more generally for a k -form ω^k we would compute the discrete version of it as

$$\omega_\sigma^k = \int_\sigma \omega^k = \langle \sigma, \omega^k \rangle$$

where σ is a k -simplex in our simplicial complex describing the underlying manifold.

Letting S denote the sampling operation and W the interpolation with Whitney forms we see that

$$S \circ W = I \quad \text{while} \quad W \circ S \approx I,$$

i.e., sampling after interpolation is the identity on the space of discrete k -forms while interpolation after sampling is at best close to the identity on the space of smooth k -forms. Trying to get the latter closer to the identity is the motivation for many attempts to design higher order k -form interpolants. One could also pursue questions of pre- and post-filtering when looking at this from a signal processing perspective. We will not do this here and simply work with PL approximations.