

Edge Subdivision Schemes and the Construction of Smooth Vector Fields

Ke Wang

Weiwei

Yiyong Tong

Mathieu Desbrun

Peter Schröder

Caltech

Abstract

Vertex- and face-based subdivision schemes are now routinely used in geometric modeling and computational science, and their primal/dual relationships are well studied. In this paper, we interpret these schemes as defining bases for *discrete differential 0- resp. 2-forms*, and complete the picture by introducing *edge-based* subdivision schemes to construct the missing bases for discrete differential 1-forms. Such subdivision schemes map scalar coefficients on edges from the coarse to the refined mesh and are *intrinsic* to the surface. Our construction is based on treating vertex-, edge-, and face-based subdivision schemes as a *joint triple* and enforcing that subdivision commutes with the topological exterior derivative. We demonstrate our construction for the case of arbitrary topology triangle meshes. Using Loop's scheme for 0-forms and generalized half-box splines for 2-forms results in a *unique* generalized spline scheme for 1-forms, easily incorporated into standard subdivision surface codes. We also provide corresponding boundary stencils. Once a metric is supplied, the scalar 1-form coefficients define a smooth tangent vector field on the underlying subdivision surface. Design of tangent vector fields is made particularly easy with this machinery as we demonstrate.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling— Curve, surface, solid, and object representations, Physically based modeling, Splines; J.6 [Computer Applications]: Computer-Aided Engineering— Computer-aided design (CAD)

Keywords: Subdivision, Discrete Exterior Calculus, Discrete Differential Geometry, vector fields, smooth surface modeling.

1 Introduction

Subdivision schemes are a broadly deployed tool in all areas of geometric modeling and computer graphics [Zorin and Schröder 2000; Warren and Weimer 2001]. Their foremost benefit is the ease with which they accommodate the construction of smooth surfaces in the arbitrary topology setting. They also offer many favorable computational properties for applications ranging from surface compression [Khodakovskiy et al. 2000] to physical modeling [Grinspun et al. 2002]. Their mathematical properties are by now well understood [Reif 1995; Reif 1999; Zorin 2000a; Zorin 2000b] and a large variety of subdivision schemes and extensions have been developed. Broadly, subdivision schemes are classified as either primal (e.g., Catmull-Clark [1978], Loop [1987], and $\sqrt{3}$ [Kobbelt 2000]) with vertices carrying the data and faces being split, or dual (e.g., Doo-Sabin [1978] and dual- $\sqrt{3}$ [Oswald and Schröder 2003]) in which data lives at faces and vertices are split.

In this paper we present a novel class of subdivision schemes which carry scalar coefficients on *edges* from the coarser mesh into scalar coefficients on edges in the refined mesh. The method can be viewed as constructing higher regularity bases for *discrete*

differential 1-forms [Desbrun et al. 2005] in the arbitrary topology 2-manifold (with boundaries) setting. Given the metric induced by an underlying surface, 1-forms then yield smooth tangent vector fields, which are useful in many applications including texture synthesis [Turk 2001], fluid simulation [Shi and Yu 2005], crowd animation [Reynolds 1999], and shading [Schlick 1994]. In particular the *design* of vector fields is greatly facilitated by the intuitive relationship between coefficients and the resulting vector field (see Figure 1).




Figure 1: Example of vector field design for use in a fur shader. The user marks selected vertices as sources (red) and sinks (blue) on the control mesh. Loop subdivision together with our novel 1-form subdivision results in a tangent vector field which interpolates the given sparse constraints over the subdivision surface (visualized on the bottom right). Such fields can be used directly to control standard fur shaders (here using Autodesk Maya).

Context While the relationship between edge-based subdivision and smooth tangent vector fields on surfaces is of interest in and of itself, our approach forms a piece of a larger picture. We treat the construction of edge-based subdivision schemes in the context of a *joint triple* of vertex-, edge-, and face-based subdivision schemes, which are linked through Stokes' theorem. This has important practical consequences: a 0-form (function), given through coefficients at vertices, has a differential, a 1-form, which can be represented *exactly* in the basis of 1-forms generated by the edge-based subdivision scheme. And similarly, the differential of a 1-form can be represented *exactly* in the basis of 2-forms generated by the face-based subdivision scheme. Additionally, these coefficients are related through local operations due to the adjoint relationship that Stokes' theorem establishes between the *exterior derivative* d and the *boundary operator* ∂ . Consequently, at the level of the mesh, the exterior derivative is *computationally realized* by taking signed sums of coefficients according to the incidence relations: computing the differential of a 0-form amounts to visiting each edge and assigning to it the difference of coefficients at its end points; similarly, the differen-

tial of a 1-form is found by visiting each triangle and assigning it the signed sum (according to orientation) of its three edge values. In this way, rather cheap and simple computations at the discrete level of the mesh correspond *exactly* to operations from classical vector calculus, e.g., taking the gradient, curl, or divergence. Since Stokes' theorem holds by definition, one also gets standard vector calculus identities, e.g., $\nabla \times \nabla = 0$, for free and *exactly*. In many applications these identities are at the core of conserved quantities and structural relationships between field variables. Ensuring that they hold at the discrete computational level avoids, with no cost overhead, spurious numerical problems in a great variety of numerical simulation scenarios ranging from standard elliptic problems to applications in fluids, electro magnetism, and elasticity [Elcott et al. 2005; Bossavit 1998; Arnold et al. 2006].

Approach and Contributions We consider the joint construction of vertex-, edge-, and face-based subdivision schemes so as to create smooth bases for 0-, 1-, and 2-forms. The linear spans of these bases, together with the exterior derivative, form a *chain complex*. Assuming convergent subdivision schemes, the chain complex property is achieved by enforcing *commutative relations* at the discrete level: given the *discrete* exterior derivative operators d^0 and d^1 (also called "co-boundary operators"), the subdivision operators S_0 (vertex), S_1 (edge), and S_2 (face) must satisfy

$$d^0 S_0 = S_1 d^0 \quad \text{and} \quad d^1 S_1 = S_2 d^1. \quad (1)$$

In words: subdividing first, followed by application of the discrete exterior derivative, should be the same as first applying the discrete exterior derivative followed by an *associated* subdivision scheme. Relations of this type are known as *formule de commutation* [Dyn et al. 1987] from the univariate setting. Our approach generalizes these relations to the bivariate, non tensor product setting.

Applying this line of reasoning to piecewise linear (PL) subdivision recovers the well known Whitney forms [1957]. Asking for smoother bases over arbitrary triangulations leads to considering Loop subdivision for S_0 . For S_2 one may then choose (a generalization of) half-box splines [Prautzsch et al. 2002]. With S_0 and S_2 fixed in this manner, the 1-form scheme S_1 that we introduce in this paper follows *uniquely* using (1). It too is (a generalization of) a piecewise polynomial spline scheme. More generally, one may begin with a desired support (stencil size) and symmetries for the 0-form (vertex-based) subdivision scheme and then derive fully parameterized families of subdivision scheme triples from (1). While we demonstrate this only in the case of Loop (and triangles), the approach applies equally well to other settings, e.g., quadrilaterals with Catmull-Clark for S_0 and Doo-Sabin for S_2 [Wang 2006].

2 Refinable Forms

After briefly fixing notation and introducing the relevant concepts from Discrete Exterior Calculus (DEC), we discuss the PL case (Section 2.1), i.e., hat functions, and their associated *Whitney* forms. (Additional details on DEC, including many of the facts we will use in this section, can be found in [Desbrun et al. 2005].) Observing that Whitney forms are refinable (Section 2.2) we construct smoother bases for DEC using standard subdivision machinery in the regular setting (Section 2.3).

Setup and Elements of DEC Consider an arbitrary topology orientable 2-manifold (possibly with boundary) simplicial complex $K = (V, E, T)$ with triangle set $T = \{t_{ijk}\}$, edge set $E = \{e_{ij} | e_{ij} \in t_{ijk} \in T\}$ and vertex set $V = \{v_i\}$, where $i, j, k = 1, \dots, n (= |V|)$ range over the vertices. Each triangle and edge is given an (arbitrary, but fixed) intrinsic orientation. (Vertices have positive

intrinsic orientation by definition.) If a particular edge e_{ij} is positively oriented then e_{ji} refers to the same edge but with negative orientation. Similarly, with t_{ijk} positively oriented any even (odd) permutation of (ijk) denotes the same triangle with positive (negative) orientation. Note that in practice, these orientations are often induced by the immersion of the mesh into \mathbb{R}^3 , but this is not a requirement. Assuming an arbitrary but fixed indexing for the triangles ($t = 1, \dots, |T|$), edges ($e = 1, \dots, |E|$) and vertices ($v = 1, \dots, |V|$) the *incidence matrices* of K now define the *boundary operators*. Let ∂^2 denote the matrix of $|E|$ rows and $|T|$ columns with $(\partial^2)_{et} = \pm 1$ if edge e is incident on triangle t and their intrinsic orientations agree/disagree (and zero otherwise). Similarly ∂^1 denotes the incidence matrix of vertices ($|V|$ rows) on edges ($|E|$ columns). In DEC the *discrete exterior derivative* mapping discrete 0-forms to discrete 1-forms is given as $d^0 = (\partial^1)^T$, while $d^1 = (\partial^2)^T$ maps discrete 1-forms to discrete 2-forms. For example, a (column) vector of coefficients associated with vertices, c_v , is mapped as $c_e = d^0 c_v$ into a (column) vector of coefficients associated with edges and similarly for a vector of edge coefficients c_e mapping to coefficients on triangles, $c_t = d^1 c_e$.

One may now ask: what are suitable interpolants, or bases, for the data on vertices, edges, and triangles such that the discrete exterior derivative applied to coefficients corresponds to the exterior derivative applied to the underlying interpolated forms? Such interpolants define a *chain map* and Whitney forms provide a PL answer to this question.

2.1 Whitney Forms

Let $\phi^v = \{\phi_i | v_i \in V\}$ be barycentric interpolation functions (taken as a row vector), one for each vertex, and supported on the triangles incident to a given vertex. In the abstract setting of simplices given as convex hulls of their incident vertices the barycentric interpolation functions are the standard PL hat functions. (Note that this observation requires neither an embedding nor a metric.) Given a vector of coefficients c_v , the inner product $f = \phi^v c_v$ defines an interpolated function f of the vertex data over the entire complex K . $f(x)$ may be evaluated at any point x by furnishing the barycentric coordinates of x . With ϕ^v given, Whitney defined the appropriate 1-form bases as

$$\phi_{ij} = \phi_i \mathbf{d} \phi_j - \phi_j \mathbf{d} \phi_i,$$

one for each edge e_{ij} and $\phi^e = \{\phi_{ij} | e_{ij} \in E\}$. (We use bold \mathbf{d} to denote the exterior derivative, while d^0 resp. d^1 are reserved for the *discrete* version acting on vectors of coefficients.) Note that orientation matters since $\phi_{ji} = -\phi_{ij}$. With these definitions one has

$$\mathbf{d} f = \mathbf{d}(\phi^v c_v) = \phi^e d^0 c_v. \quad (2)$$

To see this, consider a single triangle t_{ijk} and the interpolant formed over it by the data at the incident vertices

$$\begin{aligned} \mathbf{d}(c_i \phi_i + c_j \phi_j + c_k \phi_k) &= (c_j - c_i)(\phi_i \mathbf{d} \phi_j - \phi_j \mathbf{d} \phi_i) + \\ &\quad (c_k - c_j)(\phi_j \mathbf{d} \phi_k - \phi_k \mathbf{d} \phi_j) + \\ &\quad (c_i - c_k)(\phi_k \mathbf{d} \phi_i - \phi_i \mathbf{d} \phi_k) \end{aligned}$$

(where we used $\phi_i + \phi_j + \phi_k = 1$ and $\mathbf{d}(\phi_i + \phi_j + \phi_k) = 0$). Applying this reasoning to all triangles simultaneously we arrive at (2). Obviously $\mathbf{d} f(x)$ may be evaluated anywhere on the complex given the barycentric coordinates of x .

The Whitney 2-forms are given as

$$\begin{aligned} \phi_{ijk} &= 2(\phi_i \mathbf{d} \phi_j \wedge \mathbf{d} \phi_k + \phi_j \mathbf{d} \phi_k \wedge \mathbf{d} \phi_i + \phi_k \mathbf{d} \phi_i \wedge \mathbf{d} \phi_j) \\ &= 2 \mathbf{d} \phi_j \wedge \mathbf{d} \phi_k = 2 \mathbf{d} \phi_k \wedge \mathbf{d} \phi_i = 2 \mathbf{d} \phi_i \wedge \mathbf{d} \phi_j \end{aligned}$$

(again, using $\phi_i + \phi_j + \phi_k = 1$ and $\mathbf{d}(\phi_i + \phi_j + \phi_k) = 0$) and hence are constant in the basis of 2-forms, *no matter the metric*. They verify for an arbitrary 1-form $g = \phi^e c_e$

$$\mathbf{d}g = \mathbf{d}(\phi^e c_e) = \phi^t \mathbf{d}^1 c_e, \quad (3)$$

which can be seen by considering a single triangle

$$\begin{aligned} & \mathbf{d}(c_{ij}\phi_{ij} + c_{jk}\phi_{jk} + c_{ki}\phi_{ki}) \\ &= 2(c_{ij}\mathbf{d}\phi_i \wedge \mathbf{d}\phi_j + c_{jk}\mathbf{d}\phi_j \wedge \mathbf{d}\phi_k + c_{ki}\mathbf{d}\phi_k \wedge \mathbf{d}\phi_i) \\ &= (c_{ij} + c_{jk} + c_{ki})\phi_{ijk} \end{aligned}$$

(using $\mathbf{d}\phi_{ij} = 2\mathbf{d}\phi_i \wedge \mathbf{d}\phi_j$ and $\mathbf{d}(\phi_i + \phi_j + \phi_k) = 0$). Equation (3) then follows by summing over all triangles.

With d^0 and d^1 defined as transposes of the boundary operators (incidence matrices) and taking account of the above calculations for Whitney forms, Stokes' theorem follows by construction

$$\int_{\sigma^{k+1}} \mathbf{d}f =: \langle \sigma^{k+1}, \mathbf{d}f \rangle = \langle \partial \sigma^{k+1}, f \rangle := \int_{\partial \sigma^{k+1}} f,$$

i.e., the integral of the exterior derivative of a k -form integrated over a $(k+1)$ -simplex equals the integral of the k -form itself over the boundary of the $(k+1)$ -simplex ($k=0,1$). By linearity of integration this definition extends to any linear combination of k -simplices in K . Here we also see why the transposes of the incidence matrices, the *co-boundary operators*, play the role of the discrete exterior derivative: it is a natural requirement if we want a discrete Stokes' theorem to hold at the simplicial level. Conversely, Stokes' theorem, together with the definition of the boundary operator, becomes the definition of the discrete d . Computationally all manipulations of forms are now reduced to manipulating coefficient vectors (c_v, c_e or c_t) with sparse matrices (d^0, d^1). Only if an interpolation is required do the bases come into play.

Discussion Whitney 0-, 1-, and 2-forms are interpolating by construction. This is certainly true for 0-forms, being barycentric interpolators. In the case of 1-forms we have

$$\int_{e_{ij}} \phi_{ij} = \int_{e_{ij}} \mathbf{d}\phi_j = \int_{\partial e_{ij}} \phi_j = 1 - 0 = 1.$$

Here we used $\phi_i + \phi_j = 1$ on e_{ij} to write ϕ_{ij} as $(1 - \phi_j)\mathbf{d}\phi_j - \phi_j\mathbf{d}(1 - \phi_j)$. For all other edges the integral vanishes. This follows from the simple observation that for any other edge e at least one of i or j (say i) is not incident to e and hence ϕ_i as well as $\mathbf{d}\phi_i$ are zero on e . For the 2-forms we get similarly

$$\int_{t_{ijk}} \phi_{ijk} = \int_{t_{ijk}} 2\mathbf{d}\phi_i \wedge \mathbf{d}\phi_j = \int_{t_{ijk}} \mathbf{d}\phi_{ij} = \int_{\partial t_{ijk}} \phi_{ij} = 1.$$

Here we used $\phi_{ijk} = 2\mathbf{d}\phi_i \wedge \mathbf{d}\phi_j$, which *only* holds on t_{ijk} . Any other triangle is missing at least one of i, j , or k (say i). With both ϕ_i and $\mathbf{d}\phi_i$ vanishing on such a triangle, each summand in $\phi_{ijk} = 2(\phi_i\mathbf{d}\phi_j \wedge \mathbf{d}\phi_k + \phi_j\mathbf{d}\phi_k \wedge \mathbf{d}\phi_i + \phi_k\mathbf{d}\phi_i \wedge \mathbf{d}\phi_j)$ is identically zero, implying the claim that Whitney 2-forms are interpolating as well. (Of course, orientation matters with the value flipping sign under orientation change.) This also suggests that the coefficients c_v, c_e and c_t should be seen as *integrated* quantities. In many practical settings this is preferred in any case, since some quantities may not be well defined pointwise while their integrals are well defined (*e.g.*, mean and Gaussian curvatures on triangle meshes). Note that this will have consequences for the normalization of the subdivision schemes: 1 for 0-forms, 1/2 for 1-forms, and 1/4 for 2-forms (in contrast to standard normalizations which always enforce that stencil coefficients sum to 1).

To visualize Whitney 0-, 1-, and 2-forms we must first introduce a metric. For convenience we will assume the Euclidean metric in all our visualizations. In the case of Whitney 0-forms this gives us PL hat functions over the embedded triangles, while 2-forms are constant functions ($= \text{Area}_{ijk}^{-1}$) supported on single triangles. To visualize 1-forms as tangent vector fields we must furthermore fix a coordinate frame and will use the standard Euclidean orthonormal frame. Recall that a differential 1-form maps tangent vectors to reals, $\mathbf{d}\varphi : T_p M \rightarrow \mathbb{R}$. With a metric and coordinate frame fixed, a 1-form may be represented as a co-vector field and evaluation of the 1-form amounts to taking the inner product between this *1-form proxy* and the argument tangent vector. Further identifying vectors and co-vectors we may visualize the 1-forms as vector fields. Figure 2 (middle) shows such a visualization of the *vector proxy* $\phi_{ij} = \phi_i \nabla \phi_j - \phi_j \nabla \phi_i$ (having picked a metric, \mathbf{d} applied to a 0-form becomes the gradient operator ∇). This vector field is supported on the two triangles incident to the given edge (or one triangle if it is a boundary edge) where it varies linearly. At the common edge the vector field is not continuous. However, it is *tangentially continuous*, *i.e.*, if we consider its inner product with a vector parallel to the edge it is continuous across the shared edge.




Figure 2: Visualization of Whitney 0-, 1-, and 2-forms. The 0-forms correspond to the usual PL hat functions, while 2-forms are piecewise constant over each triangle. The 1-forms can be visualized as vector fields if we choose a metric.

So far the topology of K has played absolutely no role. Neither irregular vertices nor boundaries change anything. On the other hand piecewise linear interpolation is obviously insufficient for many applications. The goal now is to build form bases which are smoother and for this method to work elegantly in the arbitrary topology setting. Simply adding local polynomials as is customary in many finite element constructions leads to a number of compatibility conditions [Hiptmair 2001], not unlike what one finds when attempting to join individual parametric surface patches smoothly [Prautzsch et al. 2002]. Instead we go the subdivision route. The resulting subdivision bases will have larger support and (in general) will not be interpolating anymore, though they will retain the property that the action of \mathbf{d} on a form is still realized by the standard topological d^0 resp. d^1 matrices acting on the coefficient vector of the form and Stokes' theorem will continue to hold *exactly*.

2.2 Whitney Forms are Refinable

Because Whitney 0-, 1-, and 2-forms are interpolating one readily checks that they are refinable over barycentric quadrisection of triangles. For example, the barycentric interpolator ϕ_i takes on the value 1/2 at the barycentric “midpoint” of edge e_{ij} by definition. Similarly the Whitney 2-form integrated over a barycentric “quarter” triangle results in a coefficient of 1/4. (Note that these statements do not depend on the metric or an embedding.) See Figure 3 for the corresponding 0- and 2-form refinement stencils on the left resp. right. Note that orientation of e_{ij} is indicated by an arrow pointing from i to j and orientations of triangles are CCW throughout this paper.

For Whitney 1-forms (using the notation of Figure 3) we get

$$\phi_{ij} = \frac{1}{4}(2\phi_{ip} + 2\phi_{pj} + \phi_{mn} - \phi_{np} - \phi_{pm} + \phi_{rq} - \phi_{qp} - \phi_{pr}),$$




Figure 3: Left side: even (top) and odd (bottom) refinement stencils for Whitney 0-, 1-, and 2-forms; right side: notation for 1-form refinement relation.

which gives rise to the edge based refinement stencils shown in Figure 3. These coefficients follow from elementary computations, e.g., $\int_{e_{mn}} \phi_{ij} = \frac{1}{2} \int_{e_{mn}} \mathbf{d}\phi_j = \frac{1}{4}$ using the fact that $\phi_i + \phi_j = \frac{1}{2}$ on e_{mn} . Letting now S_0^W , S_1^W , and S_2^W denote the subdivision operators for Whitney 0-, 1-, and 2-forms, the refinement equations for basis forms at the coarser (Φ) and finer (ϕ) levels follow as

$$\Phi^v = \phi^v S_0^W, \quad \Phi^e = \phi^e S_1^W, \quad \text{and} \quad \Phi^t = \phi^t S_2^W. \quad (4)$$

Using (4) together with (2) and (3) we can now prove that the associated subdivision operators satisfy the commutative relations (1)

$$\phi^e d^0 S_0^W c_v = \mathbf{d}(\phi^v S_0^W c_v) = \mathbf{d}(\Phi^v c_v) = \Phi^e d^0 c_v = \phi^e S_1^W d^0 c_v$$

for S_0^W and S_1^W (and entirely analogously for S_1^W and S_2^W).

Discussion While the above calculations used Whitney forms, the underlying argument is generic. Suppose the spaces spanned by a set of bases for 0-, 1-, and 2-forms together with \mathbf{d} form a chain complex and the bases define a chain map with respect to the co-boundary operator. If these bases satisfy a refinement relation, then the refinement relation commutes with the co-boundary operator. Conversely, given refinement relations which commute with the co-boundary operator and assuming the associated subdivision schemes are convergent, then the resulting bases form a chain complex with respect to exterior differentiation and the subdivision bases yield a chain map (for \mathbf{d} and the co-boundary operator). For this reason we make the commutative relations the core of our approach.

Practically speaking, one now seeks smoother bases by looking for subdivision schemes which produce smoother limit forms while still satisfying the commutative relations. Even though the resulting schemes may have larger support and may not be interpolating anymore (as is the case in our proposed construction), the action of \mathbf{d} is still computationally realized by the simple co-boundary operators we have used all along.

2.3 Increasing the Smoothness of Whitney Forms

A well known fact from subdivision [Warren and Weimer 2001] states that, in the regular setting, a given subdivision scheme can be transformed into a subdivision scheme of higher regularity through convolution. Since we are working with triangles, convolution along the three principal directions is appropriate and we denote its discrete representation with C . With such an additional convolution Whitney 0-forms yield quartic box splines (with subdivision operator $S_0^L = CS_0^W$) while Whitney 2-forms give rise to half box splines ($S_2^H = CS_2^W$). Convolution of the Whitney 1-forms gives a new, smooth 1-form basis with subdivision scheme $S_1^E := CS_1^W$. The support of the smooth 1-form basis is contained in the 1-ring of triangles around the two triangles incident to the given edge. Over this support the 1-form is piecewise polynomial (quartic) in terms of barycentric coordinates [Wang 2006]. The subdivision stencils for S_0^L , S_1^E , and S_2^H (in the regular setting) are summarized in Figure 4.

To check that the new subdivision schemes still satisfy the commutative relations (1) we consider the subdivision process in




Figure 4: Even (top) and odd (bottom) stencils (regular setting) for smooth 0-, 1-, and 2-forms for 3-direction convolved Whitney forms.

the Fourier domain over the \mathbb{Z}^2 lattice with diagonals. For $\mathbf{z} = (z_1, z_2)$ let $S(\mathbf{z})$ denote the symbol of the subdivision scheme S as usual. The symbols of the co-boundary operators are given by

$$d^0(\mathbf{z}) = \begin{pmatrix} -1+z_1^{-1} \\ -z_1^{-1}+z_1^{-1}z_2^{-1} \\ 1-z_1^{-1}z_2^{-1} \end{pmatrix} \quad \text{and} \quad d^1(\mathbf{z}) = \begin{pmatrix} 1 & -1 & 1 \\ -1 & -z_1z_2 & -z_2 \end{pmatrix}.$$

The symbol of three direction convolution is $C(\mathbf{z}) = 1/8(1+z_1)(1+z_2)(1+z_1z_2)$. While $S_0(\mathbf{z})$ is a scalar polynomial, $S_1(\mathbf{z})$ is a 3x3 matrix valued polynomial, grouping the coefficients associated with the three edge types into a 3-vector, and $S_2(\mathbf{z})$ is a 2x2 matrix valued polynomial, with coefficients of “up”/“down” triangles gathered into a 2-vector. Importantly, $C(\mathbf{z})$ is scalar and thus “pulls through” the relevant products in the Fourier domain

$$\begin{aligned} d^0(\mathbf{z})S_0^L(\mathbf{z}) &= d^0(\mathbf{z})C(\mathbf{z})S_0^W(\mathbf{z}) = C(\mathbf{z})d^0(\mathbf{z})S_0^W(\mathbf{z}) \\ &= C(\mathbf{z})S_1^W(\mathbf{z})d^0(\mathbf{z}) = S_1^E(\mathbf{z})d^0(\mathbf{z}). \end{aligned}$$

An entirely analogous calculation shows that the commutative relation between S_1^E and S_2^H holds as well (with respect to d^1).

Unfortunately, the convolution argument does not carry over into the arbitrary topology setting where irregular vertices and boundaries need to be dealt with. To address this issue we follow the usual path of using the standard stencils in the locally regular setting and modify the weights in the irregular setting only. For the latter we will use the commutative relations directly.

3 Arbitrary Topology Setting

For 0-forms we use Loop subdivision near irregular vertices (as well as at the boundary) following [Biermann et al. 2000] (see Figure 5 [left] for the modified even stencil). For half-box splines we follow [Oswald and Schröder 2003] (see Figure 5 [right] for the modified odd stencils) near irregular vertices. (We are not aware of any boundary rules.) As usual the odd 0-form stencil and even 2-form stencil remain unchanged.

Making now the standard assumption that the 1-form odd stencil (Figure 4, bottom middle) remains unchanged one may set up a fully parameterized—in terms of α and β —even stencil for the 1-form near irregular vertices using the commutative relations (1). The resulting stencil is shown in Figure 6. The solution is unique as a function of α and β and both α and β can be varied independently.

Figure 7 visualizes a few basis 1-forms in the regular and irregular setting. To demonstrate their smoothness we show the graph of the component functions. Note that the 1-forms in the irregular setting are only smooth with respect to the parameterization induced by the 0-form characteristic map (as is to be expected).




Figure 5: Even stencil for Loop (left) and odd stencils for generalized half-box splines (right) near irregular vertices.




Figure 6: Irregular even 1-form stencil.

Discussion A natural question that arises in this context is why are the directional derivatives of the quartic box spline not a good basis for the 1-forms? Letting M_{222} be the symbol for the quartic box spline (indicating 2-fold averaging along the three principal directions of the mesh), the box splines $M_{122}/M_{212}/M_{221}$ would appear to be a natural choice of basis for the 1-forms. Unfortunately this choice leads to far lower smoothness of the 2-forms. While vector field design, strictly speaking, only requires “nice” 1-form bases, many applications simultaneously require nice 0- and 2-form bases. Examples include physical simulation and more generally any approach relying on a proper Hodge decomposition. While we do not pursue this further here, it does provide the ultimate background behind our particular choice.

That our proposed smooth 1-form bases are C^1 in the regular setting follows trivially from the fact that 3-direction convolution increases smoothness by two orders (their *tangential* continuity is even C^2). What about the smoothness at irregular vertices? Consider the eigen vector/value pairs of the three subdivision operators in light of (1).¹ For example, the leading eigen vectors of S_1 are the image under d^0 of the subdominant eigen vectors of S_0 with the same eigen value. Now consider the need to span all linear 1-forms (a necessary condition to be C^1): we need four independent eigen vectors. Three of these arise as the image (under d^0) of the “quadratic” eigen vectors of S_0 . The fourth one is the pre-image of the constant eigen vector of S_2 (this also explains why α and β are *independent* parameters in our 1-form construction). But at the irregular vertex we know that Loop’s scheme lacks a full span of “quadratic” eigen vectors making it unlikely that we can achieve C^1 smoothness at the irregular vertices so long as we fix S_0^L . The scheme as described here is only

¹The following argument is only a rough outline as we entirely neglect issues such as generalized eigen vectors and other legal but complicated degeneracies in the spectra [Zorin 2000b; Reif 1999].

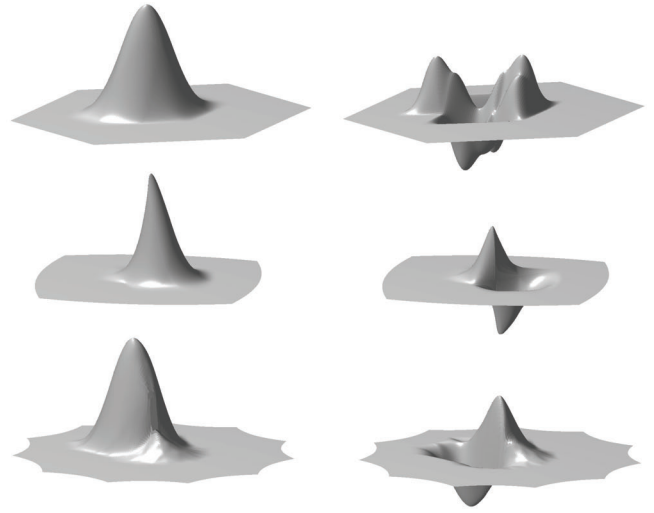


Figure 7: Visualization (x resp. y component of vector proxy) of 1-form bases in the regular ($k = 6$, top) and irregular ($k = 4$, middle; $k = 11$, bottom) setting. A single edge coefficient incident to the (ir-)regular vertex is set to 1. The edge is aligned with the x axis.

C^0 at irregular vertices of valence $k > 3$ (for $k = 3$ an additional modification is required [Wang 2006]).

3.1 Boundary Rules

The design of suitable boundary rules is more subtle since the boundary rules of Loop subdivision are fixed so that the boundary curve depends only on the boundary data. Additionally the cross boundary second derivative vanishes. Consequently we expect a reduced approximation power of the 1-forms near the boundary. To use (1) we also need to fix suitable boundary rules for the 2-form scheme.

In what follows, k is used to denote the number of triangles adjacent to vertex v on the boundary. We call a boundary vertex v regular when $k = 3$, irregular otherwise. Boundary vertices are marked in black while interior vertices are marked in white. Vertices without markers may or may not be on the boundary.




Figure 8: Odd stencil for interior edges adjacent to irregular boundary vertices (left) and the even/odd boundary rules for 0-forms (middle). On the right the 1-form boundary stencil.

3.2 0- and 2-Form Boundary Rules

For the boundary 0-form rules we follow Biermann *et al.* [2000]. In particular we assume the γ modified odd rules for interior edges incident on boundary irregular vertices (see Figure 8, left). To design boundary rules for 2-forms, we wish to preserve the maximal size of the odd stencil (see Figure 5, top middle). This simple constraint results in a 7-parameter family of 2-form stencils (with $A, B, C, D, F_0, F_1, F_2$ as parameters) near the boundary as shown in Figure 9 (here we consider only $k \geq 3$).

3.3 1-Forms Near the Boundary

The even 1-form boundary rules are simple as they must commute with the 0-form boundary rules which reproduce cubic splines along the boundary. Consequently even 1-form stencils along the boundary are the rules for quadratic splines up to a factor of $1/2$ (see Figure 8, right). With these fixed, and a fully parameterized set of 1-form stencils near the boundary, we can




Figure 12: Gallery of vector fields. Top: tet with 3 edges incident to a vertex set to +1 (left) resp. 3 edges incident to a face set to +1 (right); torus with 2 vortices and an open surface with 1 vortex.

computation of the circulation of the limit vector field on the limit edge of the surface. These can be computed analytically using the Bézier representation for a regular edge [Wang 2006] and extended to the irregular setting through a geometric series argument [Halstead et al. 1993]. Just as 0-form limit stencils are useful for approximation of a given surface with a subdivision surface (see for example [Litke et al. 2001]) the limit circulation stencils are useful for the approximation of given vector fields over a surface.

4.2 Design of Vector Fields

In surface modeling one manipulates control points to change the shape of the surface. To model a desired tangent vector field (over the underlying surface), one now manipulates *control coefficients on edges*. These coefficients have rather intuitive meaning, making design of a desired field simple. Placing a single coefficient of 1 on a given oriented edge produces a smooth vector field whose overall direction is roughly aligned with this edge and which decays smoothly to zero over the 1-ring of the two triangles incident to the edge. The magnitude of the vectors is directly controlled by the magnitude of the coefficient. Similarly, typical tasks in vector field design such as the placement of sinks, sources, and vortices becomes extremely easy. Sinks (and sources) at a vertex are where the divergence of a 1-form has negative (sink) or positive (source) value at the selected vertex. Vortices are similarly characterized by a positive (CW vortex) or negative (CCW vortex) curl (valued at a face) of a 1-form (see Figure 12).

However, setting only a few non-zero coefficients creates tangent vector fields that vanish over large parts of the surface since the basis 1-forms have finite support. In many scenarios one would like to place only a few select sources/sinks and/or CCW/CW vortices and then directly find a global vector field which satisfies these sparse features. This problem can be seen as the converse of the traditional Hodge decomposition of vector fields [Tong et al. 2003; Desbrun et al. 2005].

Hodge Composition A simple design tool that we have employed proceeds as follows. Given the coarse control mesh of a surface (assumed genus-0 for now), the user places sources and sinks at some chosen vertices (and implicitly zeros at all other vertices). A 0-form is then solved for by using this data as the rhs of Laplace’s equation over the vertices. CCW and CW vortices are placed at selected faces (and implied zeros at




Figure 13: 1-form basis forms on an open surface near the boundary.




Figure 14: Vector fields resulting from sparse interpolation. Top left: placing ± 1 at two opposing vertices results in a global (curl free) vector field with a single source and sink. Placing ± 1 at selected faces results in a global (divergence free) vector field with two opposing CCW/CW vortices. To produce a non-trivial harmonic field on a higher genus surface (here, a torus), selected edge coefficients were set to +1 (right).

all other faces). Using Laplace’s equation for 2-forms with the given data as rhs we get a corresponding global 2-form. (Any other sparse interpolation technique could of course be used as well.) If boundaries are present one may also supply Dirichlet or Neumann data there to enforce tangential or normal vector fields (or a mixture) at the boundary. Once 0-form (c_v) and 2-form data (c_t) is defined on all 0- respectively 2-simplices the final 1-form data c_e for all edges e_{ij} is set as

$$c_e = d^0 c_v + \delta^2 c_t.$$

Here δ^2 is the *co-differential* defined as follows. Letting t_{ijk} and t_{ijl} be the two triangles (with areas $A_{t_{ijk}}$ and $A_{t_{ijl}}$) incident to e_{ij} and angles \hat{k} and \hat{l} at v_k and v_l we get

$$(\delta^2 c_t)_{ij} = (c_{t_{ijk}}/A_{t_{ijk}} + c_{t_{ijl}}/A_{t_{ijl}})/(\cot \hat{k} + \cot \hat{l}).$$

Figures 14 and 15 demonstrate examples of this approach. Here we only used a low order approximation of the Laplace operator since it is sufficient for our experiments. For applications requiring high accuracy (e.g., physical modeling) one can use exact Laplace stencils with wider support [Wang 2006].

While this approach allows for the flexible placement of sources, sinks, and vortices, some vector fields, in particular when the surface has non-trivial topology, are neither the differential of a 0-form nor the co-differential of a 2-form. Figure 14 (right) shows such an example, a non-vanishing harmonic vector field

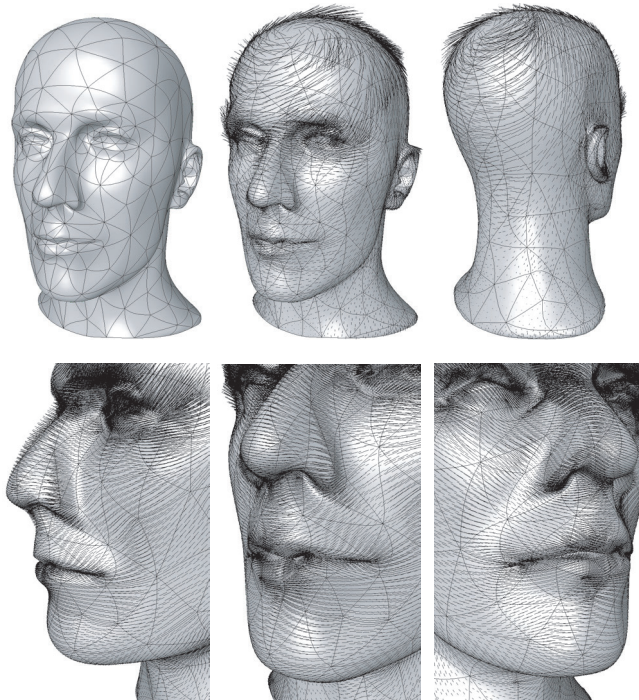


Figure 15: Fun with the mannequin head. Two vortices were placed on the head and a global vector field interpolated on the dual graph with zero Neumann boundary conditions. Note in the close ups the smooth variation of the vector field even in the presence of irregular vertices.

which is specified by setting edge coefficients directly. Importantly, for arbitrary topology surfaces any vector field can be specified with a combination of 0-, 1-, and 2-form data placed at the appropriate simplices.

5 Summary and Future Work

We have presented a novel class of subdivision schemes which carry scalar coefficients on edges from a coarse, to a refined mesh. The method constructs higher regularity bases for discrete differential 1-forms in the arbitrary topology 2-manifold setting and is intricately linked, through the commutative relations, with the associated 0- and 2-form subdivision schemes. These schemes should prove useful for the many applications which require smoothly varying tangent vector fields, and which benefit from a simple and intuitive design methodology. They should also prove useful for applications in computational science which benefit from smoother finite element type bases for divergence and curl free field variables (examples include fluids [Elcott et al. 2005] and electromagnetism [Bossavit 1998], among others), and other applications of DEC.

While we have given a concrete set of stencil weights, other choices are possible and a deeper study of the analytic properties of these new bases is required to provide further guidance in fine tuning the scheme, designing a proper Hodge decomposition, and fully understanding the underlying approximation properties to facilitate approximation of given vector fields with high accuracy. When the 1-forms are used in isolation it would be desirable to design weights which ensure C^1 continuity at irregular vertices, even if this can only be achieved by “breaking” the commutative relations. Finally, our approach applies equally well to other settings, e.g., for quadrilateral-based modeling with Catmull-Clark for S_0 and Doo-Sabin for S_2 [Wang 2006].

Acknowledgement This research has been supported in part by NSF (CCF-0528101, CCR-0133983, and ITR DMS-0453145), DOE (W-7405-ENG-

48/B341492 and DE-FG02-04ER25657), the Caltech Center for Mathematics of Information, nVidia, and Autodesk.

References

- ARNOLD, D. N., FALK, R. S., AND WINTHER, R. 2006. *Finite Element Exterior Calculus, Homological Techniques, and Applications*. *Acta Numerica* 15.
- BIERMANN, H., LEVIN, A., AND ZORIN, D. 2000. *Piecewise Smooth Subdivision Surfaces with Normal Control*. *Comp. Graphics (ACM/SIGGRAPH Proc.)*, 113–120.
- BOSSAVIT, A. 1998. *Computational Electromagnetism*. Academic Press, Boston.
- CATMULL, E., AND CLARK, J. 1978. *Recursively Generated B-Spline Surfaces on Arbitrary Topological Meshes*. *Comp. Aid. Des.* 10, 6, 350–355.
- DESBRUN, M., KANSO, E., AND TONG, Y. 2005. *Discrete Differential Forms for Computational Modeling*. In *Discrete Differential Geometry*, E. Grinspun, P. Schröder, and M. Desbrun, Eds., Course Notes. ACM SIGGRAPH.
- DOO, D., AND SABIN, M. 1978. *Analysis of the Behaviour of Recursive Division Surfaces near Extraordinary Points*. *Comp. Aid. Des.* 10, 6, 356–360.
- DYN, N., GREGORY, J. A., AND LEVIN, D. 1987. *A 4-point Interpolatory Subdivision Scheme for Curve Design*. *Comput. Aided Geom. Des.* 4, 4, 257–268.
- ELCOTT, S., TONG, Y., KANSO, E., SCHRÖDER, P., AND DESBRUN, M. 2005. *Stable, Circulation-Preserving, Simplicial Fluids*. Submitted for Publication.
- GRINSPUN, E., KRYSL, P., AND SCHRÖDER, P. 2002. *CHARMS: A Simple Framework for Adaptive Simulation*. *ACM Trans. on Graph.* 21, 3, 281–290.
- HALSTEAD, M., KASS, M., AND DEROSE, T. 1993. *Efficient, Fair Interpolation using Catmull-Clark Surfaces*. *Comp. Graphics (ACM/SIGGRAPH Proc.)*, 35–44.
- HIPTMAIR, R. 2001. *Higher-Order Whitney Forms*. *Progress in Electromagnetics Research* 32, 271–299.
- KHODAKOVSKY, A., SCHRÖDER, P., AND SWELDENS, W. 2000. *Progressive Geometry Compression*. *Comp. Graphics (ACM/SIGGRAPH Proc.)*, 271–278.
- KOBBELT, L. 2000. *$\sqrt{3}$ Subdivision*. *Comp. Graphics (ACM/SIGGRAPH Proc.)*, 103–112.
- LITKE, N., LEVIN, A., AND SCHRÖDER, P. 2001. *Fitting Subdivision Surfaces*. In *Visualization '01*, 319–324.
- LOOP, C. 1987. *Smooth Subdivision Surfaces Based on Triangles*. Master’s thesis, University of Utah, Department of Mathematics.
- OSWALD, P., AND SCHRÖDER, P. 2003. *Composite Primal/Dual $\sqrt{3}$ Subdivision Schemes*. *Comput. Aided Geom. Des.* 20, 5, 135–164.
- PRAUTZSCH, H., BOEHM, W., AND PALUSZNY, M. 2002. *Bézier and B-Spline Techniques*. Springer.
- REIF, U. 1995. *A Unified Approach to Subdivision Algorithms Near Extraordinary Points*. *Comput. Aided Geom. Des.* 12, 153–174.
- REIF, U. 1999. *Analyse und Konstruktion von Subdivisionalgorithmen für Freiformflächen beliebiger Topologie*. Shaker Verlag. Habilitationsschrift.
- REYNOLDS, C. W. 1999. *Steering Behaviors For Autonomous Characters*. In *Game Developers Conference*.
- SCHLICK, C. 1994. *An Inexpensive BRDF Model for Physically-Based Rendering*. *Comput. Graph. Forum* 13, 3, 233–246.
- SHI, L., AND YU, Y. 2005. *Taming Liquids for Rapidly Changing Targets*. In *ACM/EG Symp. on Comp. Anim.*, 229–236.
- SHIUE, L.-J., ALLIEZ, P., URSU, R., AND KETTNER, L., 2005. *A Tutorial on CGAL Polyhedron for Subdivision Algorithms*.
- STAM, J. 1998. *Exact Evaluation of Catmull-Clark Subdivision Surfaces at Arbitrary Parameter Values*. *Comp. Graphics (ACM/SIGGRAPH Proc.)*, 395–404.
- TONG, Y., LOMBAYDA, S., HIRANI, A. N., AND DESBRUN, M. 2003. *Discrete Multiscale Vector Field Decomposition*. *ACM Trans. on Graph.* 22, 3, 445–452.
- TURK, G. 2001. *Texture Synthesis on Surfaces*. *Comp. Graphics (ACM/SIGGRAPH Proc.)*, 347–354.
- WANG, K. 2006. *Online Companion to “Edge Subdivision Schemes and the Construction of Smooth Vector Fields”*. Tech. rep., Caltech.
- WARREN, J., AND WEIMER, H. 2001. *Subdivision Methods for Geometric Design: A Constructive Approach*, first ed. Morgan Kaufman Publishers.
- WHITNEY, H. 1957. *Geometric Integration Theory*. Princeton University Press.
- ZORIN, D., AND KRISTJANSSON, D. 2002. *Evaluation of Piecewise Smooth Subdivision Surfaces*. *Visual Computer* 18, 5–6, 299–315.
- ZORIN, D., AND SCHRÖDER, P., Eds. 2000. *Subdivision for Modeling and Animation*. Course Notes. ACM SIGGRAPH.
- ZORIN, D. 2000. *A Method for Analysis of C^1 -Continuity of Subdivision Surfaces*. *SIAM J. Numer. Anal.* 37, 5, 1677–1708.
- ZORIN, D. 2000. *Smoothness of Subdivision on Irregular Meshes*. *Constructive Approximation* 16, 3, 359–397.