

A Discrete Approach to Multiresolution Curves and Surfaces

Luke Olsen and Faramarz Samavati

University of Calgary

Abstract. Subdivision surfaces have been widely adopted in modeling in part because they introduce a separation between the surface and the underlying basis functions. This separation allows for simple general-topology subdivision schemes. Multiresolution representations based on subdivision, however, incongruently return to continuous functional spaces in their construction and analysis. In this paper, we propose a discrete multiresolution framework applicable to many subdivision schemes and based only on the subdivision rules. Noting that a compact representation can only afford to store a subset of the detail information, our construction enforces a constraint between adjacent detail terms. In this way, all detail information is recoverable for reconstruction, and a decomposition approach is implied by the constraint. Our framework is demonstrated with case studies in Dyn-Levin-Gregory curves and Catmull-Clark surfaces, each of which our method produces results on par with earlier methods. It is further shown that our construction can be interpreted as biorthogonal wavelet systems.

1 Introduction

Multiresolution (MR) methods have become a standard paradigm for curve and surface editing, allowing a transition between different resolutions while maintaining geometric details between edits. Subdivision curves and surfaces are supported by many popular geometric modeling programs, fitting naturally into an iterative, multi-scale design process. Since subdivision techniques naturally create a hierarchy of different resolutions, many MR approaches require the connectivity structure imposed by subdivision.

An important factor in the adoption of subdivision over parametric approaches is the simplicity that comes from the absence of basis functions: high-quality surfaces can be modeled without any direct functional evaluations. Clean, simple, affine operations are the only necessary background to learn and implement subdivision. However, the underlying theory of subdivision is based on the refinement of scaling functions [1]. An important recent movement is focused on analyzing subdivision surfaces based only on discrete operations, without explicit use of continuous functions [2].

Unlike discrete subdivision formulations, most MR schemes are defined based on wavelet theory that is heavily grounded in a functional representation. MR

settings based on wavelets have been proposed for many curve and surface subdivision types – B-splines [3], Doo-Sabin [4], Loop [5,6,7], and Catmull-Clark [8,9] surfaces. So while wavelet-based approaches can indeed be used to construct MR systems, several related questions have not been addressed. Is it possible to construct the full MR representation of a subdivision scheme only from its discrete description? Is it necessary to use continuous functions (wavelets and scalings) to create a discrete MR system? Like many problems in computer graphics, it is perhaps more natural to employ a discrete approach rather than delving into the “labyrinth of the continuum” [10].

In this paper, we present a discrete method for constructing compact and efficient multiresolution settings for a variety of subdivision schemes directly from the subdivision rules, without any direct use of basis functions (scaling or wavelets). Our construction, described in Sec. 2, is demonstrated with case studies in Dyn-Levin-Gregory curves (Sec. 3) and Catmull-Clark surfaces (Sec. 4), which are shown to offer competitively low approximation error.

1.1 Related Work

Much of the work in multiresolution is based on the theory of wavelets, although it should be noted that the two are not the same. Wavelets are a tool for “hierarchically decomposing functions” [11] into complementary bases; one basis, the coarse scaling functions, encode an approximation of the function, while the wavelet basis encodes the missing details. They are often used in geometric modeling because of their natural fit to a multi-scaling functional representation. That is, wavelet analysis undoes scaling function refinement, and can be used to derive a hierarchical multiresolution representation for curves and surfaces [12,13]. Since our paper seeks to diverge from the wavelet approach, we’ll refer the curious reader to the book of Stollnitz et al. [1] for further details.

Samavati and Bartels [14,4] investigated an alternative to wavelets for subdivision curves and Doo surfaces, constructing multiresolution systems by reversing subdivision rules. The methodology is based on discrete least-squares and can be considered to be wavelet-free, but the construction is not easily extended to general topology surfaces.

More recent multiresolution approaches for general-topology surfaces follow Sweldens’ lifting method [15], a two-stage process for constructing so-called second-generation wavelet systems. First an initial “lazy” wavelet with poor fitting properties is selected. It is then improved with respect to some criteria (eg. fitting quality or support) by one or more lifting stages. Instances of the lifting method include work on Loop [5,6] and Catmull-Clark surfaces [8,9].

Each of these methods requires a semi-regular input surface, i.e. one with subdivision connectivity. This is not a particularly restrictive condition in modeling, save for surfaces extracted from range or point-set data. Several remeshing methods have been proposed [12,16,17,18] to make semi-regular meshes from irregular ones. Remeshing comes at the expense of precision and complexity, as a remeshed surface often has more vertices and contains error relative to the original. If the original surface must be maintained, Valette and Prost [19] pro-

posed a multiresolution method for irregular meshes, but the lack of stationary subdivision rules hampers the run-time efficiency.

Some common applications of multiresolution systems are surface editing [20] and compression [21,8]. Many other representations have been considered for surface editing, such as displacement volumes [22], Laplacian encodings [23], mean value coordinates [24], and coupled prisms [25].

1.2 Notation

Let $V = \{v_1, \dots, v_n\}$ denote a set of vertices defining an object. For a curve, V represents a piecewise-linear approximation of a (usually smooth) curve. A polygon mesh M can be defined as $M = \{V, F\}$, where V defines the geometry, and F is a set of faces defining the topology. The valence n of a vertex is the number of incident edges, while the valence n^f of a face is the number of sides.

To represent relative positioning in a subdivision sequence or multiresolution hierarchy, a superscript k is used. Thus, subdivision is a process for computing a set of *fine* vertices V^{k+1} from *coarse* vertices V^k . The topological information F^{k+1} must be constructed according to F^k and the subdivision rules.

A key property of subdivision is that the subdivided position of a vertex results from a *linear* combination of a *local* neighborhood of vertices, implying that subdivision can be implemented in linear time with discrete operations. Typically a matrix \mathbf{P} is used to capture all vertex interactions during subdivision, such that $V^{k+1} = \mathbf{P}V^k$.

Decomposition is a process for converting fine data V^{k+1} to a coarse approximation V^k plus some details D^k . This should also be a linear process – let \mathbf{A} and \mathbf{B} be the associated matrices for which $V^k = \mathbf{A}V^{k+1}$ and $D^k = \mathbf{B}V^{k+1}$. Together V^k and D^k should be enough to reconstruct V^{k+1} , i.e. for some matrix \mathbf{Q} , $V^{k+1} = \mathbf{P}V^k + \mathbf{Q}D^k$. A biorthogonal MR system is one that satisfies

$$\begin{bmatrix} \mathbf{A} \\ \mathbf{B} \end{bmatrix} [\mathbf{P}|\mathbf{Q}] = \mathbf{I}.$$

In other words, decomposition and reconstruction are inverse processes. There are other orthogonality classifications [1], but anything stronger than biorthogonality is hard to satisfy in a mesh multiresolution.

Though both reconstruction and decomposition can be expressed with matrices, each matrix is sparse (due to the locality of each operation) and exhibits regularity within the columns or rows (because the same linear weights are applied to each vertex neighborhood). Thus the notion of *filters* (for curves) or *masks* (for surfaces) is useful for expressing the regular entries in a matrix. The \mathbf{P} and \mathbf{Q} matrices contain regular columns, each shifted downward by two elements from the previous column. Let $p = \{p_0, p_{\pm 1}, p_{\pm 2}, \dots\}$ and $q = \{q_0, q_{\pm 1}, q_{\pm 2}, \dots\}$ represent the non-zero entries of a regular column in \mathbf{P} and \mathbf{Q} , respectively; call these *reconstruction filters*. Similarly, \mathbf{A} and \mathbf{B} are characterized by regular entries across the rows, denoted by the *decomposition filters*, $a = \{a_0, a_{\pm 1}, \dots\}$ and $b = \{b_0, b_{\pm 1}, \dots\}$.

2 Method Overview

Our approach to constructing MR systems is based on the observation that all subdivisions impose a particular structure. New vertices are created, old vertices are (possibly) displaced, and in the mesh case, the face structure is changed (eg. one triangle split into four). Common terminology – descended from the indexing of 1D subdivision curves – refers to new vertices as *odd* and old vertices as *even*, a notation which we will use. (For vertex-splitting schemes such as Chaikin [26] or Doo-Sabin [27], one split vertex can be classified as even, the remaining as odd.)

This inherent structure recommends a decomposition strategy: even vertices are a natural choice for coarse vertices, and then, to have a compact representation, only odd vertices can be replaced with details. Notice that in addition to compactness, this approach naturally fits to the structure of subdivision.

An important question is, how can we fully reconstruct a surface if only some of the details are stored? Our idea is to place a constraint on the details. Specifically, the coarse surface is chosen such that the associated details satisfy a particular relationship, implicitly defining the missing details. Consider Fig. 1: if v^k is chosen such that its associated detail d^k is a linear combination of adjacent details – eg. $d^k = \alpha(d_l^k + d_r^k)$ – then d^k need not be stored explicitly.

In the surface case, this idea can be generalized to: choose the coarse position of even vertices such that the associated detail vector is a linear combination of the details from adjacent odd vertices. For uniform subdivisions, it is sensible for adjacent details of the same type to contribute equally to an even-vertex detail, i.e. have the same weight. Therefore, the detail constraint has a general form of

$$d^e = \alpha \sum d_j^o, \quad (1)$$

where d^e and d^o denote even and odd details respectively, and j enumerates the odd details adjacent to d^e .

Our method is as follows. First, we define a detail constraint that expresses even details in terms of adjacent odd details, and then use the subdivision rules to find an initial (or *trial*) decomposition mask that satisfies this constraint. To improve the fitting behavior, we then compute a refinement of the coarse surface by a local optimization step.




Fig. 1. (a) When a surface is decomposed, some vertices are coarsened and others are replaced with detail vectors; (b) If the details at even vertices can be computed from adjacent odd details, then a compact MR is possible

The initial masks that arise from satisfying the detail constraint represent the equivalent of a lazy wavelet decomposition. They typically have narrow support, which can lead to poor fitting performance. To increase the fitting quality, we consider a refinement of each coarse vertex v_i^k , in the form of a per-vertex displacement vector $\delta_i: v_i^k \leftarrow v_i^k + \delta_i$.

The optimal value of δ_i depends on details adjacent to v_i^k , which in turn depend on their local neighborhood of original data, and so the refinement process is equivalent to widening the support of the decomposition mask. If Δ is a vector of all such refinements, then $\Delta = \mathbf{R}D^k$ for some sparse matrix \mathbf{R} . It was noted earlier that refinement is equivalent to widening the mask support. In fact, it is equivalent to lifting, because $V^k = \mathbf{A}V^{k+1} + \mathbf{R}D^k = (\mathbf{A} + \mathbf{R}\mathbf{B})V^{k+1}$, exactly as in lifting.

Since the initial decomposition mask reverses the subdivision rules, and the size of V^k and D^k together equals the size of V^{k+1} , the initial masks are biorthogonal. According to the lifting theory, any matrix \mathbf{R} preserves biorthogonality. Therefore, the MR systems produced by our construction are biorthogonal.

This construction is best illustrated with an example. First we consider a simple curve subdivision proposed by Dyn et al. (Sec. 3). We then present a more complex surface subdivision example in Sec. 4.

3 Dyn-Levin-Gregory Curve Subdivision

Dyn et al. [28] describe an interpolating subdivision scheme based on a 4-point filter for odd points (even points are not moved), which will be referred to as *DLG* subdivision. Their construction contains a parameter w that relates to the smoothness of the limit curve. We use the typical value of $w = \frac{1}{16}$, which results in the following filters (see Fig. 2):

$$v_{2i}^{k+1} = v_i^k, \text{ and,} \tag{2}$$

$$v_{2i+1}^{k+1} = \frac{9}{16}(v_i^k + v_{i+1}^k) - \frac{1}{16}(v_{i-1}^k + v_{i+2}^k). \tag{3}$$

3.1 Trial Filter

For an interpolating scheme, there is an obvious choice for the initial filter. Since even vertices are not displaced by subdivision, a multiresolution setting can be attained by simply retaining the even vertices in V^{k+1} to form V^k , and replacing the odd vertices with details. That is, even vertices do not move:

$$\tilde{v}_i^k = v_{2i}^{k+1}. \tag{4}$$

Odd vertices can be replaced with detail vectors that capture the difference between V^{k+1} and $\mathbf{P}V^k$:

$$\begin{aligned} d_{i+1}^k &= v_{2i+1}^{k+1} - \left(\frac{9}{16} (\tilde{v}_i^k + \tilde{v}_{i+1}^k) - \frac{1}{16} (\tilde{v}_{i-1}^k + \tilde{v}_{i+2}^k) \right) \\ &= v_{2i+1}^{k+1} - \frac{9}{16} (v_{2i}^{k+1} + v_{2i+2}^{k+1}) + \frac{1}{16} (v_{2i-2}^{k+1} + v_{2i+4}^{k+1}). \end{aligned} \tag{5}$$




Fig. 2. The DLG subdivision filter for odd vertices, with $w = \frac{1}{16}$; even vertices are unmoved (Eqn. 2). Right: the corresponding subdivision matrix \mathbf{P} .

The details at even vertices are $d_i^k = \mathbf{0}$ (i.e. $\alpha = 0$ in Eqn. 1). In this case, the detail constraint is not needed to derive an initial decomposition filter.

In matrix form, the trial filters can be summarized as:

$$\tilde{\mathbf{A}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 1 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & 1 & \dots \\ \vdots & & & & & \ddots \end{bmatrix}, \tilde{\mathbf{B}} = \begin{bmatrix} \frac{-9}{16} & 1 & \frac{-9}{16} & 0 & \frac{1}{16} & 0 & 0 & 0 & \dots & \frac{1}{16} & 0 \\ \frac{1}{16} & 0 & \frac{9}{16} & 1 & \frac{9}{16} & 0 & \frac{1}{16} & 0 & \dots & 0 & 0 \\ \vdots & & & & & & & & & & \ddots \end{bmatrix}, \tilde{\mathbf{Q}} = \begin{bmatrix} 0 & 0 & 0 & 0 & \dots \\ 1 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & \dots \\ \vdots & & & & \ddots \end{bmatrix}.$$

3.2 Refinement

As Bartels and Samavati note [29], in the general case the fine data V^{k+1} will have non-zero details. In such cases, we may be able to achieve lower approximation error by choosing a non-trivial decomposition filter.

Our approach is to improve the trial filter by a refinement step, with the goal of reducing the local error, call it E , at each vertex $v_i^k \in V^k$. Based on a three-point neighborhood around v_i^k , the local error is given by the magnitude of the details. In particular, recalling that $d_i^k = \mathbf{0}$, we have

$$E = \|d_{i-1}^k\|^2 + \|d_{i+1}^k\|^2.$$

If V^k is a poor approximation of V^{k+1} , the details (and therefore the error) will be large.

After refinement, \tilde{v}_i^k will be replaced by $v_i^k = \tilde{v}_i^k + \delta_i$. This has an impact on the local error of v_i^k and neighboring vertices: d_i^k becomes $-\delta_i$, and d_{i+1}^k becomes $d_{i+1}^k - \frac{9}{16}\delta_i$ by Eqn. 5. Thus, after refinement the local error becomes

$$E(\delta) = \|-\delta_i\|^2 + \|d_{i-1}^k - \frac{9}{16}\delta_i\|^2 + \|d_{i+1}^k - \frac{9}{16}\delta_i\|^2.$$

We should choose δ_i to minimize this error.

After a bit of algebraic manipulation, $E(\delta)$ becomes

$$\begin{aligned} E(\delta) &= \frac{209}{128} \|\delta_i\|^2 - \left(\frac{9}{8} (d_{i-1}^k + d_{i+1}^k) \right) \cdot \delta_i + E \\ &= a \|\delta_i\|^2 - \mathbf{g} \cdot \delta_i + E . \end{aligned}$$

A minimal solution occurs where the derivate equals 0, or

$$2a\delta_i - \mathbf{g} = 0 \rightarrow \delta_i = \frac{a^{-1}}{2} \mathbf{g} .$$

Therefore, an optimal displacement for vertex v_i^k is

$$\delta_i = \frac{72}{209} (d_{i-1}^k + d_{i+1}^k) . \tag{6}$$

The refinement step can be incorporated into a closed-form filter by expressing d^k in terms of elements from V^{k+1} .

$$\begin{aligned} v_i^k &= \tilde{v}_i^k + \delta_i \\ &= v_{2i}^{k+1} + \frac{72}{209} (d_{i-1}^k + d_{i+1}^k) \\ &= \frac{128}{209} v_{2i}^{k+1} + \frac{72}{209} v_{2i\pm 1}^{k+1} - \frac{36}{209} v_{2i\pm 2}^{k+1} + \frac{9}{418} v_{2i\pm 4}^{k+1} . \end{aligned}$$

Our MR setting for DLG subdivision can be summarized by the following filters:

$$\begin{aligned} a &= \left\{ \frac{128}{209}, \frac{72}{209}, -\frac{36}{209}, 0, \frac{9}{418} \right\} , \\ b &= \left\{ 1, -\frac{9}{16}, 0, \frac{1}{16} \right\} , \text{ and} \\ q &= \left\{ \frac{128}{209}, -\frac{72}{209}, -\frac{36}{209}, 0, \frac{9}{418} \right\} . \end{aligned}$$

The q filter given above results from refining the trivial q filter ($q = \{1, 0, \dots, 0\}$).

That is, $\mathbf{Q} \leftarrow \mathbf{Q} - \mathbf{P}\mathbf{R}$ where \mathbf{R} is the refinement matrix defined by Eqn. 6. This set of filters allows for perfect reconstruction of a decomposed curve.

3.3 Results

To evaluate our MR construction, the work of Samavati and Bartels [29] is used as a comparison. Their local least-squares approach yields an optimal decomposition filter $a_{opt} = \left\{ \frac{107}{161}, \frac{48}{161}, -\frac{24}{161}, 0, \frac{3}{161} \right\}$. Based on the width of the filter and

the relative magnitude of the weights, we anticipate that our DLG filter will perform close to optimally.

$$\mathbf{A} = \frac{1}{418} \begin{bmatrix} 256 & 144 & -72 & 0 & 9 & 0 & 9 & 0 & -72 & 144 \\ -72 & 144 & 256 & 144 & -72 & 0 & 9 & 0 & 9 & 0 \\ 9 & 0 & -72 & 144 & 256 & 144 & -72 & 0 & 9 & 0 \\ 9 & 0 & 9 & 0 & -72 & 144 & 256 & 144 & -72 & 0 \\ -72 & 0 & 9 & 0 & 9 & 0 & -72 & 144 & 256 & 144 \end{bmatrix},$$

$$\mathbf{Q} = \frac{1}{418} \begin{bmatrix} -144 & 0 & 0 & 0 & -144 \\ 256 & -72 & 9 & 9 & -72 \\ -144 & -144 & 0 & 0 & 0 \\ -72 & 256 & -72 & 9 & 9 \\ 0 & -144 & -144 & 0 & 0 \\ 9 & -72 & 256 & -72 & 9 \\ 0 & 0 & -144 & -144 & 0 \\ 9 & 9 & -72 & 256 & -72 \\ 0 & 0 & 0 & -144 & -144 \\ -72 & 9 & 9 & -72 & 256 \end{bmatrix}$$

Results from a number of curves (shown in Fig. 3 and 4) are listed in Table 1. For each curve V^k , the decomposition filter is applied k times to get a coarse approximation V^0 of the original curve; the difference between the approximation $P^k V^0$ and the original gives an error measure $L = \|V^k - P^k V^0\|^2$ for comparing different filters. In each case, our refined filter is a substantial improvement over the trial and virtually identical to the optimal result.

To illustrate the better fitting behavior of the refined filter relative to the trial filter, consider the coastline data depicted in Fig. 4. At the southwest tip of the island, for instance, the refined filter (c) is able to approximate the original data more closely than the trial filter (b). This makes the refined filter more suitable for compression applications.




Fig. 3. Curves used to evaluate the DLG filters: (a) face; (b) car; (c) tree; (d) wolf. The original curve is shown in black, and the curve approximated by our refined filter in red. Numeric results are listed in Table 1.





Fig. 4. Comparison of our 4-point filters on real coastline data (from [30]). The original data (4096 points) is decomposed six times to obtain a 64-point approximation, which is then subdivided to the original resolution. The refined filter is able to better approximate the original data (numeric results given in Table 1).

Table 1. L_2 error relative to the original curve for the trial, refined, and optimal DLG filters (normalized against the trial error value). Shorter bars are better.



	Coast	Face	Car	Tree	Wolf
■ Optimal	0.85	0.74	0.86	0.77	0.88
□ Refined	0.86	0.75	0.86	0.78	0.89

4 Catmull-Clark Surface Subdivision

Catmull-Clark subdivision [31,32] is a popular scheme [33] for manifold surfaces of arbitrary topology. It produces a quadrilateral mesh from any base mesh by adding vertices at each edge and face, then inserting edges from each edge vertex to adjacent face vertices; see Fig. 5. The limit surface is C^2 -continuous at regular (valence-4) vertices, C^1 elsewhere. This subdivision scheme is a popular




Fig. 5. An iteration of Catmull-Clark subdivision creates new vertices at each face (f) and edge (e), while displacing old vertices (v), to create a smooth surface in the limit

alternative to triangle-mesh schemes such as Loop, providing more pleasing results when the input mesh contains mostly quadrilateral faces.

The subdivision masks for Catmull-Clark subdivision are given by:

$$f_i^{k+1} = \frac{1}{n_i^k} \left(v^k + e_i^k + e_{i+1}^k + \sum_j f_{i,j}^k \right), \tag{7}$$

$$e_i^{k+1} = \frac{1}{4} (v^k + e_i^k + f_{i-1}^{k+1} + f_i^{k+1}), \tag{8}$$

$$v^{k+1} = \frac{n-2}{n} v^k + \frac{1}{n^2} \sum_i e_i^k + \frac{1}{n^2} \sum_i f_i^{k+1}, \tag{9}$$

according to the notation illustrated in Fig. 6 (indices are computed modulo n). Note that these masks, as presented, assume that the face vertices f^{k+1} are computed first.

4.1 Trial Mask

The general form of the detail constraint (Eqn. 1) has only one free parameter. For Catmull-Clark surfaces, the heterogeneity of odd vertices necessitates two free parameters; we denote these weights α_e for edge-vertex details, and α_f for face-vertex details. The detail constraint is then

$$d^v = \alpha_e \sum d_i^e + \alpha_f \sum d_i^f, i = 1, \dots, n, \tag{10}$$

where d^v , d^e , and d^f represent the details of vertex-, edge-, and face-vertices respectively.

The vertex subdivision mask (Eqn. 9) can be rewritten so that the e^k terms are replaced with e^{k+1} terms:

$$v^{k+1} = \frac{n-3}{n} v^k + \frac{4}{n^2} \sum_i e_i^{k+1} - \frac{1}{n^2} \sum_i f_i^{k+1}.$$




Fig. 6. Notation for the neighborhood of vertex v^k . The 1-ring of v^k consists of edge neighbors e_i^k and face neighbors $f_{i,j}^k$. After subdivision, the 1-ring of v^{k+1} contains edge neighbors e_i^{k+1} and face neighbors f_i^{k+1} .

With this new form, the vertex-, edge- and face-details can be expressed by

$$d^v = v^{k+1} - \left(\frac{n-3}{n} v^k + \frac{4}{n^2} \sum_i \tilde{e}_i^{k+1} - \frac{1}{n^2} \sum_i \tilde{f}_i^{k+1} \right),$$

$$d_i^e = e_i^{k+1} - \tilde{e}_i^{k+1}, \text{ and,}$$

$$d_i^f = f_i^{k+1} - \tilde{f}_i^{k+1},$$

where \tilde{e}_i^{k+1} and \tilde{f}_i^{k+1} represent the approximations of e_i^{k+1} and f_i^{k+1} (i.e. subdivision of the coarse approximation).

To determine α_e and α_f , we must compare the left and right sides of Eqn. 10: $\sum_i \tilde{e}_i^{k+1}$ and $\sum_i \tilde{f}_i^{k+1}$ appear on both sides, with respective weights of $\frac{4}{n^2}$ and $-\frac{1}{n^2}$ on the left, and α_e and α_f on the right side. By setting $\alpha_e = \frac{4}{n^2}$ and $\alpha_f = -\frac{1}{n^2}$, those terms are eliminated, leaving

$$v^{k+1} - \frac{n-3}{n} v^k = \frac{4}{n^2} \sum_i e_i^{k+1} - \frac{1}{n^2} \sum_i f_i^{k+1}.$$

Thus, the detail constraint is satisfied if this relationship between coarse and fine vertices holds, or equivalently, if

$$v^k = \frac{1}{n-3} \left(n v^{k+1} - \frac{4}{n} \sum_i e_i^{k+1} + \frac{1}{n} \sum_i f_i^{k+1} \right). \tag{11}$$

Using this equation to decompose all even vertices v^{k+1} to coarse vertices v^k will ensure that the detail constraint is satisfied for $\alpha_e = \frac{4}{n^2}$ and $\alpha_f = -\frac{1}{n^2}$.




Fig. 7. Valence-3 vertices must be handled with special-case masks

Note that this represents an inversion of the vertex subdivision mask (Eqn. 9). That is, if a surface is the product of subdivision, then decomposition with this mask will produce a zero-error coarse approximation.

Valence-3 Vertices. The decomposition given by Eqn. 11 is undefined for valence-3 vertices, due to the $\frac{1}{n-3}$ term. Fortunately, a decomposition mask that works for most valence-3 vertices can be found by an alternate method.

Consider a valence-3 vertex v with at least one non-valence-3 coarse neighbor v_n (see Fig. 7). Furthermore, let e be the edge vertex between v and v_n , and f_a and f_b be the face-vertices adjacent to e . According to Eqn. 8, these vertices are related by $e \approx \frac{1}{4}(v^k + v_n^k + f_a + f_b)$.

Since v_n is not valence-3, it can be decomposed via Eqn. 11 to v_n^k , leaving v^k as the only unknown. Therefore,

$$v^k \approx 4e - v_n^k - f_a - f_b . \tag{12}$$

Although it was assumed that v_n is not valence-3, Eqn. 12 allows v to be decomposed as long as v_n has already been decomposed, regardless of v_n 's valence. Therefore, isolated valence-3 vertices can be decomposed by cascading inwards from non-valence-3 vertices. In practice there is rarely a need for this approach, though: because all edge vertices are valence-4, two valence-3 vertices can only be adjacent in the base mesh.

One downside to this approach is that the detail constraint is not satisfied, meaning that a detail term must be explicitly stored for each of these vertices. Another drawback is that the position of v^k does not depend on v^{k+1} ; in cases where v^{k+1} deviates largely from a smooth position, Eqn. 12 will produce a poor approximation. To alleviate this issue somewhat, each already-decomposed neighbor of v^{k+1} (a maximum of 3) can nominate a ‘‘candidate’’ position for v^k based on Eqn. 12. The final position of v^k is then taken to be the average of all candidates.

4.2 Refinement

The initial decomposition mask has small support, considering that a coarse vertex depends on only one level of vertex neighbors. Multiresolution systems

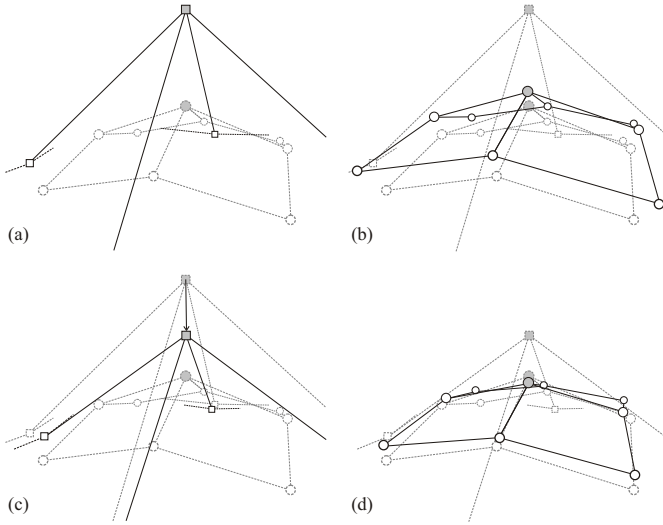


Fig. 8. Refinement reduces the error in coarse data: (a) initial coarse approximation; (b) subdivision of (a); (c) coarse surface after refinement; (d) subdivision of (c). The surface in (d) is “closer” to the original surface than (b).

generally exhibit greater stability and better fitting properties when the decomposition mask has wider support [29].

We address the stability and fitting properties from another direction, recalling the geometric interpretation of detail vectors. A detail vector captures the difference between a surface and the subdivision of its coarse approximation, i.e the *error* in the coarse surface. The larger the error, the larger the magnitude of the detail vectors.

We can look at error on local and global scales. Since efficiency is a primary concern in our construction, we restrict ourselves to a local examination. If the local error can be reduced everywhere in a mesh, then the global error is also reduced. The local error of a vertex is represented by the magnitude of a k -ring of details. Here we consider the 1-ring, in which case the local error of v^k is

$$E = \|d^v\|^2 + \sum \|d_i^e\|^2 + \sum \|d_i^f\|^2 .$$

Figure 8a-b depicts the error introduced by decomposition.

If the error is non-zero, there should be a vector δ for which $v^k + \delta$ is a more optimal position than v^k . To find δ – recalling that the error is a measure of how far a subdivided surface is from the original – we must determine how the displacement of v^k impacts \tilde{v}^{k+1} , \tilde{e}_i^{k+1} , and \tilde{f}_i^{k+1} . If chosen correctly, we should see a reduction in the error (Fig. 8c-d).

From the subdivision rules, it can be seen that the weight applied to v^k is r for \tilde{v}^{k+1} , s_i for \tilde{e}_i^{k+1} , and t_i for \tilde{f}_i^{k+1} , where

$$r = \frac{n-2}{n} + \frac{1}{n^2} \sum_i \frac{1}{n_i^f},$$

$$s_i = \frac{1}{4} \left(1 + \frac{1}{n_i^f} + \frac{1}{n_{i-1}^f} \right), \text{ and}$$

$$t_i = \frac{1}{n_i^f}.$$

When v^k is displaced, the error changes. For instance, \tilde{v}^{k+1} becomes $\tilde{v}^{k+1} + r\delta$, and therefore $d^v = v^{k+1} - \tilde{v}^{k+1}$ becomes $v^{k+1} - (\tilde{v}^{k+1} + r\delta) = d^v - r\delta$; similarly, d_i^e and d_i^f are impacted according to s_i and t_i . Thus, after displacement the local error becomes

$$E(\delta) = \|d^v - r\delta\|^2 + \sum \|d_i^e - s_i\delta\|^2 - \sum \|d_i^f - t_i\delta\|^2.$$

This expression expands to

$$E(\delta) = a\|\delta\|^2 - \mathbf{g} \cdot \delta + E, \tag{13}$$

where

$$a = r^2 + \sum_i s_i^2 + \sum_i t_i^2, \text{ and}$$

$$\mathbf{g} = 2 \left(rd^v + \sum_i s_i d_i^e + \sum_i t_i d_i^f \right).$$

Equation 13 is quadratic in δ , so the minimizing value can be found analytically by differentiating and finding a zero-crossing.

$$E'(\delta) = 2a\delta - \mathbf{g} = 0 \rightarrow \delta = \frac{\mathbf{g}}{2a}.$$

This is a minimum, because $E''(\delta) = 2a > 0$. Therefore

$$\delta = \frac{\mathbf{g}}{2a} = \frac{\sum(r\alpha_e + s_i)d_i^e + \sum(r\alpha_f + t_i)d_i^f}{r^2 + \sum s_i^2 + \sum t_i^2},$$

where the detail constraint (Eqn. 10) is used to replace the d^v term. (Because the detail constraint does not apply to valence-3 vertices, the d^v term should be retained when refining them.)

Due to the interrelationships between vertices (eg. coarse vertices sharing an edge each contribute to the edge vertex), this refinement is only optimal when applied to a single vertex. That is, if a neighbor of v^k has been displaced by refinement, the details of any shared face- and edge-vertices have already been reduced, leading to an over-refinement of v^k .

We have observed [7] that scaling the displacements by the central weight of the subdivision mask adequately accounts for these interrelationships. Here, the central weight (i.e. the contribution of v^k to v^{k+1}) is r . Scaling δ by r yields

$$\delta = \frac{\sum r(r\alpha_e + s_i)d_i^e + \sum r(r\alpha_f + t_i)d_i^f}{r^2 + \sum s_i^2 + \sum t_i^2} . \tag{14}$$

4.3 Boundaries and Creases

Because Catmull-Clark surfaces are a generalization of cubic B-spline curves, curve subdivision is typically used along boundary vertices and edges [33]. Sharp features can also be accommodated with boundary masks, increasing the versatility of the representation. Similarly, a multiresolution system can employ a B-spline curve multiresolution (such as [3,14,7]) when decomposing a mesh with boundaries and sharp features. As illustrated in Fig. 9, using boundary masks can provide much better results after decomposition, introducing less error and also making the mesh easier to edit.

4.4 Results

To evaluate the fitting quality of our construction, we use a similar approach as the previous section. Each of the models in Figs. 9–12 was decomposed k times, then subdivided without details back to the original resolution. For comparison purposes, the Catmull-Clark MR system described by Bertram et al. [8] was implemented.

Table 2 summarizes the L_2 -norm error of the approximation relative to the original. The approximation error introduced by our method is comparable to Bertram’s method; in some instances ours outperforms theirs and vice versa, depending on the characteristics of the particular model. Based on visual inspection and the numeric results, our method provides good fitting behavior suitable for use in compression and mesh editing applications.

A benefit of subdivision-based multiresolution systems is their run-time efficiency; decomposition and reconstruction are linear, if designed well. For the




Fig. 9. Sharp features can be handled with B-spline curve masks, providing a better coarse approximation of the original data. In this example, the boundary between the silver and blue sections was automatically tagged as a sharp feature.




Fig. 10. De-noising of a double-torus with our refined masks, from level 3 (left) to level 0 (right)





Fig. 11. Decomposition of a terrain section, from level 3 (left) to level 0 (right). (Data from USGS[34].)

Table 2. L_2 error relative to the original surface (normalized against the trial error). Shorter bars are better.






Fig. 12. Multiresolution editing: geometric details are retained after performing an edit at a lower resolution

Table 3. Time required to reconstruct a surface from the previous level, i.e. level 1 refers to the reconstruction of level 1 from level 0 (reported in seconds)

Model	Reconstruction time (s)			
	M^1	M^2	M^3	M^4
Double torus	0.62	0.63	1.47	3.50
Terrain	0.15	0.31	0.36	1.21
Bullet	0.10	0.15	0.17	–
Dog	0.17	0.53	1.22	–



Fig. 13. Decomposition of a complex surface: with only 1/16th of the geometric detail, the rightmost surface is a faithful representation of the original

models in Figs. 9–12, we measured the time required to reconstruct the model from the lowest resolution (level 0) to the highest; these results are summarized in Table 3. The system was implemented in a high-level language (C#), so performance is less than optimal. Despite the overhead, however, the efficiency is suitable for interactive editing of complex surfaces.

For a visual analysis of the fitting quality, consider Fig. 13. One level of decomposition reduces the geometric complexity by a factor of 4, yet has very little effect on the visual quality of the object. Further decomposition results in a surface with 1/16th of the complexity, yet the coarse model is still a faithful representation of the original.

5 Conclusion

We have proposed a wavelet-free multiresolution construction that is applicable to a variety of subdivision surfaces. By constraining the details and then choosing the coarse surface to satisfy a detail constraint, it is possible to store only a subset of the detail terms and compute the rest. To have high fitting quality, it is necessary to perform a local optimization on the coarse surface.

Though our construction can be viewed through the lens of wavelet analysis, and in fact is an instance of the lifting method, no knowledge of wavelets or scaling functions is required to understand or apply the method. Thus, the simplicity of subdivision is carried over to multiresolution systems, in both description and implementation.

Our construction was illustrated by constructing MR systems for Dyn-Levin-Gregory subdivision curves and Catmull-Clark surfaces, including a consideration of boundary masks. The resulting systems are stable and provide fitting behavior that is comparable to other methods, showing that our constructions are suitable for editing and compression applications.

Acknowledgments

This research was supported in part by the National Science and Engineering Research Council of Canada.

References

1. Stollnitz, E., DeRose, T., Salesin, D.: *Wavelets for Computer Graphics: Theory and Applications*, pp. 152–159. Morgan Kaufmann Publishers, Inc., San Francisco (1996)
2. Warren, J., Weimer, H.: *Subdivision Methods for Geometric Design: A Constructive Approach*. Morgan Kaufmann, San Francisco (2001)
3. Finkelstein, A., Salesin, D.: Multiresolution curves. In: *Proc. of SIGGRAPH 1994*, pp. 261–268. ACM Press, New York (1994)
4. Samavati, F., Mahdavi-Amiri, N., Bartels, R.: Multiresolution surfaces having arbitrary topologies by a reverse doo subdivision method. *Computer Graphics Forum* 21(2), 121–136 (2002)
5. Bertram, M.: Biorthogonal Loop-Subdivision Wavelets. *Computing* 72(1-2), 29–39 (2004)
6. Li, D., Qin, K., Sun, H.: Unlifted loop subdivision wavelets. In: *12th Pacific Conference on Computer Graphics and Applications* (2004)
7. Olsen, L., Samavati, F., Bartels, R.: Multiresolution for curves and surfaces based on constraining wavelets. *Computers and Graphics* (2007)
8. Bertram, M., Duchaineau, M., Hamann, B., Joy, K.: Generalized b-spline subdivision-surface wavelets for geometry compression. *IEEE Transactions on Visualization and Computer Graphics* 10, 326–338 (2004)
9. Wang, H., Qin, K.H., Tang, K.: Efficient wavelet construction with catmull-clark subdivision. *The Visual Computer* 22, 874–884 (2006)

10. Leibniz, G.: *The Labyrinth of the Continuum: Writings on the Continuum Problem*, pp. 1672–1686. Yale University Press, New Haven (2001)
11. Stollnitz, E., DeRose, T., Salesin, D.: Wavelets for computer graphics: A primer, part 1. *IEEE Computer Graphics and Applications* 15(3), 76–84 (1995)
12. Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., Stuetzle, W.: Multiresolution analysis of arbitrary meshes. In: *Proc. of SIGGRAPH 1995*, pp. 173–182 (1995)
13. Lounsbery, M., DeRose, T.D., Warren, J.: Multiresolution analysis for surfaces of arbitrary topological type. *ACM Trans. Graph.* 16(1), 34–73 (1994)
14. Samavati, F., Bartels, R.: Multiresolution curve and surface representation by reversing subdivision rules. *Computer Graphics Forum* 18(2), 97–120 (1999)
15. Sweldens, W.: The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.* 29(2), 511–546 (1997)
16. Lee, A., Sweldens, W., Schröder, P., Cowsar, L., Dobkin, D.: Maps: multiresolution adaptive parameterization of surfaces. In: *Proc. of SIGGRAPH 1998*, pp. 95–104. ACM Press, New York (1998)
17. Litke, N., Levin, A., Schröder, P.: Fitting subdivision surfaces. In: *IEEE Visualization 2001*, pp. 319–324 (2001)
18. Boier-Martin, I., Rushmeier, H., Jin, J.: Parameterization of triangle meshes over quadrilateral domains. In: *Proc. of ACM Symposium on Geometry Processing (SGP 2004)*, pp. 193–203 (2004)
19. Valette, S., Prost, R.: Wavelet-based multiresolution analysis of irregular surface meshes. *IEEE Transactions on Visualization and Computer Graphics* 10(2), 113–122 (2004)
20. Zorin, D., Schröder, P., Sweldens, W.: Interactive multiresolution mesh editing. In: *Proc. of SIGGRAPH 1997*, pp. 259–268. ACM Press, New York (1997)
21. Gioia, P.: Reducing the number of wavelet coefficients by geometric partitioning. *Computational Geometry: Theory and applications* 14(1-3), 25–48 (1999)
22. Botsch, M., Kobbelt, L.: Multiresolution surface representation based on displacement volumes. *Computer Graphics Forum* 22, 483–491 (2003)
23. Sorkine, O., Cohen-Or, D., Lipman, Y., Alexa, M., Rössl, C., Seidel, H.P.: Laplacian surface editing. In: *Proc. of ACM Symposium on Geometry Processing (SGP 2004)*, pp. 175–184 (2004)
24. Ju, T., Schaefer, S., Warren, J.: Mean value coordinates for closed triangular meshes. In: *Proc. of SIGGRAPH 2005*, pp. 561–566 (2005)
25. Botsch, M., Pauly, M., Gross, M., Kobbelt, L.: Primo: Coupled prisms for intuitive surface modeling. In: *Proc. of ACM Symposium on Geometry Processing (SGP 2006)*, pp. 11–20 (2006)
26. Chaikin, G.: An algorithm for high speed curve generation. *Computer Graphics and Image Processing* 3(4), 346–349 (1974)
27. Doo, D., Sabin, M.: Behaviour of recursive subdivision surfaces near extraordinary points. *Computer-Aided Design* 10(6), 356–260 (1978)
28. Dyn, N., Levine, D., Gregory, J.: A 4-point interpolatory subdivision scheme for curve design. *CAGD* 4, 257–268 (1987)
29. Bartels, R., Samavati, F.: Reversing subdivision rules: Local linear conditions and observations on inner products. *Journal of Computational and Applied Mathematics* 119, 29–67 (2000)
30. National Geophysical Data Center: Coastline extractor (2007), <http://rimmer.ngdc.noaa.gov/mgg/coast/getcoast.html>

31. Catmull, E., Clark, J.: Recursively generated b-spline surfaces on arbitrary topological surfaces. *Computer-Aided Design* 10(6), 350–355 (1978)
32. Zorin, D., Schröder, P.: Subdivision for modeling and animation. In: *SIGGRAPH 2000 Course Notes*. ACM Press, New York (2000)
33. DeRose, T., Kass, M., Truong, T.: Subdivision surfaces in character animation. In: *Proc. of SIGGRAPH 1998*, pp. 85–94 (1998)
34. United States Geological Survey: Seamless data distribution system (2004), <http://seamless.usgs.gov/website/seamless>