

A smooth subdivision scheme for hexahedral meshes

Chandrajit Bajaj

Department of Computer Sciences, University of Texas

Joe Warren

Department of Computer Science, Rice University

Guoliang Xu

State Key Laboratory of Scientific and Engineering Computing,

Chinese Academy of Sciences, Beijing

Abstract

In a landmark paper, Catmull and Clark described a simple generalization of the subdivision rules for bi-cubic B-splines to arbitrary quadrilateral surface meshes. This smooth subdivision scheme has become a mainstay of surface modeling systems. Joy and MacCracken described a generalization of this surface scheme to volume meshes. Unfortunately, little is known about the smoothness and regularity of this scheme due to the complexity of the subdivision rules. This paper presents an alternative subdivision scheme for hexahedral volume meshes that consists of a simple split and average algorithm. Along extraordinary edges of the volume mesh, the scheme provably converges to a smooth limit volume. At extraordinary vertices, the authors supply strong experimental evidence that the scheme also converges to a smooth limit volume. The scheme automatically produces reasonable rules for non-manifold topology and can easily be extended to incorporate boundaries and embedded creases expressed as Catmull-Clark surfaces and B-spline curves.

1 Introduction

Subdivision has evolved into a fundamental technique for describing geometric shape, mainly due to its simplicity. Given a coarse polygonal mesh \mathcal{p}_0 , subdivision produces a sequence of increasingly dense polygonal meshes \mathcal{p}_i with each successive mesh related by the linear transformation $\mathcal{p}_i = S_{i-1} \mathcal{p}_{i-1}$. If these subdivision matrices S_{i-1} are chosen appropriately, then this sequence of polygonal meshes converge to a smooth limit shape \mathcal{p}_∞ that approximates the initial coarse mesh \mathcal{p}_0 . Figure 1 is an example of this process. An initial coarse mesh is repeatedly refined into a sequence of increasingly dense meshes. The limit of this process is a smooth surface that follows the initial mesh.




Figure 1

Typically, the \mathcal{p}_i are treated as vectors that store the positions of each vertex of the mesh. The topological connectivity of the mesh is specified via an auxiliary data structure T_i that consists of a list of faces of the mesh. In this framework, each round of subdivision consists of two separate steps: topological splitting of each face in T_{i-1} to produce T_i , then computing new vertex position \mathcal{p}_i via the subdivision transformation $S_{i-1} \mathcal{p}_{i-1}$. The beauty of this arrangement is that there are very simple rules for generating reasonable transformations S_{i-1} (i.e. ones whose limit is well-behaved) directly from the topology T_{i-1} . The example of figure 1 uses the subdivision rules for bi-cubic B-splines. Each quadrilateral is split into four quadrilaterals. Each new vertex of \mathcal{p}_i is a convex combination of its nearby neighbors on \mathcal{p}_{i-1} . One of three different combinations is used depending on whether the new vertex lies on a face, edge or vertex of T_{i-1} .

Unfortunately, these rules only apply to quadrilateral meshes with tensor product topology (i.e. each interior vertex is shared by four quadrilaterals). This restriction greatly reduces the applicability of these rules. Standard shapes such as a sphere require meshes with non-tensor product topology. In a landmark paper, Catmull and Clark [2] developed a generalization of the rules for bi-cubic B-splines that apply to any type of quadrilateral mesh. In particular, they developed a rule

for repositioning those vertices that are shared by n quadrilaterals. This new rule was designed to produce smooth surfaces in the limit even in the case of extraordinary vertices (i.e. $n \neq 4$).

The development of this scheme ignited the growth of subdivision into a practical surface modeling tool. Major corporations such as Pixar have adopted CatmullClark surfaces as their basic surface modeling primitive. For volume modeling, the use of subdivision is still in its infancy. Joy and MacCracken [8] introduced the first generalization of the rules for Catmull-Clark subdivision to volume meshes. Unfortunately, the complexity of these rules makes analyzing the smoothness of the resulting volume scheme very difficult and has in all likelihood hampered the use of the scheme in practice. In this paper, we suggest an alternative set of subdivision rules that have a much simpler description in terms of splitting and averaging. In addition, we provide a theoretical analysis that yields strong evidence that this new scheme produces smooth volume meshes everywhere.

2 The scheme

To understand the new scheme, we first consider the subdivision rules for bi-cubic B-splines. Initially, the bi-cubic rules were described in terms of topological splitting followed by vertex positioning. One difficulty with this approach was the need for three types of positioning rules: one for each possible topological position of a new vertex in T_i with respect to T_{i-1} (i.e. on a face of T_{i-1} , on an edge of T_{i-1} , at a vertex of T_{i-1}). One way to avoid this problem is to recast the bi-cubic rules into a slightly different framework.

The geometric positioning rules embodied by the subdivision S_{i-1} can be factored into two separate transformations: bi-linear subdivision followed by a simple averaging operation. Figure 2 illustrates this factorization for the first step of figure 1. A coarse mesh is first split using bi-linear subdivision. The resulting mesh is then smoothed using a simple averaging operator supported over the nearby neighbors of each vertex.




Figure 2

There are two superficial advantages of this approach. First, bi-linear subdivision can be carried out during topological splitting. Second, the averaging operator for tensor product meshes employs only a single type of averaging rule. Similarly, the rule for tri-cubic subdivision can be factored into tri-linear subdivision followed by averaging. The true advantage of this decomposition is that the problem of developing a variant of the tri-cubic rule for arbitrary hexahedral meshes can be focused solely on the averaging operation. This approach allows us to develop a single, simple averaging rule for arbitrary meshes (as opposed to the four required under the Joy and MacCracken approach). Our task is to develop a generalization of this averaging rule that reproduces the tensor product rule on tensor product meshes and yields smooth limit volumes for arbitrary hexahedral meshes. This generalization, developed in 2D by Morin et al. [10], Zorin and Schröder [20] and Stam [15] simultaneously, is remarkably simple.

2.1 Multi-linear subdivision

Before proceeding to the problem of averaging, we consider the problem of implementing multi-linear subdivision. At first glance, this might appear to be an imposing task. However, with the correct choice of data structures, implementing multi-linear subdivision is straightforward. (Note that we consider the general d -dimensional case since the three dimensional solution allows for little simplification.) Essentially, multi-linear subdivision consist of splitting a topological d -cube into 2^d subcubes and positioning the new vertices using multi-linear interpolation. To implement the topological split, we must first settle on a representation for the topology of a d -dimensional cube (a d -cube). We suggest a simple, recursive representation with a 0-cube consisting of a vertex index and a d -cube ($d > 0$) consisting of a list of two $(d - 1)$ -cubes. For example, a square consists of a list of two line segments and a cube consist of a list of two squares. Given this representation, an algorithm for splitting a d -cube C into its 2^d subcubes is as follow:

Given a d -cube C , recursively compute the multi-linear subdivision of two $(d-1)$ -cubes comprising and call the two resulting lists of 2^{d-1} $(d-1)$ -cubes top and bottom , respectively. top and bottom are splits of the "top" and "bottom" faces of C . Next, use linear interpolation to compute a list of 2^{d-1} $(d-1)$ -cubes called middle that lie halfway between top and bottom . Finally, return 2^{d-1} d -cubes from corresponding pairs of cubes in top and middle and 2^{d-1} d -cubes from corresponding pairs of cubes in middle and bottom .

Given this topological representation for a single d -cube, the topology T_1 for a mesh of d -cubes is simply a list of such d -cubes. To ensure topological consistency of this representation, each new vertex that is common to several subcubes must use the same index for each subcube. To maintain these indices, we suggest using a hash table. The same hash table can be used to store the new vertex locations provided by multi-linear interpolation.

2.2 Cell averaging

The key to finding a suitable averaging operator for non-tensor product meshes is to understand the structure of the averaging operator in the tensor product case. In the univariate case, the subdivision rules for cubic B-splines can be expressed as linear subdivision followed by averaging with the mask $(\frac{1}{4} \quad \frac{1}{2} \quad \frac{1}{4})$. Specifically, the subdivision matrix S_{k-1} can be decomposed as follows:

$$\begin{pmatrix} \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 & 0 & \cdot \\ \cdot & 0 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & \cdot \\ \cdot & 0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & 0 & \cdot \\ \cdot & 0 & 0 & \frac{1}{2} & \frac{1}{2} & 0 & \cdot \\ \cdot & 0 & 0 & \frac{1}{8} & \frac{3}{4} & \frac{1}{8} & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{pmatrix} = \begin{pmatrix} \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \frac{1}{4} & \frac{1}{2} & \frac{1}{4} & 0 & 0 & \cdot \\ \cdot & 0 & \frac{1}{4} & \frac{1}{2} & \frac{1}{4} & 0 & \cdot \\ \cdot & 0 & 0 & \frac{1}{4} & \frac{1}{2} & \frac{1}{4} & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \end{pmatrix} \cdot \begin{pmatrix} \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & 1 & 0 & 0 & \cdot \\ \cdot & \frac{1}{2} & \frac{1}{2} & 0 & \cdot \\ \cdot & 0 & 1 & 0 & \cdot \\ \cdot & 0 & \frac{1}{2} & \frac{1}{2} & \cdot \\ \cdot & 0 & 0 & 1 & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \end{pmatrix}.$$

Note that application of the mask $(\frac{1}{4} \quad \frac{1}{2} \quad \frac{1}{4})$ has a simple geometric interpretation: reposition a vertex at the midpoint of the midpoints of the two segments that contain the vertex. In the bivariate case, bi-cubic subdivision is equivalent to bi-linear subdivision followed by averaging with the tensor product of the univariate mask with itself. Plotted in 2D, the coefficients of this averaging mask have the form $\begin{pmatrix} \frac{1}{16} & \frac{1}{8} & \frac{1}{16} \\ \frac{1}{8} & \frac{1}{4} & \frac{1}{8} \\ \frac{1}{16} & \frac{1}{8} & \frac{1}{16} \end{pmatrix}$. Note that this mask can be decomposed into the sum of four submasks,

$$\begin{pmatrix} \frac{1}{16} & \frac{1}{8} & \frac{1}{16} \\ \frac{1}{8} & \frac{1}{4} & \frac{1}{8} \\ \frac{1}{16} & \frac{1}{8} & \frac{1}{16} \end{pmatrix} = \frac{1}{4} \left(\begin{pmatrix} 0 & 0 & 0 \\ \frac{1}{4} & \frac{1}{4} & 0 \\ \frac{1}{4} & \frac{1}{4} & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{1}{4} & \frac{1}{4} \\ 0 & \frac{1}{4} & \frac{1}{4} \end{pmatrix} + \begin{pmatrix} 0 & \frac{1}{4} & \frac{1}{4} \\ 0 & \frac{1}{4} & \frac{1}{4} \\ 0 & 0 & 0 \end{pmatrix} + \begin{pmatrix} \frac{1}{4} & \frac{1}{4} & 0 \\ \frac{1}{4} & \frac{1}{4} & 0 \\ 0 & 0 & 0 \end{pmatrix} \right).$$

This mask has a geometric interpretation analogous to the univariate case: reposition a vertex at the centroid of the centroids of the four quadrilaterals that contain the vertex. The averaging mask for tri-cubic B-splines (the tensor of the mask $(\frac{1}{4} \ \frac{1}{2} \ \frac{1}{4})$ with itself three times) can be expressed decomposed in a similar manner. Again, this mask has a geometric interpretation analogous to the univariate case: reposition a vertex at the centroid of the centroids of the eight hexahedra that contain the vertex. These geometric interpretations of the tensor product rules lead to our general rule for smoothing a mesh of topological d -cubes,

Cell averaging: Given a vertex v , compute the centroids of those topological d -cubes that contain v . Reposition v at the centroid of these centroids.

Due to its simplicity, cell averaging can be implemented in a very straightforward manner with a minimal amount of topological computation. Given a mesh of the form $\{T, p_{old}\}$, first compute $val[v]$, the number of d -cubes in T that contain the vertex v . This quantity can easily be computed during a single pass through T . Next, initialize a table of new vertex positions p_{new} to be zero. Finally, make a second pass through T . For each d -cube in T , compute the centroid $cent$ of its vertices and update the position of each vertex v of the d -cube via:

$$p_{new}[v] += \frac{cent}{val[v]}. \quad (1)$$

Since there are exactly $val[v]$ d -cubes containing v , $p_{new}[v]$ accumulates the centroid of the $val[v]$ centroids.

Combining multi-linear subdivision with cell averaging yields a subdivision scheme that we refer to as multi-linear cell averaging (MLCA) subdivision. Figure 3 shows the effect of applying MLCA subdivision to an initial surface mesh consisting of six squares forming a cube. The top row shows the effect of bi-linear subdivision. The bottom row shows the effect of cell averaging. As we shall see, the resulting limit surface is smooth even at vertices of valence three.




Figure 3

2.3 Composite subdivision rules

In some situations (such as analyzing the smoothness of the scheme), it is preferable to have a vertex positioning rule that accounts for the effect of multi-linear subdivision and cell averaging simultaneously (as done in Catmull-Clark). We next derive such a rule for a mesh of d -cubes sharing a single common vertex (i.e. the subdivision rule at an extraordinary vertex). To aid in the derivation, we define several relevant functions for a mesh of topological d -cubes:

Let $\text{ring}[v_1, v_2, \dots]$ be the set of d -cubes containing all of the vertices v_1, v_2, \dots ,

Let $\text{val}[v_1, v_2, \dots]$ be the number of d -cubes in $\text{ring}[v_1, v_2, \dots]$,

Let $\text{dim}[v_1, v_2, \dots]$ be the dimension of the smallest facet containing all of v_1, v_2, \dots (e.g. 0 if they coincide, 1 if they lie on a common edge, 2 if they lie on a common face, etc.)

Given a coarse (old) $\{T_{\text{old}}, P_{\text{old}}\}$, let $\{T_{\text{new}}, P_{\text{new}}\}$ be the fine (new) produced by MLCA subdivision. If v is a vertex of the old, then its position in the new is a weighted sum of the vertices u in $\text{ring}[v]$ on the old. In particular, we claim that the weight associated with u has the form:

$$\frac{1}{4^d \text{val}[v]} 3^{d-\dim[u,v]} \text{val}[u, v] \quad (2)$$

For a one-dimensional curve ($d = 1$), $\text{val}[v]$ is two for all vertices. Therefore, by equation 2, the subdivision at vertex of T_{01d} has the form $(\frac{3^{1-1} * 1}{4^{1-2}}, \frac{3^{1-0} * 2}{4^{1-2}}, \frac{3^{1-1} * 1}{4^{1-2}})$ which is equal to $(\frac{1}{8}, \frac{3}{4}, \frac{1}{8})$. In the surface case ($d = 2$), the subdivision rule at an extraordinary vertex of valence n has the weights shown in figure 4. One interesting feature of the equation 2 is that the weight assigned to the central vertex v by this rule is always $(\frac{3}{4})^d$ since $\text{val}[v] = \text{val}[v, v]$ and $\dim[v, v] = 0$.




Figure 4

We conclude this section by proving the correctness of equation 2. To this end, let C be a cell of T_{new} containing v and whose parent also contains u . After multi-linear subdivision, but prior to cell averaging, the vertices of C involve weighted combinations of u due to repeated linear interpolation. If u and v share a facet of dimension $\dim[u, v]$ in T_{01d} , then these weights appear on a complementary sub-facet of C of dimension $(d - \dim[u, v])$. The sum of these weights in u is given by the expression

$$\sum_{i=0}^{d-\dim[u,v]} \binom{d-\dim[u,v]}{i} 2^{i-d}. \quad (3)$$

where $\binom{a}{b}$ is the binomial coefficient $\frac{a!}{b!(a-b)!}$. This expression can be simplified via *Mathematica* to $2^{-d} 3^{d-\dim[u,v]}$. Thus, the total contribution of u to the centroid of C during cell averaging is $2^{-d} 3^{d-\dim[u,v]}$ divided by the number of vertices in the cell 2^d . Thus, the total contribution of u to the new position of v from the cell C is $\frac{3^{d-\dim[u,v]}}{4^d \text{val}[v]}$. Since there are $\text{val}[u, v]$ cells that contain u and v , the total contribution over all cells must be as in equation 2. Since both

multi-linear subdivision and cell averaging reproduce constants, the weights of equation 2 must form a partition of unity, i.e.

$$\sum_{u \in \text{ring}[v]} 3^{d-\dim[u,v]} \text{val}[u, v] = 4^d \text{val}[v]. \quad (4)$$

3 Extensions

3.1 Non-manifold meshes

One particular advantage of MLCA subdivision is that it makes no restrictions on the local topology of the mesh. In particular, the scheme can be applied without change to meshes with non-manifold topology. In the univariate case, non-manifold topology gives rise to curve networks with $n > 2$ segments meeting at a vertex. The vertex positioning rule at such a non-manifold vertex takes $\frac{3}{4}$ of the original vertex plus $\frac{1}{4n}$ times each of its n neighbors. This rule is very similar to the variational rule for curve networks developed in Warren and Weimer [18]. Figure 5 show an example of MLCAaveraging subdivision applied to an butterfly shaped curve network. Figure 6 shows an example of an MLCA subdivision applied to a quadrilateral mesh consisting of 12 squares. The final surface mesh inherits the non-manifold curve present in the initial mesh.




Figure 5




Figure 6

3.2 Boundaries and embedded creases

Hoppe et al. [5] introduced a useful variant of the Catmull-Clark subdivision rules that allowed the introduction of B-spline "crease" curves on a Catmull-Clark surface. The benefit of these crease curves were two-fold: they allowed the controlled introduction of tangent plane discontinuities on a smooth Catmull-Clark surface and they allowed for direct control of boundary curves. By making a small modification to the rule for cell averaging, MLCA subdivision can be adjusted to introduce "creases" of any dimension into a smooth limit shape.

To introduce creases, we allow the topology T of a mesh to contain topological cubes of differing dimensions. The topology list T_0 for a mesh of quadrilaterals may also contain line segments and vertices. These lower dimensional cells define crease curves and crease vertices in the resulting surface mesh. For example, the initial topology T_0 for the bounded quad surface patch of figure 1 consists of 4 faces, 8 boundary edges and 4 corner vertices. These extra cells were added to the mesh to ensure that the boundary curves of the limit surface were defined by the boundary of the initial mesh. If the

vertices in this 3×3 grid are numbered $\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix}$, then the initial topology T_0 is a list of the form

$$\begin{aligned} & \{\{1, 2, 5, 4\}, \{4, 5, 8, 7\}, \{2, 3, 6, 5\}, \{5, 6, 9, 8\}, \\ & \{1, 4\}, \{4, 7\}, \{1, 2\}, \{2, 3\}, \{3, 6\}, \{6, 9\}, \{7, 8\}, \{8, 9\}, \\ & \{1\}, \{3\}, \{7\}, \{9\}\}. \end{aligned}$$

Given this new representation, multi-linear subdivision is performed as usual with each d -cube split into 2^d subcubes. However, cell averaging is adjusted as follows: for each vertex v , we compute $\dim[v]$, the dimension of the lowest dimension cell in T that contains v . In the example above, $\dim[v]$ has the values

$$\begin{array}{rcccccccc} v & : & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ \dim[v] & : & 0 & 1 & 0 & 1 & 2 & 1 & 0 & 1 & 0 \end{array}$$

$\dim[v]$ as well as $\text{val}[v]$ (now the number of cells containing v of dimension $\dim[v]$) can be computed in a single pass through T . Given these two quantities, the averaging rule of equation 1 can be reused with the following adjustment:

Cell averaging with creases : Given a vertex v , compute the centroids of those cells of dimension $\dim[v]$ that contain v . Reposition v at the centroid of these centroids.

This modification ensures that the subdivision rule for vertices on a crease is influenced only by those vertices that give rise to the crease. For example, if T contains a topological 0-cube (i.e. a vertex), cell averaging simply leaves the position of the vertex unchanged and MLCA subdivision interpolates the vertex. Figure 1 is an example of MLCA subdivision applied to a mesh incorporating creases along its boundaries. Note that boundaries of this patch are themselves B-spline curves. Figure 7 shows a cross-section of figure 3. The surface mesh in figure 3 is actually the boundary for a volume mesh starting from a single cube.




Figure 7

Figure 8 shows two examples of objects modeled using MLCA subdivision. Each mesh is a volume mesh whose boundary is made to conform to the shape of a ring and an axe, respectively, using crease surface, crease edges and crease vertices.




Figure 8

A ring and an axe modeled using MLCA subdivision with creases. (Thanks to Scott Schaefer for his help with this figure)

Finally, we note that the ability to define smooth volume meshes with embedded crease surface should be particularly useful. 3D application such as geological modeling require precise representation of embedded faults and fractures. We intend to pursue applications of MLCA subdivision in this area.

3.3 Non-hexahedral initial meshes

Any polygonal surface mesh can be converted into a quadrilateral mesh by splitting each face n -gon into n quadrilateral faces. Each of these n quadrilaterals consist of a vertex, the midpoints of its two incident edges and the centroid of the polygon. Figure 9 shows three examples of this split. Ideally, a similar method would allow one to split any convex polyhedron into a mesh of hexahedra. Unfortunately, such a split is possible only for those polyhedra for which all vertices have valence three. Given such a polyhedra with n vertices, it can be split into n hexahedra where each hexahedra consists of a vertex, 3 adjacent edge midpoints, 3 adjacent face centers and the centroid of the original polyhedron. Note that the split on the faces of the polyhedron is exactly the 2D split for a polygon. Figure 10 shows two examples of this 3D split for a tetrahedron and a triangular prism.





Figure 9

Given this observation, we suggest two possible approaches to converting a polyhedral mesh to a hexahedral mesh. Since all the vertices of a tetrahedron have valence three, one approach is to triangulate a non-hexahedral volume mesh and then split each tetrahedron into four hexahedra. The main drawback of this approach is that it introduces an extraordinary vertex of valence 4 into the center of every tetrahedron and destroys any kind of tensor product structure that might have existed in the initial mesh. Given a set of points in space, a second approach is to construct the 3D Voronoi diagram for the point set. In general position, the cells of this diagram are convex polyhedra whose vertices have valence three. Each of these polyhedra can then be partitioned into a collection of hexahedra as described above.



Figure

10

We are also considering other approaches for generating initial hexahedral meshes from point data (see [4], [9], [14] and [16] for examples of such methods). Our results will be summarized in a forthcoming paper [1].

4 Smoothness analysis

The MLCA subdivision scheme is applicable to meshes of both quadrilaterals in two dimensions and meshes of hexahedra in three dimensions. In this section, we analyze the smoothness of MLCA scheme in each case. In two dimension, we show that the scheme produces limit surfaces that are smooth everywhere. In three dimension, we prove that MLCA subdivision converges to a smooth limit volume along extraordinary edges of the mesh. At extraordinary vertices of the 3D mesh, we examine the characteristic map for MLCA subdivision and provide strong evidence that MLCA subdivision most likely converges to smooth limit volume at these vertices. The key to this analysis is the spectral technique described in Reif [13], Prutzsch [12], Peters and Reif [11] and Warren [17].

4.1 The surface case

For quadrilateral meshes, topological subdivision produces quadrilateral meshes that are locally tensor product almost everywhere. In particular, the meshes are locally tensor product everywhere except at extraordinary points (i.e. those vertices of the original mesh whose valence is not four). Since MLCA subdivision was designed to reproduce the subdivision rule for bi-cubic B-splines on tensor product meshes, the limit surfaces produced by the scheme are C^2 everywhere except at these extraordinary points. Thus, our task is to determine the smoothness of MLCA subdivision at an extraordinary vertices. Luckily, much of the analysis has been done in previous work by Reif and Peters [11] on the smoothness of Catmull-Clark subdivision at an extraordinary vertex v .

Consider an extraordinary point v shared by n quadrilaterals. In the neighborhood of v , repeated topological subdivision yields a topological mesh T consisting of n quadrants of a tensor product mesh surrounding v (see figure 11). In this framework, the behavior of the subdivision scheme at v is characterized by the spectral structure of the subdivision matrix S associated with T . For a valence five extraordinary vertex v , the matrix S (restricted to $\text{ring}[v]$) has the form:

$$\begin{pmatrix} 1 - na & 6a & 6a & 6a & 6a & 6a & a & a & a & a & a \\ \frac{3}{8} & \frac{3}{8} & \frac{1}{16} & 0 & 0 & \frac{1}{16} & \frac{1}{16} & 0 & 0 & 0 & \frac{1}{16} \\ \frac{3}{8} & \frac{1}{16} & \frac{3}{8} & \frac{1}{16} & 0 & 0 & \frac{1}{16} & \frac{1}{16} & 0 & 0 & 0 \\ \frac{3}{8} & 0 & \frac{1}{16} & \frac{3}{8} & \frac{1}{16} & 0 & 0 & \frac{1}{16} & \frac{1}{16} & 0 & 0 \\ \frac{3}{8} & 0 & 0 & \frac{1}{16} & \frac{3}{8} & \frac{1}{16} & 0 & 0 & \frac{1}{16} & \frac{1}{16} & 0 \\ \frac{3}{8} & \frac{1}{16} & 0 & 0 & \frac{1}{16} & \frac{3}{8} & 0 & 0 & 0 & \frac{1}{16} & \frac{1}{16} \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & \frac{1}{4} & 0 & 0 & 0 & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & \frac{1}{4} & 0 & 0 & 0 \\ \frac{1}{4} & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & \frac{1}{4} & 0 & 0 \\ \frac{1}{4} & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & \frac{1}{4} & 0 \\ \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & \frac{1}{4} & 0 & 0 & 0 & 0 & \frac{1}{4} \end{pmatrix} \quad (5)$$

Note that we have not completely specified the position rule at the extraordinary vertex. The variable a is free and can be chosen as a function of n .

If the eigenvalues λ_i of S and their corresponding eigenvectors z_i are indexed in order of decreasing magnitude, a necessary condition for the scheme to smooth at extraordinary vertex v is that these eigenvalues must satisfy $1 = \lambda_0 > |\lambda_1| \geq |\lambda_2| > |\lambda_3| \geq \dots$ (see [12]). Based on the block circulant structure of S , the eigenvalues and eigenvectors of S can be computed directly using the discrete Fourier transform. All but two of the eigenvalues of S (and

their corresponding eigenvectors) are independent of the choice of a . The remaining two eigenvalues (and their corresponding eigenvectors) depend on the choice of a and are eigenvalues of the reduced matrix

$$\begin{pmatrix} 1 - 7na & 6na & na \\ \frac{3}{8} & \frac{1}{2} & \frac{1}{8} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \end{pmatrix}. \quad (6)$$

Now, Catmull-Clark suggests a vertex rule with $a = \frac{1}{4n^2}$. Using some basic calculus, Reif and Peters [11] show that the eigenvalues of the resulting S are real and satisfy $1 > \lambda_1 = \lambda_2 > \lambda_3 \dots$ where $\lambda_2 > \frac{1}{4}$. On the other hand, the subdivision matrix S for MLCA subdivision has a vertex rule with $a = \frac{1}{16n}$. For this choice of a , the two eigenvalues of the reduced matrix 6 are $\frac{1}{4}$ and $\frac{1}{16}$. Since the remaining eigenvalues of S are unchanged, the spectrum of S continues to satisfy $1 > \lambda_1 = \lambda_2 > \lambda_3 \dots$ for MLCA subdivision.

Unfortunately, this condition alone is not sufficient to guarantee that the scheme is smooth at the extraordinary vertex v . For schemes with a spectrum of the form $\lambda_0 = 1 > |\lambda_1| \geq |\lambda_2| > \dots$, the eigenvectors z_1 and z_2 corresponding to the subdominant eigenvalues λ_1 and λ_2 determine whether the limit surface forms a manifold in the neighborhood of v . Given an initial mesh of the form $\{T, p_0\}$ with $p_0 = \{z_1, z_2\}^T$, Reif [13] defines the **characteristic map** ψ associated with the eigenvectors z_1 and z_2 to be the limit of the subdivision process $p_k = S p_{k-1}$ (see figure 11). Now, if the characteristic map ψ is **regular** (i.e. it is 1-1 and onto everywhere), then subdivision scheme associated with S is guaranteed to be smooth at v . For the subdivision matrix of equation 5, Reif and Peters [11] show that the characteristic map is regular for all valences $n \geq 3$. In other words, Catmull-Clark subdivision and MLCA subdivision share the same characteristic map. Thus, MLCA subdivision produces smooth limit surfaces at extraordinary vertices.

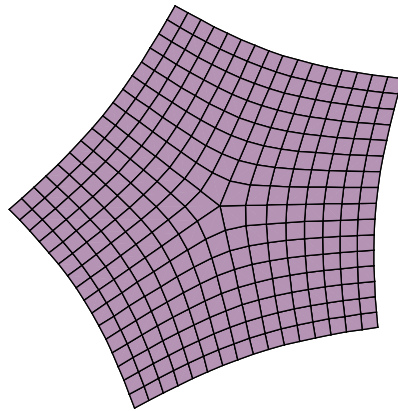


Figure 11

4.2 The volume case

To analyze the smoothness of MLCA subdivision on hexahedral meshes, we first consider the effect of topological subdivision on a hexahedral mesh. Repeated subdivision of a single hexahedron produces a tensor product volume mesh on the interior of the hexahedra. This tensor product structures extends across the faces of the hexahedra to face adjacent hexahedra. Since MLCA subdivision reproduces the subdivision rule for tri-cubic B-splines on tensor product meshes, the limiting volumes are C^2 on both the interiors of the hexahedra and the interiors of their faces.

For n hexahedra sharing a common edge, repeated topological subdivision yields a volume mesh whose topology is the tensor product of uniform quad mesh surrounding an extraordinary vertex of a valence n extraordinary point and a curve mesh. Figure 12 shows local examples of such meshes for $n = 3$ and $n = 5$. Using equation 2, one can verify that the subdivision rules produced by MLCA along this "extraordinary" edge are the tensor product of the bivariate MLCA rule for a valence n vertex and the univariate rule for cubic B-splines. Since these rules deliver C^1 surfaces and C^2 curves respectively, the tensor product rules produces C^1 volumes. *Note that volume scheme of Joy and MacCracken does not have this tensor product property.* As a result, analyzing the smoothness of Joy and MacCracken scheme along extraordinary edges appears to be impossible with current analysis techniques.




Figure 12


All that remains is to analyze the smoothness of MLCA subdivision at extraordinary vertices of the original mesh. Given a vertex v of the original hexahedral mesh, repeated topological subdivision produces an infinite topological mesh T that is invariant under topological subdivision and whose structure is determined by $\text{ring}[v]$. As in the 2D case, the spectral structure of the subdivision matrix S associated with T governs the behavior of the subdivision scheme at v . In particular, a similar set of conditions on the eigenvalues and eigenvectors of S is sufficient to show that the scheme is smooth at v . These conditions can be summarized as follows:

1. Let the eigenvalues λ_i of S be indexed in order of decreasing magnitude. These eigenvalues must satisfy $1 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \lambda_3 > \lambda_4 \dots$

2. Let z_1, z_2 and z_3 be the (infinite) eigenvectors associated with eigenvalues λ_1, λ_2 and λ_3 , respectively. The characteristic map ψ defined as the limit of the subdivision process on $\{\mathbb{T}, \{z_1, z_2, z_3\}^{\mathbb{T}}\}$ must be regular.

Unfortunately, in three dimensions, the topological mesh \mathbb{T} lacks the rotational symmetries that made it possible to use the discrete Fourier transform to compute the spectral structure of S in two dimensions. In general, the spectral structure of S depends on the topological structure of $\text{ring}[\mathbb{v}]$ in a way that appears to defy symbolic computation. Instead, we take a brute force approach and simply enumerate possible topologies for $\text{ring}[\mathbb{v}]$. For each distinct topology, we compute the eigenvalues and eigenvectors of S restricted to the two-ring of \mathbb{v} and attempt to verify the two conditions above.

In practice, there are two limiting factors on this approach. First, the number of distinct topologies \mathbb{T} grows exponentially in $\text{val}[\mathbb{v}]$. Thus, we analyzed every distinct topology \mathbb{T} for which $\text{val}[\mathbb{v}] \leq 12$ and also analyzed randomly generated topologies with values of $\text{val}[\mathbb{v}]$ of up to 50. Second, determining whether the characteristic map associated with the eigenvectors z_1, z_2 and z_3 is regular is a difficult computation. In practice, we restricted these eigenvectors to the two-ring of \mathbb{v} and tested whether the resulting hexahedral mesh is regular under multi-linear interpolation (as opposed to MLCA subdivision). Note this test is not conclusive and, in theory, could report that the mesh is regular when, in fact, it is not. Developing a precise test for regularity is possible (although complicated by the presence of extraordinary edges). This topic will be the subject of a future paper by the authors. Figure 13 shows the two-rings of the characteristic map for nine distinct low valence topologies (i.e. all configurations for which the extraordinary vertex \mathbb{v} has 7 or fewer edge adjacent neighbors).



Figure

13

For all of the numerous topologies tested, the corresponding subdivision matrix S satisfied the eigenvalue condition $1 = \lambda_0 > \lambda_1 \geq \lambda_2 \geq \lambda_3 > \lambda_4$ and had a characteristic map whose two-ring was regular under multi-linear interpolation. We view this as strong evidence that MLCA produces a smooth volume at an extraordinary vertex v for all possible one-rings. We hypothesize that this behavior is due to the use of cell averaging in our scheme. This averaging reduces the elastic energy of the volume mesh and thus should tend to avoid introducing the type of self-folds that cause the characteristic map to be non-regular. Again, this topic is a subject for further research.

5 Conclusions

MLCA subdivision has many attractive features. It is fast, simple to implement, handles non-manifold meshes and easily incorporates boundaries/creases. Our intention is to explore both the theory and application of MLCA subdivision in future work. Potential application domains include 3D mesh modeling for physical simulations, representing smooth implicit surfaces and constructing smooth deformations for volume textures. We are also currently in the process of implementing an exact test for regularity of meshes created through MLCA subdivision. This test can be used to answer the question of smoothness at extraordinary vertices posed in the previous section as well as for mesh modeling applications where regularity is important.

6 References

- [1] Bajaj, C., Warren, J. and Xu, G.: *Adaptive Hexahedral Finite Element Meshing of Solids with Embedded Particles*, manuscript, 2001.
- [2] Catmull, E. and Clark, J.: *Recursively generated B-spline surfaces on arbitrary topological meshes*. Computer Aided Design, vol. 16, no. 6, pp. 350—355, 1978.
- [3] Chaikin, G.: *An Algorithm for High Speed Curve Generation*. Computer Graphics and Image Processing, 3, pp. 346-349, 1974.
- [4] Eppstein, D.: *Linear Complexity Hexahedral Mesh Generation*, Proceedings of the Twelfth Annual Symposium on Computational Geometry, pp. 58—67, 1996.
- [5] Hoppe, Hugues and DeRose, Tony and Duchamp, Tom and Halstead, Mark and Jin, Hubert and McDonald, John and Schweitzer, Jean and Stuetzle, Werner: *Piecewise smooth surface reconstruction*. In SIGGRAPH94 proceedings, annual conference series, ACM press, pp. 295—302, 1994.
- [6] Lane, J. M. and Riesenfeld, R. F.: *A theoretical development for the computer generation and display of piecewise polynomial functions*. Transactions on Pattern Analysis and Machine Intelligence, vol. 2, no. 1, pp. 35—46, 1980.

-
- [7] Loop, Charles T.: *Smooth subdivision surfaces based on triangles*. Master's Thesis, Department of Mathematics, University of Utah, August 1987.
- [8] MacCracken, R. and Joy, K.: *Free-form deformations of solid primitives with constraints*. Proceedings of SIGGRAPH 1996, Computer Graphics Proceedings, Annual Conference Series, pp. 181—188 (July 1996). ACM Press / ACM SIGGRAPH / Addison Wesley Longman. 1996.
- [9] Mitchell, S.: *Hexahedral mesh generation via the dual*, 11th ACM Symposium on Computational Geometry, pp. C4—C5, 1995.
- [10] Morin, Géraldine, Warren, Joe and Weimer, Henrik: *A Subdivision Scheme for Surfaces of Revolution*, to appear in a special issue on subdivision of Computer Aided Geometric Design.
- [11] Peters, Jorg and Reif, Ulrich: *Analysis of Generalized B-Spline Subdivision Algorithms*. SIAM Journal of Numerical Analysis 35, pp. 728—748, 1998.
- [12] Prautzsch, Hartmut: *Smoothness of subdivision surfaces at extraordinary points*. Advances in Computational Mathematics, 9, pp. 377—389, 1998.
- [13] Reif, Ulrich: *A unified approach to subdivision algorithms near extraordinary points*. Computer Aided Geometric Design 12, pp. 153—174, 1995.
- [14] Schneiders, R. and Bunten, R., Automatic Generation of Hexahedral Finite Element Meshes, Computer Aided Geometric Design, vol. 12, pp. 693—707, 1995.
- [15] Stam, Jos: *Subdivision as repeated averaging*. Submitted to CAGD special issue on subdivision, 2001.
- [16] Tautges, T., Blacker, T. and Mitchell, S.: *The Whisker Weaving Algorithm: A Connectivity-Based Method for Constructing All-Hexahedral Finite Element Meshes*, International Journal for Numerical Methods in Engineering, vol. 39, pp. 3327-3349, 1996.
- [17] Warren, Joe: *Subdivision Methods for Geometric Design*, Manuscript available online at <http://www.cs.rice.edu/~jwarren/>, 1995.
- [18] Weimer, Henrik and Warren, Joe: *Subdivision Schemes for Thin Plate Splines*. Proceedings of Eurographics 1998, Computer Graphics Forum, vol. 17, no. 3, pp. 303—313 & 392, 1998.

[19] Zorin, Denis, Schröder, Peter and Sweldens, Wim: *Interpolating subdivision for meshes with arbitrary topology*. In SIGGRAPH96 proceedings, annual conference series, ACM Press, pp. 189—192, 1996.

[20] Zorin, Denis and Schröder, Peter: *A unified framework for primal/dual quadrilateral subdivision schemes*. To appear in CAGD, 2001.