Non-Uniform Recursive Subdivision Surfaces

Thomas W. Sederberg[1]
Jianmin Zheng[2]
Brigham Young University

David Sewell[3]
Sewell Development
Provo, Utah

Malcolm Sabin[4]
Numerical Geometry Ltd.
Cambridge, UK

Abstract

Doo-Sabin and Catmull-Clark subdivision surfaces are based on the notion of repeated knot insertion of uniform tensor product B-spline surfaces. This paper develops rules for non-uniform Doo-Sabin and Catmull-Clark surfaces that generalize non-uniform tensor product B-spline surfaces to arbitrary topologies. This added flexibility allows, among other things, the natural introduction of features such as cusps, creases, and darts, while elsewhere maintaining the same order of continuity as their uniform counterparts.

Categories and Subject Descriptors: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling–surfaces and object representations.

Additional Key Words and Phrases: B-splines, Doo-Sabin surfaces, Catmull-Clark surfaces.

1 INTRODUCTION

Tensor product non-uniform rational B-spline surfaces have become an industry standard in computer graphics, as well as in CAD/CAM systems. Because surfaces of arbitrary topological genus cannot be represented using a single B-spline surface, there has been considerable interest in the generalization, based on knot insertion, called 'recursive subdivision,' which removes this limitation.

However, despite being based on knot insertion, the recursive subdivision techniques published so far are the analogues of equal interval, *uniform* B-splines rather than of *non-uniform* B-splines.

[1] tom@byu.edu
[2] zheng@cs.byu.edu (On leave from Zhejiang University)
[3] dave@sewelld.com
[4] malcolm@geometry.demon.co.uk

This paper explores the possibility of achieving the extra flexibility of unequal knot intervals in a recursive subdivision scheme including, for example, the ability to express features such as creases and darts by simply setting some of the knot intervals to zero. Schemes are presented for achieving non-uniform Doo-Sabin and Catmull-Clark surfaces. We will refer to these collectively as Non-Uniform Recursive Subdivision Surfaces (NURSSes).

Figure 1 (left) shows a Doo-Sabin surface, and Figure 1 (right) shows an example of a non-uniform Doo-Sabin surface in which the knot spacings along certain control edges have been set to zero (as labeled), thereby creating a G^0 discontinuity along the oval edge on the left. Figure 2 shows
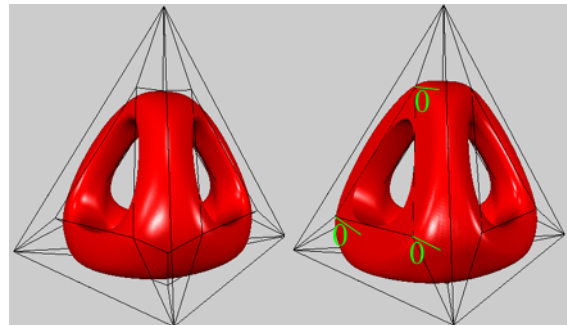

Figure 1: Uniform and non-uniform Doo-Sabin surfaces.

two non-uniform Catmull-Clark surfaces. The one on the left contains a dart formed by setting two pairs of control-edge knot spacings to zero. The one on the right shows shape modification induced by changing the knot spacing along the top edges to 10 and along the center horizontal edges to 0.1. The control net used here has the topology of a B-spline control net, but these shapes cannot be obtained using NURBS or uniform Catmull-Clark surfaces.

1.1 Background

The concept of image space subdivision as a graphics technique had been around for a long time when recursive subdivision appeared as an object definition technique. The first
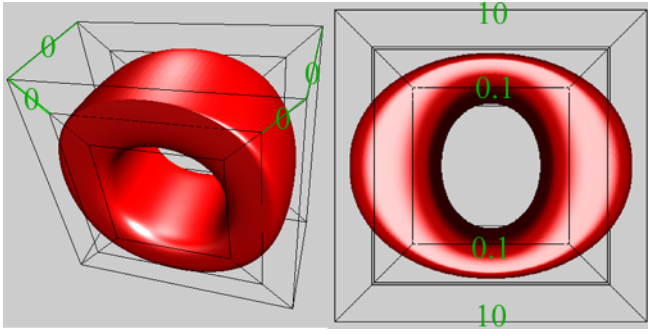
Figure 2: Non-uniform Catmull-Clark surfaces.

relevant result in parametric space subdivision of sculptured surfaces was the de Casteljau algorithm, which both evaluated a point on a Bézier curve and provided the control points for the parts of the curve meeting there. This was generalized to B-splines in the form of the Oslo algorithm [6] and Boehm subdivision [3] and used as a basis for interrogation methods applying parametric space subdivision.

However, what sparked the imagination of the graphics and modeling communities in 1975 was a much more specific subdivision of a quadratic B-spline, proposed by Chaikin as a curve rendering technique [5], and recognized for what it was by Forrest [11] and by Riesenfeld [26]. It later turned out that the concept of a curve being the limit of a polygon under the operation of cutting off the corners had been explored by de Rham in the 1940s and 50s. His results were translated into modern terminology by de Boor [7].

It was quickly appreciated that the curve ideas could give surface techniques just by applying the concept of tensor product, but the important key concept, that subdivision could overcome the rigid rectangular partitioning of the parametric domain — one of the major limitations of tensor products — was reached more or less simultaneously by Catmull and Clark [4] and by Doo and Sabin [8]. Since then there have been five major directions of development:

1. The analysis of what happens near an extraordinary point, started in the Doo-Sabin paper [8], was taken up by Ball and Storry. This led to an optimization of the coefficients for the cubic case [1, 2] which unfortunately missed one of the possible variations, and the task was completed by Sabin [28], in a paper which also identified that a cubic construction could never give full G^2 continuity at the singular points, and that continuity at such points was in fact a much more complicated question than had been assumed. Further analysis was carried out by Reif [25]. The nature of the behavior around the extraordinary points is now well understood.

2. Constructions based on box-splines, rather than on tensor products, were explored by Farin [10] and by Loop [15], and a collection of possible constructions was assembled by Sabin [27].

3. Constructions that interpolate the control points were explored by Dyn, Gregory and Levin [9], and an improved scheme was derived by Zorin, Schröder and Sweldens [31]. Kobbelt [14] proposed an alternative for quadrilateral nets with arbitrary topology. The simpler ones in this category can be viewed as duals of quadratic B-spline constructions.

4. Nasri [17] studied the problems of efficient implementation and practical edge-conditions and extended this to modifications of the basic technique to achieve various interpolation conditions [18, 19, 20]. Halstead, Kass, and DeRose showed that a fairness norm could be computed exactly for Catmull-Clark surfaces [12], enabling the determination of more fair limit surfaces.

5. The idea of using just a small number of subdivision steps, and then using n-sided combinations of patches to fill in a configuration made more regular in some sense by those steps, was explored by Loop [16], Peters [21, 22, 23] and Prautzsch [24]. Ball and Storry [29] took the opposite line, of using subdivision to define an n-sided patch.

What was not explored until now was that the general topology subdivision schemes were as rigid as the equal interval splines from which they were derived.

1.2 Overview

Section 2 reviews knot-doubling for non-uniform B-spline curves of degree two and three and introduces a simple approach for labeling the knot intervals on the control polygon— an idea that is crucial for the extension to subdivision surfaces. Section 3 then gives the corresponding expression for knot doubling of non-uniform tensor product B-spline surfaces. Section 4 proposes subdivision rules for non-uniform Doo-Sabin and Catmull-Clark surfaces, which reduce to non-uniform B-spline surfaces when the control net is a rectangular grid and when all knot intervals along every given row and column are the same. A continuity analysis is given in section 5, showing that non-uniform Doo-Sabin surfaces are G^1 and non-uniform Catmull-Clark surfaces are generally G^2, but G^1 at certain points. Section 6 makes some observations on NURSSes, and offers a conclusion.

2 CURVE KNOT DOUBLING

For a quadratic periodic B-spline curve, each vertex of the control polygon corresponds to a single quadratic curve segment. It is convenient then to express the knot vector by writing the knot interval d_i of each curve segment next to its corresponding control vertex $\mathbf{P}_i$. If a new knot is inserted at the midpoint of each current knot interval, the resulting control polygon has twice as many control points, and their

coordinates $\mathbf{Q}_k$ are:

$$\mathbf{Q}_{2i} = \frac{(d_i + 2d_{i+1})\mathbf{P}_i + d_i\mathbf{P}_{i+1}}{2(d_i + d_{i+1})}$$

$$\mathbf{Q}_{2i+1} = \frac{d_{i+1}\mathbf{P}_i + (2d_i + d_{i+1})\mathbf{P}_{i+1}}{2(d_i + d_{i+1})} \quad (1)$$
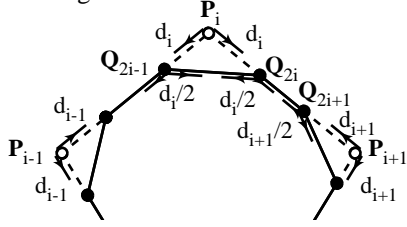
as illustrated in Figure 3.

Figure 3: Non-uniform quadratic B-spline curve.

For cubic periodic B-spline curves, each *edge* of the control polygon corresponds to a single cubic curve segment, and so we write the knot intervals adjacent to each edge of the control polygon. The equations for the new control points $\mathbf{Q}_k$ generated upon inserting a knot midway through each knot interval are:

$$\mathbf{Q}_{2i+1} = \frac{(d_i + 2d_{i+1})\mathbf{P}_i + (d_i + 2d_{i-1})\mathbf{P}_{i+1}}{2(d_{i-1} + d_i + d_{i+1})} \quad (2)$$

$$\mathbf{Q}_{2i} = \frac{d_i\mathbf{Q}_{2i-1} + (d_{i-1} + d_i)\mathbf{P}_i + d_{i-1}\mathbf{Q}_{2i+1}}{2(d_{i-1} + d_i)} \quad (3)$$
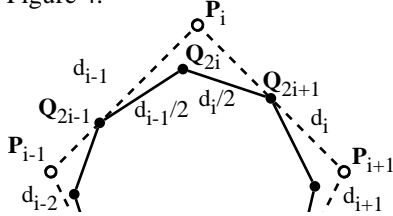
as shown in Figure 4.

Figure 4: Non-uniform cubic B-spline curve.

3 SURFACE KNOT DOUBLING

The knot-doubling formulae for B-spline curves extend easily to surfaces. Non-uniform B-spline surfaces are defined in terms of a control net that is topologically a rectangular grid, for which all horizontal knot vectors are scales of each other, and all vertical knot vectors are scales of each other.

3.1 Quadratic Case

The formulae for the new control points $\mathbf{F}_A$ can be written in Doo-Sabin form, which is significant because in this form the new control points are seen as being in groups, creating a new face in each old face, and the vertices of each such new face are in 1:1 correspondence with the vertices of the old,

whereas under the tensor product form we merely see all the new vertices as forming a new regular array (see Figure 5).

$$\mathbf{F}_A = \frac{\mathbf{V} + \mathbf{A}}{2} + \frac{ac(\mathbf{B} + \mathbf{C} - \mathbf{A} - \mathbf{D})}{4(ad + ac + bc + bd)}, \quad (4)$$

where

$$\mathbf{V} = \frac{bd\mathbf{A} + ad\mathbf{B} + bc\mathbf{C} + ac\mathbf{D}}{bd + ad + bc + ac}. \quad (5)$$
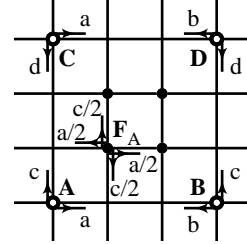

Figure 5: Knot doubling, quadratic B-spline.

3.2 Cubic Case

For non-uniform cubic B-spline surfaces, the refinement rules can be written as follows (see Figure 6). First, each face is replaced with a new vertex F_i. For example,

$$\mathbf{F}_1 = [(e_3 + 2e_4)(d_2 + 2d_1)\mathbf{P}_0 + (e_3 + 2e_4)(d_2 + 2d_3)\mathbf{P}_1$$

$$+ (e_3 + 2e_2)(d_2 + 2d_3)\mathbf{P}_5 + (e_3 + 2e_2)(d_2 + 2d_1)\mathbf{P}_2]$$

$$/[4(e_2 + e_3 + e_4)(d_1 + d_2 + d_3)]. \quad (6)$$

Then, each edge is split with an edge vertex E_i, e.g.

$$\mathbf{E}_1 = \frac{e_2\mathbf{F}_1 + e_3\mathbf{F}_4 + (e_2 + e_3)\mathbf{M}_1}{2(e_2 + e_3)}, \quad (7)$$

where

$$\mathbf{M}_1 = \frac{(2d_1 + d_2)\mathbf{P}_0 + (d_2 + 2d_3)\mathbf{P}_1}{2(d_1 + d_2 + d_3)}. \quad (8)$$

Finally, each original control point is replaced with a vertex point $\mathbf{V}$

$$\mathbf{V} = \frac{\mathbf{P}_0}{4} + \frac{d_3e_2\mathbf{F}_1 + d_2e_2\mathbf{F}_2 + d_2e_3\mathbf{F}_3 + d_3e_3\mathbf{F}_4}{4(d_2 + d_3)(e_2 + e_3)} +$$

$$[d_3(e_2 + e_3)\mathbf{M}_1 + e_2(d_2 + d_3)\mathbf{M}_2 + d_2(e_2 + e_3)\mathbf{M}_3$$

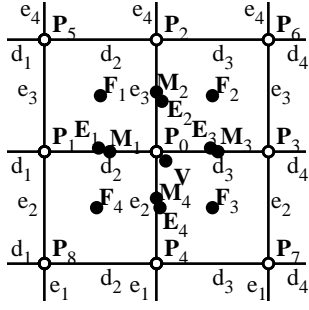$$+ e_3(d_2 + d_3)\mathbf{M}_4]/[4(d_2 + d_3)(e_2 + e_3)]. \quad (9)$$

Figure 6: Face, edge and vertex points.

4 NURSS REFINEMENT

In uniform B-spline surfaces, all knot spacings are the same. Doo-Sabin and Catmull-Clark proposed generalizations of uniform B-spline surface schemes that allow for vertices of the control mesh to have valence other than four, and the faces of the control mesh to have other than four sides. Their subdivision rules were designed such that when the control mesh happens to be a rectangular grid, the subdivision rules are equivalent to knot doubling of uniform B-spline surfaces. The subdivision surfaces are then defined as the limit of the control meshes when these subdivision rules are applied an infinite number of times.

We here define generalizations of non-uniform B-spline surfaces. As in the cubic curve case, each edge in the control polyhedron of a non-uniform Catmull-Clark surface is assigned a knot spacing. For a non-uniform Doo-Sabin surface, each vertex is assigned a knot spacing (possibly different) for each edge radiating from it. Our objective is to devise a set of refinement rules for NURSSes such that if all knot intervals are equal, the quadratic NURSS reduces to Doo-Sabin and the cubic NURSS reduces to Catmull-Clark. There are actually two distinct rules to be devised. First, we need to revise the Doo-Sabin and Catmull-Clark rules for the new point coordinates, taking the knot spacings into account. Second, we need rules for determining the new knot spacings.

Note that "NURSS" could just as well stand for "Non-Uniform *Rational* Subdivision Surfaces," because it is a simple matter to first project rational control points to 4-D, then apply our rules, and finally to project back to 3-D.

In this section, bold capital letters stand for points, and non-bold typeface for knot spacings. The indices for knot spacing d_{ij}^k indicate that the spacing pertains to an edge with $\mathbf{P}_i$ as one endpoint. Referring to Figure 8, the notation d_{ij}^0 indicates the knot spacing for edge $\mathbf{P}_i$–$\mathbf{P}_j$. Rotating counter-clockwise about $\mathbf{P}_i$, d_{ij}^1 denotes the knot spacing for the first edge encountered, d_{ij}^2 indicates that of the second edge, etc. For the cubic case, each edge has a single knot spacing, so $d_{ij}^0 = d_{ji}^0$.

4.1 Quadratic Case

In the quadratic case, refinement proceeds in a manner identical to Doo-Sabin subdivision: A polyhedron spawns a refined polyhedron for which new faces (of type F, type E and type V respectively) are created for each face, edge, and vertex of the previous polyhedron. During the subdivision step, each face is replaced by a new face connected across the old edges and across the old vertices by other new faces. In such refinement schemes, the extraordinary points are at the "center" of n-sided faces with $n \neq 4$. After one iteration, every vertex of the new polyhedron will have valence four, and the number of faces with other than four sides will remain constant.

Refer to Figure 7 for labels. The new vertex $\bar{\mathbf{P}}_i$ is computed:

$$\bar{\mathbf{P}}_i = \frac{\mathbf{V} + \mathbf{P}_i}{2} + (d_{i+1,i+2}^0 d_{i+3,i+2}^0 + d_{i-1,i-2}^0 d_{i-3,i-2}^0)$$

$$\times \frac{-n\mathbf{P}_i + \sum_{j=1}^n \left(1 + 2\cos\left(\frac{2\pi|i-j|}{n}\right)\right)\mathbf{P}_j}{8\sum_{k=1}^n d_{k-1,k}^0 d_{k+1,k}^0} \quad (10)$$

where

$$\mathbf{V} = \frac{\sum_{k=1}^n d_{k-1,k}^0 d_{k+1,k}^0 \mathbf{P}_k}{\sum_{k=1}^n d_{k-1,k}^0 d_{k+1,k}^0}.$$
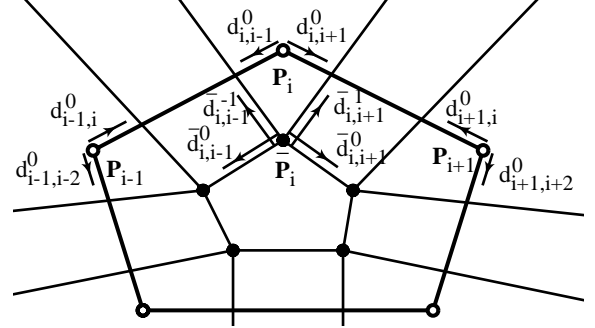

Figure 7: Quadratic refinement rules.

4.1.1 New knot spacings

New knot spacings $\bar{d}_{ij}^k$ can be specified in numerous ways. Here are two straightforward options:

$$\bar{d}_{i,i+1}^0 = \bar{d}_{i,i-1}^{-1} = d_{i,i+1}^0/2$$
$$\bar{d}_{i,i-1}^0 = \bar{d}_{i,i+1}^1 = d_{i,i-1}^0/2$$

or

$$\bar{d}_{i,i+1}^0 = d_{i,i+1}^0/2, \; \bar{d}_{i,i-1}^{-1} = (d_{i,i+1}^0 + d_{i,i-1}^{-1})/4$$

$$\bar{d}_{i,i-1}^0 = d_{i,i-1}^0/2, \; \bar{d}_{i,i+1}^1 = (d_{i,i-1}^0 + d_{i,i+1}^1)/4$$

The former allows the refinement matrix to remain constant after a few iterations. The latter seems to produce more satisfactory shapes.

4.2 Cubic Case

Our development parallels that for Catmull-Clark surfaces. As shown in Figure 8, the face point for a face with n sides is computed as

$$\mathbf{F} = \frac{\sum_{i=0}^{n-1} w_i \mathbf{P}_i}{\sum_{i=0}^{n-1} w_i},$$

(11)

where

$$w_i = (d_{i+1,i}^0 + d_{i+1,i}^2 + d_{i+1,i}^{-2} + d_{i-2,i-1}^0 + d_{i-2,i-1}^2 + d_{i-2,i-1}^{-2})$$
$$\times (d_{i-1,i}^0 + d_{i-1,i}^2 + d_{i-1,i}^{-2} + d_{i+2,i+1}^0 + d_{i+2,i+1}^2 + d_{i+2,i+1}^{-2})$$
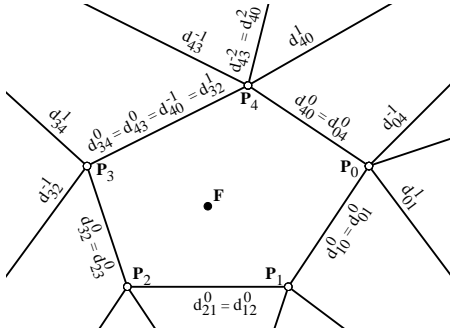
(12)

Figure 8: Face point.

The edge point is computed (see Figure 9a):

$$\mathbf{E} = (1 - \alpha_{ij} - \alpha_{ji})\mathbf{M} + \alpha_{ij}\mathbf{F}_{ij} + \alpha_{ji}\mathbf{F}_{ji},$$

(13)

where

$$\alpha_{ij} = \frac{d_{ji}^1 + d_{ij}^{-1}}{2(d_{ji}^1 + d_{ij}^{-1} + d_{ji}^{-1} + d_{ij}^1)}$$

(14)

if $d_{ji}^1 + d_{ij}^{-1} + d_{ji}^{-1} + d_{ij}^1 \neq 0$ and $\alpha_{ij} = 0$ otherwise.

$$\mathbf{M} = \frac{(d_{ji}^0 + d_{ji}^2 + d_{ji}^{-2})\mathbf{P}_i + (d_{ij}^0 + d_{ij}^2 + d_{ij}^{-2})\mathbf{P}_j}{d_{ji}^0 + d_{ji}^2 + d_{ji}^{-2} + d_{ij}^0 + d_{ij}^2 + d_{ij}^{-2}}$$
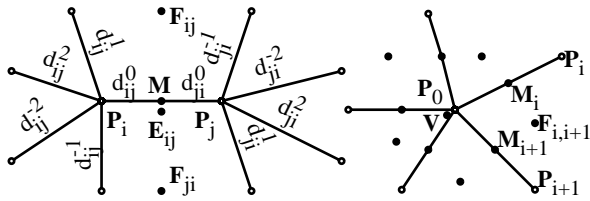
(15)

Figure 9: a) Edge point.   b) Vertex point.

The vertex point for a point of valence n is expressed (see Figure 9b):

$$\mathbf{V} = c\mathbf{P}_0 + \frac{3 \sum_{i=1}^{n} (m_i \mathbf{M}_i + f_{i,i+1} \mathbf{F}_{i,i+1})}{n \sum_{i=1}^{n} (m_i + f_{i,i+1})},$$

(16)

where $\mathbf{M}_i$ are defined as (15), $\mathbf{F}_{i,i+1}$ as (11), and

$$m_i = (d_{0i}^1 + d_{0i}^{-1})(d_{0i}^2 + d_{0i}^{-2})/2$$

(17)

$$f_{ij} = d_{0i}^1 d_{0j}^{-1}$$

(18)

$$c = \frac{n - 3}{n} \quad \text{if } \sum_{i=1}^{n} (m_i + f_{i,i+1}) \neq 0$$

(19)

otherwise, $c = 1$.

4.2.1 New knot spacings

Each n-sided face is split into n four sided faces, whose knot spacings are determined as shown in Figure 10.
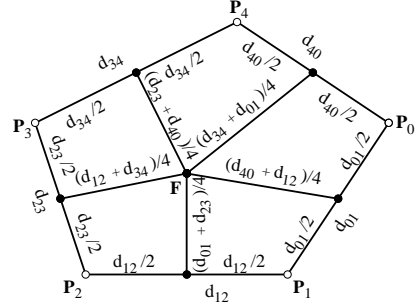

Figure 10: New knot spacings.

5 CONTINUITY ANALYSIS

For each construction, we consider the behavior of the limit surface for all the knot spacings being positive.

5.1 Quadratic case

5.1.1 Limit surface structure

After one iteration, all type V and type E faces are four-sided, with the property that the intervals crossing the original edges are equal as seen from the two sides. Thus each such type V face reduces to a mesh that is equivalent to a uniform biquadratic B-spline, while each type E face reduces to a mesh that is equivalent to a non-uniform biquadratic B-spline. This leaves only the regression in the type F faces to analyze for continuity.

The limit surface for a face consists of patches in the structure as shown in Figure 11a. The face is four-sided in this example.
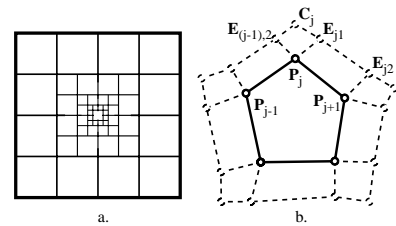

Figure 11: a) Sequence of quadratic polynomial pieces for a four-sided face; b) Configuration surrounding type F face.

5.1.2 Continuity at face-centers

It is convenient to consider four-sided faces alongside the more general n-sided faces. After at most two subdivisions, the configuration surrounding a type F face may be represented as in Figure 11b. In the center lies the face (of type F) $P_1 P_2 \ldots P_n$. Each edge of this face (for example $P_j P_{j+1}$) is adjacent to a four-sided face of type E with vertices $P_j P_{j+1} E_{j2} E_{j1}$. The neighborhood of each vertex P_j is completed by a four-sided face of type V with vertices $P_j E_{j1} C_j E_{(j-1)2}$.

Let the configuration around this type F face be represented by the vector of points

$$M = [P_1, \ldots, P_n, E_{11}, E_{12}, \ldots, E_{n1}, E_{n2}, C_1, \ldots, C_n]^T,$$

and $\bar{M}$ be the corresponding configuration after subdivision. Then $\bar{M} = S_n M$, where S_n is a $4n \times 4n$ matrix called the refinement matrix. Here only consider the first option for the new knot spacings in which case S_n remains constant through all the subsequent subdivision steps. Thus we can use the eigenstructure of S_n to analyze continuity.

We carried out an algebraic eigenanalysis for orders 3 to 8 based on the discrete Fourier transform technique in an exercise reported in [30]. This leads us to

Theorem 1 *For orders 3 to 8, if all knot spacings $d_{i,j} > 0$, then the refinement matrix S_n is not defective and its eigenvalues are*

$$\lambda_1 = 1 > \lambda_2 = \lambda_3 = \frac{1}{2} > |\lambda_4|, |\lambda_5|, \cdots, |\lambda_{4n}|.$$

By an argument similar to one used in [12] we can conclude that, provided that all the knot spacings are greater than zero, the limit surface generated by the non-uniform Doo-Sabin scheme is G^1 continuous both at all ordinary points and at extraordinary points of valence less than 9.

5.2 Cubic Case

5.2.1 Limit surface structure

After one subdivision, every face is four-sided, and after two more, the pattern of knot intervals over the group of 16 subfaces replacing each original face is as shown in Figure 12a. The h_i are in arithmetic sequence as are the v_i.

When these values are substituted into (11), (13) and (16), the positions of the new vertices in, or on the boundary of, the innermost four sub-faces are exactly the same as if all the horizontal intervals had been equal and all the vertical intervals equal likewise. The innermost four subfaces therefore converge towards uniform bicubic B-splines, giving a pattern of bicubic pieces as shown in Figure 12b, where the largest square represents the face of a non-uniform Catmull-Clark net with different knot intervals along each edge. The pattern of smaller squares shows schematically the infinite progression of Bézier patches that make up the limit surface. The interior of the limit surface of every such face is therefore G^2. We need only concern ourselves with the continuity at the edges and at the vertices, where there is a regression.
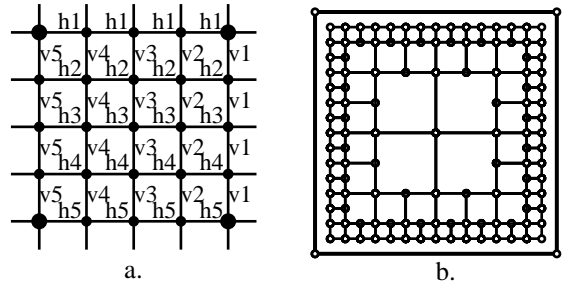

Figure 12: a) Knot intervals after three subdivisions; b) Sequence of cubic polynomial pieces.

5.2.2 Continuity at vertices

In this section we consider continuity at vertex points of valence ≥ 3 (i.e., exceptional points, as well as vertex points of valence four). Unfortunately, since the refinement matrix changes at each iteration, it is difficult to perform an eigenanalysis to determine if non-uniform Catmull-Clark surfaces are G^1 at vertex points, except for simple numerical cases. One of the few cases that yield a constant refinement matrix is the valence three vertex in Figure 13 (right). In this case, the second and the third eigenvalues are generally *different*.

To find out what is going on in the neighborhood of this point, we performed a numerical study, chosing a cube as a control polyhedron with various knot spacings. Figure 13
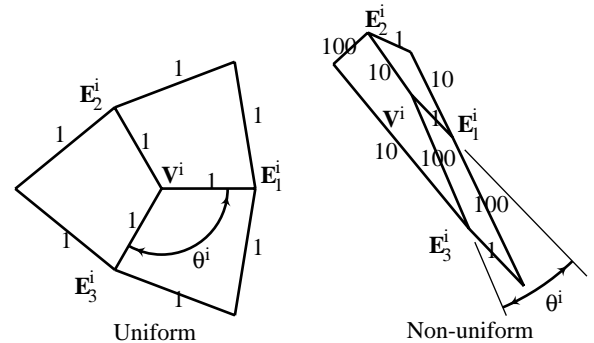

Figure 13: Neighborhood of valence 3 vertex point after 25 iterations.

shows the neighborhood of a vertex point $\mathbf{V}^i$ (that began as a corner of the cube) after 25 iterations. The figure on the left shows the uniform knot spacing; everything is symmetric as expected. The figure on the right came from the same vertex on the same cube, only with knot intervals of value 100 assigned to one set of four parallel edges on the cube; of value 10 assigned to another set of four parallel edges, and of 1 assigned to the remaining four edges, again after $i = 25$ iterations. Notice that the angles are no longer equal. In fact, it turns out that angle $\theta^i = \angle \mathbf{E}_1^i - \mathbf{V}^i - \mathbf{E}_3^i$ tends to zero as $i \to \infty$. However, the three faces become coplanar at a *much* faster rate, as shown in the following table. Here, $\Delta \mathbf{N}^i$ refers to the maximum angle between the planes $\mathbf{E}_1^i - \mathbf{V}^i - \mathbf{E}_2^i$, $\mathbf{E}_1^i - \mathbf{V}^i - \mathbf{E}_3^i$, and $\mathbf{E}_2^i - \mathbf{V}^i - \mathbf{E}_3^i$.

i	Uniform		Non-uniform	
	$\Delta\mathbf{N}^i$ π radians	θ^i π radians	$\Delta\mathbf{N}^i$ π radians	θ^i π radians
0	5×10^{-1}	0.5	5×10^{-1}	0.5
5	5×10^{-3}	0.6665	2×10^{-2}	0.37
25	7×10^{-10}	0.666667	5×10^{-8}	0.15
50	1×10^{-20}	0.666667	9×10^{-15}	0.043
75	2×10^{-30}	0.666667	1×10^{-21}	0.013
100	6×10^{-40}	0.666667	2×10^{-28}	0.004

The normals are becoming parallel at a rate that is roughly 10^7 times faster than the rate at which θ^i is approaching zero. After 100 iterations, the configuration is as close to G^1 as anyone could possibly have need for. The facets are 25 orders of magnitude smaller than they need to be for any practical use (five iterations are plenty for most graphics applications).

We also did a similar study on valence four, using widely varying knot spacings, and again observed a fast convergence of normal vectors, but *no* tendency of any face angles to tend to zero. Hence, we are confident that valence four points are G^1. Preliminary experiments with $n > 4$ indicate very similar behavior.

5.2.3 Continuity across edges

Across the interior of an edge, the situation may be regarded as a standard non-uniform B-spline with a perturbation due to the original variation between the knot intervals. This perturbation tends to zero with a convergence rate of $O\left(2^{-i}\right)$. Note that the non-uniform B-spline is C^2 and the surrounding vertices converge to a plane configuration with a rate of $O\left(4^{-i}\right)$. Therefore, in the limit, the non-uniform Catmull-Clark surface is G^1 across every edge.

6 DISCUSSION

Figure 14 shows the effect that knot spacing can have on a surface. In the grid at the left, all edges are assigned knot spacing of 1, except for the four edges labeled with a 0. Two steps of non-uniform Catmull-Clark subdivision result in the meshes shown (minus a few outer layers of quadrilaterals). This configuration of knot spacings causes the limit surface to interpolate the center point with G^0 continuity.
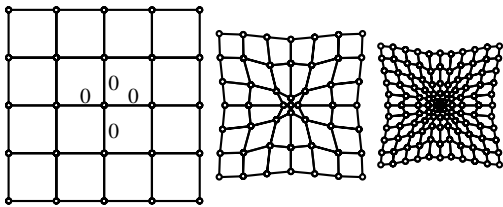

Figure 14: Effect of non-uniform knot spacing.

In the absence of non-uniform subdivision surfaces, Hoppe et. al. proposed a scheme for imposing features such as creases, corners, and darts on an otherwise G^1 subdivision surface that uses special-case "masks" [13]. NURSSes can provide for such features without the need for special masks. For example, Figure 1 shows a crease imposed on a Doo-Sabin surface by setting three knot spacings to zero. Figure 2 (left) shows a dart on a non-uniform Catmull-Clark surface, created by setting four knot intervals to zero.

Figures 16–19 show a variety of shapes that can be attained using NURSSes, but not using uniform subdivision surfaces. The initial control polyhedra are shown in wireframe. Sharp features can be imposed by setting to zero the knot spacing of appropriate edges.

Another use for knot spacing is in shape modification. Figure 2 (right) shows the effect of altering the knot spacing on several control polygon edges of a torus-shaped Catmull-Clark surface. Figure 17 shows a uniform Doo-Sabin surface (on the left) and two non-uniform counterparts, formed by choosing different knot spacings as shown. The sphere in Figure 18 cannot be expressed exactly using uniform subdivision surfaces, rational or otherwise. Figure 19c,e,h are other examples of non-uniform Catmull-Clark surfaces.

In summary,

- this method extends the known general topology methods by permitting unequal knot intervals, thus allowing a single surface description the strengths of both the standard non-uniform tensor products and the uniform recursive subdivision surfaces in one representation.

- Even in the situation where there are no extraordinary points, this theory extends current capability by giving a G^1 surface when the knot intervals are chosen individually for every edge in the control polygon, not constrained to support a tensor product structure.

- This scheme provides a lot of freedom to adjust the shape of the surface. In particular, it can model sharp features by properly setting certain knot spacings to zero.

Future work will design a convenient modeling interface for the interactive purpose and determine how to use knot spacing to best advantage.

ACKNOWLEDGEMENTS

REFERENCES

[1] A A Ball and D J T Storry. Recursively Generated B-spline Surfaces. *Proc. CAD84*, pages 112–119, 1984. ISBN 0408 01 4407.

[2] A A Ball and D J T Storry. Conditions For Tangent Plane Continuity Over Recursively Generated B-spline Surfaces. *ACM ToG*, 7:83–102, 1988.

[3] W Boehm. Inserting New Knots Into B-spline Curves. *Computer-Aided Design*, 12:199–201, 1980.

[4] E Catmull and J Clark. Recursively Generated B-spline Surfaces On Arbitrary Topological Meshes. *Computer-Aided Design*, 10:350–355, 1978.

[5] G Chaikin. An Algorithm For High-speed Curve Generation. *Computer Graphics and Image Processing*, 3:346–349, 1974.

[6] E Cohen, T Lyche, and R F Riesenfeld. Discrete B-splines And Subdivision Techniques In Computer Aided Design And Computer Graphics. *Computer Graphics and Image Processing*, 14:87–111, 1980.

[7] C de Boor. Cutting Corners Always Works. *Computer Aided Geometric Design*, 4:125–131, 1987.

[8] D Doo and M Sabin. Behaviour Of Recursive Division Surfaces Near Extraordinary Points. *Computer-Aided Design*, 10:356–360, 1978.

[9] N Dyn, D Levin, and J A Gregory. A 4-point Interpolatory Subdivision Scheme For Curve Design. *Computer Aided Geometric Design*, 4:257–268, 1987.

[10] G Farin. Designing C^1 Surfaces Consisting Of Triangular Cubic Patches. *Computer-Aided Design*, 14:253–256, 1982.

[11] A R Forrest. Notes On Chaikin's Algorithm. Technical Report Memo CGP74/1, University of East Anglia, Norwich, UK, 1974.

[12] M Halstead, M Kass, and T DeRose. Efficient, Fair Interpolation Using Catmull-Clark Surfaces. *Computer Graphics (SIGGRAPH 93 Conference Proceedings)*, 27:35–44, 1993.

[13] H Hoppe, T DeRose, T Duchamp, M Halstead, H Jin, J McDonald, J Schweitzer, and W Stuetzle. Piecewise Smooth Surface Reconstruction. *Computer Graphics (SIGGRAPH 94 Conference Proceedings)*, 28:295–302, 1994.

[14] L Kobbelt. Interpolatory Subdivision On Open Quadrilateral Nets With Arbitrary Topology. *Computer Graphics Forum (Eurographics 96)*, 1996.

[15] C Loop. *Smooth Subdivision Surfaces Based On Triangles*. Master's thesis, University of Utah, Dept. of Mathematics, 1987.

[16] C Loop. Smooth Spline Surfaces Over Irregular Meshes. *Computer Graphics (SIGGRAPH 94 Conference Proceedings)*, 28:303–310, 1994.

[17] A Nasri. *Polyhedral Subdivision Methods For Freeform Surfaces*. PhD thesis, University of East Anglia, 1984.

[18] A Nasri. Polyhedral Subdivision Methods For Freeform Surfaces. *ACM ToG*, 6:29–73, 1987.

[19] A Nasri. Surface Interpolation On Irregular Network With Normal Conditions. *Computer Aided Geometric Design*, 8:89–96, 1991.

[20] A Nasri. Curve Interpolation In Recursively Generated B-spline Surfaces Over Arbitrary Topology. *Computer Aided Geometric Design*, 14:13–30, 1997.

[21] J Peters. Joining Smooth Patches Around A Vertex To Form A C^k Surface. *Computer Aided Geometric Design*, 9:387–411, 1992.

[22] J Peters. Smooth Free-form Surfaces Over Irregular Meshes Generalizing Quadratic Splines. *Computer Aided Geometric Design*, 10:347–361, 1993.

[23] J Peters. C^1 Surface Splines. *SIAM J Num Anal*, 32:645–666, 1995.

[24] H Prautzsch. Freeform Splines. *Computer Aided Geometric Design*, 14:201–206, 1997.

[25] U. Reif. A Unified Approach To Subdivision Algorithms Near Extraordinary Vertices. *Computer Aided Geometric Design*, 12:153–174, 1995.

[26] R F Riesenfeld. On Chaikin's Algorithm. *Computer Graphics and Image Processing*, 4:304–310, 1975.

[27] M A Sabin. Recursive Division. In J Gregory, editor, *The Mathematics of Surfaces*, pages 269–282. Clarendon Press, Oxford, 1986. ISBN 0 19 853609 7.

[28] M A Sabin. Cubic Recursive Division With Bounded Curvature. In P J Laurent, A le Mehaute, and L L Schumaker, editors, *Curves and Surfaces*, pages 411–414. Academic Press, 1991. ISBN 0 12 438660 1.

[29] D J T Storry and A A Ball. Design Of An N-sided Surface Patch From Hermite Boundary Data. *Computer Aided Geometric Design*, 6:111–120, 1989.

[30] J Zheng, T Sederberg, and M A Sabin. Eigenanalysis Of Non-Uniform Doo-Sabin Surfaces. Technical report, Brigham Young University, Department of Computer Science (appears as an appendix in the electronic version of this paper), November 1997.

[31] D Zorin, P Schröder, and W Sweldens. Interpolating Subdivision For Meshes With Arbitrary Topology. *Computer Graphics (SIGGRAPH 96 Conference Proceedings)*, 30:189–192, 1996.

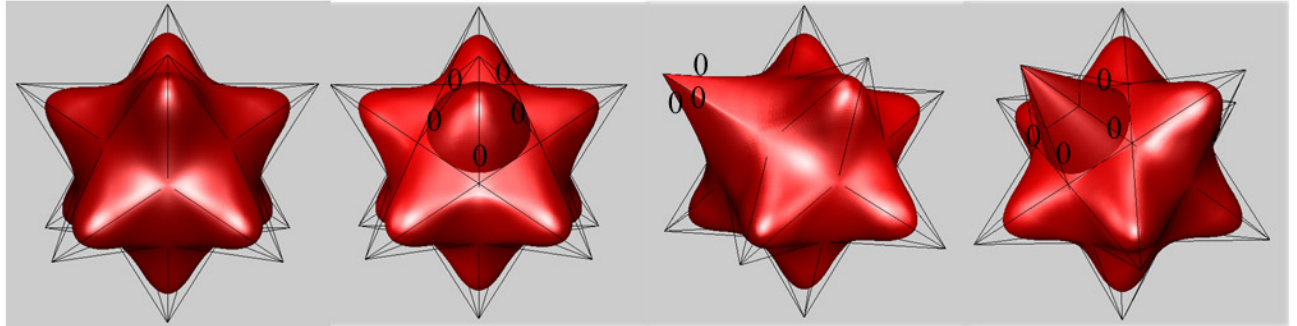# Non-Uniform Recursive Subdivision Surfaces

Figure 16a
Doo-Sabin surface

Figure 16b   Figure 16c   Figure 16d
Non-Uniform Doo-Sabin surfaces showing various features
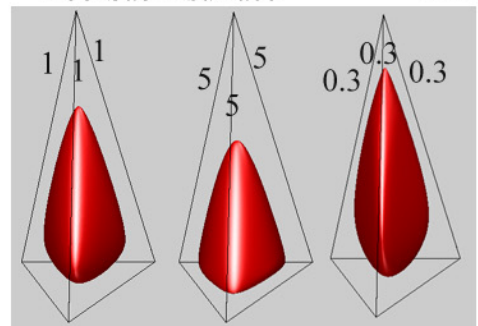

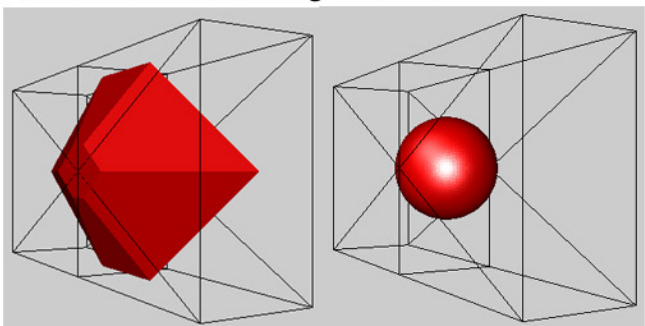Figure 17: NU Doo-Sabin
with various knot spacings

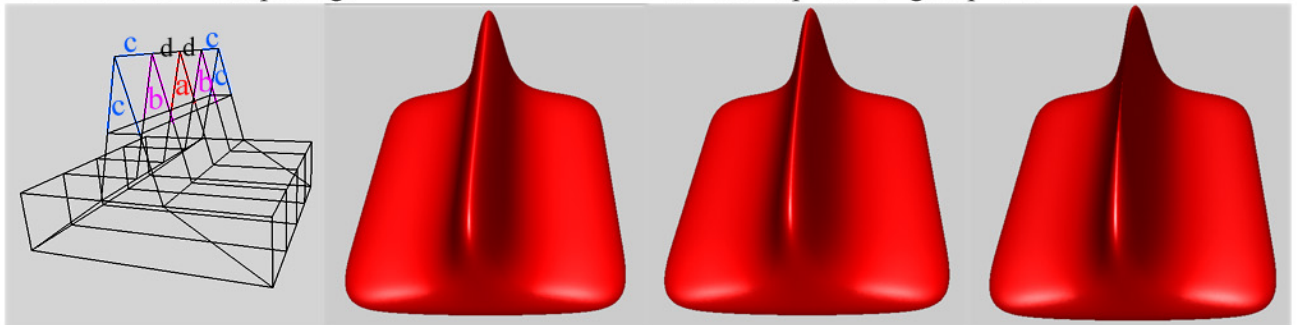Figure 18: Non-uniform rational Catmull-Clark
surface representing a sphere

Figure 19a
Initial polyhedron with
labels for knot spacings

Figure 19b
Uniform Catmull-
Clark surface

Figure 19c
a = b = 0.2
c = d = 1

Figure 19d
a = b = 0
c = d = 1
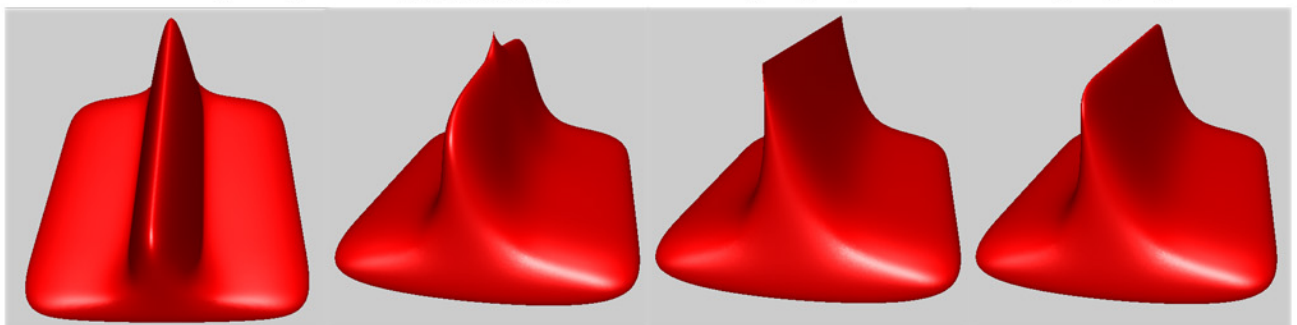

Figure 19e
a=b=20, c=d=1

Figure 19f
a=d=0, b=c=1

Figure 19g
a=b=c=0, d=1

Figure 19h
a=b=c=0.2, d=1

Non-uniform Catmull-Clark surfaces with various knot spacings