On Subdivision Schemes Generalizing Uniform B-spline Surfaces of Arbitrary Degree

Jos Stam

Alias | wavefront

1218 Third Ave, 8th Floor, Seattle, WA 98101, USA

Abstract

We introduce a new class of subdivision surfaces which generalize uniform tensor product B-spline surfaces of any bi-degree to meshes of arbitrary topology. Surprisingly, this can be done using subdivision rules that involve direct neighbors only. Consequently, our schemes are very easy to implement, regardless of degree. The famous Catmull-Clark scheme is a special case. Similarly we show that triangular box splines of total degree $3m+1$ can be generalized to arbitrary triangulations. Loop subdivision surfaces are a special case when $m=1$. Our new schemes should be of interest to the high-end design market where surfaces of bi-degree up to 7 are common.

Key words:
CAD, Curves & Surfaces, Solid Modeling, Subdivision Surfaces
PACS:

1 Introduction

In this paper we present a new class of subdivision surfaces which generalize uniform tensor product B-spline surfaces of any bi-degree to meshes of arbitrary topology. The generalizations for bi-degrees 2 and 3 were published in 1978 simultaneously by Catmull and Clark [1] and by Doo and Sabin [3]. Until recently [9], it seems no attempt has been successful in extending this work to surfaces of higher bi-degrees. This is partly due to the erroneous belief that such schemes require large subdivision masks and, consequently, would be difficult to implement. In this paper we show that the subdivision rules for uniform B-spline surfaces of any bi-degree can be generalized to meshes of arbitrary topology using operations which involve direct neighbors only. Consequently, our schemes are very easy to implement, regardless of degree. In fact, one version of our schemes requires only minimal modifications to an existing implementation of the Catmull-Clark scheme. The key idea behind our method is a generalization of the recurrence that computes binomial coefficients, viz. the Pascal Triangle.

Preprint submitted to Elsevier Preprint

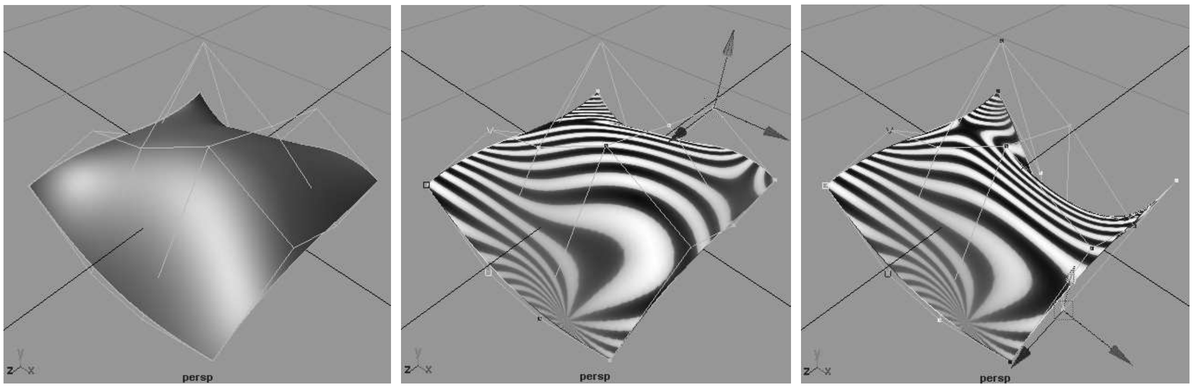

Fig. 1. This figure shows the reflection lines on a NURBS patch of bi-degree 5. By varying the control vertices a designer has enough degrees of freedom to control the variation of the reflection lines.

This fact was first observed by Lane and Riesenfeld in the curve setting [6]. However, they did not generalize this property to arbitrary meshes.

The main motivation behind this work is to bring subdivision surfaces to the wider design market, especially the high-end design market, such as the automotive industry where surfaces of bi-degree 5 and 7 are frequently used. Such surfaces are important because they give designers enough degrees of freedom to directly control how a car body reflects its environment. A standard environment is often composed of a set of parallel lines. The smoothness of the surfaces is then evaluated by observing how these lines are reflected off of the surface. The resulting patterns are known as “reflection lines.” In order to control the reflection lines directly a designer has to control the variation of the curvature. Consequently, surfaces of bi-degree higher than 3 are needed. Figure 1 shows the reflection lines on a NURBS patch of bi-degree 5.

The main problem designers currently face concerns regions of irregular topology. For example, corners are typically handled by trimming 3 NURBS surfaces of high bi-degree. Trimming is a time-consuming process that yields only approximate continuity. Indeed, at the trim, the surface may not even be C^0 continuous. Consequently, designers are forced to fine-tune many arcane parameters for the best fit. Our schemes, on the other hand, are naturally curvature-continuous everywhere except at a finite number of points where the curvature generally diverges—just like the Catmull-Clark scheme. Furthermore, by using higher-degree surfaces, we can reduce the rate of divergence.

Given our technique’s simplicity, ease of implementation, and practicality, we were surprised that these techniques weren’t well known in the field of computer aided design. We are aware of only one reference which states such a technique prior to our work. This is the “mid point” scheme of Prautzsch [9]. Unfortunately he did not discuss any practical applications and extensions of his scheme. We suppose that lower-degree surfaces are sufficient for character animation and free-form modeling, which have up to now been the only applications of subdivision surfaces in industry. By contrast, the high-end design industry has only recently begun to consider the use of subdivision surfaces. In addition to its practical importance, our work is also of theoretical interest. Indeed, our new schemes elegantly generalize the theory of uniform B-splines to arbitrary meshes. We show that uniform B-spline surfaces can be viewed as a special case of a simple smoothing algorithm on general meshes. Also, some of our schemes can be interpreted as single steps in a multigrid solver. They therefore tie in nicely with the variational approaches to subdivision introduced by Kobbelt [5] and Weimer and Warren [15]. We believe that the different point of view described in this paper has the potential to lead to further insights and developments in the theory of subdivision and surface modeling in general.

After we finished this research, we became aware of related work done independently by both Zorin and Schröder [16] and Warren and Weimer [14]. Their generalizations of the Lane-Riesenfeld algorithm to arbitrary meshes are essentially identical to our `Simple` algorithm and the mid-point scheme of Prautzsch [9]. Zorin and Schröder call their algorithm "repeated averaging" and describe an elegant implementation using quad trees. Their technique also extends to multi-resolution meshes. Warren and Weimer use their algorithm to define subdivision schemes in higher dimensions than 2. An important example is subdivision volumes defined by repeatedly averaging a three-dimensional mesh. We also mention that Velho and Zorin have recently proposed a new set of subdivision schemes for $4-8$ meshes that also relies on a factorization of the subdivision rules [13].

The rest of this paper is organized as follows. In the next section we briefly review subdivision for uniform B-spline curves and a recursive algorithm to compute them. These results are then easily extended to tensor-product B-spline surfaces of any degree. In Section 3, we present several possible generalizations of these rules to meshes of arbitrary topology. We discuss the continuity of these schemes in Section 4. Finally, in Section 5 we present some examples of surfaces generated with our new algorithm and conclude in Section 6.

2 The Curve Case

The subdivision rules for uniform B-spline curves of any degree are well known and are related to the binomial coefficients. Lane and Riesenfeld first showed how to subdivide efficiently using only local averages [6]. This result is proven elegantly using the discrete Fourier transform as shown in Chui's monograph [2]. We do not repeat the proof here. We start by describing the Lane-Riesenfeld algorithm for doubling the number of control vertices for a given B-spline curve. The algorithm subdivides a set of initial control vertices

$$\cdots, P^0_{-2}, P^0_{-1}, P^0_0, P^0_1, P^0_2, \cdots$$

of a degree d uniform B-spline curve in d steps. First, the set of control vertices is linearly subdivided:

$$P^1_n = \begin{cases} P^0_{n/2} & n \ \text{even} \\ \left(P^0_{(n-1)/2} + P^0_{(n+1)/2}\right)/2 & n \ \text{odd.} \end{cases} \tag{1}$$

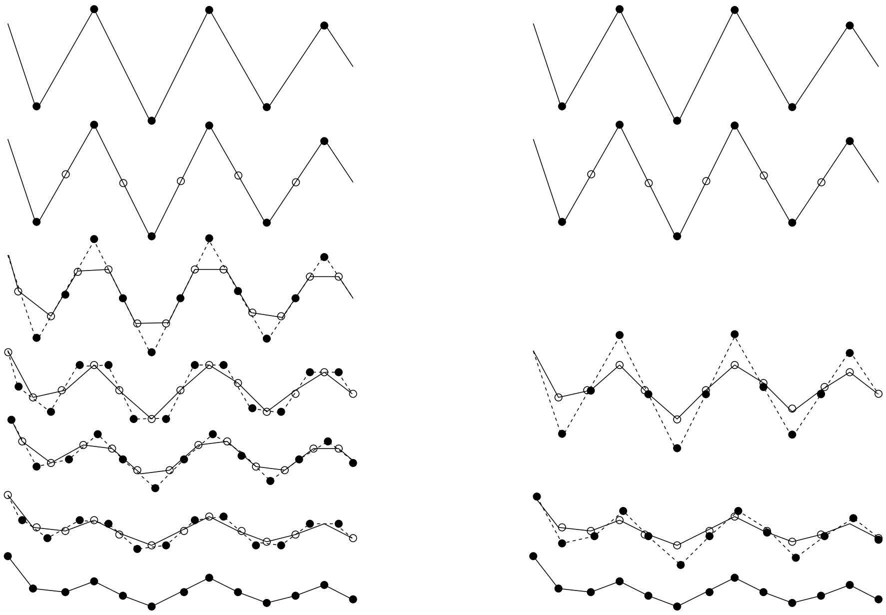

Fig. 2. On the left, Lane-Riesenfeld subdivision is applied to the control vertices of a uniform B-spline curve of degree 5. On the right, applying 2 averaging steps at a time allows the control vertices to be computed in place.

Subsequently, new control vertices are generated using $d-1$ averaging steps:

$$P_n^k = \frac{1}{2}\left(P_n^{k-1} + P_{n+1}^{k-1}\right) \qquad k = 2, \cdots, d. \tag{2}$$

See the left half of Figure 2 for a visual depiction of this algorithm. Notice that in the first step the number of vertices is doubled, while in the averaging steps it remains the same. Each averaging step raises by one the degree for which the original and new control vertices produce the same curve.

We now propose two modifications of this scheme that are both formally equivalent to it. We mention these schemes here because their generalizations to meshes are easier to implement and can handle boundaries and creases more gracefully. As is evident from the left side of Figure 2, the vertices before and after an averaging step are staggered, while the control vertices before and after two averaging steps are "in place." The latter is more desirable when implementing this algorithm. Therefore, we propose an algorithm equivalent to the Lane-Riesenfeld scheme which leaves the vertices in place by performing two averaging steps at a time. We obtain an algorithm for uniform B-splines of odd degree $d = 2m+1$ by performing one linear subdivision step followed by m "smoothing steps":

$$P_n^{2k+1} = \frac{1}{4}\, P_n^{2k-1} + \frac{1}{2}\, P_{n+1}^{2k-1} + \frac{1}{4}\, P_{n+2}^{2k-1} \qquad (k = 1, \cdots, m). \tag{3}$$

Each smoothing step effectively elevates the degree by 2. The effect of the algorithm is depicted in the right half of Figure 2. A similar algorithm which keeps the vertices in place can be designed for uniform B-splines of even degree $d = 2m+2$ as follows. We start with one linear subdivision step (Equation 1) and one averaging step (Equation 2) followed by

m smoothing steps, as in Equation 3 with $2k+1$ and $2k-1$ replaced by $2k+2$ and $2k$, respectively.

The generalization of these schemes to tensor-product uniform B-spline surfaces of arbitrary bi-degree is straightforward—simply apply the algorithm twice: once for each direction of the control vertex mesh (possibly using different degrees in each of the two directions).

3 The Surface Case

We now introduce our new surface subdivision schemes. The input to our schemes is an arbitrary (manifold) mesh `M0`. A mesh is defined by 3 arrays that store the vertices, edges and faces. Each vertex contains its position as well as the indices of the neighboring edges and faces. Each edge contains the indices of its end points and the indices of the two adjacent faces. Similarly, each face contains the indices of its neighboring faces, edges and vertices.

Each one of our schemes starts by linearly subdividing the mesh `M0`:

```
M1 = LinSubdivide ( M0 ).
```

In this step the vertices of the original mesh are unaffected and new vertices are added at the midpoint of each edge and at the centroid of each face. Edges are added connecting face centroids with each of the surrounding edge midpoints, guaranteeing that the new mesh consists of quadrilateral faces. If we stop here, the subdivision scheme is a generalization of a uniform B-spline of bi-degree 1. Starting from this linearly subdivided mesh, there are different means of obtaining generalizations of uniform B-spline surfaces of arbitrary degree. Before we proceed, let us review some standard nomenclature. To each vertex we assign a *valence* which is equal to the number of edges emanating from the vertex. A vertex is called *regular* if it has valence four; otherwise it is called *extraordinary*. The bi-degree of the B-spline surface we wish to generalize is denoted by `d` in this section.

3.1 The Simplest Scheme

Our simplest scheme is a direct generalization of the Lane-Riesenfeld algorithm. We apply `d-1` averaging steps to the linearly subdivided mesh `M1`:

```
Simple ( M0, d )
   M1 = LinSubdivide ( M0 )
   for k=1 to d-1 do
      M1 = Dual ( M1 )
   end for
   return M1.
```

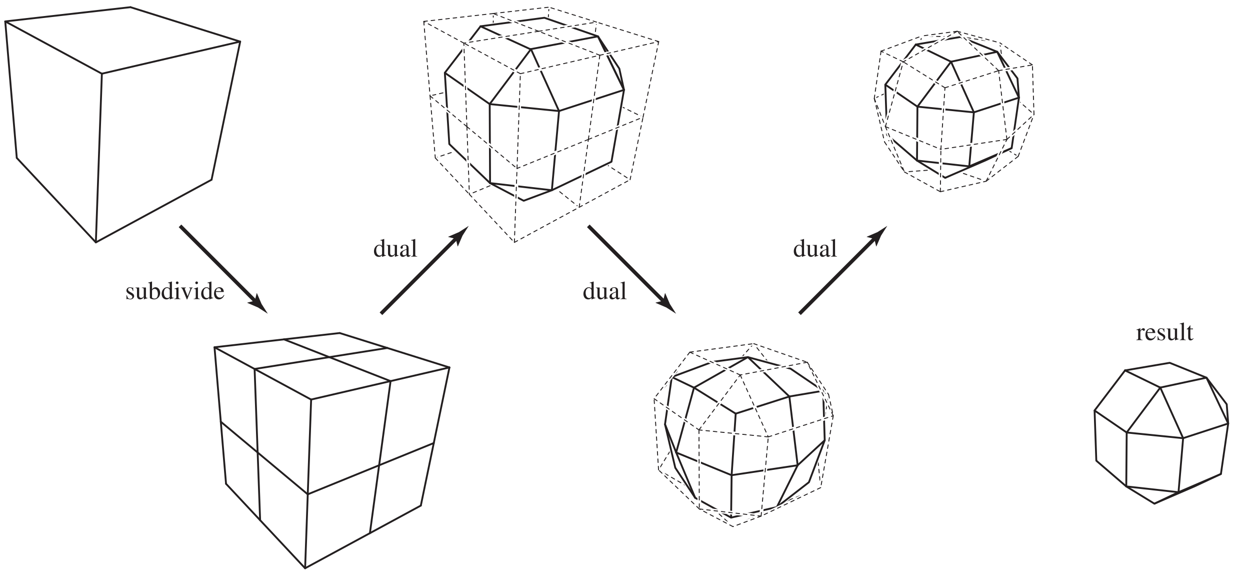

Fig. 3. One application of the `Simple` algorithm. The bi-degree is 4 in this example.

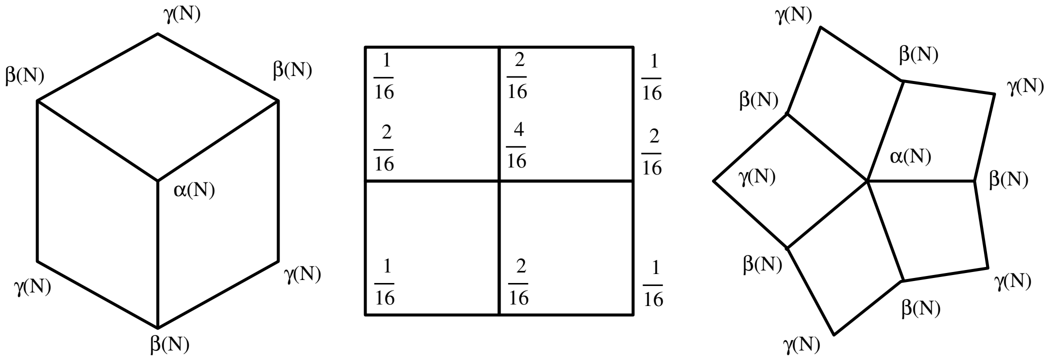

Fig. 4. Smoothing masks for our odd degree subdivision scheme applied to vertices of valences 3, 4 and 5, respectively.

The `Dual` function generalizes the averaging step of Equation 2 to arbitrary meshes. The dual of a mesh `M1` is a new mesh whose vertices are the centroids of `M1`'s faces and whose edges join centroids of faces that share a common edge in `M1`. Figure 3 illustrates the linear subdivision step and several dual steps. Clearly, each step requires only that we find the direct neighbors of vertices, edges, and faces.

Although this scheme is conceptually very simple, it is inefficient because the `Dual` routine modifies the topology of the mesh. This is the main reason we introduce two other generalizations next. See, however, the work of Zorin and Schröder, who developed an elegant implementation of this scheme using quad trees [16].

3.2 *The Odd Degree Scheme*

As we suggested in Section 2, we can define a more efficient subdivision scheme that generalizes uniform B-spline surfaces of odd degree. This scheme generalizes the smoothing steps of Equation 3:

```
Odd ( M0, d )
   M1 = LinSubdivide ( M0 )
   for k=1 to (d-1)/2 do
      OddSmooth ( M1 )
   end for
   return M1.
```

The smoothing function `OddSmooth` updates the vertices of the mesh. The smoothing of a vertex `v` involves only the vertices of the faces adjacent to vertex `v`. The vertices are updated in a "Jacobi manner" and require extra storage to cache the position of each vertex before a smoothing step. For regular vertices, the smoothing mask is simply the tensor-product version of Equation 3, whose mask consists of $3 \times 3 = 9$ numbers, as shown in the middle of Figure 4. At an extraordinary vertex `v`, the mask is more complicated and includes `v` itself as well as N *edge vertices* and N *face vertices*, where N is the valence of `v`. We use $\alpha(N)$ to denote the coefficient associated with `v`. For reasons of symmetry, all the coefficients associated with `v`'s edge vertices have to be equal; call them $\beta(N)$. Similarly, all the coefficients for `v`'s face vertices are equal; call them $\gamma(N)$. See Figure 4 for a depiction of the masks for extraordinary vertices of valence 3 and 5. There are many possible choices for the values of the coefficients, but they should at least give rise to an affine invariant subdivision scheme:

$$\alpha(N) + N\,\beta(N) + N\,\gamma(N) = 1.$$

It is also desirable to have coefficients that are positive to ensure stability and to guarantee the convex hull property. One choice of parameters that has all the desired properties is as follows:

$$\alpha(N) = \frac{N-3}{N}, \quad \beta(N) = \frac{2}{N^2} \quad \text{and} \quad \gamma(N) = \frac{1}{N^2}.$$

A single smoothing step with these masks is equivalent to the original Catmull-Clark scheme [1]. It is, however, different from the `Simple` algorithm run with an odd degree. Also note that existing implementations of the Catmull-Clark scheme can easily be modified to handle surfaces of arbitrary odd degree. Simply follow a Catmull-Clark subdivision step with `(d-3)/2` odd smoothing steps.

3.3 The Even Degree Scheme

To construct a scheme for even degrees, we apply one averaging step (the `Dual` function) to the linearly subdivided mesh first, followed by `(d-2)/2` even smoothing steps. (Alternatively, we could start with a mesh subdivided using the Doo-Sabin scheme, for example [3].) Here is our algorithm:

```
Even ( M0, d )
   M1 = LinSubdivide ( M0 )
   M1 = Dual ( M1 )
   for k=1 to (d-2)/2 do
      EvenSmooth ( M1 )
   end for
   return M1.
```

Before smoothing, each vertex of the mesh is regular, and therefore we can define `EvenSmooth` as follows. Replace each vertex with the average of the centroids of the four adjacent faces. This operation is local and has the desired properties of affine invariance and positive coefficients. In fact, using this scheme is identical to using `Simple` with an even degree. However, the `Even` function is more easily implemented and more efficient than `Simple` because it does not replace `M1` with a dual mesh whose topology differs in every smoothing step.

3.4 Boundaries, Corners and Creases

Sharp boundaries and creases [4] are easily incorporated in our odd subdivision schemes. For bi-degree 1 surfaces all edges are creased and all boundaries are sharp. Therefore, the first subdivision step is the same as the one called in the routine `Odd`. Only the smoothing step has to be modified. For vertices not on the boundary or on a crease, we apply the same mask as in Section 3.2. For a vertex on a crease or on a boundary, we simply use the curve smoothing mask defined by Equation 3 along the crease or the boundary. When more than two crease edges meet at a vertex, we consider the vertex a cusp or corner and perform no smoothing. We believe this implementation to be particularly elegant since it requires only minor modifications to the routine `Odd`. It is not at all clear how to add boundary and crease rules to even degree generalizations of uniform B-splines. This is a well-known problem with dual schemes such as Doo-Sabin. When creases and sharp boundaries are needed, we recommended the use of the modified `Odd` scheme.

3.5 Triangular Schemes

Our generalizations of uniform tensor-product B-splines can also be extended to handle triangular subdivision surfaces. A scheme generalizing the subdivision rules for triangular box spline surfaces of total degree 4 was first proposed by Loop in 1987 [7]. Similarly, the

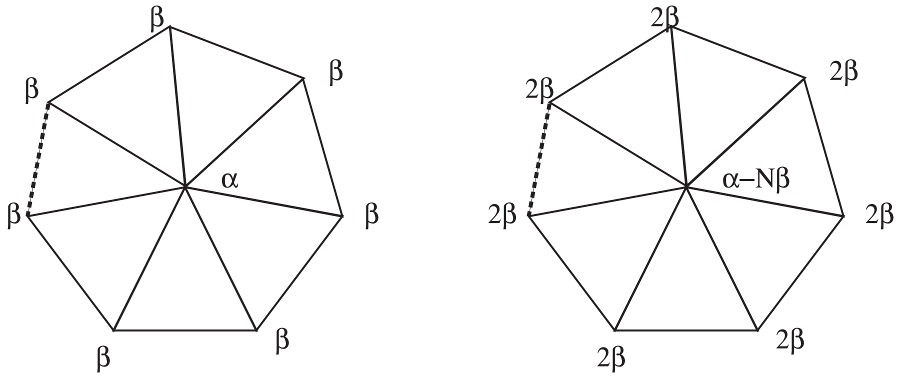

Fig. 5. On the left, vertex subdivision mask for Loop subdivision scheme. On the right, corresponding smoothing mask for this vertex in our scheme.

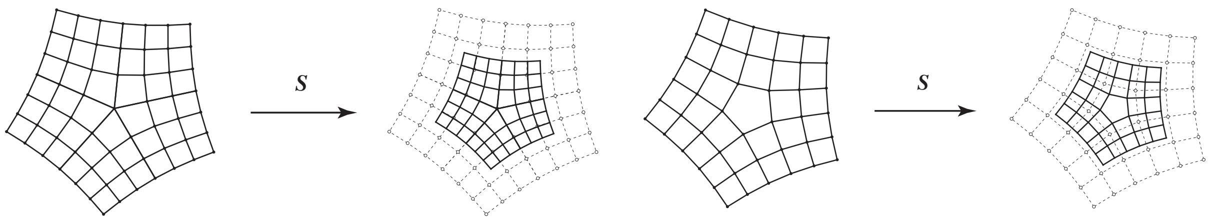

Fig. 6. The subdivision matrix relates the control vertices surrounding an irregular point after one subdivision step.

subdivision masks for triangular box splines of higher degree can be related to lower degree ones. We construct generalizations of triangular box splines of total degree $d = 3m + 1$ by a straightforward extension of our `Odd` scheme. First we add new vertices at the midpoint of each edge of the triangulation, next we perform m smoothing steps of the vertices. The smoothing mask involves direct neighbors only and is related to the vertex subdivision mask of Loop subdivision surfaces [7]. Let α and β be the coefficients of the vertex mask as shown on the left of Figure 5. Then the coefficients of the smoothing mask are given by $\alpha' = \alpha - N\beta$ and $\beta' = 2\beta$, where N is the valence of the vertex. For a regular vertex having valence 6 we have $\alpha' = 1/4$ and $\beta' = 1/8$. In this new formulation there is no need to distinguish between vertex and edge rules. Generalizations to other degrees are possible but schemes analogous to the `Simple` one for B-splines are trickier since the dual of a regular triangular mesh is hexagonal, whereas for regular B-spline meshes the dual remains a quad mesh.

4 Remarks on Smoothness

The surfaces obtained using our new schemes are naturally C^{d-1}-continuous except at a finite number of *irregular* points. When d is odd these points correspond to the extraordinary vertices, while for even degrees they correspond to the centroids of the non-quadrilateral faces. This is obvious since away from the irregular points, our subdivision rules are identical to the ones for tensor-product B-splines of bi-degree d. Each irregular region shrinks at every subdivision step and approaches a point in the limit of infinite subdivision.

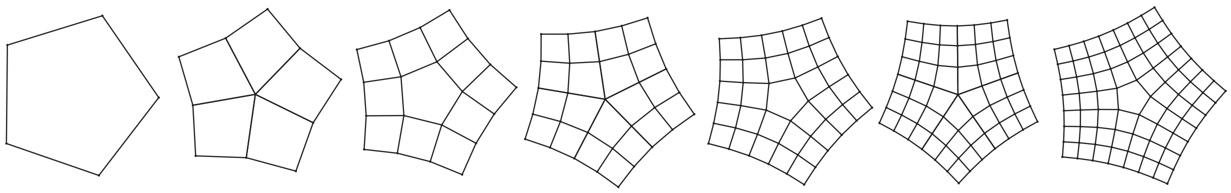

Fig. 7. Characteristic meshes of the `Odd` and `Even` schemes at an irregular point of valence 5. The bi-degrees range from 2 (left) to 8 (right).

The smoothness analysis at the irregular points is more subtle and has only recently been fully clarified [8,10,17]. Of crucial importance is the eigenstructure of the subdivision matrix. To every irregular point we associate a set of neighboring vertices that influence its limit position on the surface. For odd degrees this set consists of the extraordinary vertex and $(d-1)/2$ rings of vertices surrounding the extraordinary vertex. The left half of Figure 6 shows these vertices for valence $N = 5$ and bi-degree $d = 7$. For even degrees the set includes $d/2$ rings of vertices surrounding the face. See the right half of Figure 6 for an example with $N = 5$ and $d = 6$. The subdivision matrix S specifies how this set of vertices is transformed into a topologically similar set of vertices by one iteration of subdivision.

We have computed the eigenstructure of subdivision matrices for a wide range of different valences and bi-degrees. Due to the affine invariance of the subdivision rules, the largest eigenvalue is always equal to one. The subdominant (next largest) eigenvalue $\lambda < 1$ is of geometric multiplicity two. The two eigenvectors corresponding to λ are crucial in the smoothness analysis. In particular these eigenvectors form a planar mesh which, when subdivided, yields a surface called the *characteristic map* [10]. When this map is injective and regular, the corresponding subdivision scheme is C^1. Figure 7 shows the meshes of the characteristic map of the `Odd` and `Even` schemes for an irregular point of valence 5 and bi-degrees ranging from 2 to 8. The surfaces obtained by subdiving these meshes seem to be injective. We therefore believe that these schemes are C^1 at the irregular points. Zorin and Schröder have rigorous C^1 proofs for the `Simple` scheme for degrees up to 9 [16]. Finding rigorous proofs for higher bi-degrees remains an open problem, however.

More interesting is the analysis of higher-order smoothness. A *necessary* condition for curvature continuity at the irregular points is that the eigenvalue μ next largest after λ satisfies $\mu = \lambda^2$ [3]. Unfortunately our schemes never satisfy this condition at the irregular points. Therefore none of our schemes are C^2 at the irregular points. However, the smoothness improves with higher degree as μ approaches λ^2. In Table 1 we list the ratio $\delta = \log \mu / \log \lambda$ for several valences and bi-degrees. [1] The ideal ratio is 2 as in the regular case when the valence is 4. When $\delta < 2$ the curvature diverges. When $\delta > 2$ the curvature is zero, indicating that the limit surface is locally flat. It is interesting to note that for valence 3, additional smoothing actually increases the curvature divergence. For all other valences, the rate of divergence decreases as the degree goes up, which is desirable. We note that satisfying the condition $\mu = \lambda^2$ guarantees bounded curvature but does not guarantee continuous curvature. See Zorin [17] for more details.

[1] There is no entry for $d = 2$ and $N = 3$ because we have only 3 eigenvalues and the ratio is therefore not defined.

`Simple` scheme

degree \ valence	3	4	5	6	7
2	—	2.000	2.205	2.087	1.923
3	1.555	2.000	1.804	1.635	1.505
4	1.498	2.000	1.967	1.890	1.791
5	1.519	2.000	1.900	1.774	1.655
6	1.502	2.000	1.962	1.889	1.799
7	1.510	2.000	1.932	1.833	1.727
8	1.503	2.000	1.965	1.898	1.815

`Odd` scheme

degree \ valence	3	4	5	6	7
3	2.010	2.000	1.804	1.635	1.505
5	1.608	2.000	1.900	1.774	1.655
7	1.578	2.000	1.932	1.833	1.727
9	1.556	2.000	1.948	1.865	1.769
11	1.544	2.000	1.958	1.885	1.797
13	1.536	2.000	1.964	1.899	1.818
15	1.531	2.000	1.969	1.909	1.833

Table 1

Ratio $\delta = \log \mu / \log \lambda$ for different valences and bi-degrees. Results for the `Simple` scheme left and `Odd` scheme right.

The lack of C^2 continuity in our schemes of degree $d < 6$ comes as no surprise; Reif proved that any C^2 subdivision scheme generalizing polynomial surfaces must be of at least bi-degree 6 [11].

5 Results

We have implemented our new schemes as a plugin in our modeling and animation software package MAYA. Fortunately, we were able to reuse many existing routines already available in our software. Also our new `Odd` scheme with creases for bi-degree 3 turned out to be a much simpler implementation than our current implementation of Catmull-Clark schemes with creases. In Figure 8 (top) we depict our `Simple` scheme applied to a simple cube. The bi-degrees range from $d = 1$ to $d = 5$. Below it we depict surfaces generated using our implementation of the `Even` and `Odd` schemes. The surfaces generated using the `Even` scheme are identical to the corresponding ones generated using `Simple`. It is interesting to note that the "nicest shape" in this case is obtained with the `Odd` with a degree $d = 3$, i.e., the Catmull-Clark scheme. This is consistent with our smoothness analysis in Section 4: the curvature divergence near a valence 3 vertex is slowest when $d = 3$.

In Figure 9 we show the reflection lines on a surface with a vertex of valence 3 (top) and of valence 7 (bottom). The bi-degrees of the surfaces are $d = 3, 5$ and 7 from left to right. As noted above, the behavior of the reflection lines worsens with degree when the valence is three while it improves with degree for higher valences.

Figure 10 depicts results from our implementation of the triangular scheme. The case $d = 4$ is equivalent to Loop's scheme.

Finally Figure 11 is an example of the type of surfaces often encountered in the automotive

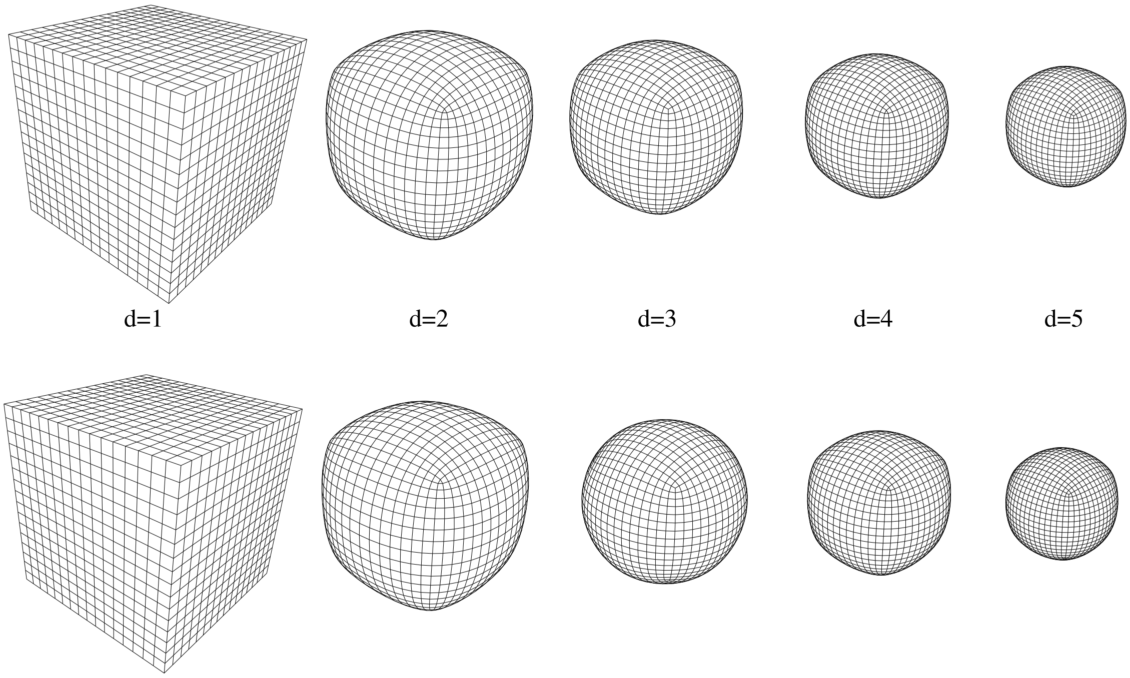

Fig. 8. Our new subdivision schemes applied to a cube. We used the `Simple` scheme (top) and the `Even` and `Odd` schemes (below).

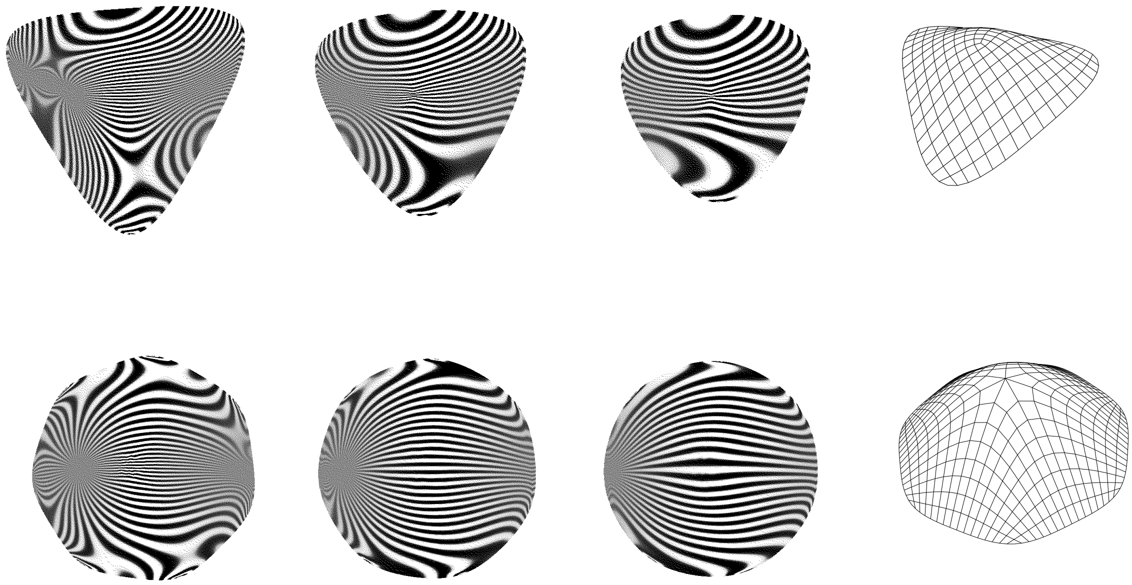

Fig. 9. Reflection lines on surfaces generated using our `Odd` schemes for bi-degree 3, 5, and 7 (from left to right). The top surfaces have a vertex of valence 3 in the center, while the bottom ones have a vertex of valence 7 in the center. The shape of the surfaces are shown in the right most column of the figure.

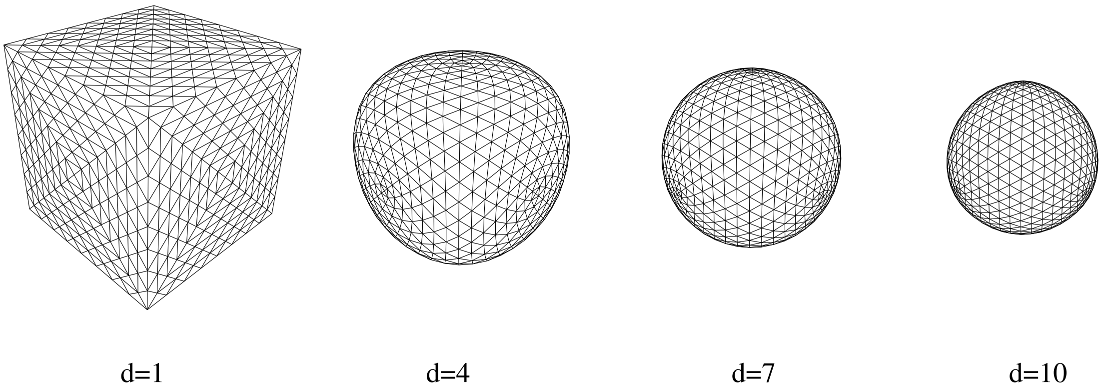

Fig. 10. Our new triangular schemes applied to a cube. Total degree of the surface ranges from 1 to 10. The case $d = 4$ is equivalent to Loop's scheme.

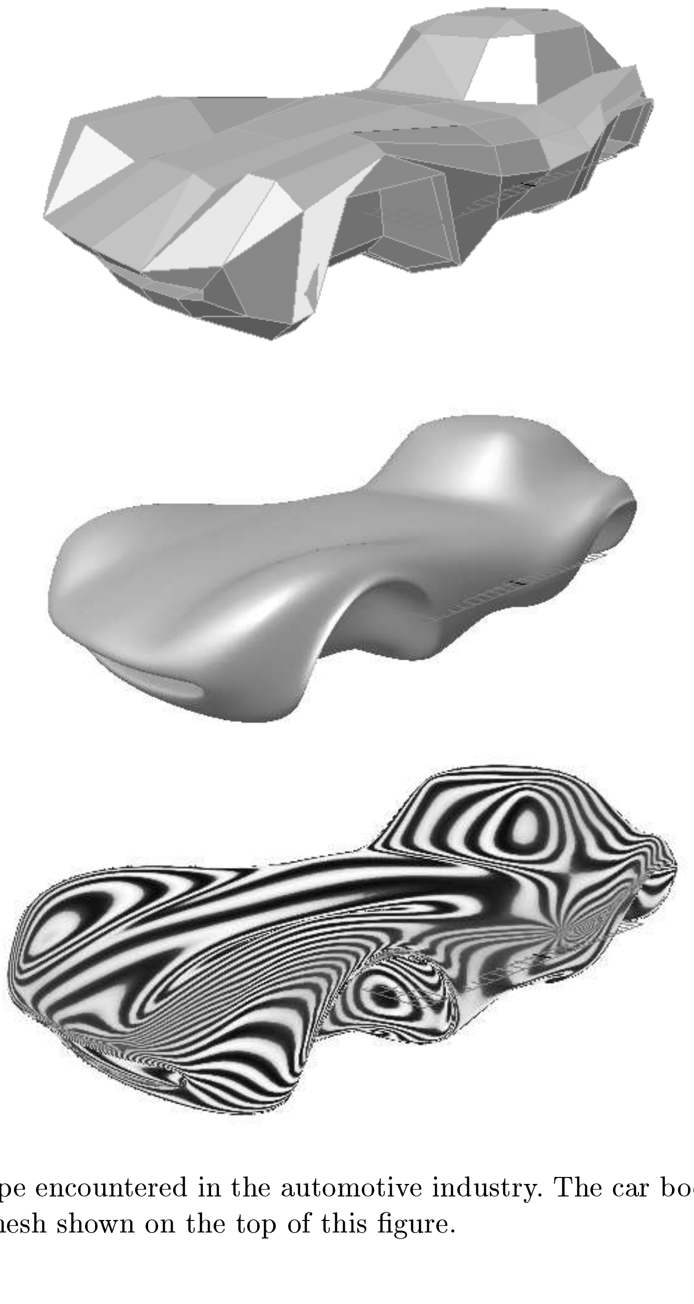

Fig. 11. A typical shape encountered in the automotive industry. The car body is entirely modeled using the polygonal mesh shown on the top of this figure.

industry. The entire car body was modeled using a single polygonal mesh shown on the top of Figure 11. To model a similar surface with trimmed NURBS alone would have been much more time consuming. In Figure 11 we show both a shaded version of the model (middle) and its reflection lines (bottom). The reflection lines are very well behaved on this model. The surfaces shown in Figure 11 were generated using the `Odd` scheme of bi-degree 5. In general the amount of smoothing could be a user specified attribute for each vertex. For example, near valence 3 vertices less smoothing might be desirable. These surfaces could be a possible candidate to replace high order NURBS in the high-end design market.

6 Conclusions and Future Work

In this paper we have introduced a new class of subdivision schemes generalizing tensor-product B-splines of any bi-degree to arbitrary meshes. Although these generalizations are not C^2 everywhere, they provide useful alternatives to the current practice in the high end design market. Our schemes offer C^{d-1} continuity everywhere except at the irregular points, where they seem to be C^1. This continuity compares favorably with trimmed NURBS, which are often not even C^0. While the curvature of our schemes in general diverges at the irregular points, it does so very slowly. Our schemes can also be evaluated exactly everywhere, even near irregular points, using a straightforward extension of Stam's work on evaluating Catmull-Clark surfaces [12]. The eigenbasis functions are easily precomputed for a set of valences and bi-degrees. In addition, the initial mesh has to be subdivided a sufficient number of times to isolate the extraordinary points.

Our approach should be of general interest to the surface modeling community as it sheds new light on the problem of handling arbitrary topologies. We believe our solution to be particularly simple and elegant.

We also hope that our approach will lead someone to settle a long-standing question: are there locally supported subdivision schemes with the convex hull property that are C^2 at irregular points?

Acknowledgments

I wish to thank Dick Rice for helpful discussions in the early stages of this work. Thanks to Michael Lounsbery, Eric Stollnitz and Pamela Jackson for carefully proofreading the paper. Thanks to Jérôme Maillot for writing the MAYA plugin and to Brad Clarkson for letting me use his car model. I also wish to thank the organizers of the Dagstuhl Workshop on Subdivision for inviting me to present and discuss this work with the leading experts in subdivision.

References

[1] E. Catmull and J. Clark. Recursively Generated B-Spline Surfaces On Arbitrary Topological Meshes. *Computer Aided Design*, 10(6):350–355, 1978.

[2] C. K. Chui. *Multivariate Splines.* CBMS-NSF Regional Conference Series in Applied Mathematics. SIAM, Philadelphia, PA, 1988.

[3] D. Doo and M. A. Sabin. Behaviour Of Recursive Subdivision Surfaces Near Extraordinary Points. *Computer Aided Design*, 10(6):356–360, 1978.

[4] H. Hoppe, T. DeRose, T. Duchamp, M. Halstead, H. Jin, J. McDonald, J. Schweitzer, and W. Stuetzle. Piecewise smooth surface reconstruction. In *Computer Graphics Proceedings, Annual Conference Series, 1994*, pages 295–302, July 1994.

[5] L. Kobbelt. A Varitational Approach to Subdivision. *Computer Aided Geometric Design*, 13:743–761, 1996.

[6] J. M. Lane and R. F. Riesenfeld. A Theoretical Development For the Computer Generation adn Display of Piecewise Polynomial Surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2(1):35–46, January 1980.

[7] C. T. Loop. *Smooth Subdivision Surfaces Based on Triangles.* M.S. Thesis, Department of Mathematics, University of Utah, August 1987.

[8] J. Peters and U. Reif. Analysis Of Generalized B-Splines Subdivision Algorithms. To appear in SIAM Journal of Numerical Analysis.

[9] H. Prautzsch. Smoothness of subdivision surfaces at extraordinary points. *Advances in Computational Mathematics*, 9:377–389, 1998.

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

[11] U. Reif. A degree estimate for subdivision surfaces of higher regularity. *Proceedings of the American Mathematical Society*, 124:2167–2174, 1996.

[12] J. Stam. Exact Evaluation of Catmull-Clark Subdivision Surfaces at Arbitrary Parameter Values. In *Computer Graphics Proceedings, Annual Conference Series, 1998*, pages 395–404, July 1998.

[13] L. Velho and D. Zorin. 4-8 Subdivision. *Computer Aided Geometric Design. Special issue on Subdivision Surfaces.*, 18:(this issue), 2001.

[14] J. Warren and H. Weimer. Subdivision for Geometric Design. To appear, 2000.

[15] H. Weimer and J. Warren. Subdivision Schemes for Thin Plate Splines. *Proceedings of Eurographics 1998, Computer Graphics Forum*, 17(3):303–313, 1998.

[16] D. Zorin and P. Schröder. A Unified Framework for Primal/Dual Quadrilateral Subdivision Schememes. *Computer Aided Geometric Design. Special issue on Subdivision Surfaces.*, 18:(this issue), 2001.

[17] D. N. Zorin. *Subdivision and Multiresolution Surface Representations.* PhD thesis, Caltech, Pasadena, California, 1997.