Volume xx (200y), Number z, pp. 1–5

Quad/Triangle Subdivision

Jos Stam[1] and Charles Loop[2]

[1] Alias | wavefront, Seattle, WA, USA.
[2] Microsoft Research, Redmond, WA, USA.

Abstract

In this paper we introduce a new subdivision operator that unifies triangular and quadrilateral subdivision schemes. Designers often want the added flexibility of having both quads and triangles in their models. It is also well known that triangle meshes generate poor limit surfaces when using a quad scheme, while quad-only meshes behave poorly with triangular schemes. Our new scheme is a generalization of the well known Catmull-Clark and Loop subdivision algorithms. We show that our surfaces are C^1 everywhere and provide a proof that it is impossible to construct a C^2 scheme at the quad/triangle boundary. However, we provide rules that produce surfaces with bounded curvature at the regular quad/triangle boundary and provide optimal masks that minimize the curvature divergence elsewhere. We demonstrate the visual quality of our surfaces with several examples.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Curve, surface, solid, and object representations

1. Introduction

Subdivision surfaces are currently one of the most powerful surface representations used to model smooth shapes. Unlike regular surface splines, such as NURBS, subdivision surfaces can handle shapes of arbitrary topology in a unified framework. Also, unlike polygonal meshes, subdivision surfaces generate smooth surfaces, which is important in designing aesthetically pleasing shapes. Subdivision surfaces were introduced in 1978 by both Catmull and Clark [2] and Doo and Sabin [3]. They both generalized tensor product B-splines of bi-degree three and two, respectively, to arbitrary topologies by extending the refinement rules to irregular parts of the control mesh. Later, in 1987 Loop generalized triangular Box splines of total degree four to arbitrary triangular meshes [5]. These subdivision schemes have the desirable property that they admit a polynomial representation on the regular part of the mesh. Consequently, these surfaces are curvature continuous almost everywhere [7, 12] and can be evaluated explicitly anywhere [8].

The visual quality of a subdivision surface depends in a crucial way on the initial, or base, mesh of control vertices. For general shapes designers often want to model certain regions with triangle patches and others with quad patches. Unfortunately, both Catmull-Clark and Loop surfaces require that all patches be quadrilateral or triangular, respectively. In theory this is not a problem, since any triangle can be converted into three quads and any quad can be tesselated. However, as illustrated in Figure 1, Catmull-Clark surfaces behave very poorly on triangle-only base meshes: the resulting surface exhibits annoying undulating artifacts. Similarly, Loop schemes do not perform well on quad-only meshes. More importantly, designers often want to preserve quad patches on regular areas of the surface where there are two "natural" directions. Consequently it is often desirable to have surfaces that have a hybrid quad/triangle patch structure.

To the best of our knowledge only one previous hybrid quad/triangle subdivision scheme has been proposed in the literature [6]. This scheme reproduces Catmull-Clark on the quad region but does not reproduce Loop on the triangular part. The triangular scheme on the **regular** part of the mesh is actually non-polynomial and is not even curvature continuous. Our scheme on the other hand reproduces Loop surfaces on the triangular part of the mesh. Also, no analysis of the limit surface was provided in [6]. Our hybrid subdivision algorithm is surprisingly simple and is based on the fact that both Loop and Catmull-Clark surfaces can be generated by a linear subdivision step followed by vertex-smoothing [9]. This

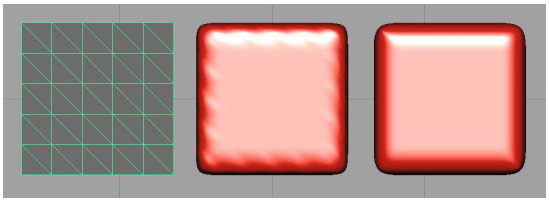

Figure 1: *A regular triangular mesh (left) behaves poorly with Catmull-Clark (middle) and behaves nicely with Loop.*

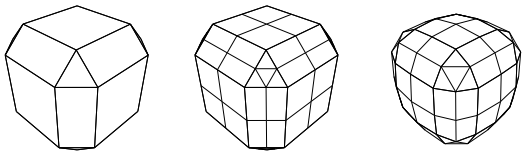

Figure 2: *Our subdivision scheme comprises two steps. A base mesh (left) is first linearly subdivided (middle) and then followed by an averaging of the vertices (right).*

algorithm is related to the "repeated averaging" schemes of Zorin and Schröder [11] and Warren and Weimer [10] in that two averaging steps are taken at once.

2. Quad/Triangle Subdivision

2.1. The Generalization

Our new surfaces are modeled by a base mesh that is formed of quads and triangles only. We assume that polygons of the base mesh other than quads and triangles are quadrangulated by placing a vertex at the centroid of the polygon; as in step one of the Catmull-Clark algorithm. Optionally, these faces could be triangulated. Figure 2 (left) shows a typical base mesh. Our subdivision rules comprise two steps. In the first step we evenly split each edge into two, each quad into four and each triangle into three, as shown in Figure 2 (middle). Then in a second step we replace each vertex with a linear combination of itself and its direct neighbors resulting in the mesh shown in Figure 2 (right). When a vertex is entirely surrounded by triangles or quads we use the first two masks shown in Figure 4, otherwise when the mask is surrounded by both triangles and quads we introduce the new third mask. We first consider the *regular* case when a vertex is surrounded by two adjacent quads and three adjacent triangles. In this case the obvious choice is the mask shown in Figure 3 (right), which is a simple average of the regular quad and triangle masks also shown in Figure 3. To derive a general mask at an *irregular vertex* we introduce some notations: let n_e be the number of edges emanating from the vertex and let n_q be the number of quads surrounding the vertex. As shown in Figure 4, we denote by α the weights associated with the irregular vertex and let β and γ be the weights associated with the neighboring edge and face vertices, respectively. A natural generalization of the regular case is to let the weights be equal to

$$\beta = \frac{\alpha}{2} \quad \textit{and} \quad \gamma = \frac{\alpha}{4}.$$

The value of α is then determined from the condition that the rules should be affine invariant:

$$\alpha + n_e \frac{\alpha}{2} + n_q \frac{\alpha}{4} = 1.$$

This equation is readily solved for α:

$$\alpha = \frac{1}{1 + \frac{n_e}{2} + \frac{n_q}{4}}. \tag{1}$$

The choice for our mask is quite arbitrary as there are potentially many other affine invariant generalizations. Ideally we want a mask which gives the most aesthetically pleasing surfaces. One way to formalize this requirement is to make the curvature well behaved at each vertex. In fact, for $n_e < 5$ our masks are far from optimal. Before we discuss the modification to our masks we have to recall some well known results from the theory of subdivision surfaces.

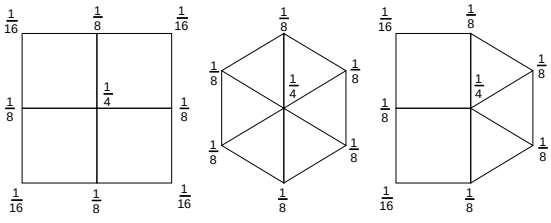

Figure 3: *Averaging masks for regular quads (left), regular triangles (middle) and regular quad-triangles (right).*

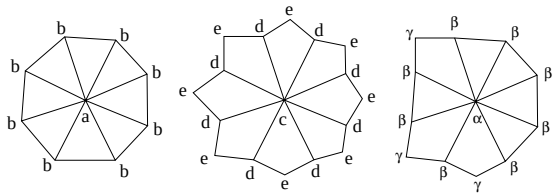

Figure 4: *Smoothing masks for Loop, Catmull-Clark and our new scheme. The weights are* $a = 2(3/8 + 1/4\cos(2\pi/n_e))^2 - 1/4$, $b = (1-a)/n_e$, $c = (n_e-3)/n_e$, $d = 2/n_e^2$ *and* $e = 1/n_e^2$. *In our example* $n_e = 8$ *and* $n_q = 3$.

2.2. Eigen-Analysis of Subdivision

Of crucial importance in the theory of subdivision surfaces is the subdivision matrix S. To every vertex of the control mesh we associate a one-ring of neighboring vertices as depicted in Figure 4. The subdivision matrix specifies how this set

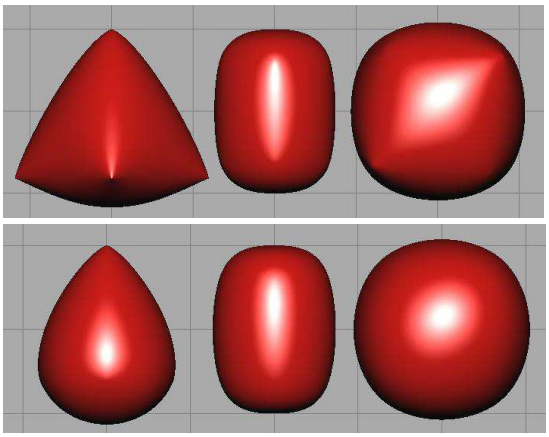

Figure 5: *Surface generated using our new subdivision scheme: without vertex correction (top) after vertex correction (bottom).*

of vertices is transformed into a similar set of vertices after one iteration of the subdivision rules. This matrix is of size $K \times K$, where $K = 1 + n_e + n_q$. As first pointed out in [3] the eigenstructure of this matrix is important in the analysis of the limit behavior of the surface at the central vertex. Due to the property of affine invariance, the matrix S always has a maximum eigenvalue equal to one. The next five eigenvalues in order of magnitude:

$$1 > \lambda \geq \lambda_2 > \mu \geq \mu_2 \geq \mu_3 > \cdots$$

are important in characterizing the behavior of the tangent plane and the behavior of the curvature at the central vertex. In particular, the two left eigenvectors corresponding to λ and λ_2 can be used to compute the normal (if it exists) of the limit surface at the central vertex. When the surface curvature is continuous, $\mu = \lambda^2$. Consequently, to measure the quality of the curvature we propose the following *ratio* $\rho = \mu/\lambda^2$. Ideally, this ratio should be equal to 1. When $\rho < 1$ the curvature is zero and the surface has a flat spot, which is undesirable. On the other hand, when $\rho > 1$ the curvature diverges which results in surfaces which appear to be "pinched" at the central vertex.

2.3. Vertex Correction

Refer to Figure 5 (top) depicting three surfaces created with our generalized masks. Clearly the corners of these surfaces appear to be overly pinched. Somehow we want the limit surface at the corners to be drawn more inward. One simple technique to achieve this is to follow the averaging step with a *correction step*. This idea was first introduced (implicitly) by Catmull and Clark to improve the behavior of the surface at a vertex where three quads meet: a corner. This fact was first pointed out in [6]. The idea behind our correction step is to translate the smoothed vertex along the direction defined by the difference between its current position v_1 and its previous position v_0 by an amount η:

$$v_2 = v_1 + \eta(v_1 - v_0).$$

To make the corners more rounded we let $\eta > 0$, so that the vertex is drawn more inward. We numerically determine the optimal value of η which makes the ratio $\rho = 1$. These values are reported in Table 1. When $n_e = 3$ the optimal value $\rho = 1$ cannot be achieved because larger values of η result in a subdivision matrix having a negative eigenvalue, which results in undesirable oscillations. When this happens we compute the value of η which sets the smallest eigenvalue exactly to zero. Figure 5 (bottom) shows the effect of the correction step. Notice how the pinched behavior has almost completely disappeared and the surfaces have a rounder appearance.

n_e	n_q	η	ρ
2	1	-0.20505	1.000
3	1	0.80597	1.227
3	2	0.61539	1.242
4	1	0.34792	1.000
4	2	0.21380	1.000
4	3	0.10550	1.000

Table 1: *Values of the correction factor η for different configurations. Also given is the corresponding value of the ratio ρ.*

When $n_e \geq 5$ we did not notice any improvement in the ratio ρ when varying the correction factor η. Nor did we notice any differences in the appearance of the resulting surfaces. Therefore we simply ignore the vertex correction step when $n_e \geq 5$. It is possible to improve on the behavior of the curvature in these regions by modifying the original mask given by Equation 1. We did not explore this line of enquiry further as the surfaces we experimented with had no apparently bad artifacts in those regions.

2.4. Creases, Boundaries and Hierarchy

Creases and boundaries are easily incorporated in our scheme following [9]. Smoothing along boundaries and creases is achieved using the $[\frac{1}{4}\ \frac{1}{2}\ \frac{1}{4}]$ cubic B-spline averaging mask along the boundary/crease. Alternatively, more clever rules could be devised to handle pathological boundary cases as shown in [1]. On the quad and triangle regions of the mesh we can use any of the well known hierarchical schemes.

3. Smoothness

3.1. C^1-Continuity

Our surfaces are naturally curvature continuous on both the regular quad and the triangle regions of the mesh since each

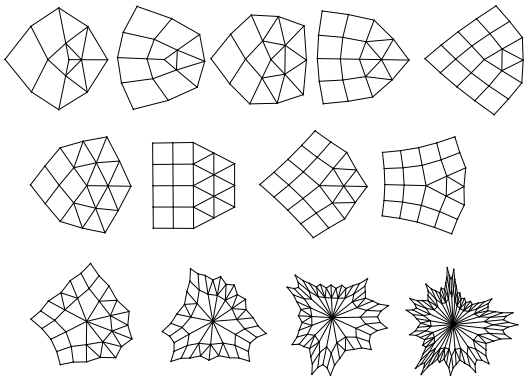

Figure 6: *Control nets defining the characteristic map for different values of n_e and n_q.*

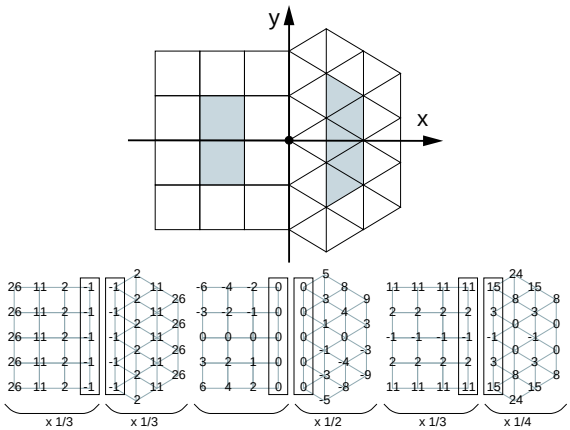

Figure 7: *Coefficients of the Box splines corresponding to the shaded area (top) that reproduce the three quadratics: x^2, xy and y^2 (bottom).*

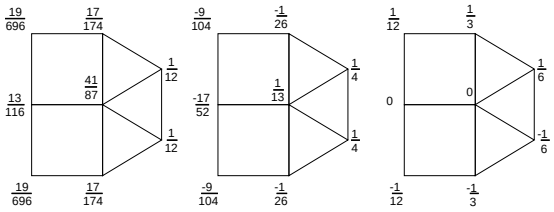

Figure 8: *Limit masks for position (left) and normal (middle/right) for a regular quad/triangle vertex.*

agree exactly with regular bi-cubic B-splines and triangular Box splines of total degree four, respectively. At the irregular quad and triangle regions the surface is C^1 but not curvature continuous. We refer the reader to the literature for a proof of this fact [7, 12]. The crucial tool used in these proofs is the *characteristic map* first introduced by Reif [7]. This map is the surface defined by the two-dimensional mesh formed by the two eigenvectors corresponding to λ and λ_2 (see Section 2.2) of the larger subdivision matrix S', which includes two rings of neighboring vertices. Figure 6 shows the aforementioned meshes for various configurations of quads and triangles around a central vertex. A fundamental theorem of subdivision surfaces states that when the characteristic map is both regular and injective, the surface is C^1 at the central vertex. In Figure 6 we depict several control meshes for characteristic maps corresponding to different values of n_e and n_q. It seems reasonable to infer that the corresponding surfaces are injective. In fact we can verify injectivity for these meshes up to machine precision. This argument does not constitute a formal proof of C^1 for all possible cases. However in practice what matters most is that we can compute a limit normal from the left two eigenvectors corresponding to λ and λ_2 [4].

3.2. Non-C^2-Continuity

We now show that our surfaces cannot be C^2 along the regular quad/triangle boundary. In order for the surface to be C^2 it must be able to reproduce the three quadratics x^2, xy and y^2 at any patch along the boundary. After one subdivision step, one half of the patches at the boundary are exact polynomials (shown in grey at the top of Figure 7). On these patches we can exactly represent any of the quadratics in terms of the Box spline basis functions. This completely constrains the values of the control vertices on both sides of the quad/triangle boundary. The values of these vertices are readily computed and are given on the bottom of Figure 7. For the surface to be C^2 these coefficients have to agree at the quad/triangle boundary. Unfortunately this is not the case for the coefficients of y^2 and consequently the surface cannot possibly be C^2. Note that our proof carries through even if we use a different configuration than the one shown in Figure 7. Other configurations change the coefficients we present. These may change the quadratics spanned, but no choice captures all three.

However, for the rules we have chosen, the surface at the quad/triangle boundary has bounded curvature since it has the following regular set of eigenvalues:

$$1, \frac{1}{2}, \frac{1}{2}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{8}, \frac{3}{32},$$

such that $\mu = \lambda^2$.

4. Implementation

We have implemented our scheme as a MAYA shape plugin node. Our plugin takes as an input any MAYA-modeled polygonal mesh comprised of triangles and quads and displays a shaded polygonal mesh using the limit positions and normals of the vertices at a given level of subdivision. To

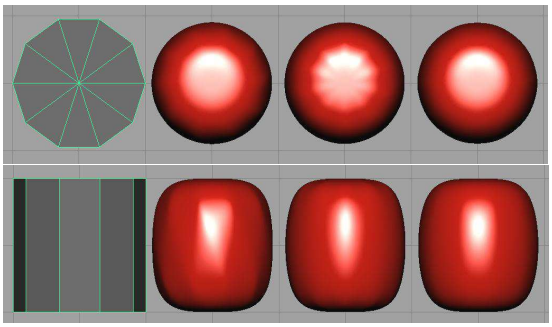

Figure 9: *Comparison of Loop (left), Catmull-Clark (middle) and our new scheme (right).*

compute limit positions and limit normals we use the left eigenvectors of the corresponding subdivision matrix (see Section 2.2). In the interior of the faces these limit masks are regular and well known. At irregular quad and triangle vertices we use the masks given in [4] and [5]. At the vertices that share triangles and quads we compute these eigenvectors numerically only once when reading in the base mesh. Figure 8 provides the limit masks for the regular quad/triangle boundary vertices.

Figure 9 demonstrates that our new scheme performs better than Loop and Catmull-Clark on a cylinder-like polygonal model. Figure 10 depicts different surfaces created using our new subdivision scheme. Note that the right-most pictures show a reflection-line plot of the surface, which provides an excellent visualization of the curvature.

5. Conclusions and Future Work

In this paper we have proposed a novel hybrid quad/triangle scheme which unifies Catmull-Clark and Loop surfaces in a single framework. We obtained our subdivision rules by decomposing the subdivision process into separate steps: linear subdivision followed by vertex-averaging followed by a vertex correction. We have shown that our surfaces are tangent plane continuous but not curvature continuous. However, we have introduced masks which optimize the behavior of the curvature. In general, for hybrid quad/triangle meshes our new scheme produces nicer surfaces than previous schemes.

In the future we intend to find an exact evaluation procedure for our surfaces similar to [8]. The evaluation depends on the ability to evaluate exactly along the regular quad/triangle boundary. Once this is established evaluation everywhere follows directly. We are also searching for a formal proof of the fact that our surfaces are C^1, perhaps by recasting the problem into the known quad or triangle framework or by inventing a new method of proof.

Finally, even though our schemes work well on hybrid quad/triangle meshes, they still perform poorly when the base mesh contains faces with a very bad aspect ratio. This is a more general problem that plagues uniform stationary subdivision. A topic for future research could be to come up with new rules which can handle any type of base mesh and still produce aesthetically pleasing looking shapes.

References

1. H. Biermann, A. Levin, and D. Zorin. Piecewise Smooth Subdivision with Normal Control. In *SIGGRAPH 2000 Conference Proceedings, Annual Conference Series*, pages 113–120, July 2000.
2. E. Catmull and J. Clark. Recursively Generated B-Spline Surfaces On Arbitrary Topological Meshes. *Computer Aided Design*, 10(6):350–355, 1978.
3. D. Doo and M. A. Sabin. Behaviour Of Recursive Subdivision Surfaces Near Extraordinary Points. *Computer Aided Design*, 10(6):356–360, 1978.
4. M. Halstead, M. Kass, and T. DeRose. Efficient, Fair Interpolation Using Catmull-Clark Surfaces. In *Proceedings of SIGGRAPH '93*, pages 35–44. Addison-Wesley Publishing Company, August 1993.
5. C. T. Loop. *Smooth Subdivision Surfaces Based on Triangles*. M.S. Thesis, Department of Mathematics, University of Utah, August 1987.
6. J. Maillot and J. Stam. A Unified Subdivision Scheme for Polygonal Modeling. *Proceedings of Eurographics 2001, Computer Graphics Forum*, 20(3):471–479, 2001.
7. U. Reif. A Unified Approach To Subdivision Algorithms Near Extraordinary Vertices. *Computer Aided Geometric Design*, 12:153–174, 1995.
8. 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.
9. J. Stam. On subdivision schemes generalizing uniform b-spline surfaces of arbitrary degree. *Computer Aided Geometric Design. Special Edition on Subdivision Surfaces*, 18:383–396, 2001.
10. J. Warren and H. Weimer. *Subdivision Methods For Geometric Design: A Constructive Approach*. Morgan Kaufmann, 2001.
11. D. Zorin and P. Schröder. A Unified Framework for Primal/Dual Quadrilateral Subdivision Schememes. *Computer Aided Geometric Design. Special issue on Subdivision Surfaces.*, 18:429–454, 2001.
12. D. N. Zorin. *Subdivision and Multiresolution Surface Representations*. PhD thesis, Caltech, Pasadena, California, 1997.

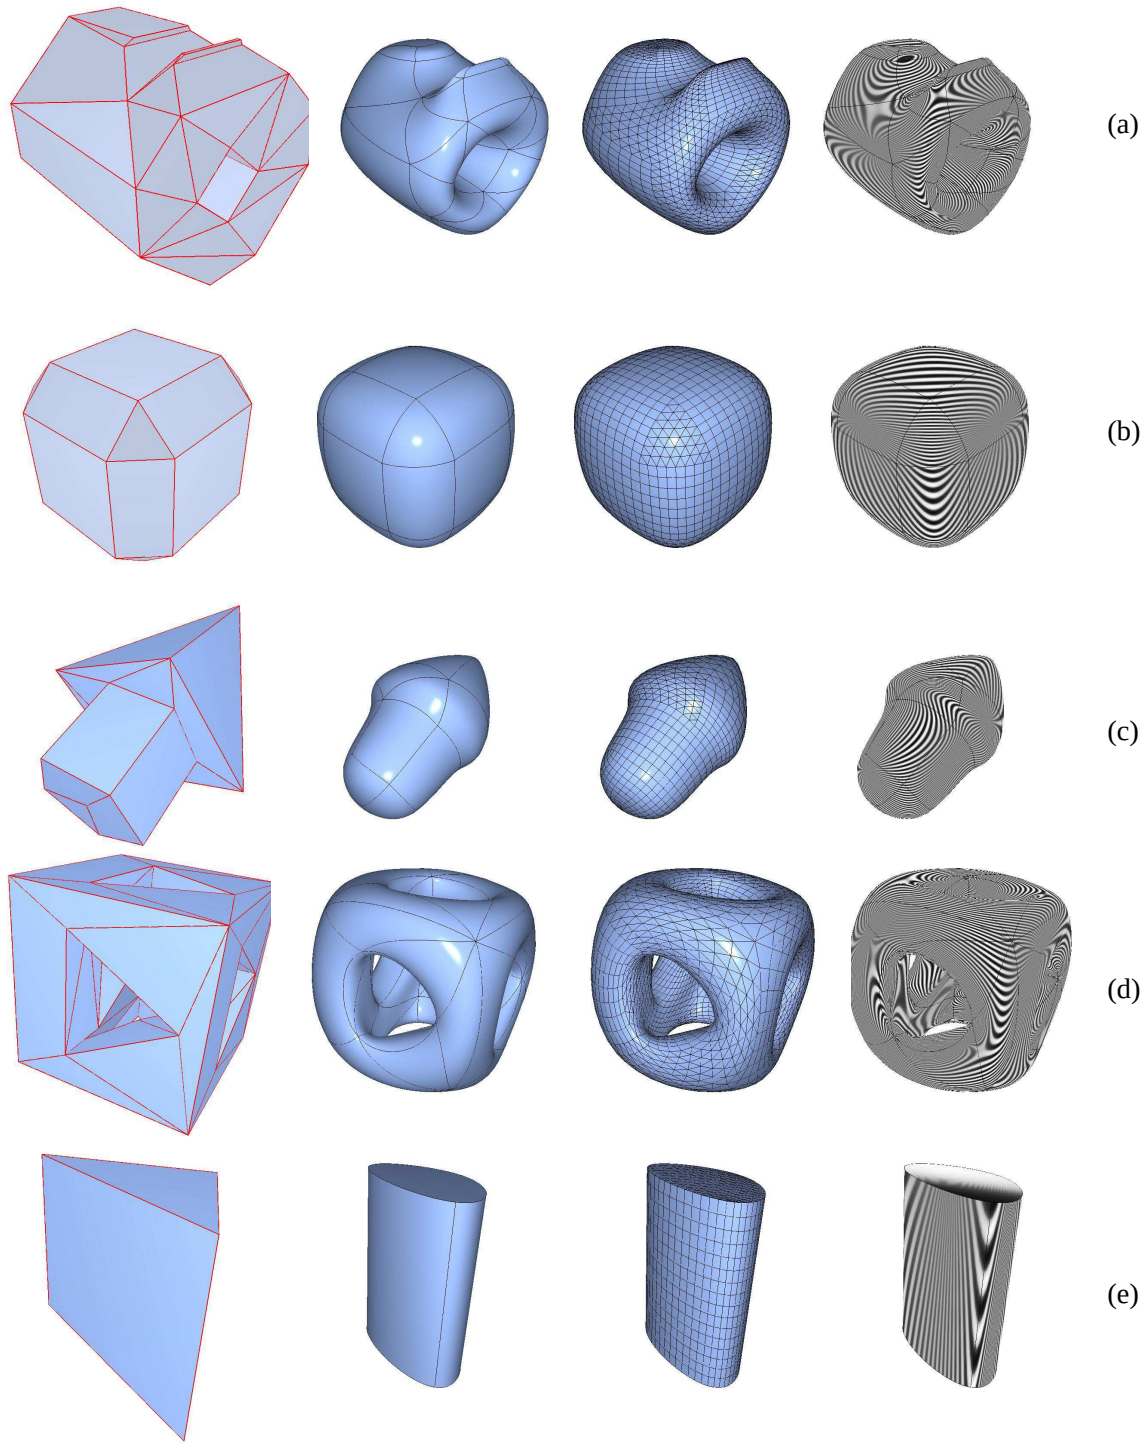

Figure 10: *Different surfaces generated using our new scheme.*