

Beyond Catmull–Clark?

A survey of advances in subdivision surface methods

Thomas J. Cashman

University of Lugano

Abstract

Subdivision surfaces allow smooth freeform surface modelling without topological constraints. They have become a fundamental representation for smooth geometry, particularly in the animation and entertainment industries. This survey summarizes research on subdivision surfaces over the last fifteen years in three major strands: analysis, integration into existing systems, and the development of new schemes. We also examine the reason for the low adoption of new schemes with theoretical advantages, explain why Catmull–Clark surfaces have become a de facto standard in geometric modelling, and conclude by identifying directions for future research.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Curve, surface, solid, and object representations; Splines

1. Introduction

Subdivision surfaces are smooth freeform surfaces which are generated using recursive rules (see Figure 1). The surface is specified using a coarse control mesh and, as a key advantage of the representation, this control mesh is not required to have

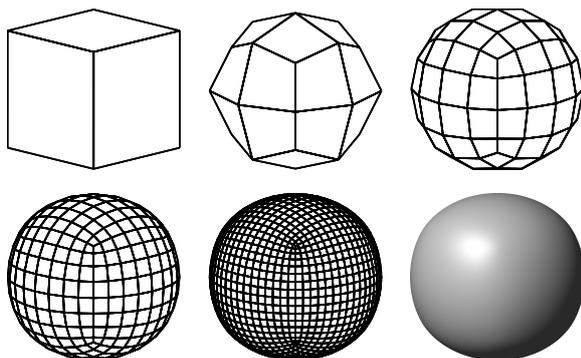


Figure 1: A subdivision scheme acting on a cube: the control mesh, the result of four subdivision steps, and the smooth limit surface. Extraordinary vertices (which are all of valency three, in this example) are preserved by each subdivision step, leading to singularities surrounded by regular surface.

a regular graph structure. This frees subdivision surfaces from topological constraints, and grants users a large amount of design freedom.

The first papers on subdivision surfaces were famously published by Catmull and Clark [CC78] and Doo and Sabin [Doo78, DS78] over thirty years ago. Subdivision surfaces have since become invaluable in entertainment applications such as animated films [DKT98] and special effects [NCP*09], and are increasingly important in real-time applications such as computer games [KMDZ09] as well. In the twenty years following their invention, a rich variety of subdivision schemes was developed, each one in the spirit of Catmull–Clark and Doo–Sabin in that they create a smooth surface by recursive application of simple rules. We review some of these basic schemes in Section 2.1, and also recommend the survey of the ‘subdivision zoo’ by Zorin et al. [ZSD*00].

In the last fifteen years, researchers have steadily addressed and overcome many of the limitations of these early schemes. This survey summarizes work in three major strands:

- theory for analysing subdivision surfaces,
- tools for integrating subdivision schemes into existing applications, and
- the development of surfaces with improved properties such as fairness, compatibility, or the removal of artifacts.

The two latter themes are those with practical applications, yet it is only the work on integrating existing schemes that has found widespread use, to the author’s knowledge. This raises the question of why early subdivision schemes are so enduring in applications where they have already been accepted, and continue to find relatively little use in applications where they have not. We might expect the development of schemes with theoretical improvements to lead to an increasing diversity of schemes, used in a wider range of applications. Instead we observe that Catmull–Clark (the very first subdivision scheme for surfaces) and Loop [Loo87] (the first subdivision scheme based on triangular patches) are overwhelmingly the most popular in practice. We conclude the survey by examining the reasons for this strong preference, and considering directions for future research.

2. Background

This paper complements earlier surveys of research in subdivision. Zorin et al. [ZSD*00] give a survey primarily aimed at computer graphics practitioners. Dyn and Levin [DL02] summarize the most important schemes with accompanying analyses for convergence and smoothness, while Sabin [Sab05] gives a comprehensive view, not just for subdivision surfaces but for univariate and trivariate schemes as well. In this paper we focus on research in the last fifteen years; this section gives a brief background to early work, but we refer the reader to other surveys for details.

2.1. Subdivision schemes

We can understand subdivision schemes in terms of two main properties:

- the rules that are used to insert points where the control mesh is regular,
- how those rules are generalized to allow for meshes with irregular connectivity.

We consider the first of these properties in this section and review important tools for understanding the second in Section 2.2. Here the definition of regular connectivity depends on the type of mesh refinement; in Figure 2, for example, all vertices and faces are regular. In general, regular regions have the same graph structure as a regular tiling of the plane, where the refinement pattern defines the polygons used in the relevant tiling.

Subdivision schemes were first discovered by generalizing knot insertion rules for spline surfaces on a regular grid. Knot insertion reproduces a given spline function on a refined grid, and subdivision rules can be extracted as the affine combinations of coefficients which are taken to form coefficients of the new basis. For example knot insertion on tensor-product B-splines of odd and even degree leads to the patterns of refinement shown in Figures 2a and 2b respectively. These were the patterns used by the Catmull–Clark and

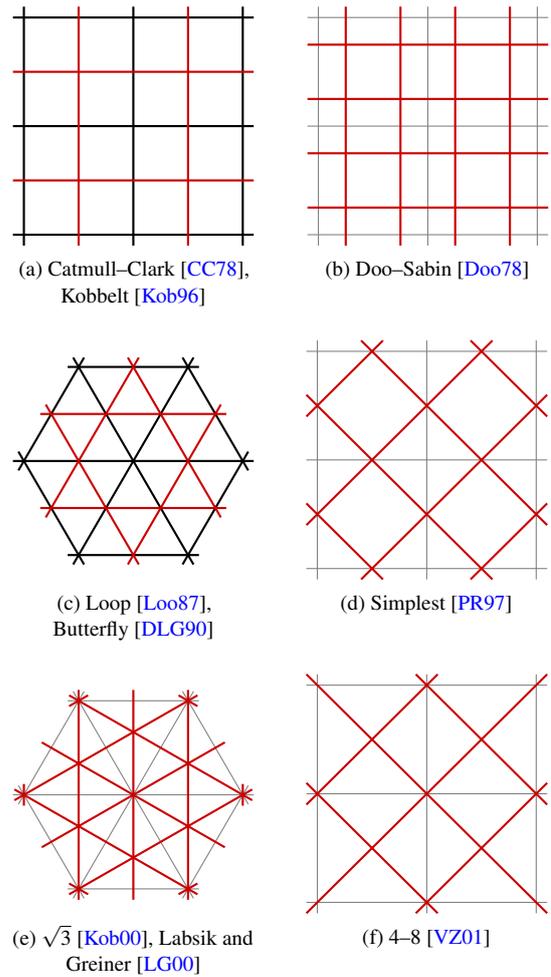


Figure 2: Regular refinement patterns for important subdivision schemes. New edges are drawn in red, and original edges appear in black where they form part of the refined pattern, or in grey where they do not.

Doo–Sabin schemes; the tensor-product structure means that both schemes operate on quadrilateral meshes.

Nearly ten years after these first schemes were introduced, Loop [Loo87] filled a natural gap for a subdivision scheme which uses triangular control meshes instead. The refinement pattern for Loop’s scheme is shown in Figure 2c, and Dyn et al. [DLG90] used the same type of refinement to create the first surface subdivision scheme where the limit surface interpolates a control mesh rather than approximating it. Kobbelt [Kob96] developed a similar scheme for quadrilateral meshes. Researchers have also found subdivision schemes with refinement patterns that incorporate a rotation of the grid directions. In particular, this includes the Simplest scheme by Peters and Reif [PR97], the $\sqrt{3}$ schemes by Kobbelt [Kob00]

and Labsik and Greiner [LG00] and the 4–8 scheme by Velho and Zorin [VZ01]. Comprehensive classifications of refinement patterns are provided by Han [Han03] and Ivrisimtzis et al. [IDS04].

While the Catmull–Clark and Doo–Sabin schemes generalize tensor-product B-splines, the schemes named in Figures 2c to 2f are based on a variety of surface types. Loop, Simplest and 4–8 all generalize knot insertion rules for box splines. The interpolatory Butterfly and $\sqrt{3}$ schemes sample cubic interpolants, in the same way as the four-point subdivision scheme for curves [DLG87]. In regular regions Kobbelt’s $\sqrt{3}$ scheme creates a non-polynomial surface with C^2 continuity but a fractal support for each basis function [ISD04].

There are also schemes which use the flexibility of subdivision to allow a mix of refinement patterns in the same mesh. Stam and Loop [SL03] and Schaefer and Warren [SW05] present schemes which preserve both quadrilateral and triangular faces: the Catmull–Clark scheme is used on the quadrilateral regions and Loop’s scheme on the triangular parts. These schemes therefore mix the refinement patterns shown in Figures 2a and 2c. The ‘4–3’ scheme by Peters and Shiue [PS04] also combines these two patterns, although using a box-spline-based subdivision scheme for the quadrilateral regions rather than Catmull–Clark.

2.2. Analysis at singularities

The second defining property of a subdivision surface scheme concerns the rules that are used around irregularities in the control mesh. Such an irregularity could be an *extraordinary vertex* (see examples in Figures 3a and 3c) or an *extraordinary face* (see Figures 3b and 3d). In both cases the *valency* of an extraordinary element is the number of incident edges. Irregularities are preserved by each subdivision step (see Figure 1), and therefore generate *singularities* surrounded by regular limit surface [PR08]. In this section we review the analysis of subdivision surfaces around such singularities.

The analysis depends on writing the linear map computed by a subdivision step as a matrix. If we consider the action of a subdivision step on the whole control mesh, then this matrix is taller than it is wide, as each step increases the density of the mesh. However, the subdivision rules operate locally, so at each step we can consider a fixed number of vertices around an extraordinary element to obtain a square subdivision matrix \mathbf{S} . If a subdivision scheme is *stationary*, then \mathbf{S} is constant for each subdivision step. For an initial control mesh \mathbf{Q} , the position of the singularity is therefore given by $\mathbf{S}^\infty \mathbf{Q}$. Doo and Sabin [DS78] observed that we can infer properties of \mathbf{S}^∞ using a diagonalization of \mathbf{S} , assuming that it exists. If a subdivision scheme is also *uniform* then the same subdivision rules apply in every part of the mesh, and around an extraordinary element of valency n , the action of a subdivision step therefore has a rotational symmetry of order n . Doo and Sabin showed that in this case, a Discrete

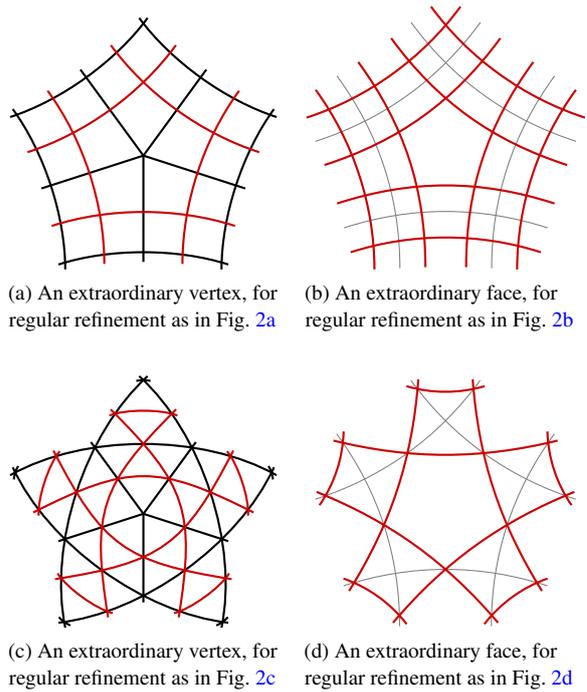


Figure 3: Example refinement patterns for the region around irregularities in the control mesh. As in Figure 2, original edges are drawn in black where they form part of the refined pattern, or in grey where they do not.

Fourier Transform (DFT) can simplify the analysis further, as the DFT $\hat{\mathbf{S}}$ of \mathbf{S} is a block diagonal matrix. We can therefore decompose a subdivision step into blocks $\hat{\mathbf{S}}_\omega$, each of which acts on the ω -th Fourier component of the input data. All of the schemes discussed in Section 2.1 are both stationary and uniform, and in this survey we assume that both properties apply, explicitly mentioning where a scheme is either non-stationary or non-uniform.

In order to arrive at a fixed non-zero limit surface, subdivision rules must be affine combinations (i.e. take weighted means). Each row of \mathbf{S} therefore sums to one, and so the vector of ones, $\mathbf{1}$, is an eigenvector with eigenvalue 1 (i.e. $\mathbf{S}\mathbf{1} = \mathbf{1}$). This eigenvector appears in $\hat{\mathbf{S}}_0$, as $\mathbf{1}$ is a constant (zero frequency) vector, and so we say that it has *Fourier index* 0. An eigenfunction is the limit of applying a subdivision scheme to an eigenvector, and the eigenfunction corresponding to this unit eigenvalue is the constant unit function.

To analyse a scheme further, we need some constraints on \mathbf{S} . In this survey we are interested in schemes with the following properties:

- The unit eigenvalue is dominant. The limit position $\mathbf{S}^\infty \mathbf{Q}$ is therefore given by taking this dominant component after representing \mathbf{Q} in the eigenbasis of \mathbf{S} . With the *left eigen-*

vector \mathbf{V} for the unit eigenvalue (which satisfies $\mathbf{V}\mathbf{S} = \mathbf{V}$), the limiting position is given by $\mathbf{V}\mathbf{Q}$ [HKD93]; \mathbf{V} therefore gives weights which find the singular point on the limit surface as an affine combination of the input points.

- After the dominant unit eigenvalue, the next largest eigenvalue is real and double. This corresponds to the fact that the space of bivariate linear functions has dimension two.
- The subdominant eigenvalue λ has Fourier index ± 1 ; the rotational symmetry which allowed us to use the DFT means that λ appears with equal value in both $\hat{\mathbf{S}}_1$ and $\hat{\mathbf{S}}_{-1}$.
- The space of bivariate quadratic functions has dimension three, and there are therefore three eigencomponents that form a corresponding basis. These eigencomponents capture quadratic properties of the limit surface at the singularity $\mathbf{S}^\infty \mathbf{Q}$, and have real eigenvalues.
- One of these eigencomponents has Fourier index 0. The associated eigenvalue, which we shall write as μ_0 , is subdominant in $\hat{\mathbf{S}}_0$ (recall that the *dominant* eigenvalue in $\hat{\mathbf{S}}_0$ is the unit eigenvalue).
- The other two quadratic eigencomponents correspond to the dominant eigenvalues in $\hat{\mathbf{S}}_{\pm 2}$; we shall write their value as μ_2 . As for λ , this eigenvalue is double as a result of rotational symmetry.
- All other eigenvalues are strictly less than the eigenvalues μ_0 and μ_2 in absolute value.

To summarize, in the discussion below we assume that a subdivision scheme is stationary and uniform, and that the square subdivision matrix \mathbf{S} has one of the following spectra:

$$1 > \lambda = \lambda > \mu_0 \geq \mu_2 = \mu_2 > \text{other eigenvalues}$$

$$1 > \lambda = \lambda > \mu_2 = \mu_2 \geq \mu_0 > \text{other eigenvalues}$$

Not all subdivision schemes fit this description. For example, although the Simplest scheme [PR97] is stationary and uniform, it has a subdominant eigenvalue which is eightfold rather than double, for an extraordinary face with valency 3. In this survey we are only interested in the analysis of schemes where the above properties hold, however, and it simplifies our discussion to assume that they do. For a more thorough mathematical treatment, we recommend the monograph by Peters and Reif [PR08], who derive the analyses below without our simplifying assumptions.

3. Analysis of subdivision surfaces

The first theme of work in this survey encompasses tools and theory for analysing subdivision surfaces. While the analysis in Section 2.2 was first presented by Doo and Sabin in 1978, recent work has continued to develop our understanding of subdivision surfaces and their properties.

3.1. C^1 regularity

Reif [Rei95] provides the first comprehensive analysis of subdivision surface smoothness by defining the *characteristic map* ψ that evaluates the two eigenfunctions corresponding

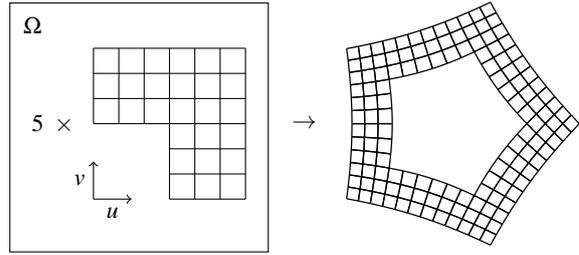


Figure 4: An example characteristic ring derived from a subdivision matrix for valency 5.

to λ . Since it is built from eigenvectors of \mathbf{S} , the characteristic map is composed of an infinite union of scaled *characteristic rings* [PR08] such as the one shown in Figure 4. Each characteristic ring is a map $\Omega \rightarrow \mathbb{R}^2$ from a parametric domain Ω ; in Figure 4, for example, Ω is a union of L-shaped pieces. For subdivision surfaces which are sufficiently smooth (C^1 and regular), an affine transform of the characteristic map provides a first-order Taylor approximation to the surface around a singularity, where the affine transform is given by the representation of \mathbf{Q} in the eigenbasis of \mathbf{S} . As in Section 2.2, we can find this representation using left eigenvectors of \mathbf{S} ; in this case the relevant left eigenvectors are those corresponding to λ .

Reif shows that a subdivision scheme generates C^1 regular surfaces if ψ is regular and injective, in addition to conditions on λ which are satisfied by our assumptions in Section 2.2. Zorin [Zor00b] extends this analysis for more general cases which do not satisfy those assumptions. Peters and Reif [PR98] use this theory to prove C^1 regularity for the Catmull–Clark and Doo–Sabin subdivision schemes, and Umlauf [Uml00] does the same for the Loop scheme. Both Zorin [Zor00a] and Peters and Reif [PR08] provide conditions which make it easier to verify regularity and injectivity of ψ for other subdivision schemes.

3.2. Higher regularity

Using the characteristic map, Prautzsch and Reif [PR99] are able to provide an important lower bound on the degree of *r-flexible* stationary polynomial subdivision schemes. They define a subdivision scheme as *r-flexible* if, at each singularity, there is some setting for \mathbf{Q} such that the surface has any given partial derivatives up to order r . This means that for a scheme to be *r-flexible*, it must be at least C^r . Zorin [Zor06] applies this definition of *r-flexibility* to a whole parametric surface by considering every point on the surface in the same way.

The important observation by Prautzsch and Reif is that for a subdivision scheme to be *r-flexible*, its eigenfunctions must span all r -degree polynomials of ψ . The characteristic map is itself generated by the subdivision scheme, and if it is

a C^k function then for general n it must have polynomial degree at least $k + 1$. The Catmull–Clark scheme, for example, has $k = 2$, since it generates curvature-continuous surfaces in regular regions, including the regular regions of the characteristic map. Now r -degree polynomials of this map have minimum degree $r(k + 1) = rk + r$, and this is therefore a lower bound on the degree of polynomials used in the subdivision scheme. We discuss schemes which meet this bound exactly in Section 5.6. Reif [Rei96] shows that this means stationary curvature-flexible polynomial schemes ($r = 2$) require at least degree 6, and Prautzsch and Reif [PR99] generalize the result for arbitrary r . By considering $r = 2$, this analysis also shows that for a subdivision scheme to be 2-flexible, we need restrictions on the eigenfunctions corresponding to μ_0 and μ_2 . This is analogous to the result that for a scheme to be C^1 -continuous, there are restrictions on the eigenfunctions corresponding to λ (i.e. the characteristic map; see Section 3.1).

These degree estimates are a consequence of the assumption that a subdivision scheme is uniform and stationary. Zulti et al. [ZLLT06] give an example scheme that overcomes the degree barrier by breaking these assumptions, generalizing the same quartic box spline as Loop [Loo87] yet still managing to be C^2 and 2-flexible. Their scheme is non-uniform, allows only a single extraordinary vertex in the mesh, and uses special subdivision rules along every chain of edges emanating from an extraordinary vertex. Peters and Karčiauskas [PK10] also discuss low-degree 2-flexible subdivision schemes by breaking the assumption on stationarity instead. They gain more flexibility by using an *accelerated* refinement, where a greater number of spline patches are introduced at every subdivision step. Myles and Peters [MP09] use this idea to generate a low-degree 2-flexible polar subdivision scheme that we discuss further in Section 5.4. In all of these cases, the combination of high smoothness and low degree is only possible by making the scheme either non-uniform or non-stationary.

To analyse their non-uniform scheme, Zulti et al. [ZLLT06] use the *joint spectral radius* for establishing Hölder continuity along the extraordinary edges, in the same way as Levin and Levin’s analysis [LL03] of Stam and Loop’s quad/triangle subdivision [SL03]. The joint spectral radius is a powerful tool for analysis, as it is able to give the exact Hölder regularity of a scheme, not just the number of continuous derivatives [Sab10]. Unfortunately, it is also very hard to compute an exact joint spectral radius in general [TB97], and so where this tool is invoked, we often gain only upper- and lower-bounds on Hölder regularity. The joint spectral radius is also only applicable for the analysis of subdivision *curves* or extraordinary edges: the technique cannot be used to analyse point singularities of subdivision surfaces.

In order to analyse the curvature properties of those singularities in the same way as the analysis for position and normals, Reif [Rei07] introduces the *embedded Weingarten*

map \mathbf{W} . For a C^1 subdivision surface, the position of a singularity is given as the limit of $\mathbf{S}^m \mathbf{Q}$ as $m \rightarrow \infty$, and the normal is defined in a similar way, by examining the convergence of a sequence of normal vectors. However studying curvature at the singularity is more difficult. One could try using the sequence of principal curvatures and directions, but the principal directions are undefined at umbilic points of the surface, where the principal curvatures are equal. It is therefore not necessary for the principal directions to converge for a surface to be curvature continuous, whereas convergence of \mathbf{W} , a 3×3 matrix, is both necessary and sufficient. The name ‘embedded Weingarten map’ is by analogy to the Weingarten map, which has principal directions and curvatures as its eigenvectors and eigenvalues respectively. However \mathbf{W} refers to coordinates in the embedding space instead of the surface tangent space, which makes it particularly well-suited for curvature analysis at singularities of subdivision surfaces.

3.3. Curvature of low-degree stationary schemes

For low-degree, uniform and stationary schemes like Catmull–Clark, the work described in Sections 3.1 and 3.2 means that the question of C^1 continuity is settled, and the possibility of C^2 , 2-flexible singularities is ruled out. However the second derivative can still exhibit a range of possible behaviours at singularities. Sabin et al. [SDHI03] summarize the most important options, including the situation where a subdivision matrix has the spectrum:

$$1 > \lambda = \lambda > \overbrace{\lambda^2 = \lambda^2 = \lambda^2}^{\mu_0 = \mu_2 = \lambda^2} > \text{other eigenvalues}$$

This gives the subdivision scheme a property known as *bounded curvature*, which is a necessary condition for a 2-flexible subdivision scheme [PR99]. It is not sufficient, because of the eigenfunction conditions described in Section 3.2. Subdivision schemes with bounded curvature do, however, preserve curvatures in all three of the quadratic eigencomponents through subdivision. This avoids several alternative outcomes, all undesirable:

- If $\mu_0 < \lambda^2$ and $\mu_2 < \lambda^2$, then the surface has a flat spot, as the quadratic components shrink faster than the square of the linear components. Prautzsch and Umlauf [PU98] use this spectrum as a way of forcing trivial curvature continuity, but the resulting surfaces are not 2-flexible, and the artifacts created using this enforced flatness are too severe for most practical purposes.
- If $\mu_0 > \lambda^2$ or $\mu_2 > \lambda^2$, then the surface has divergent curvature [DS78].
- If $\mu_0 > \mu_2$, then the surface has prescribed positive Gaussian curvature for almost all initial control meshes [PR04].
- If $\mu_2 > \mu_0$, then the surface has prescribed negative Gaussian curvature for almost all initial control meshes [PR04].

Bounded-curvature schemes, by contrast, allow extraordinary regions to hold an arbitrary non-zero curvature, just as in regular regions. However, if the eigenfunctions for the quadratic

components are not quadratic functions of the characteristic map (as is the case for any modifications to the Catmull–Clark subdivision rules where $n \neq 4$), then curvature at the singularity $\mathbf{S}^\infty \mathbf{Q}$ is undefined. Peters and Umlauf [PU01] show that the resulting curvature is bounded between values that depend on both \mathbf{S} and \mathbf{Q} .

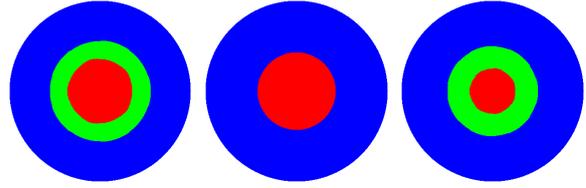
Karčiauskas et al. [KPR04] show that the Catmull–Clark scheme has $\lambda^2 < \mu_0 < \mu_2$ at all valencies greater than four, leading to a limit surface with divergent curvature, and hyperbolic shape for almost every possible control mesh. Reif and Schröder [RS01] show that despite these deficiencies, the principal curvatures are square-integrable for any scheme satisfying our assumptions in Section 2.2.

3.4. Shape analysis

Peters and Reif [PR04] analyse second-order properties further by means of a surface ring they call the *central surface* and define for each input \mathbf{Q} . This is the dominant term in the eigenbasis expansion of the surface around a singularity when written in a local coordinate system: the origin is placed at the limiting position $\mathbf{S}^\infty \mathbf{Q}$, and the coordinate xy -plane is set as the corresponding tangent plane. Therefore these terms (corresponding to the eigenvalues 1 and λ respectively) play no part in the z -value of the central surface, which is then determined entirely by the subsubdominant eigenvector(s). In our notation, the corresponding subsubdominant eigenvalue is either μ_0 , μ_2 , or the common value if they are equal.

If $\mu_0 = \mu_2$, Karčiauskas et al. [KPR04] consider the central surface for the complete range of inputs with quadratic shape, to build a *shape-in-the-limit chart*. For each input in the chart, they record whether the Gaussian curvature at the singularity is positive, negative, or *hybrid*, where the corresponding central surface is neither entirely elliptic nor hyperbolic. A subdivision surface with hybrid curvature therefore contains both positive and negative Gaussian curvature in every infinitesimal region of a singularity. This could be the outcome of a bounded curvature scheme, for example, if the bounded range for a particular input includes both positive and negative curvature. Augsdörfer et al. [ADS06] adapt shape-in-the-limit charts to additionally plot the range of Gaussian curvature in the central surface, and they propose a polar rather than barycentric layout for the chart (see Figure 5). Ginkel and Umlauf [GU08] make it easier to compute the resulting charts by analysing their rotational symmetries.

Another strand of work considers *artifacts* in a subdivision surface, defined as features of a surface that cannot be removed by modifications to the control mesh. Sabin and Barthe [SB03] categorize a wide range of artifacts and their sources, some of which are features of subdivision surfaces even on a completely regular grid. Augsdörfer et al. quantify this type of artifact for subdivision schemes on quadrilateral [ADS11a] and triangular [ADS11b] grids.



© Eurographics Association 2006

Figure 5: Ginkel and Umlauf [GU06] show shape-in-the-limit charts for a bounded curvature variant of Loop subdivision, in the polar layout suggested by Augsdörfer et al. [ADS06]. Red points show where the central surface has elliptic shape, blue indicates hyperbolic, and green points have hybrid curvature, for valency 5 (left), 6 (middle) and 7 (right).

3.5. Non-linear analysis

Although there is a growing body of work on non-linear and geometry-sensitive subdivision schemes for curves, at present there are few schemes of this type for surfaces (see Section 5.3 for some early examples). Nevertheless an important tool for analysing schemes of this type is the theory of *proximity* developed by Wallner and Dyn [WD05] for studying manifold-valued subdivision schemes for curves. The idea is that a non-linear scheme may converge, in the limit, to a linear scheme that we know how to analyse. If the convergence rate is high enough, then the two schemes can be shown to share the same continuity. Myles and Peters [MP09] use this type of argument to show C^2 continuity for a non-stationary surface scheme (see Section 5.4).

3.6. Approximations and proxy surfaces

To support new techniques in rendering, which we discuss in Section 4.1, we need to understand the properties of *approximations* to subdivision surfaces. Approximations are necessary because the only explicit form for a subdivision surface with control mesh irregularities involves an infinite number of surface patches. A graphics card can never render the exact geometry of such a surface, and so always uses an approximation instead. This may not be observable, as the approximation might be sufficiently accurate that there is no difference to the true subdivision surface in a rendered image. Since an approximation is used, however, we could always find a view of the same geometry such that a difference becomes apparent. Good rendering methods avoid all visible error by making the approximation view-dependent, but this just means that many different approximations are required rather than one.

The simplest available approximation is a triangulation, and it is common to render (possibly a triangulation of) the control mesh after a small number of subdivision steps. Several researchers have tried to find, a priori, the required num-

ber of Catmull–Clark subdivision steps for a given approximation error. Zeng and Chen [ZC06] provide estimates based on the first differences of the control mesh, but these are over-conservative and lead to a large prediction for the number of required steps. Cheng et al. [CY06, CCY06, CC06] make improved estimates based on second differences instead. Huang and Wang [HW07b] find optimal convergence rates for these second differences, but still predict a large number of subdivision steps compared to typical use in practice. They therefore consider the alternative approximation where each vertex of the control mesh is projected to its corresponding limit position, to give a *limit mesh*. We saw in Section 2.2 that the limit position of a vertex is given by a weighted combination of its surrounding vertices, using weights in the dominant left eigenvector \mathbf{V} . Huang et al. then provide subdivision depth estimates for this limit mesh approximation for both Catmull–Clark [HW07a, HDW08] and Loop [HW08] surfaces.

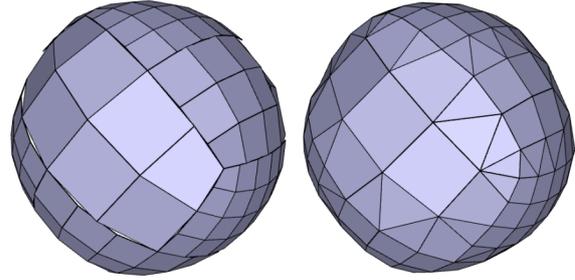
Peters and Wu [PW09] consider the problem for general subdivision surfaces, and use a reparametrization to show that the approximation error is proportional to $\max\{\mu_0, \mu_2, \frac{1}{4}\}^m$ for m subdivision steps. They also consider a posteriori estimates, based on measuring the error after subdivision, and recommend using a priori estimates only after one or two local subdivision steps. More generally, Peters and Reif [PR08] develop a theory of *proxy surfaces*, encompassing not only triangulations but also the higher-order approximations we discuss in Section 4.1.3. They draw the distinction between *parametric* distance, the distance between a subdivision surface and its proxy at common parameter values, and *geometric* distance, for example the Hausdorff distance between a surface and its approximation. Geometric distance removes the influence of the parametrization chosen for the proxy surface, and their results confirm that the convergence rate of geometric distance depends on the subsubdominant eigenvalue (μ_0 or μ_2).

4. Integrating subdivision surfaces into applications

The second theme of work in this survey covers tools for incorporating subdivision surfaces into existing hardware and software. The main example of this work is hardware rendering, which we discuss in Section 4.1. However there has also been work on providing subdivision surfaces with the full programming interface familiar from regular spline surfaces, including differentiable parametrizations, and tools such as Boolean operations and trimming curves. These we discuss in Section 4.2.

4.1. Improvements to rendering methods

Subdivision surface approximations can be categorized into two main classes. Some methods evaluate the exact value of the surface at given parameter values, while others approximate a subdivision surface over larger patches. One might



© IEEE Publications 2007

Figure 6: Schaefer and Warren [SW07] show the result of using separate linear approximations on each surface patch (left), and their dual tessellation procedure for generating a watertight approximation (right).

assume that exact computation results in higher rendering accuracy, but positions that are evaluated discretely must still be connected into a surface before being drawn on screen. In the process, it is possible to introduce greater error than a method which computes an approximation directly [NCP*09]. Therefore we cannot assume that exact evaluation gives a surface approximation with lower error. In this section we look first at methods for exact evaluation on the CPU (Section 4.1.1) and GPU (Section 4.1.2), followed by work that computes non-linear approximations directly (Section 4.1.3).

4.1.1. Evaluation

For subdivision schemes which generate a polynomial spline in regular regions, such as Loop and Catmull–Clark, we can evaluate the surface exactly if we have polynomial coefficients for the relevant surface patch. Near a singularity, this can involve an arbitrarily large number of subdivision steps, but Stam [Sta98] gives the first approach that can evaluate any point of the surface with a constant bound on computation time. He proposes using the same eigenbasis that we considered in Section 2.2, by projecting the input \mathbf{Q} into the eigenbasis of \mathbf{S} . In this basis, multiplication by \mathbf{S} is the result of simply multiplying coefficients by their corresponding eigenvalues. Zorin and Kristjansson [ZK02] develop this theme by pointing out that the same technique can be used on other bases; any block-diagonal form for \mathbf{S} with triangular blocks no larger than 3×3 still allows evaluation with constant-time complexity. They use this approach to evaluate the subdivision surfaces with creases introduced by Biermann et al. [BLZ00], where computing an eigenbasis is not always possible due to mesh parameters and tags which modify \mathbf{S} .

For schemes such as Kobbelt’s [Kob96] non-polynomial interpolating scheme, there is no closed regular form, and so the techniques of Stam [Sta98] and Zorin and Kristjansson [ZK02] are not applicable. To fill this gap, Schaefer and Warren [SW07] present evaluation procedures that also work for non-polynomial schemes. Their approach relies on the

scaling relation satisfied by refinement rules, to tabulate basis functions at rational parameter values. This allows them to calculate surface limit points using *stencils*, which take weighted combinations of control mesh vertices. They follow exact evaluation with a dual tessellation approach that gives a watertight linear approximation to the subdivision surface, even if adjacent patches are evaluated at different densities (see Figure 6). In a later paper [SW08] they also include the details for evaluating tangents to the surface. Bolz and Schröder [BS02] use similar pretabulated basis functions for high-performance evaluation, by making careful use of the CPU cache and data-level parallelism.

4.1.2. Evaluation on the GPU

The increasing power of graphics hardware has made it possible to render subdivision surfaces directly on a graphics card, transferring only the coarse control mesh from the CPU. This frees up a large amount of system memory for other purposes and makes it possible to consider subdivision surfaces in real-time applications such as computer games [KMDZ09]. The first techniques for hardware rendering mirror the exact rendering discussed in Section 4.1.1, but now implemented on the GPU instead.

Pulli and Segal [PS96] present an early hardware implementation of Loop subdivision for SGI geometry engines. They propose a compact index for mesh vertices that uses a simple array lookup for access to neighbouring regular vertices; this is a precursor to similar indexing schemes in all hardware implementations. Shiu et al. [SJP05] calculate recursive subdivision steps on the GPU by processing a surface *fragment* surrounding each vertex, while Bunnell [Bun05] renders Catmull–Clark surfaces by handling each *patch* of the surface separately instead. Like Schaefer and Warren [SW07], Bunnell gives a scheme for adaptively subdividing patches while maintaining a watertight tessellation of the surface, but using triangle fans rather than the dual approach shown in Figure 6.

Shiu et al. and Bunnell use a *fragment shader* to compute a fine tessellation of the surface, and this requires the graphics card to operate in two passes: one to calculate the tessellation, and another to render the generated surface to screen. Modern graphics cards are equipped with a *geometry shader* which allows this process to operate in just one pass; Kazakov [Kaz07] describes an implementation of Catmull–Clark evaluation using this more capable hardware.

4.1.3. Higher-order approximations

With the exception of Pulli and Segal [PS96], each of the hardware evaluation methods in Section 4.1.2 shares a need to subdivide a mesh once or twice on the CPU before sending it to the graphics card. This brings the set of possible mesh configurations down to a manageable level by ensuring that each patch is incident on at most one extraordinary vertex. However, the cost of this simplification is that the

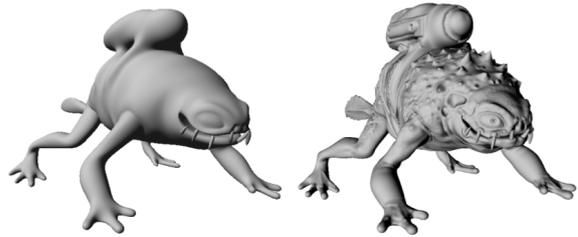


Figure 7: Loop and Schaefer [LS08a] demonstrate that their approximation to a Catmull–Clark subdivision surface (left) can be used for displacement mapping to obtain a more detailed surface (right).

mesh increases in size by 4 (subdividing once) or 16 times (subdividing twice) for Catmull–Clark and Loop subdivision. This, coupled with the necessity to approximate before rendering (see Section 3.6), has resulted in work to compute efficient *substitutes* for subdivision surfaces that give high-performance rendering without the complexity of exact evaluation. There are a wide range of approaches to higher-order approximation; we review the most important ideas here and Ni et al. [NCP*09] provide a fuller survey.

As the most popular subdivision scheme in practice, Catmull–Clark has received by far the most attention for work in this area. Peters [Pet00] converts each quadrilateral face of a Catmull–Clark mesh into a bicubic non-uniform B-spline patch, but may require one or two subdivision steps before doing so. Like Vlachos et al. [VPBM01], Loop and Schaefer [LS08a] take advantage of the fact that graphics hardware uses separate geometry and normal channels by providing different approximations to each channel. Their approximation uses bicubic Bézier patches for geometry and patches of degree 3×2 for each of two tangent fields. These fields can be used to generate a surface normal, or for displacement mapping [LMH00] (see Figure 7). Their work has been extended by Kovacs et al. [KMDZ09] to handle surface creases. Although not intended for high-performance rendering, Alexa and Boubekur [AB08] also present a similar idea. They create surfaces with smoother shading by replacing the true normal channel of a subdivision surface with one created using the same subdivision rules as those used for the geometry.

Another class of substitutes for Catmull–Clark surfaces use a true C^1 approximation, so the normal channel is simply filled with normals evaluated from the approximating surface. Myles et al. [MYP08] achieve this using biquintic patches around extraordinary vertices, expressed as biquintic perturbations of a bicubic surface. Ni et al. [NYM*08] are able to use lower degree by using *composite patches* instead. Both these approaches require a mesh with only quadrilateral faces, but Loop et al. [LSNCn09] provide an approximation which can also incorporate triangular faces, with the option

of using approximate normals instead of the true normal field evaluated from the surface. Myles et al. [MNP08] go further, extending their previous work to handle both triangular and pentagonal faces.

4.2. Adapting spline tools for subdivision

With many good approaches for evaluation (Section 4.1), subdivision surfaces can be treated like any other parametric spline surface. However, meshes with irregular connectivity create new problems, and some new opportunities, for several of the tools which are familiar from regular spline surfaces.

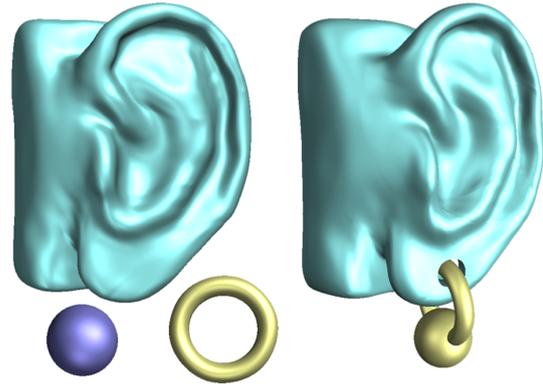
Litke et al. [LLS01] address the application of *trimming* subdivision surfaces. For regular spline surfaces it is necessary to compute the parametric *preimage* of a trim curve so that the relevant part of the surface can be excluded for evaluation; this can be difficult and lack robustness. Litke et al. point out that with Levin's [Lev99] combined subdivision scheme, which can satisfy boundary constraints, it is possible to meet a desired trim curve exactly without computing an exact preimage. The compromise they make is a modification (within a specified tolerance) of the surface near the trim curve.

Boier-Martin and Zorin [BMZ04] address parametrization of subdivision surfaces. Although we know that subdivision surfaces like Catmull–Clark are C^1 regular surfaces, and therefore possess a C^1 parametrization (Section 3.1), the *natural parametrization* from scaled copies of Ω (see Figure 4) may not even be differentiable. Boier-Martin and Zorin point out that

$$\frac{\partial \psi}{\partial u} \left(\frac{u}{2}, \frac{v}{2} \right) = 2\lambda \frac{\partial \psi}{\partial u} (u, v)$$

and likewise for partial derivatives in the v direction. Since the subdivision surface is a weighted sum of its eigenfunctions (including ψ), and for Catmull–Clark $\lambda > \frac{1}{2}$ wherever the valency $n > 4$, this means that the natural parametrization has divergent derivatives at every extraordinary vertex with valency greater than 4. They propose two alternative parametrizations which solve this problem: one with vanishing derivatives and the other, based on inverting the characteristic map ψ , which has a non-degenerate Jacobian. For computing a *data-dependent* parametrization with minimal distortion (for use in texturing, for example), He et al. [HSH10] also give a method specific to subdivision surfaces.

There are a sequence of related problems in *interference* detection and *intersection* calculation. In simulation, it may be important to know whether two surfaces intersect but unnecessary to calculate the exact intersection curve. This is the problem of interference detection, which is solved approximately by DeRose et al. [DKT98] for Catmull–Clark and exactly by Grinspun and Schröder [GS01] for Loop subdivision surfaces. Wu and Peters [WP04] give a more efficient approach which also applies to a larger class of subdivision surfaces.



© Association for Computing Machinery, Inc. 2001

Figure 8: An example, from Biermann et al. [BKZ01], of Boolean operations on subdivision surfaces. They calculate a union of the sphere and the torus, and an enlarged version of the torus is subtracted from the ear.

Lanquetin et al. [LFKN03] propose an algorithm for the situation where an exact intersection curve is required. Their approach may not detect all intersections, however, and Severn and Samavati [SS06] repair this defect. Once an exact intersection curve has been calculated, it can be used to carry out Boolean operations: Biermann et al. [BKZ01] give the first details of this procedure for subdivision surfaces (see Figure 8). Later work attempted to improve robustness by using voxelization [LC07] or by operating on the *limit mesh* [JS09] (see Section 3.6 for a definition).

5. New subdivision surface schemes

The third and final theme in this survey covers the construction of new subdivision schemes. The new schemes satisfy a variety of requirements and constraints, and they appear in this section roughly ordered by similarity to the basic schemes in Section 2.1. The work ranges from bounded curvature schemes (Section 5.1), which only require small modifications to the weights used as part of a subdivision step, to high-continuity schemes (Section 5.6), which can have a significantly different algorithmic structure.

5.1. Bounded curvature and tuning

The Doo–Sabin scheme was constructed with bounded curvature (see Section 3.3) from the beginning [DS78], but the Catmull–Clark scheme was not, and Loop's scheme satisfies $\lambda^2 = \mu_0$ but not $\lambda^2 = \mu_2$, so it is not a bounded-curvature scheme either. For the Catmull–Clark scheme, Sabin [Sab91] presents the first bounded-curvature variant, and Holt [Hol96] does the same for Loop's scheme. Loop [Loo02] presents another bounded-curvature variant of his own scheme that retains the *convex hull property*: the property that the surface is contained entirely within the convex hull of the control

mesh. This is guaranteed if the affine subdivision weights are non-negative.

From the artifact and shape analysis in Section 3.4 there arise a wide range of criteria that we might want a subdivision scheme to satisfy. Barthe and Kobbelt [BK04] treat subdivision weights as degrees of freedom in a nonlinear optimization, and manipulate the subdivision matrix \mathbf{S} towards desirable properties for its eigenstructure, including bounded curvature. They coin the name *tuning* for this kind of subdivision rule optimization.

A bounded-curvature scheme prevents the collapse or explosion of curvature near singularities, but the curvature can still have large oscillations, as Peters [Pet02] observes. Augsdörfer et al. [ADS06] explicitly minimize these oscillations by tuning bounded-curvature variants of several subdivision schemes. They minimize the variation of Gaussian curvature in the central surface (see Section 3.4) for a sampling of quadratic shapes; the end result is still just a modification of the weights used in each subdivision step. Ginkel and Umlauf [GU06] also use shape-in-the-limit charts for tuning Loop's scheme, but they try to eliminate hybrid shapes altogether by correcting input data that would lead to a hybrid limit surface. Unfortunately this is not always possible in general: for very high valencies, Ginkel and Umlauf [GU08] show that an entire shape-in-the-limit chart can be hybrid, leaving no possible correction that would lead to non-hybrid shape.

5.2. Modified for C^2 by blending or finite patching

The analysis in Section 3.2 shows that no tuning of the Catmull–Clark or Loop subdivision weights can give surfaces which are 2-flexible at singularities, as the polynomial degree is too low to achieve this using a stationary subdivision scheme. Given this impossibility, several researchers consider alternative modifications instead. Levin [Lev06] and Zorin [Zor06] both present methods to give 2-flexible surfaces by smoothly blending the subdivision surface with another, best-fit C^2 surface. Levin's technique is described for Catmull–Clark, while Zorin's uses the example of Loop's scheme.

In Section 4.1.3 we considered finite approximations to subdivision surfaces for the purposes of efficient rendering, but it is also possible to create a finite collection of patches that have a higher continuity than a given subdivision scheme. Loop [Loo04] presents a curvature-continuous modification for Catmull–Clark surfaces by using n bisepic patches around each singularity of valency n . These extraordinary patches mimic the shape of the target Catmull–Clark surface. This method again shows that the cost of high flexibility at singularities is usually high polynomial degree, also in the finite setting. Loop and Schaefer [LS08b] improve the approximation by making use of unconstrained degrees of freedom to give fairer finite patchings, and also to handle surfaces with boundary.

5.3. Interpolatory subdivision

Subdivision surface schemes either *approximate* a control mesh or *interpolate* it. Early interpolating schemes include Butterfly [DLG90] and Kobbelt's [Kob96] scheme, both of which are C^1 in regular regions. Zorin et al. [ZSS96] also present a modified version of the Butterfly scheme with the same continuity class but higher visual smoothness.

For approximating schemes such as Catmull–Clark and Loop, the limit surface does not pass exactly through the control mesh, but roughly follows the same shape. However several authors adapt or modify approximating schemes to interpolate a control mesh instead. Li and Ma [LM07] give schemes which are a blend between a given approximating scheme and an interpolating version, by interpolating the differences between a control mesh and its refinement. As examples they give interpolating variants of Catmull–Clark, Loop, and $\sqrt{3}$ subdivision. Other approaches find a limit surface that interpolates the input data by modifying the control mesh instead, thus gaining interpolating surfaces with curvature continuity. Halstead et al. [HKD93] find interpolating Catmull–Clark surfaces by solving a global linear system which additionally minimizes a quadratic fairness measure. Maekawa et al. [MMN07] also construct interpolating surfaces from approximating schemes, but using iterative corrections based on closest-point computations instead.

Schaefer and Warren [SW03] present an interpolating scheme for quadrilateral meshes by factoring the four-point scheme [DLG87] into differencing and averaging passes. This is similar to the *refine and smooth* factorizations that are used to create general-degree subdivision surfaces, described in Section 5.5. Schaefer [Sch03] also extends this work to produce interpolating surfaces of revolution, based on the approximating scheme by Morin et al. [MWW01] for surfaces of revolution.

Several non-linear schemes have appeared which generate interpolating surfaces using geometric constructions. Dyn et al. [DLL92] use a geometric construction to create smooth convexity-preserving surfaces. Korbacher et al. [KSH00] and Yang [Yan05] present non-linear schemes for the refinement pattern shown in Figure 2c, and Dodgson et al. [DSS07] unsuccessfully propose a $\sqrt{3}$ scheme based on sampling spheres. Although sensitivity to local geometry may be the key to simple constructions with high fairness, so far these non-linear schemes have not been able to consistently outperform their linear counterparts.

5.4. Polar subdivision

Karčiauskas and Peters [KP07c] identify a *polar* control mesh configuration which is not supported by classical subdivision schemes (see Figure 9). They present an extension to Catmull–Clark subdivision [KP07a] which adds support for polar configurations and gives bounded curvature at polar vertices. Myles and Peters [MP09] go further by creating a

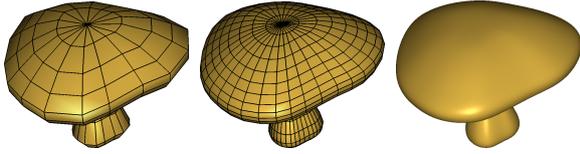


Figure 9: Myles and Peters [MP09] demonstrate the benefits of polar subdivision using this mushroom model: control mesh (left), one subdivision step (middle) and the limit surface (right). The example polar configuration lies at the top of the mushroom, where many isoparameter lines converge in an extraordinary vertex surrounded by triangular faces.

modification of Catmull–Clark with C^2 and 2-flexible polar singularities. They manage to circumvent the high degree estimate we encountered in Section 3.2, by using subdivision rules which increase the valency of a polar vertex at every subdivision step (as shown in Figure 9). This gives a non-stationary scheme which cannot be analysed using an eigenbasis as in Section 2.2. Instead, Myles and Peters prove their scheme is C^2 using the techniques of proximity [WD05], adapted for their scenario where the scheme is non-stationary in the mesh connectivity, rather than in the subdivision weights.

5.5. Schemes generalizing tensor-product B-splines

The polynomial subdivision schemes we have considered so far all generalize B-splines or box splines of relatively low degree. However several researchers realized that the Lane-Riesenfeld refine and smooth algorithm [LR80] for subdividing *regular* arbitrary-degree B-splines could be generalized for *irregular* meshes as well. Prautzsch [Pra98] and Warren and Weimer [WW01] describe the natural generalization where each smoothing stage replaces a face with its barycenter, and Zorin and Schröder [ZS01] show that the resulting subdivision surfaces are C^1 at singularities for degrees ≤ 9 . Stam [Sta01] and Stewart and Foisy [SF04] address some practical considerations by describing variants where the topology of the mesh is invariant under smoothing, while Prautzsch and Chen [PC11] prove C^1 continuity at all degrees ≥ 2 . In regular regions, all of these schemes generate tensor-product B-splines of any specified degree d , and are therefore C^{d-1} . At singularities, we know from the Prautzsch and Reif degree estimates (Section 3.2) that all are only C^1 . Oswald and Schröder [OS03] extend the same refine and smooth technique to non-polynomial schemes with a framework that incorporates general subdivision rules, such as the $\sqrt{3}$ subdivision operator shown in Figure 2e.

These schemes all generalize surfaces with *uniform* parametrizations, where each surface patch has equal edge lengths in parameter space. Sederberg et al. [SZSS98] describe the first *non-uniform* subdivision surfaces, which they call NURSS. Their knot insertion rules specialize to the

Doo–Sabin and Catmull–Clark rules in the uniform case, but also allow the surface to represent non-uniform biquadratic or bicubic B-splines exactly. Generalizing non-uniform B-splines brings all the benefits of non-uniform parametrizations, such as boundaries that meet a given B-spline curve while retaining cross-boundary tangent and curvature control, or selective reduction of continuity in the interior of a surface. However, Qin and Wang [QW99] show that for an extraordinary face with high valency and large variation in knot spacing, the biquadratic NURSS schemes may not even converge to a limit point.

Müller et al. [MRF06] give a different variant of Catmull–Clark which generalizes non-uniform bicubic B-splines and makes it possible to evaluate the limit surface at any given vertex. Müller et al. [MFR*10] extend this work to give a stationary subdivision matrix around control mesh vertices, making it possible to apply the eigenanalysis in Section 2.2. Sederberg et al. [SZBN03] also develop their non-uniform construction further by allowing a control mesh to contain T-junctions; they called the resulting surfaces *T-splines*. Karčiauskas and Peters [KP09a] analyse a different non-uniform variant of Catmull–Clark, which generalizes only uniform B-splines, but allows for non-uniform ‘adjustable speed’ spline rings around a singularity.

Subdivision surfaces which generalize non-uniform and arbitrary-degree B-splines bring us closer to Non-Uniform Rational B-Splines (NURBS), the standard freeform surface representation for Computer-Aided Design [PT87]. NURBS have no constraint on polynomial degree or parametrization, and so a NURBS-compatible subdivision scheme must be both non-uniform and arbitrary-degree. Cashman et al. [CAD09] present the first scheme of this type for odd degrees, with restrictions on multiple knots near singularities. These restrictions allow them to selectively insert knots to create a uniform configuration around extraordinary vertices. Although their scheme has been analysed only for the uniform case, this uniformization procedure allows their analysis to apply to surfaces with more general starting parametrizations. Except for the modifications that they introduce to give bounded curvature, this NURBS-compatible scheme also has an even-degree counterpart [Cas10]. Figure 10 shows how the schemes discussed in this section relate to (subsets of) NURBS.

5.6. High-continuity schemes

The analysis in Section 3.2 raises the question of whether it is possible to find schemes that have exactly degree $rk + r$: the minimum degree for a stationary, polynomial, r -flexible subdivision scheme which is C^k in regular regions. Both Prautzsch [Pra97] and Reif [Rei98] answer this question in the affirmative, with a particular focus on curvature-flexible schemes with C^2 continuity. Their degree estimates mean that such schemes use patches which are at least bisextic. Both schemes are finite constructions, although they allow for sub-

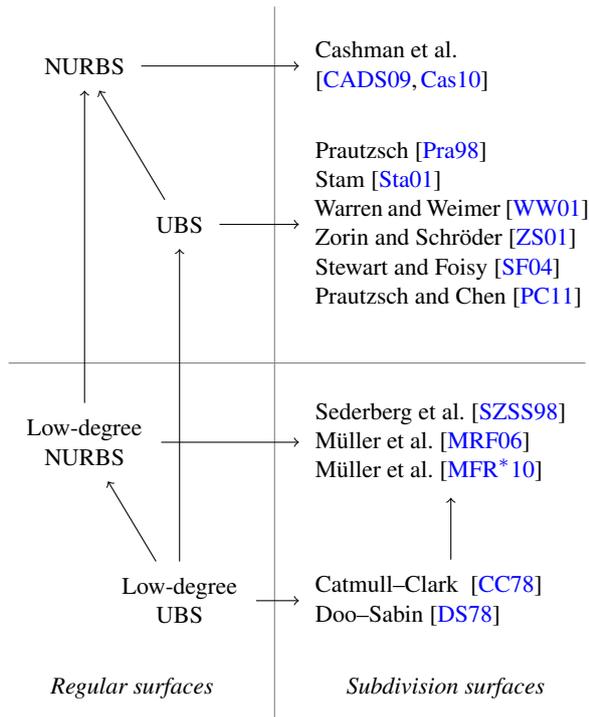
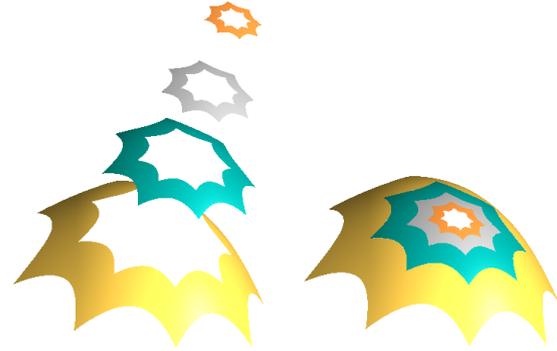


Figure 10: An overview of subdivision surfaces based on B-splines. This diagram shows classes of surfaces with subset relations between them (where \rightarrow represents \subset). The acronym NURBS is shortened to UBS for ‘Uniform B-Splines’. Surfaces above the horizontal line are based on general-degree B-splines; those below it are limited to biquadratic and bicubic degrees.

division algorithms as well. They are therefore concerned with filling an n -sided hole with n bisextic patches. Both also choose a parametrization, essentially turning the characteristic map into an input rather than a result of analysis. Freeform splines, the scheme by Prautzsch [Pra97], simply chooses any C^2 planar ring, such as the image $\psi(\Omega)$ of the map shown in Figure 4, to act as the parametrization of the input ring. Reif [Rei98] uses a particular singular parametrization instead.

As an example, we now focus exclusively on freeform splines, as it is the simpler of the two schemes to describe. With the set-up outlined so far, Prautzsch’s construction now chooses any quadratic polynomial \mathbf{p} for the region around the singularity, and composes \mathbf{p} with the parametrization to gain bisextic patches. The final step is a modification of the input ring so that the new patches join smoothly to the input: Prautzsch recommends a least-squares fit for choosing \mathbf{p} so that the required modifications are small. As a complete process, we can view this construction as building a custom stationary subdivision scheme \mathbf{S} for each input \mathbf{Q} .



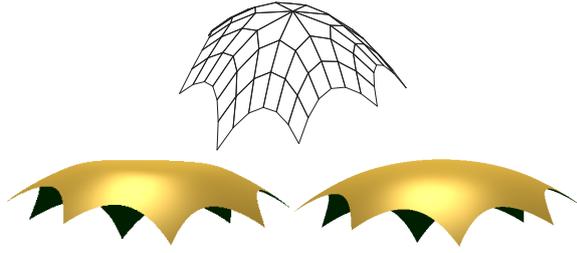
© Elsevier B. V. 2007

Figure 11: Karčiauskas and Peters [KP07b] show a sequence of guided surface rings (left) and the resulting guided surface (right).

Peters [Pet02] shows that although this gives a perfect, freely-chosen quadratic surface at the singularity, curvature artifacts tend to move outwards to the *transition layer* between the extraordinary and regular part of the surface. In response Karčiauskas and Peters [KP07b, KP08] develop *guided surfaces* (see Figure 11), where an arbitrary shape \mathbf{p} is chosen for the extraordinary region in a similar way, but this surface is only sampled: it does not form part of the limit surface directly. In particular, they propose a scheme which samples \mathbf{p} at the corners of the infinite collection of parametric rings which are scaled copies of Ω . These rings become progressively denser as the subdivision surface moves towards a singularity, which means that the limiting behaviour is given entirely by the guide surface. However low-frequency sampling in the outer rings allows a fairer join with the input. Although this framework requires bisextic surfaces for curvature continuity, Karčiauskas and Peters [KP09b] also show that the same algorithm with lower-degree surfaces can give results which are visually very similar.

The PTER framework that appears in the monograph by Peters and Reif [PR08] is based on the guided surface approach, and is intended as a general way to construct C^k subdivision schemes. Given a map ψ such as the one shown in Figure 4 (which need not be a characteristic ring, but must have a similar structure), the idea is to compose four operators, in order:

- R** Reparametrize the input ring so that instead of a map from Ω to \mathbb{R}^3 , it becomes a map from $\psi(\Omega)$ to \mathbb{R}^3 .
- E** Extend the resulting function inwards from $\psi(\Omega)$ to $\lambda\psi(\Omega)$, where λ is the scaling factor corresponding to the map ψ .
- T** Turn-back so that we have a new ring as a function from $\frac{1}{2}\Omega$ to \mathbb{R}^3 .
- P** Project the resulting ring into the space of rings which join C^k with the input.



© Elsevier B. V. 2008

Figure 12: Karčiauskas and Peters [KP08] compare two bicubic subdivision surfaces defined by the same control mesh (top): a bounded-curvature variant of Catmull-Clark [ADS06] (bottom left), and a bicubic guided surface (bottom right).

Peters and Reif show that guided subdivision fits almost trivially into this framework, where the operator \mathbf{E} just yields the guide surface on the scaled domain $\lambda\psi(\Omega)$, and the operator \mathbf{P} samples the guide, and its derivatives, in a way that gives C^2 -connecting rings. Naturally this process can be iterated to gain a sequence of rings which fill an n -sided hole in the same way as a subdivision surface, as shown in Figure 11.

6. Subdivision schemes in practice

Despite the large body of research described in Section 5, at present the dominant schemes in subdivision surface modelling remain Loop [Loo87] and Catmull-Clark [CC78]. This is reflected for Catmull-Clark, in particular, by the large amount of work in this survey which applies specifically to Catmull-Clark subdivision, in all three areas of analysis, integration, and modification to create new schemes. Perhaps surprisingly, subdivision schemes used in practice almost always appear in their original, unmodified forms, even though these early schemes feature the avoidable curvature problems discussed in Section 3.3.

First we consider why Catmull-Clark is not replaced with superior alternatives in industries which have already adopted subdivision surfaces. Some non-technical reasons include:

- that the amount of work specialized particularly for Catmull-Clark creates an ecosystem (and large amounts of highly-optimized code) which may be difficult to change;
- that as a de facto standard, Catmull-Clark subdivision creates an expectation for how a subdivision surface should behave.

We find an example of the latter in Karčiauskas and Peter's comparison [KP08, Fig. 11] between bicubic guided surfaces and bounded-curvature subdivision (see Figure 12). The guided surface in this comparison passes through the same limit point as a Catmull-Clark surface, and given a certain control mesh, they show that a tuned bounded-curvature variant of Catmull-Clark [ADS06] creates a surface that is

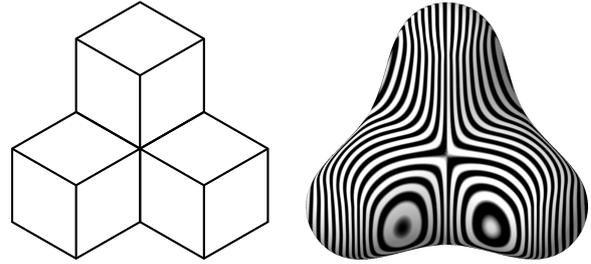


Figure 13: Catmull-Clark surface defined by the control mesh shown left and drawn with reflection lines (right).

considerably flatter than an untuned surface. This is not a surface artifact (Section 3.4), as the flatness can be ameliorated by modifying the control mesh, but the bounded-curvature variant may appear deficient simply by behaving differently. Alternatively we could argue that the bounded-curvature variant is less desirable because an extraordinary vertex is treated differently to regular vertices: here it seems to have a smaller effect on the limit surface. However the same is true of any subdivision scheme, so this argument may be more subjective than it first appears.

The most obvious *technical* reason for the dominance of Catmull-Clark is

- that it is good enough for purpose, and the theoretical shape problems discussed in Section 3.3 rarely appear in practice.

When used in animation, careful control mesh design eliminates high-valency singularities, or at least restricts them to flat or hidden parts of the surface [SAP01]. This means that the most important valencies are 3 and 5, and here the values of λ^2 , μ_0 and μ_2 differ by less than 0.7% and 12.5% respectively. For these close-to-regular valencies, curvature is therefore not too far from bounded, and at the scale of Figure 13, for example, the unbounded and generically zero curvatures are not visible at all, even when examining the surface with reflection lines. One final explanation for Catmull-Clark's longevity is its

- simplicity and generality.

The alternatives may be perceived as considerably more complex (e.g. guided subdivision; Section 5.6), to have limited applicability (e.g. C^2 polar subdivision; Section 5.4), or to offer dubious benefits (e.g. bounded curvature schemes, in Section 5.1, which can suffer from curvature oscillations [Pet02]).

These arguments apply equally to Loop subdivision, since it is the de facto standard for subdivision surfaces on a triangular mesh, and there is also a large investment into tools and code specifically for Loop surfaces. Loop even has bounded elliptic curvature at singularities, since $\lambda^2 = \mu_0$. Although the value of μ_2 leads to curvature that is generally *unbounded* at

valencies greater than 6 [KPR04], for close-to-regular valencies of 5 and 7, μ_2 differs from λ^2 by 18.4% and 13.3% respectively. So here too, curvature is not too far from bounded, and Loop subdivision certainly shares the simplicity and generality of Catmull–Clark.

For industries where subdivision surfaces do not yet have a foothold, the question is why none of the more recent schemes have been adopted. Inertia may again play a role; as an example the Computer-Aided Design (CAD) industry supports a large body of highly robust code which does not easily accommodate new representations [SFL*08]. However there is clearly a demand for more flexible modelling in engineering: Dassault Systèmes introduce subdivision surfaces in their CATIA product *Imagine and Shape*, and T-splines [SZBN03] make subdivision surfaces available to CAD packages such as Rhinoceros 3D and SolidWorks. One barrier to adoption is the high-level interface these modellers present to users: tools such as intersections, offsets, fillets and blends [IGPP01] leave a typical CAD model as the union of many approximating trimmed NURBS surfaces [SFL*08]. There is nothing to gain from replacing many trimmed NURBS surfaces with many trimmed subdivision surfaces, and it is not clear that subdivision surfaces make it any easier to provide these high-level operators. Nevertheless we do not believe that the potential of subdivision surfaces has been exhausted for all possible applications. In particular high-fairness schemes such as guided surfaces [KP08] may prove to be a powerful solution to future problems in surface modelling.

7. Conclusion

Subdivision surfaces are now a mature technology, and linear subdivision schemes are well understood. Possible future directions for research include:

- high-fairness, non-linear geometry-sensitive subdivision schemes with simple rules, extending the work on non-linear interpolatory schemes discussed in Section 5.3;
- a unified analysis for these non-linear schemes, extending the work in Section 3.5;
- schemes which fully realize the promise of arbitrary-topology NURBS (Section 5.5) with at least bounded curvature, and without restrictions to odd degrees or certain knot configurations;
- standardization and recommendations for the large number of choices that must be made to implement a scheme in the PTER or freeform spline frameworks (Section 5.6).

Even if geometry-sensitive schemes start to fulfil their promise, or high-continuity schemes become more widely adopted, it seems clear that it will be impossible to match the unique combination offered by Catmull–Clark surfaces: a remarkably simple and predictable construction, with highly-efficient implementations, and which generalizes an equally simple and efficient closed-form regular surface. There are many cases where Catmull–Clark surfaces are smooth

enough, and with increasing real-time applications, we conclude that Catmull and Clark [CC78] present an algorithm which is surprisingly close to optimal, for one of the first papers on subdivision surfaces.

Acknowledgements

This paper was written with the support of the SNF under project number 200021-134639. I am grateful for the assistance of Kai Hormann and the anonymous reviewers, in providing useful feedback during the preparation of this survey. I also thank the copyright holders and authors listed below, all of whom gave permission for their work to be reproduced.

Figure 5 appears with the kind permission of Ingo Ginkel, Georg Umlauf and the Eurographics Association.

Figure 6 was provided by Scott Schaefer and Joe Warren, with the permission of the IEEE.

Figure 7 appears courtesy of Charles Loop, Scott Schaefer and Bay Raitt of Valve Software, who provided the ‘Monsterfrog’ model shown in this figure.

Figure 8 is reproduced with the permission of Denis Zorin and the ACM.

Figure 9 was kindly provided by Ashish Myles and Jörg Peters.

Figures 11 and 12 appear with the permission of Kęstutis Karčiauskas, Jörg Peters and Elsevier.

References

- [AB08] ALEXA M., BOUBEKEUR T.: Subdivision shading. *ACM Transactions on Graphics* 27, 5 (2008), #142:1–4. 8
- [ADS06] AUGSDÖRFER U. H., DODGSON N. A., SABIN M. A.: Tuning subdivision by minimising Gaussian curvature variation near extraordinary vertices. *Computer Graphics Forum* 25, 3 (2006), 263–272. 6, 10, 13
- [ADS11a] AUGSDÖRFER U. H., DODGSON N. A., SABIN M. A.: Artifact analysis on B-splines, box-splines and other surfaces defined by quadrilateral polyhedra. *Computer Aided Geometric Design* 28, 3 (2011), 177–197. 6
- [ADS11b] AUGSDÖRFER U. H., DODGSON N. A., SABIN M. A.: Artifact analysis on triangular box-splines and subdivision surfaces defined by triangular polyhedra. *Computer Aided Geometric Design* 28, 3 (2011), 198–211. 6
- [BK04] BARTHE L., KOBBELT L.: Subdivision scheme tuning around extraordinary vertices. *Computer Aided Geometric Design* 21, 6 (2004), 561–583. 10
- [BKZ01] BIERMANN H., KRISTJANSSON D., ZORIN D.: Approximate Boolean operations on free-form solids. In *Proceedings of SIGGRAPH (2001)*, Fiume E., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 185–194. 9
- [BLZ00] BIERMANN H., LEVIN A., ZORIN D.: Piecewise smooth subdivision surfaces with normal control. In *Proceedings of SIGGRAPH (2000)*, Akeley K., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 113–120. 7

- [BMZ04] BOIER-MARTIN I., ZORIN D.: Differentiable parameterization of Catmull-Clark subdivision surfaces. In *Proceedings of the Symposium on Geometry Processing* (2004), Eurographics, pp. 155–164. [9](#)
- [BS02] BOLZ J., SCHRÖDER P.: Rapid evaluation of Catmull-Clark subdivision surfaces. In *Proceedings of the Web3D Symposium* (2002), ACM, pp. 11–17. [8](#)
- [Bun05] BUNNELL M.: Adaptive tessellation of subdivision surfaces with displacement mapping. In *GPU Gems 2: Programming Techniques for High-Performance Graphics and General-Purpose Computation*, Pharr M., Fernando R., (Eds.). Addison-Wesley, 2005. [8](#)
- [CAD09] CASHMAN T. J., AUGSDÖRFER U. H., DODGSON N. A., SABIN M. A.: NURBS with extraordinary points: High-degree, non-uniform, rational subdivision schemes. *ACM Transactions on Graphics* 28, 3 (2009), #46:1–9. [11](#), [12](#)
- [Cas10] CASHMAN T. J.: *NURBS-compatible subdivision surfaces*. PhD thesis, University of Cambridge, 2010. Technical report number UCAM-CL-TR-773. [11](#), [12](#)
- [CC78] CATMULL E., CLARK J.: Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design* 10, 6 (1978), 350–355. [1](#), [2](#), [12](#), [13](#), [14](#)
- [CC06] CHEN G., CHENG F.: Matrix based subdivision depth computation for extra-ordinary Catmull-Clark subdivision surface patches. In *Geometric Modeling and Processing – GMP 2006*, Kim M.-S., Shimada K., (Eds.), vol. 4077 of *Lecture Notes in Computer Science*. Springer, 2006, pp. 545–552. [7](#)
- [CCY06] CHENG F., CHEN G., YONG J.: Subdivision depth computation for extra-ordinary Catmull-Clark subdivision surface patches. *Advances in Computer Graphics* (2006), 404–416. [7](#)
- [CY06] CHENG F., YONG J.: Subdivision depth computation for Catmull-Clark subdivision surfaces. *Computer Aided Design & Applications* 3, 1–4 (2006), 485–494. [7](#)
- [DKT98] DEROSE T., KASS M., TRUONG T.: Subdivision surfaces in character animation. In *Proceedings of SIGGRAPH* (1998), Cohen M., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 85–94. [1](#), [9](#)
- [DL02] DYN N., LEVIN D.: Subdivision schemes in geometric modelling. *Acta Numerica* 11 (2002), 73–144. [2](#)
- [DLG87] DYN N., LEVIN D., GREGORY J. A.: A 4-point interpolatory subdivision scheme for curve design. *Computer Aided Geometric Design* 4, 4 (1987), 257–268. [3](#), [10](#)
- [DLG90] DYN N., LEVIN D., GREGORY J. A.: A butterfly subdivision scheme for surface interpolation with tension control. *ACM Transactions on Graphics* 9, 2 (1990), 160–169. [2](#), [10](#)
- [DLL92] DYN N., LEVIN D., LIU D.: Interpolatory convexity-preserving subdivision schemes for curves and surfaces. *Computer-Aided Design* 24, 4 (1992), 211–216. [10](#)
- [Doo78] DOO D. W. H.: A subdivision algorithm for smoothing down irregular shaped polyhedrons. In *Proceedings: Interactive Techniques in Computer Aided Design, Bologna, Italy* (1978), IEEE, pp. 157–165. [1](#), [2](#)
- [DS78] DOO D., SABIN M. A.: Behaviour of recursive division surfaces near extraordinary points. *Computer-Aided Design* 10, 6 (1978), 356–360. [1](#), [3](#), [5](#), [9](#), [12](#)
- [DSS07] DODGSON N., SABIN M., SOUTHERN R.: *Preconditions on geometrically sensitive subdivision schemes*. Tech. Rep. UCAM-CL-TR-691, University of Cambridge, 2007. [10](#)
- [GS01] GRINSPUN E., SCHRÖDER P.: Normal bounds for subdivision-surface interference detection. In *Proceedings of the conference on Visualization '01* (2001), IEEE, pp. 333–340. [9](#)
- [GU06] GINKEL I., UMLAUF G.: Loop subdivision with curvature control. In *Proceedings of the Symposium on Geometry Processing* (2006), Polthier K., Sheffer A., (Eds.), Eurographics, pp. 163–171. [6](#), [10](#)
- [GU08] GINKEL I., UMLAUF G.: Symmetry of shape charts. *Computer Aided Geometric Design* 25, 3 (2008), 131–136. [6](#), [10](#)
- [Han03] HAN B.: Classification and construction of bivariate subdivision schemes. In *Curve and Surface Fitting: Saint-Malo 2002*, Cohen A., Merrien J.-L., Schumaker L. L., (Eds.). Nashboro Press, 2003, pp. 187–197. [3](#)
- [HDW08] HUANG Z., DENG J., WANG G.: A bound on the approximation of a Catmull-Clark subdivision surface by its limit mesh. *Computer Aided Geometric Design* 25, 7 (2008), 457–469. [7](#)
- [HKD93] HALSTEAD M., KASS M., DEROSE T.: Efficient, fair interpolation using Catmull-Clark surfaces. In *Proceedings of SIGGRAPH* (1993), Kajiji J. T., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 35–44. [4](#), [10](#)
- [Hol96] HOLT F.: Toward a curvature-continuous stationary subdivision algorithm. *Zeitschrift für angewandte Mathematik und Mechanik* 76 (1996), 423–424. [9](#)
- [HSH10] HE L., SCHAEFER S., HORMANN K.: Parameterizing subdivision surfaces. *ACM Transactions on Graphics* 29, 4 (2010), #120:1–6. [9](#)
- [HW07a] HUANG Z., WANG G.: Distance between a Catmull-Clark subdivision surface and its limit mesh. In *Proceedings of the ACM symposium on Solid and Physical Modeling* (2007), pp. 233–240. [7](#)
- [HW07b] HUANG Z., WANG G.: Improved error estimate for extraordinary Catmull-Clark subdivision surface patches. *The Visual Computer* 23, 12 (2007), 1005–1014. [7](#)
- [HW08] HUANG Z., WANG G.: Bounding the distance between a Loop subdivision surface and its limit mesh. In *Advances in Geometric Modeling and Processing*, Chen F., Jüttler B., (Eds.), vol. 4975 of *Lecture Notes in Computer Science*. Springer, 2008, pp. 33–46. [7](#)
- [IDS04] IVRISIMTZIS I., DODGSON N., SABIN M.: A generative classification of mesh refinement rules with lattice transformations. *Computer Aided Geometric Design* 21, 1 (2004), 99–109. [3](#)
- [IGPP01] IGLESIAS A., GÁLVEZ A., PUIG-PEY J.: Computational methods for geometric processing. Applications to industry. *Computational Science – ICCS* (2001), 698–707. [14](#)
- [ISD04] IVRISIMTZIS I. P., SABIN M. A., DODGSON N. A.: On the support of recursive subdivision. *ACM Transactions on Graphics* 23, 4 (2004), 1043–1060. [3](#)
- [JS09] JIANG D., STEWART N. F.: Robustness of Boolean operations on subdivision-surface models. In *Numerical Validation in Current Hardware Architectures*, Cuyt A., Krämer W., Luther W., Markstein P., (Eds.). Springer-Verlag, 2009, pp. 161–174. [9](#)
- [Kaz07] KAZAKOV M.: Catmull-Clark subdivision for geometry shaders. In *Proceedings of AFRIGRAPH* (2007), Slay H., Spencer S. N., Bangay S., (Eds.), ACM, pp. 77–84. [8](#)
- [KMDZ09] KOVACS D., MITCHELL J., DRONE S., ZORIN D.: Real-time creased approximate subdivision surfaces. In *Proceedings of the Symposium on Interactive 3D Graphics and Games* (2009), pp. 155–160. [1](#), [8](#)
- [Kob96] KOBBELT L.: Interpolatory subdivision on open quadrilateral nets with arbitrary topology. *Computer Graphics Forum* 15, 3 (1996), 409–420. [2](#), [7](#), [10](#)

- [Kob00] KOBBELT L.: $\sqrt{3}$ -subdivision. In *Proceedings of SIGGRAPH* (2000), Akeley K., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 103–112. 2
- [KP07a] KARČIAUSKAS K., PETERS J.: Bicubic polar subdivision. *ACM Transactions on Graphics* 26, 4 (2007), #14:1–6. 10
- [KP07b] KARČIAUSKAS K., PETERS J.: Concentric tessellation maps and curvature continuous guided surfaces. *Computer Aided Geometric Design* 24, 2 (2007), 99–111. 12
- [KP07c] KARČIAUSKAS K., PETERS J.: Surfaces with polar structure. *Computing* 79, 2 (2007), 309–315. 10
- [KP08] KARČIAUSKAS K., PETERS J.: On the curvature of guided surfaces. *Computer Aided Geometric Design* 25, 2 (2008), 69–79. 12, 13, 14
- [KP09a] KARČIAUSKAS K., PETERS J.: Adjustable speed surface subdivision. *Computer Aided Geometric Design* 26, 9 (2009), 962–969. 11
- [KP09b] KARČIAUSKAS K., PETERS J.: Guided spline surfaces. *Computer Aided Geometric Design* 26, 1 (2009), 105–116. 12
- [KPR04] KARČIAUSKAS K., PETERS J., REIF U.: Shape characterization of subdivision surfaces—case studies. *Computer Aided Geometric Design* 21, 6 (2004), 601–614. 6, 14
- [KSH00] KARBACHER S., SEEGER S., HÄUSLER G.: A non-linear subdivision scheme for triangle meshes. In *Vision, Modeling and Visualization* (2000), pp. 163–170. 10
- [LC07] LAI S., CHENG F.: Robust and error controllable Boolean operations on free-form solids represented by Catmull-Clark subdivision surfaces. *Computer Aided Design and Applications* 4, 1–4 (2007), 487–496. 9
- [Lev99] LEVIN A.: Combined subdivision schemes for the design of surfaces satisfying boundary conditions. *Computer Aided Geometric Design* 16, 5 (1999), 345–354. 9
- [Lev06] LEVIN A.: Modified subdivision surfaces with continuous curvature. *ACM Transactions on Graphics* 25, 3 (2006), 1035–1040. 10
- [LFKN03] LANQUETIN S., FOUFOU S., KHEDDOUCI H., NEVEU M.: A graph based algorithm for intersection of subdivision surfaces. *Computational Science and Its Applications – ICCSA* (2003), 387–396. 9
- [LG00] LABSIK U., GREINER G.: Interpolatory $\sqrt{3}$ -subdivision. *Computer Graphics Forum* 19, 3 (2000), 131–138. 2, 3
- [LL03] LEVIN A., LEVIN D.: Analysis of quasi-uniform subdivision. *Applied and Computational Harmonic Analysis* 15, 1 (2003), 18–32. 5
- [LLS01] LITKE N., LEVIN A., SCHRÖDER P.: Trimming for subdivision surfaces. *Computer Aided Geometric Design* 18, 5 (2001), 463–482. 9
- [LM07] LI G., MA W.: A method for constructing interpolatory subdivision schemes and blending subdivisions. *Computer Graphics Forum* 26, 2 (2007), 185–201. 10
- [LMH00] LEE A., MORETON H., HOPPE H.: Displaced subdivision surfaces. In *Proceedings of SIGGRAPH* (2000), Akeley K., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 85–94. 8
- [Loo87] LOOP C. T.: Smooth subdivision surfaces based on triangles. *Master's thesis, University of Utah, Department of Mathematics* (Aug. 1987). 2, 5, 13
- [Loo02] LOOP C.: Bounded curvature triangle mesh subdivision with the convex hull property. *The Visual Computer* 18 (2002), 316–325. 9
- [Loo04] LOOP C.: Second order smoothness over extraordinary vertices. In *Proceedings of the Symposium on Geometry Processing* (2004), pp. 165–174. 10
- [LR80] LANE J. M., RIESENFELD R. F.: A theoretical development for the computer generation and display of piecewise polynomial surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 2, 1 (1980), 35–46. 11
- [LS08a] LOOP C., SCHAEFER S.: Approximating Catmull-Clark subdivision surfaces with bicubic patches. *ACM Transactions on Graphics* 27, 1 (2008), #8:1–11. 8
- [LS08b] LOOP C., SCHAEFER S.: G^2 tensor product splines over extraordinary vertices. In *Proceedings of the Symposium on Geometry Processing* (2008), pp. 1373–1382. 10
- [LSNCn09] LOOP C., SCHAEFER S., NI T., CASTAÑO I.: Approximating subdivision surfaces with Gregory patches for hardware tessellation. *ACM Transactions on Graphics* 28 (2009), #15:1–9. 8
- [MFR*10] MÜLLER K., FÜNFIG C., REUSCHE L., HANSFORD D., FARIN G. E., HAGEN H.: Dinus: Double insertion, nonuniform, stationary subdivision surfaces. *ACM Transactions on Graphics* 29 (2010), #25:1–21. 11, 12
- [MMN07] MAEKAWA T., MATSUMOTO Y., NAMIKI K.: Interpolation by geometric algorithm. *Computer-Aided Design* 39, 4 (2007), 313–323. 10
- [MNP08] MYLES A., NI T., PETERS J.: Fast parallel construction of smooth surfaces from meshes with tri/quad/pent facets. *Computer Graphics Forum* 27, 5 (2008), 1365–1372. 9
- [MP09] MYLES A., PETERS J.: Bi-3 C^2 polar subdivision. *ACM Transactions on Graphics* 28, 3 (2009), #48:1–12. 5, 6, 10, 11
- [MRF06] MÜLLER K., REUSCHE L., FELLNER D.: Extended subdivision surfaces: Building a bridge between NURBS and Catmull-Clark surfaces. *ACM Transactions on Graphics* 25, 2 (2006), 268–292. 11, 12
- [MWW01] MORIN G., WARREN J., WEIMER H.: A subdivision scheme for surfaces of revolution. *Computer Aided Geometric Design* 18, 5 (2001), 483–502. 10
- [MYP08] MYLES A., YEO Y. I., PETERS J.: GPU conversion of quad meshes to smooth surfaces. In *Proceedings of the ACM symposium on Solid and Physical Modeling* (2008), pp. 321–326. 8
- [NCP*09] NI T., CASTAÑO I., PETERS J., MITCHELL J., SCHNEIDER P., VERMA V.: Efficient substitutes for subdivision surfaces. *ACM SIGGRAPH Courses* (2009), #13:1–107. 1, 7, 8
- [NYM*08] NI T., YEO Y., MYLES A., GOEL V., PETERS J.: GPU smoothing of quad meshes. In *Proceedings of the International Conference on Shape Modeling and Applications* (2008), IEEE, pp. 3–9. 8
- [OS03] OSWALD P., SCHRÖDER P.: Composite primal/dual $\sqrt{3}$ -subdivision schemes. *Computer Aided Geometric Design* 20, 3 (2003), 135–164. 11
- [PC11] PRAUTZSCH H., CHEN Q.: Analyzing midpoint subdivision. *Computer Aided Geometric Design* 28, 7 (2011), 407–419. 11, 12
- [Pet00] PETERS J.: Patching Catmull-Clark meshes. In *Proceedings of SIGGRAPH* (2000), Akeley K., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 255–258. 8
- [Pet02] PETERS J.: Smoothness, fairness and the need for better multi-sided patches. In *Topics in algebraic geometry and geometric modeling: Workshop on Algebraic Geometry and Geometric Modeling, Vilnius University, Lithuania* (2002), vol. 334, p. 55. 10, 12, 13

- [PK10] PETERS J., KARČIAUSKAS K.: An introduction to guided and polar surfacing. In *Mathematical Methods for Curves and Surfaces*, Dæhlen M., Floater M., Lyche T., Merrien J.-L., Mørken K., Schumaker L., (Eds.), vol. 5862 of *Lecture Notes in Computer Science*. Springer, 2010, pp. 299–315. [5](#)
- [PR97] PETERS J., REIF U.: The simplest subdivision scheme for smoothing polyhedra. *ACM Transactions on Graphics* 16, 4 (1997), 420–431. [2](#), [4](#)
- [PR98] PETERS J., REIF U.: Analysis of algorithms generalizing B-spline subdivision. *SIAM Journal on Numerical Analysis* 35, 2 (1998), 728–748. [4](#)
- [PR99] PRAUTZSCH H., REIF U.: Degree estimates for C^k -piecewise polynomial subdivision surfaces. *Advances in Computational Mathematics* 10, 2 (1999), 209–217. [4](#), [5](#)
- [PR04] PETERS J., REIF U.: Shape characterization of subdivision surfaces—basic principles. *Computer Aided Geometric Design* 21, 6 (2004), 585–599. [5](#), [6](#)
- [PR08] PETERS J., REIF U.: *Subdivision Surfaces*, vol. 3 of *Geometry and Computing*. Springer, 2008. [3](#), [4](#), [7](#), [12](#)
- [Pra97] PRAUTZSCH H.: Freeform splines. *Computer Aided Geometric Design* 14, 3 (1997), 201–206. [11](#), [12](#)
- [Pra98] PRAUTZSCH H.: Smoothness of subdivision surfaces at extraordinary points. *Advances in Computational Mathematics* 9, 3 (1998), 377–389. [11](#), [12](#)
- [PS96] PULLI K., SEGAL M.: Fast rendering of subdivision surfaces. In *Proceedings of the Eurographics Workshop on Rendering Techniques* (1996), Springer-Verlag, pp. 61–70. [8](#)
- [PS04] PETERS J., SHIUE L. J.: Combining 4- and 3-direction subdivision. *ACM Transactions on Graphics* 23, 4 (2004), 980–1003. [3](#)
- [PT87] PIEGL L., TILLER W.: Curve and surface constructions using rational B-splines. *Computer-Aided Design* 19, 9 (1987), 485–498. [11](#)
- [PU98] PRAUTZSCH H., UMLAUF G.: A G^2 -subdivision algorithm. In *Geometric Modelling, Dagstuhl, 1996* (1998), Farin G., Bieri H., Brunnert G., DeRose T., (Eds.), vol. 13 of *Computing Supplement*, Springer, pp. 217–224. [5](#)
- [PU01] PETERS J., UMLAUF G.: Computing curvature bounds for bounded curvature subdivision. *Computer Aided Geometric Design* 18, 5 (2001), 455–461. [6](#)
- [PW09] PETERS J., WU X.: The distance of a subdivision surface to its control polyhedron. *Journal of Approximation Theory* 161, 2 (2009), 491–507. [7](#)
- [QW99] QIN K., WANG H.: Eigenanalysis and continuity of non-uniform Doo-Sabin surfaces. In *Proceedings of Pacific Conference on Computer Graphics and Applications* (1999), IEEE, pp. 179–186. [11](#)
- [Rei95] REIF U.: A unified approach to subdivision algorithms near extraordinary vertices. *Computer Aided Geometric Design* 12, 2 (1995), 153–174. [4](#)
- [Rei96] REIF U.: A degree estimate for subdivision surfaces of higher regularity. *Proceedings of the American Mathematical Society* 124, 7 (1996), 2167–2174. [5](#)
- [Rei98] REIF U.: TURBS—Topologically unrestricted rational B-splines. *Constructive Approximation* 14, 1 (1998), 57–77. [11](#), [12](#)
- [Rei07] REIF U.: An appropriate geometric invariant for the C^2 -analysis of subdivision surfaces. In *12th IMA Conference on the Mathematics of Surfaces* (2007), Martin R. R., Sabin M. A., Winkler J. R., (Eds.), vol. 4647 of *Lecture Notes in Computer Science*, Springer, pp. 364–377. [5](#)
- [RS01] REIF U., SCHRÖDER P.: Curvature integrability of subdivision surfaces. *Advances in Computational Mathematics* 14, 2 (2001), 157–174. [6](#)
- [Sab91] SABIN M. A.: Cubic recursive division with bounded curvature. In *Curves and surfaces* (1991), Laurent P. J., Le Méhauté A., Schumaker L. L., (Eds.), Academic Press Professional, Inc., pp. 411–414. [9](#)
- [Sab05] SABIN M. A.: Recent progress in subdivision: a survey. In *Advances in Multiresolution for Geometric Modelling*, Dodgson N. A., Floater M. S., Sabin M. A., (Eds.). Springer, 2005, pp. 203–230. [2](#)
- [Sab10] SABIN M.: *Analysis and Design of Univariate Subdivision Schemes*, vol. 6 of *Geometry and Computing*. Springer, 2010. [5](#)
- [SAP01] SKARIA S., AKLEMAN E., PARKE F.: Modeling subdivision control meshes for creating cartoon faces. In *Proceedings of the International Conference on Shape Modeling and Applications* (2001), IEEE, pp. 216–225. [13](#)
- [SB03] SABIN M. A., BARTHE L.: Artifacts in recursive subdivision surfaces. In *Curve and Surface Fitting: Saint-Malo 2002*, Cohen A., Merrien J.-L., Schumaker L. L., (Eds.). Nashboro Press, 2003, pp. 353–362. [6](#)
- [Sch03] SCHAEFER S.: *A Factored Interpolatory Subdivision Scheme for Surfaces of Revolution*. Master's thesis, Rice University, 2003. [10](#)
- [SDHI03] SABIN M. A., DODGSON N. A., HASSAN M. F., IVRISSIMTZIS I. P.: Curvature behaviours at extraordinary points of subdivision surfaces. *Computer-Aided Design* 35, 11 (2003), 1047–1051. [5](#)
- [SF04] STEWART I. F., FOISY A. R.: Arbitrary-degree subdivision with creases and attributes. *Journal of Graphics Tools* 9, 4 (2004), 3–17. [11](#), [12](#)
- [SFL*08] SEDERBERG T. W., FINNIGAN G. T., LI X., LIN H., IPSON H.: Watertight trimmed NURBS. *ACM Transactions on Graphics* 27, 3 (2008), #79, 1–8. [14](#)
- [SJP05] SHIUE L. J., JONES I., PETERS J.: A realtime GPU subdivision kernel. *ACM Transactions on Graphics* 24, 3 (2005), 1010–1015. [8](#)
- [SL03] STAM J., LOOP C.: Quad/triangle subdivision. *Computer Graphics Forum* 22, 1 (2003), 79–85. [3](#), [5](#)
- [SS06] SEVERN A., SAMAVATI F.: Fast intersections for subdivision surfaces. *Computational Science and its Applications – ICCSA* (2006), 91–100. [9](#)
- [Sta98] STAM J.: Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values. In *Proceedings of SIGGRAPH* (1998), Cohen M., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 395–404. [7](#)
- [Sta01] STAM J.: On subdivision schemes generalizing uniform B-spline surfaces of arbitrary degree. *Computer Aided Geometric Design* 18, 5 (2001), 383–396. [11](#), [12](#)
- [SW03] SCHAEFER S., WARREN J.: A factored interpolatory subdivision scheme for quadrilateral surfaces. In *Curve and Surface Fitting: Saint-Malo 2002*, Cohen A., Merrien J.-L., Schumaker L. L., (Eds.). Nashboro Press, 2003, pp. 373–382. [10](#)
- [SW05] SCHAEFER S., WARREN J.: On C^2 triangle/quad subdivision. *ACM Transactions on Graphics* 24 (2005), 28–36. [3](#)
- [SW07] SCHAEFER S., WARREN J.: Exact evaluation of non-polynomial subdivision schemes at rational parameter values. In *Proceedings of Pacific Conference on Computer Graphics and Applications* (2007), IEEE, pp. 321–330. [7](#), [8](#)
- [SW08] SCHAEFER S., WARREN J.: Exact evaluation of limits

- and tangents for non-polynomial subdivision schemes. *Computer Aided Geometric Design* 25, 8 (2008), 607–620. 8
- [SZBN03] SEDERBERG T. W., ZHENG J., BAKENOV A., NASRI A.: T-splines and T-NURCCs. *ACM Transactions on Graphics* 22, 3 (2003), 477–484. 11, 14
- [SZSS98] SEDERBERG T. W., ZHENG J., SEWELL D., SABIN M.: Non-uniform recursive subdivision surfaces. In *Proceedings of SIGGRAPH* (1998), Cohen M., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 387–394. 11, 12
- [TB97] TSITSIKLIS J. N., BLONDEL V. D.: The Lyapunov exponent and joint spectral radius of pairs of matrices are hard—when not impossible—to compute and to approximate. *Mathematics of Control, Signals, and Systems* 10, 1 (1997), 31–40. 5
- [Uml00] UMLAUF G.: Analyzing the characteristic map of triangular subdivision schemes. *Constructive Approximation* 16, 1 (2000), 145–155. 4
- [VPBM01] VLACHOS A., PETERS J., BOYD C., MITCHELL J.: Curved PN triangles. In *Proceedings of the Symposium on Interactive 3D graphics* (2001), ACM, pp. 159–166. 8
- [VZ01] VELHO L., ZORIN D.: 4–8 Subdivision. *Computer Aided Geometric Design* 18, 5 (2001), 397–427. 2, 3
- [WD05] WALLNER J., DYN N.: Convergence and C^1 analysis of subdivision schemes on manifolds by proximity. *Computer Aided Geometric Design* 22, 7 (2005), 593–622. 6, 11
- [WP04] WU X., PETERS J.: Interference detection for subdivision surfaces. *Computer Graphics Forum* 23, 3 (2004), 577–584. 9
- [WW01] WARREN J., WEIMER H.: *Subdivision Methods for Geometric Design*. Morgan Kaufmann, 2001. 11, 12
- [Yan05] YANG X.: Surface interpolation of meshes by geometric subdivision. *Computer-Aided Design* 37, 5 (2005), 497–508. 10
- [ZC06] ZENG X., CHEN X.: Computational formula of depth for Catmull-Clark subdivision surfaces. *Journal of Computational and Applied Mathematics* 195, 1–2 (2006), 252–262. 7
- [ZK02] ZORIN D., KRISTJANSSON D.: Evaluation of piecewise smooth subdivision surfaces. *The Visual Computer* 18, 5 (2002), 299–315. 7
- [ZLLT06] ZULTI A., LEVIN A., LEVIN D., TEICHER M.: C^2 subdivision over triangulations with one extraordinary point. *Computer Aided Geometric Design* 23, 2 (2006), 157–178. 5
- [Zor00a] ZORIN D.: A method for analysis of C^1 -continuity of subdivision surfaces. *SIAM Journal on Numerical Analysis* 37, 5 (2000), 1677–1708. 4
- [Zor00b] ZORIN D.: Smoothness of stationary subdivision on irregular meshes. *Constructive Approximation* 16, 3 (2000), 359–398. 4
- [Zor06] ZORIN D.: Constructing curvature-continuous surfaces by blending. In *Proceedings of the Symposium on Geometry Processing* (2006), Polthier K., Sheffer A., (Eds.), Eurographics, pp. 31–40. 4, 10
- [ZS01] ZORIN D., SCHRÖDER P.: A unified framework for primal/dual quadrilateral subdivision schemes. *Computer Aided Geometric Design* 18, 5 (2001), 429–454. 11, 12
- [ZSD*00] ZORIN D., SCHRÖDER P., DEROSE T., KOBBELT L., LEVIN A., SWELDENS W.: Subdivision for modeling and animation. *ACM SIGGRAPH Courses* (2000), #23:1–194. 1, 2
- [ZSS96] ZORIN D., SCHRÖDER P., SWELDENS W.: Interpolating subdivision for meshes with arbitrary topology. In *Proceedings of SIGGRAPH* (1996), Rushmeier H., (Ed.), Computer Graphics Proceedings, Annual Conference Series, ACM, pp. 189–192. 10