

Tuning Subdivision by Minimising Gaussian Curvature Variation Near Extraordinary Vertices

U. H. Augsdörfer¹, N. A. Dodgson¹ and M. A. Sabin^{1,2}

¹Computer Laboratory, University of Cambridge, 15 J.J. Thomson Ave, Cambridge CB3 0FD, England

²Numerical Geometry Ltd, 26 Abbey Lane, Lode, Cambridge CB5 9EP, England

Abstract

We present a method for tuning primal stationary subdivision schemes to give the best possible behaviour near extraordinary vertices with respect to curvature variation.

Current schemes lead to a limit surface around extraordinary vertices for which the Gaussian curvature diverges, as demonstrated by Karčiauskas et al. [KPR04]. Even when coefficients are chosen such that the subsubdominant eigenvalues, μ , equal the square of the subdominant eigenvalue, λ , of the subdivision matrix [DS78] there is still variation in the curvature of the subdivision surface around the extraordinary vertex as shown in recent work by Peters and Reif [PR04] illustrated by Karčiauskas et al. [KPR04].

In our tuning method we optimise within the space of subdivision schemes with bounded curvature to minimise this variation in curvature around the extraordinary vertex. To demonstrate our method we present results for the Catmull-Clark [CC78], 4-8 [Vel01, VZ01] and 4-3 [PS03] subdivision schemes. We compare our results to previous work on the tuning of these schemes and show that the coefficients derived with this method give a significantly smaller curvature variation around extraordinary vertices.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling, I.3.6 [Computer Graphics]: Methodology and Techniques

1. Introduction

Subdivision is an algorithmic technique to generate smooth surfaces as the limit of a sequence of successively refined polyhedral meshes. By repeatedly applying the subdivision procedure to the initial manifold mesh of vertices joined by faces, we generate a sequence of meshes that converges to a resulting smooth surface.

Subdivision surface schemes have been around since 1978 [CC78]. Even then the question of how to choose the coefficients around extraordinary vertices was identified as a key issue [DS78]. Extraordinary points are associated with vertices and faces that do not have the ordinary number of neighbours, i.e. 4 neighbours for quadrilateral, or 6 neighbours for triangular subdivision schemes. Problems observed at and near such extraordinary points can include unbounded increase of curvature, erratic fluctuation of curvature or flat spots [SB03, PR04].

Although these misbehaviours are not serious enough to prevent use of subdivision surfaces in graphics, they are critical in CAD/CAM where Class A surfaces are required. They may be one reason why subdivision surfaces have not replaced NURBS in CAD/CAM software. Improving the behaviour of subdivision surfaces around extraordinary points in the limit is therefore critical to the success of subdivision schemes in applications requiring Class A surfaces.

In this paper we consider primal schemes. For such schemes we need to consider tuning for extraordinary vertices (EVs) only. After one subdivision step, such schemes will have no extraordinary faces and the number of EVs will stay constant.

It has been proven that, for many of the schemes, the goal of C^2 at an extraordinary vertex is unattainable [Rei96]. We instead optimise a subdivision scheme against a quality criterion, which is designed to ensure that the variation in Gaussian curvature around an EV is kept to a minimum.

The method presented can be applied to any primal subdivision scheme on any base tiling. We focus on schemes which work on quadrilateral lattices, since these are most frequently used in engineering applications. We demonstrate our method by presenting results for three approximating quadrilateral schemes: Velho’s 4-8 scheme [Vel01, VZ01], the Catmull-Clark scheme [CC78], and Peters and Shiue’s 4-3 scheme [PS03].

1.1. Tuning Subdivision Schemes

A subdivision operation creates a new finer mesh by constructing new vertices as linear combinations of old ones. The coefficient or weights of the linear combinations can be depicted either as a *mask* or as *stencils*.

The *mask* documents the coefficients by which a given old vertex influences the neighbouring new ones. It can be depicted in the form of a diagram where coefficients by which a given old vertex influences the surrounding new ones are laid out in the same pattern as the new vertices. Each scheme has exactly one mask for each valency. *Stencils* are often described using a set of diagrams where the coefficients by which nearby old vertices influence a given new one are laid out in the same pattern as those old vertices.

The purpose of tuning the coefficients of a subdivision scheme is to make the limit surface behave as well as possible around extraordinary points, in particular with respect to continuity and curvature behaviour. It is possible to choose different levels of tuning depending on how many coefficients are taken to be variable.

The simplest level of tuning, used by Loop [Loo87], Velho and Zorin [VZ01], Catmull and Clark [CC78] and Peters and Shiue [PS03], is to tune only the influence of an extraordinary vertex on its image in the next refinement.

More freedom can be gained by tuning some or all of the coefficients in the mask of an extraordinary point. This requires that the stencils used for points covered by that mask should be renormalised, but the ratios between parent vertices other than the extraordinary one remain unchanged.

Even more freedom can be obtained by tuning independently some or all of the coefficients in any stencil in which the extraordinary vertex makes a contribution. This is implied in recent work by Barthe and Kobbelt [BK04], where they describe the process of tuning two triangular grid schemes. They optimise a linear combination of a number of quality measures by selecting the coefficients in the stencils around an extraordinary point.

One can go even further, by allowing the extraordinary vertex to extend the reach of nearby ordinary ones. This was used by Zorin, Schröder and Sweldens [ZSS96] in their tuning of the Butterfly schemes. An extreme version was used by Zulti *et al.* [ZLLT06] in a C^2 scheme where all new vertices on rays propagating from an extraordinary point have non-standard stencils.

We choose to limit ourselves to the second level of tuning and ascertain what quality of tuning can be achieved at this level. We choose the second level because it is clear that the first level of tuning does not usually give enough freedoms to achieve bounded non-zero curvature. Having more freedoms to optimise can be expected to give better results. However, tuning stencils requires detailed work to cover the cases where a stencil has more than one extraordinary vertex in its footprint. Even though such cases only occur in the very first steps of subdivision, they still demand considerable implementation work and must be handled well because artifacts are known to be introduced at the first subdivision step.

2. Outline

In our tuning approach, rather than optimising within the coefficient space, which involves imposing non-linear constraints, we use the subdominant eigenvalue, λ , as a freedom and impose the condition that the subsubdominant eigenvalues, μ , are its square. We thus ensure that the condition $\lambda^2 = \mu$, necessary for bounded curvature around extraordinary vertices [DS78], is always satisfied. This also reduces the number of freedoms to two fewer than the number of coefficients in the mask and thus simplifies the optimisation.

By setting $\lambda^2 = \mu$ for all μ , the curvatures around the EV neither grow nor shrink. However, even with these eigenvalue relations, variation in the curvature of the subdivision surface around the EV is still observed as shown in recent work by Peters and Reif [PR04] subsequently illustrated by Karčiauskas *et al.* [KPR04]. Figure 4 shows two examples of the Gaussian curvature around the EV. The initial configuration was the same, but the coefficients used for the subdivision process around the EV are different. Our method aims to minimise the variation in curvature around the EV.

The optimisation process presented here involves looping through a wide range of subdominant eigenvalues, λ . For each λ we derive the corresponding optimal mask coefficients. The coefficients in the mask are optimised by sampling a representative set of input meshes and calculating variation in curvature. The optimal coefficients are those for which variation in curvature is a minimum.

2.1. The Derivation of the Mask Coefficients

We determine the mask coefficients, which determine the influence of the EV on new vertices next to it, by requiring that $\lambda^2 = \mu$, and solving the determinant equations for the first and second Fourier component (see Sabin [Sab02] for details of Fourier analysis of subdivision):

$$\det(\hat{S}_1 - \lambda I) = 0 \quad (1)$$

$$\det(\hat{S}_2 - \mu_2 I) = \det(\hat{S}_2 - \lambda^2 I) = 0 \quad (2)$$

For all schemes listed here the condition that both these determinant equations be fulfilled is sufficient to solve for

the mask coefficients of the points closest to the EV, β and γ (see Figures 3, 11, and 16 for the labelling of mask coefficients). Because solving both determinant equations simultaneously results in a quadratic equation for the β and γ pair we get two solutions for β and γ . It depends on the scheme which solution to choose to get a sensible and, with respect to λ , continuous answer.

A real solution for β and γ cannot be achieved for all possible values of λ . While for the Velho $\sqrt{2}$ scheme larger λ result in an imaginary solution for β and γ , it is smaller λ which give rise to imaginary solutions for the binary Catmull-Clark and 4-3 scheme. We consider these solutions not feasible and exclude these eigenvalues from the space in which we seek our optimum.

Once β and γ are chosen, coefficient α can be derived from the determinant equation corresponding to the \hat{S}_0 Fourier subdivision matrix. In the case of Velho's 4-8 subdivision scheme, this approach reduces an optimisation problem in three dimensions to one dimension, because all three mask coefficients are fixed by the input λ . For schemes with larger masks, like the Catmull-Clark and 4-3 scheme, coefficients which determine the influence of the EV on new vertices further away from the EV are retained as freedoms in the optimisation. We use an optimising algorithm based on the simplex search method. It is a direct search method that does not use numerical or analytic gradients. We start the search from an initial estimate, which we take to be the standard coefficients.

2.2. The Quality Measure

Because it is impossible to achieve C^2 around an EV [Rei96] we aim to improve the quality of the surface around an EV by minimising its variation in curvature. This will minimise undesirable rippling around EVs, which is unacceptable in CAD/CAM applications.

Subdivision algorithms do not generate additional extraordinary vertices, so that the number of EVs stays constant for the sequence of polyhedral meshes generated. After a few subdivision steps each EV is separated from all other EVs by growing regular regions. In the vicinity of an EV the subdivision surface generated can be regarded as the union of the EV and a sequence of spline rings. To analyse the shape at an EV where the limit surface is not necessarily C^2 , Peters and Reif [PR04] analyse the spline ring of the characteristic map around an EV. They refer to the spline ring for a characteristic map as the *central surface*. The central surfaces at different subdivision levels coincide. The size of the spline ring is different for each scheme. The spline rings for the three schemes discussed here are shown in Figure 1.

2.2.1. The Gaussian Curvature

Peters and Reif [PR04] showed that the Gaussian curvature within a spline ring around the extraordinary vertex can be




Figure 1: The mesh is drawn as a 90 degree sector for ease of drawing, but is symbolic of any valency. The EV, highlighted in grey, is in the lower left hand corner. The new vertices after one subdivision step are shown as dots. From left to right: The red dots show the vertices within the spline rings of the 4-8, Catmull-Clark and 4-3 scheme, which are used for curvature analysis.

used to detect early when a mesh will lead to undesirable curvature behaviour. We use this measure of surface quality in our work.

Because we are examining the limit case, relevant to CAD/CAM, we determine the Gaussian curvature, K , using the expression

$$K = \frac{\partial^2 z}{\partial x^2} \frac{\partial^2 z}{\partial y^2} - \left(\frac{\partial^2 z}{\partial x \partial y} \right)^2 \quad (3)$$

where x , y , and z are rectangular coordinates aligned with the tangent plane at the EV.

To determine the first and second derivative at vertices within the spline ring we can apply derivative stencils. These stencils are derived from the regular case row eigenvectors corresponding to the eigenvalues λ and μ . The derivatives determined in this way are derivatives with respect to local coordinates u and v . Since we are interested in the curvature of many points in the spline ring around the EV, we need to relate all points to a common x , y , z coordinate system to be able to use Equation 3 and we use the Jacobian matrix of the characteristic to do this.

Because we determine the Gaussian curvature at each of a dense grid of vertices in the spline ring, Equation 3 is a discrete curvature estimator only. Although it would be more accurate to have some measure which applied to the whole spline ring rather than just to sample points, it is reasonable to expect the curvature to vary smoothly over the spline ring, because these schemes are known to be C^2 in the regular region. The spline ring is sampled densely enough to capture the variation in curvature. We also analysed denser spline rings around the EV, i.e. the black region in Figure 4 and the green or blue region in Figure 12, but found that, while the computation time increases, the overall result did not change significantly. We therefore restrict the analysis here to the red spline ring shown in Figures 4 and 12.

Peters and Reif [PR04] stress the fact that hybrid curvature behaviour in the EV neighbourhood should be avoided, and initially we aimed at minimising the set of hybrid cases. As a result of this work, however, we came to believe that it

is more important to minimise the variation in Gaussian curvature, regardless of its sign. We therefore aim to determine coefficients such that curvature variation around an EV is kept to a minimum for as many initial configurations as possible. Indeed, our method subsumes and improves on Peters and Reif's because it reduces the number of configurations for which hybrid Gaussian curvature can be observed, i.e. where the minimum and maximum of the Gaussian curvature have different signs.

2.2.2. The Characteristic Mesh

In order to achieve minimised curvature variation over as many initial configurations as possible, we need to sample the range of input meshes. Any input mesh configuration has a local curvature behaviour around the EV which is some linear combination of the three second order characteristic maps in z , taken over the x, y configuration from the natural configuration.

To set up the natural configuration [BS88], we use eigenvectors calculated from the subdivision matrix corresponding to the sub- and subsubdominant eigenvalues. Once all the coefficients are known, the eigenvectors corresponding to the λ and μ eigenvalues respectively are determined from the $\omega = 0, \pm 1, \pm 2$ Fourier subdivision matrices [Sab02].

The column eigenvector corresponding to the subdominant eigenvalue λ , obtained from the $\omega = \pm 1$ Fourier subdivision matrix, gives the natural configuration, the way in which the neighbourhood of the extraordinary vertex is laid out within the tangent plane [BS88].

The column eigenvectors from the $\omega = 0$ and $\omega = \pm 2$ Fourier subdivision matrices corresponding to the subsubdominant eigenvector, μ , give the so-called cup and saddle components perpendicular to the tangent plane [Sab02].

This gives a relatively coarse local polyhedron for the cup-shape, m_c , the saddle shape, m_{s1} , with its axes aligned with the axes of the coordinate system, and the other saddle shape, m_{s2} , the axes of which are rotated by 45° . Each of these defines a second order characteristic map, which we refer to as the characteristic mesh. By combining the three characteristic meshes, m_c , m_{s1} and m_{s2} , we can derive any central surface.

2.2.3. The Representative Set of Shapes

We built up a representative set of central surfaces by linearly combining the three characteristic meshes according to

$$m_i = (1 - r)m_c + r \cos(\phi)m_{s1} + r \sin(\phi)m_{s2} \quad (4)$$

where the radius, r , ranges from $[0, 1]$, using a stepsize of 0.005 and the angle ϕ ranges from $[0, 2\pi]$, with a stepsize of $\pi/48$. We have thus a representative set of 19201 input meshes, m_i . The various configurations which are made up




Figure 2: Information about a representative set of quadratic shapes, built up according to Equation 4 can be mapped into a disk. The shape on the top left is a combination, m_i , of the three characteristic meshes, m_c , m_{s1} and m_{s2} , which are also shown. To obtain it, the shapes have been combined using Equation 4 with $r = 0.6$ and $\phi = 0.78$.

of certain proportion of cup and saddle shapes of different orientations can therefore all map into a disk parameterised on r and ϕ . This representative set was sampled densely enough to see that characteristics of the meshes vary smoothly over the disk representation. Because the central surfaces coincide at different subdivision levels, it is sufficient to determine the Gaussian curvature of only one spline ring around the EV.

Figure 2 shows how quadratic shapes are located in the disk. To accomplish a symmetric plot from which the shapes can be read off intuitively, we display the shapes at angles $\Phi = \phi/2$ and mirror the information on the bottom half of the disk. The pure cup is at its centre, and the pure saddles are points around the circumference, as shown in Figure 2. Configurations which are some mix of cup and saddle are placed in the disk on a ray whose direction depends on the orientation of the saddle and at a radius from the centre depending on the proportions of saddle.

2.2.4. The Shape-In-The-Limit Chart

We are now able to plot a quality measure against positions on the disk, a convention introduced by Karčiauskas *et al.* [KPR04]. This is referred to as the *shape-in-the-limit chart*.

Peters and Reif [PR04] define different notions of shape in the vicinity of an extraordinary vertex where the surface is not necessarily C^2 . The one we use measures the Gaussian curvature, K_c , of the subdivision surface in the central surface around the extraordinary vertex. If $K_c > 0$ or $K_c < 0$ for each point in the spline ring around the extraordinary vertex the limit surface is elliptic or hyperbolic respectively. If K_c changes sign within the spline ring, it is referred to as hybrid. If the sign of K_c is the only information we want to convey, we can use colour to represent the curvature behaviour within the spline ring for each mesh in our representative set: white for elliptic, grey for hyperbolic and black to denote hybrid behaviour of the Gaussian curvature.

The disk representation of the shape-in-the-limit chart is




Figure 3: Unnormalised mask coefficients around regular (left) and extraordinary (right) vertices for the 4-8 subdivision scheme.

our variant of the basic idea introduced by Karčiauskas *et al.* [KPR04], who treat the components more abstractly as barycentric coordinates, which maps one sector of our disk into a triangle without loss of information.

We extend this idea by plotting the variation in Gaussian curvature in the spline ring for each shape as a z-dimension. For better visualisation we use colour to indicate the size of variation in Gaussian curvature and highlight, in black, shapes for which a change in Gaussian curvature direction occurs.

2.2.5. Scheme Quality

Our metric for a ‘good’ scheme is that for each valency it minimises the mean, over the shape in the limit chart, of the variation of Gaussian curvature over the spline ring.

3. 4-8 scheme

This approximating scheme, which generalises the four-directional box spline, was first introduced by Velho [Vel01] in 2001. It is C^4 continuous everywhere except at extraordinary vertices, where it is C^1 .

3.1. Mask coefficients

The scheme makes it possible to achieve high smoothness while using small stencils, which is equivalent to having a small mask. The mask of this scheme has three coefficients, and is shown in Figure 3 for the regular and the extraordinary case.

The three coefficients in the mask can be determined using the subdominant eigenvalue, λ , as an input which fixes α , β and γ by solving the determinant equations for the first and second Fourier component as described in Section 2.1.

We implement the subdivision by means of standard stencils everywhere except for vertices around the EV, which require non-standard stencils, using values α , β and γ to weight the EV. The stencils are then renormalised so that the coefficients sum to 1.

The initial control mesh needs to be subdivided until the mesh is dense enough to apply the derivative stencil and calculate the Gaussian curvature for each vertex within the




Figure 4: The Gaussian curvature around an EV with valency $n = 5$ for the characteristic mesh m_c ($\omega = 0$ eigenvector) of the 4-8 scheme. The spline ring, which is used for analysis, is highlighted in red. Left: The variation of the characteristic mesh for $\lambda = 0.72$ has a curvature variation from 0.15 to 0.25. Right: The characteristic mesh for $\lambda = 0.7751$ has a curvature variation from 0.12 to 0.14. The latter value for λ yields the optimum coefficients with respect to curvature variation.




Figure 5: The 2D and 3D shape-in-the-limit charts for an extraordinary vertex with valency $n = 5$ are shown together with the natural configuration for three different subdominant eigenvalues λ for the 4-8 scheme. For both charts black is used to highlight undesirable hybrid cases. The subdominant eigenvalue λ has been increased from top to bottom $\lambda = 0.75, 0.7751$ and 0.79 .

spline ring around the EV. Figure 4 shows the Gaussian curvature within the spline ring for the cup-shaped characteristic mesh, m_c , around the EV corresponding to two different subdominant eigenvalues. The red parts of the data show the spline ring which is used for analysis in the optimisation process.

Three examples of shape-in-the-limit charts are shown in Figure 5 for $n = 5$. The subdominant eigenvalue λ , is increased from the top down. The charts are shown together with the natural configuration corresponding to each λ . The




Figure 6: The minimum and maximum Gaussian curvature, K_C , as a function of radius for all shapes along $\phi = 0$ in the chart for two different subdominant eigenvalues for the 4-8 scheme. Left: $\lambda = 0.75$ and right: $\lambda = 0.7751$.




Figure 7: The mean curvature variation in a chart corresponding to λ is shown for valencies $n = 5$ to 8 for the 4-8 scheme. The optimum λ is the one which results in a configuration which yields the minimum variation for a given sample set of central surfaces.

λ which yields a natural configuration for which the mean variation of Gaussian curvature is smallest is $\lambda = 0.7751$.

In Figure 6 we plot the minimum and maximum Gaussian curvature observed within the spline ring for all configurations along the $\phi = 0$ line of the chart as a function of radius. We conjecture from these graphs that minimising the variation in curvature around the EV has the added benefit of bringing us close to the minimum of the number of configuration for which the Gaussian curvature is hybrid.

3.2. Results

In Figure 7 we plot the variation in curvature as a function of input λ . The optimum coefficients are derived using the subdominant eigenvalue for which the variation is a minimum.

n	5	6	7	8
λ	0.7751	0.8281	0.8663	0.8939
α	11.479	20.848	33.552	50.274
β	1.5217	1.1492	0.8733	0.6838
γ	0.7146	0.4581	0.3175	0.2236

Table 1: The optimised subdominant eigenvalue and coefficients for the 4-8 Velho subdivision scheme, for which the mean variation in Gaussian curvature is a minimum.




Figure 8: The generic mesh where the x and y values are the standard Catmull-Clark eigenvectors corresponding to the subdominant eigenvector and $z = 1 - x^2 - 2y^2$, as used by Karčiauskas et al. [KPR04] to illustrate their results. It serves as a good example of a convex mesh, which is not a perfect cup. We use valency eight for ease of comparison with Karčiauskas et al. [KPR04].

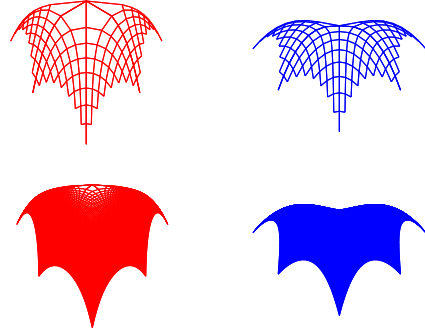


Figure 9: The figures show the limit surface of the subdivided generic mesh with a valency of eight. Top: The mesh has been subdivided six times. Bottom: The mesh has been subdivided thirteen times. Two different sets of coefficients have been used. Left: coefficients optimised using the new method with respect to curvature variation. Right: the coefficients suggested by Velho and Zorin [VZ01].

The optimum coefficients derived in this way are listed in Table 1.

We illustrate the results by subdividing the convex generic mesh, shown in Figure 8. This mesh is designed to highlight difficulties in the immediate neighbourhood of an EV which may not be easily visible in a real-world graphics example. We use this mesh to compare the optimised coefficients with those suggested by Velho and Zorin [VZ01].

Around EVs Velho and Zorin [VZ01] propose the modification of the coefficient in the stencil which determines the new position of the EV. While the weight of the old EV is unaltered, the surrounding vertices get coefficients depending on the valency of the EV around which they are located, $\beta_s = 1/(2n)$.

We subdivide a number of times before determining the limit shape using the standard limit stencil around the EV and the limit stencil derived from the row-eigenvector cor-




Figure 10: The Gaussian curvature within the spline ring after six and thirteen subdivision steps. The Gaussian curvature of the spline ring corresponding to the Velho-Zorin subdivided mesh (blue) is compared to the Gaussian curvature achieved with the optimised coefficients (red).

responding to the $\lambda = 1$ eigenvalue of the tuned subdivision matrix. Figure 9 shows how the surface develops around an EV after six and thirteen subdivision steps. Their corresponding Gaussian curvature within the spline ring around the EV are shown in Figure 10. While for the standard 4-8 scheme the curvature diverges and even becomes negative after a number of subdivision steps, our method has a small variation in curvature and does not change sign.

4. Catmull-Clark

Catmull-Clark subdivision [CC78], a binary scheme, is the most widely used subdivision scheme working on a quadrilateral mesh. Because Catmull-Clark subdivision is a generalisation of a bi-cubic B-spline it generates surfaces that reduce to a standard B-spline surface except at extraordinary points. Therefore, everywhere except at extraordinary vertices, the surface is continuous in tangent and curvature (C^2). At extraordinary points the surface is at least continuous in tangent.

4.1. Mask Coefficients

The mask of this scheme has six coefficients and is shown in Figure 11 for the regular and irregular case.

The mask coefficients α , β and γ can be solved for by using the condition $\mu = \lambda^2$, since the block structure of the Fourier matrix allows us to determine these coefficients from the first block only.




Figure 11: The unnormalised mask of the Catmull-Clark scheme for the regular and extraordinary case.




Figure 12: The Gaussian curvature for the first three spline rings, shown in three different colours, for the characteristic mesh m_{s2} , the eigenvector corresponding to the μ_2 eigenvalue of the Catmull-Clark scheme. Left: $\lambda = 0.6105$. Right: $\lambda = 0.6155$.

While the coefficients α , β , γ , are fixed for a given input λ , the coefficients, δ , ϵ and ϕ , cannot be determined using the subdivision matrix. For each input λ we optimise these coefficients with respect to the mean variation in Gaussian curvature within the first spline ring.

Two examples of Gaussian curvature around a EV are shown in Figure 12 for a choice of sub-optimal and optimal coefficients.

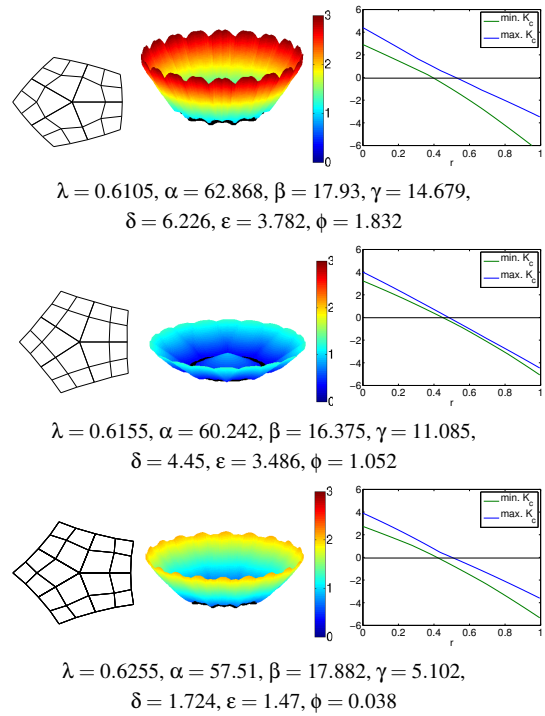


Figure 13: The 3D shape-in-the-limit chart for a mesh with an extraordinary vertex of valency $n = 5$ are shown together with the natural configuration for three different values of λ for the Catmull-Clark scheme. Black is used to highlight hybrid cases. The right hand side shows the minimum and maximum Gaussian curvature, K_c , for all shapes in the chart at $\phi = 0$ as a function of chart radius.

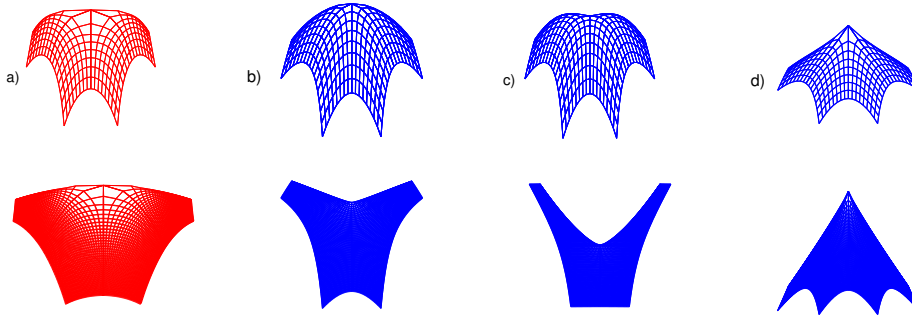


Figure 15: The figures show the limit surface of the subdivided generic mesh with a valency of eight. The mesh has been subdivided using four different variations of the Catmull-Clark algorithm. Top: The mesh has been subdivided three times. Bottom: The mesh has been subdivided 16 times. Four different sets of coefficients have been applied. a) coefficients optimised with respect to curvature variation. b) the coefficients suggested by Catmull and Clark, p. 184 of [CC78]. c) the coefficients first applied by Catmull and Clark, p. 186 of [CC78]. d) the coefficients suggested by Karčiauskas et al., p.11 of [KPR04].




Figure 14: The mean curvature variation in a chart corresponding to λ is shown for valencies $n = 5$ to 8 for the Catmull-Clark scheme. The optimum λ is the one which results in a configuration which yields the minimum variation for a given sample set of central surfaces.

Examples for the 3D shape-in-the-limit chart are shown in Figure 13. While the centre example shows the result of an optimal solution for the mask coefficients around an EV with valency $n = 5$, the two others are examples of sub-optimal choices with respect to curvature variation.

4.2. Results

The mean curvature variation for a chart is shown as a function of λ in Figure 14. The subdominant eigenvalue, λ , which yields the minimum mean variation within a chart can be read off as the optimum choice for the mask of an EV with the given valency. The mask coefficients optimised with respect to curvature variation are shown in Table 2.

Other researchers have attempted to improve the behaviour of Catmull-Clark subdivision surfaces around EVs [CC78, KPR04]. To compare our results to previous work

n	5	6	7	8
λ	0.615	0.703	0.768	0.815
α	60.242	95.021	143.090	207.678
β	16.375	11.386	8.577	6.688
γ	11.085	8.378	5.359	3.573
δ	3.610	2.510	1.731	1.382
ϵ	3.186	2.376	1.744	1.199
ϕ	0.786	0.816	0.773	0.663

Table 2: The optimum coefficients found for the Catmull-Clark scheme.

we use the same generic mesh, shown in Figure 8, as Karčiauskas et al. [KPR04] to illustrate our results. Results are shown in Figure 15. Although the same mesh has been subdivided in all figures, it is obvious that all but our coefficients cause the subdivided mesh to change shape. While b) to c) become less cup shaped and develop to a saddle shape, for Figure d) it is the cup shape which dominates.

5. The 4-3 scheme

4-3 subdivision is an alternative to Catmull-Clark subdivision on quadrilateral meshes. In regular mesh regions, 4-3 surfaces are C^2 , otherwise C^1 . 4-3 has a smaller mask than Catmull-Clark. The mask for this scheme has five coefficients. It is shown for the regular case in Figure 16 together with the extraordinary mask.

5.1. Results

The results for the optimised 4-3 scheme are obtained very similarly to those of the Catmull-Clark and are simply read off Figure 17, which shows how the mean variation in curvature for a shape-in-the-limit chart varies with λ . The λ for




Figure 16: The unnormalised mask of the 4-3 scheme for the regular and extraordinary case.




Figure 17: The mean curvature variation for a shape-in-the-limit chart of the 4-3 scheme for every λ is plotted as a function of λ . The curve is smooth with one clear minimum.

which the variation is a minimum yields the optimised coefficients for the extraordinary mask. The optimised coefficients are listed in Table 3.

In Figure 18 we compare the subdivided generic mesh using the coefficients optimised with respect to curvature variations in the spline ring to those suggested by Peters and Shiue [PS03]. Peters and Shiue suggested a non-standard stencil to be used to determine the new EV. Again, after a few subdivision steps the initial convex mesh develops a saddle-like configuration around the EV. The coefficients determined with the method presented here prevent such a development. The curvature stays positive throughout the subdivision process.

6. Discussion

We limit our degrees of freedom for optimisation by tuning the mask coefficients of the scheme. Tuning with respect to more degrees of freedoms, as has been explored in [Loo02, BK04, ZSS96, ZLLT06], can be expected to improve the achievable results. However, choosing more degrees of freedom is associated with an increase in implementation difficulties, especially as it must handle all cases where two EVs fall close together.

By treating the subdominant eigenvalue, λ , as a freedom we are able always to guarantee that the subsubdominant eigenvalue be the square of the subdominant eigenvalue, $\mu = \lambda^2$, a condition necessary for bounded curvature.

We also guarantee that no shape of the central surface pre-

n	5	6	7	8
λ	0.6185	0.705	0.769	0.816
α	14.8475	25.297	39.088	57.139
β	4.2348	2.972	2.22	1.7628
γ	2.2121	1.608	1.0576	0.5643
δ	1.0326	0.585	0.3973	0.3053
ϵ	0.6965	0.528	0.3656	0.2059

Table 3: The optimum coefficients found for the 4-3 scheme.

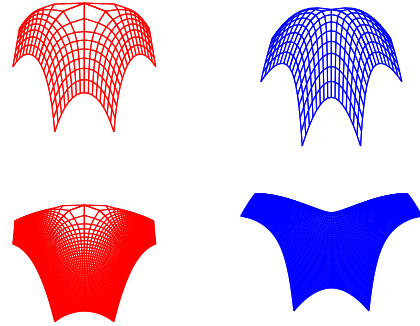


Figure 18: The figures show the limit surface of the subdivided generic mesh with a valency of eight using different variations of the 4-3 subdivision scheme. Top: The mesh has been subdivided three times. Bottom: The mesh has been subdivided eight times. Two different sets of coefficients have been applied. Left: Our optimised coefficients. Right: the coefficients suggested by Peters and Shiue [PS03].

vails over the other, by ensuring the subsubdominant eigenvalues are all equal. It has been shown that the subdivision surface of the classical algorithms, such as Catmull-Clark, cannot model certain basic shapes [KPR04], because of one central surface prevailing over another due to subsubdominant eigenvalues not being equal.

By introducing a 3D shape-in-the-limit chart as a measure of goodness of a scheme, we are able to choose mask coefficients and subdominant eigenvalue around an EV such that the variation in Gaussian curvature around an EV is a minimum for all possible configurations around an EV. Other measures of goodness could be considered in the optimisation process and it would be useful to explore these. However, we believe that minimal variation in curvature is a key requirement to ensure good behaviour around an EV.

Although we do not restrict our search to positive values, our optimisation algorithm finds only positive coefficients and thus ensures that the limit surface lies in the convex hull of the control mesh.

For large valencies this method may lead to large subdominant eigenvalues, a situation associated with the *polar artifact* [SB03]. Thus, if the application only wishes to subdi-

vide a small number of times (say 4 or 5), our results presented here may not be 'optimal'. The polar artifact is most visible when the rendering is done by merely drawing the polyhedron instead of a faceting of the limit surface itself. The polar artifact problem is in fact not a surface problem: it is a rendering problem, which is solved by simply refining further around an EV. The 4-8 scheme in particular supports this very well. Also, vertices of very high valencies are not required in order to achieve the modelling of surfaces of arbitrary genus, nor to follow the principle that edges in the mesh should run along features. Large values of λ are therefore not a practical problem.

It is clear from Figures 7, 14, and 17 that any coefficients giving significantly different λ values will give larger curvature variation. Also, the natural configuration will differ from the nice layout observed for optimised coefficients (see Figures 5 and 13).

We demonstrate our algorithm on subdivision schemes working on a quadrilateral mesh, because these are mainly used in CAD/CAM applications. Our method can clearly be applied to any primal quadrilateral subdivision scheme. However, it would be straightforward to modify it to work on primal triangular schemes.

7. Summary and Conclusion

We have presented a new method for tuning a variety of subdivision schemes around EVs. The method ensures bounded curvature and aims at minimising the curvature variation observed at and around EVs.

Our intent was to get the best possible limit surface by tuning the masks of extraordinary vertices. Applications which require only a small number of steps, rather than going to the limit surface, may be better tuned by other criteria. Tuning stencils, while more difficult, may produce even better results.

The method was demonstrated for three primal quadrilateral subdivision schemes: 4-8, Catmull-Clark, and the 4-3 scheme. For these three schemes the mask coefficients optimised with respect to curvature variations are listed in tables for EVs with a valency $5 \leq n \leq 8$. We compared our results against earlier suggestions to improve the behaviour around EVs for these schemes by subdividing a convex surface with high valency ($n = 8$). We illustrated that the limit surface is visibly better when using coefficients determined with the method presented here.

We show that subdivision schemes may not only be tuned to obtain bounded curvature, but that they can also be tuned to achieve smallest possible curvature variation around extraordinary points. This is so important for producing good limit surfaces that we believe such tuning should always be done.

References

- [BK04] BARTHE L., KOBBELT L.: Subdivision scheme tuning around extraordinary vertices. *Computer Aided Geometric Design* 21 (2004), 561 – 583.
- [BS88] BALL A., STORRY D.: Conditions for tangent plane continuity over recursively generated B-spline surfaces. *ACM Transactions on Graphics* 7, 2 (1988), 83–108.
- [CC78] CATMULL E., CLARK J.: Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer Aided Design* 10, 6 (1978), 183–188.
- [DS78] DOO D., SABIN M.: Behaviour of recursive division surfaces near extraordinary points. *Computer Aided Design* 10, 6 (1978), 177–181.
- [KPR04] KARCIAUSKAS K., PETERS J., REIF U.: Shape characterization of subdivision surfaces – case studies. *Computer Aided Geometric Design* 21, 6 (2004), 601–614.
- [Loo87] LOOP C.: *Smooth subdivision surfaces based on triangles*. Master's thesis, Department of Mathematics, University of Utah, 1987.
- [Loo02] LOOP C.: Bounded curvature triangle mesh subdivision with the convex hull property. *The Visual Computer* 18, 5-6 (2002), 316–325.
- [PR04] PETERS J., REIF U.: Shape characterization of subdivision surfaces: basic principles. *Computer Aided Geometric Design* 21, 6 (2004), 585–599.
- [PS03] PETERS J., SHIUE L.-J.: 4-3 directionally ripple-free subdivision. *ACM Transactions on Graphics* (2003).
- [Rei96] REIF U.: A degree estimate for subdivision surfaces of higher regularity. *Proc AMS* 124,7 124, 7 (1996), 2167–2174.
- [Sab02] SABIN M.: Eigenanalysis and artifacts of subdivision curves and surfaces. In *Primus Workshop Proceedings* (2002), pp. 69–92.
- [SB03] SABIN M., BARTHE L.: Artifacts in recursive subdivision surfaces. In *St Malo 2002* (2003), pp. 353–362.
- [Vel01] VELHO L.: Quasi 4-8 subdivision. *Computer Aided Geometric Design* 18, 4 (2001), 345–358.
- [VZ01] VELHO L., ZORIN D.: 4-8 subdivision. *Computer Aided Geometric Design* 18, 5 (2001), 397–427.
- [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.
- [ZSS96] ZORIN D., SCHRÖDER P., SWELDENS W.: Interpolatory subdivision for meshes with arbitrary topology. In *Proc ACM SIGGRAPH 1996* (1996), pp. 189–192.