Interpolating Subdivision for Meshes with Arbitrary Topology

Denis Zorin[†] Peter Schröder[†] Wim Sweldens[*]

Abstract

Subdivision is a powerful paradigm for the generation of surfaces of arbitrary topology. Given an initial triangular mesh the goal is to produce a smooth and visually pleasing surface whose shape is controlled by the initial mesh. Of particular interest are interpolating schemes since they match the original data exactly, and are crucial for fast multiresolution and wavelet techniques. Dyn, Gregory, and Levin introduced the Butterfly scheme [17], which yields C^1 surfaces in the topologically regular setting. Unfortunately it exhibits undesirable artifacts in the case of an irregular topology. We examine these failures and derive an improved scheme, which retains the simplicity of the Butterfly scheme, is interpolating, and results in smoother surfaces.

1 Introduction

Modeling the geometry of surfaces of arbitrary topology is an important area of research in computer graphics and approximation theory. A powerful paradigm for the construction of such surfaces is *subdivision*. Beginning with an input mesh a sequence of meshes is defined whereby new vertices are inserted as, preferably, simple local affine combinations of neighboring vertices. An attractive feature of these schemes is that they are *local*, i.e., no global system of equations needs to be solved. For example, classical spline constructions fit into this category since the resulting surface can be evaluated with the de Casteljau algorithm. These schemes are generally not interpolating, although modifications are possible to recover interpolating schemes (see below). Another class of schemes is based on *interpolating* subdivision. Perhaps the most common of these is piecewise linear interpolation. Unfortunately this is not smooth enough for many applications. A scheme that achieves C^1 continuity, at least in the topologically regular setting, was pioneered by Dyn, Gregory, and Levin [17, 18] and has been applied to the construction of smooth surfaces.

The mathematical analysis of the surfaces resulting from subdivision is not always straightforward (see for example Reif [28]). However, the simplicity of the algorithms and associated data structures makes them attractive for large datasets and interactive applications where speed is of the essence.

Recently interest in *interpolating* subdivision has increased since it allows fast multiresolution and wavelet decompositions of complex geometry [24, 19]. Interpolating subdivisions present the only currently known way to build smooth *finite* analysis and synthesis filters for wavelet algorithms on general manifolds [30]. These wavelet schemes can be thought of as instances of the so-called *lifting scheme* [31] which in a fundamental way is related to interpolating subdivision. Furthermore, the ability to build *adaptive* subdivisions relies on the interpolating nature of the subdivision rule. Multiresolution decomposition algorithms are of importance for compression, progressive display and transmission, multiresolution editing, and for multigrid/wavelet based numerical methods.

While the Butterfly scheme of Dyn, Gregory and Levin can be used to generate smooth surfaces over triangular meshes in which every vertex is of valence 6, it exhibits degeneracies when applied in a topologically irregular setting: smoothness can be lost at vertices of valence not equal to 6. Figure 1 shows an example of such a failure for a vertex of valence 3. The left shows the control mesh, in the middle the resulting

[†] Department of Computer Science, California Institute of Technology, Pasadena CA 91125.

[*] Lucent Technologies, Bell Labs Innovations, Rm. 2C-371, 600 Mountain Avenue, Murray Hill, NJ 07974; Department of Computer Science, Katholieke Universiteit Leuven, Celestijnenlaan 200 A, B 3001 Leuven, Belgium. (On leave as Senior Research Assistant of the National Fund of Scientific Research - NFWO)

dzorin@gg.caltech.edu, ps@cs.caltech.edu, wim@research.att.com

mesh after 3 levels of applying the Butterfly scheme, and on the right the result achieved with our modified scheme, which does not exhibit the cusp-like loss of smoothness.

Motivated by these observations we consider in this paper the construction of subdivision schemes under the following constraints:

- **Interpolation:** The original mesh vertices are interpolated and all newly generated vertices at any step during the subdivision are on the limit surface. In particular this implies that points will not be moved once they have been generated.
- **Locality:** The neighborhood used to define new vertex positions from old ones should be as small as possible to enable fast algorithms.
- **Symmetry:** The scheme should exhibit the same type of symmetries as the local mesh topology.
- **Generality:** The scheme should work over triangulations which are not topologically restricted, including the proper handling of boundaries.
- **Smoothness:** We require the resulting scheme to reproduce polynomials up to some power—a necessary but not sufficient condition for higher order continuity.
- **Simplicity:** The scheme should only require simple topological and flexible data structures so that adaptive level of detail, for example, is easily incorporated into an implementation.

Since the Butterfly scheme satisfies these requirements except for topological generality, we make it the starting point of our investigation.

The main result of our work is a simple modification of the Butterfly scheme around vertices of valence not equal to 6. It combats the cusp like artifacts exhibited by the unmodified scheme in those circumstances. We use eigen analysis and Fourier transform techniques [12, 1] to *synthesize* new interpolating subdivision schemes.

Before describing the main result in detail sufficient for implementation, we briefly review related work and discuss a number of issues which are relevant to smooth subdivision schemes in general. This is followed by a section discussing the results of our modification. We conclude with a discussion and outlook towards future research issues. For purposes of exposition, the mathematical derivations have been moved to the appendix.

2 Related Work

In this section we briefly review a number of algorithms which attempt to build smooth surfaces over arbitrary topology control meshes through the use of subdivision. These come in two principal varieties, approximating and interpolating. The former are typically based on generalizations of spline patch based schemes, while the latter are related to 1D interpolating schemes [15, 13, 8, 9, 10].

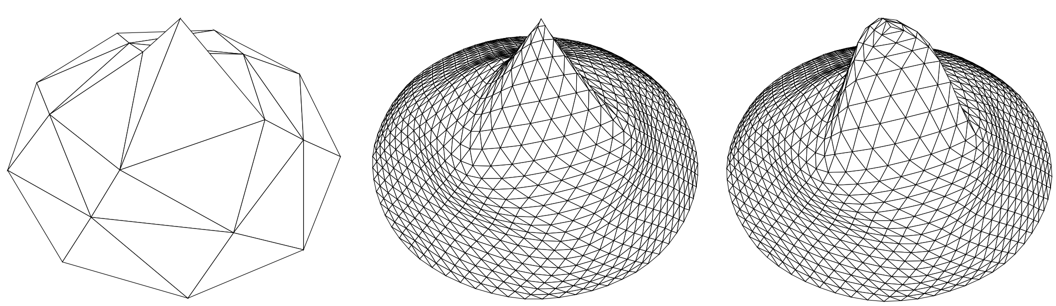

Figure 1: On the left an input control mesh with a vertex of valence 3. In the middle the result after several levels of subdivision using the Butterfly scheme. The surface loses smoothness around the vertex of valence 3. On the right the result achieved in the same situation with our modified scheme. The behavior around the extraordinary vertex remains smooth and no cusp is formed.

2.1 Approximating Subdivision Schemes

The first subdivision algorithms for meshes of arbitrary topology were given by Doo and Sabin [11, 12, 29] and Catmull and Clark [2]. These were based on generalizations of quadratic and cubic B-spline subdivision for meshes consisting of quadrilaterals. The behavior around extraordinary vertices was first analyzed by Doo and Sabin [12] using Fourier transforms and an eigen analysis of the subdivision process. This work was expanded upon by Ball and Storry [1] who derived some necessary conditions. More recently Warren [32] showed that polynomial reproduction is a necessary condition for smoothness of planar subdivision schemes. The behavior around extraordinary vertices was also studied by Reif [28], who derived necessary and sufficient conditions for C^1 continuity. A first scheme for arbitrary meshes consisting of triangles was given by Loop [22]. It is based on a generalization of quartic triangular B-splines.

Another class of approaches to arbitrary topology meshes attempts to directly derive a set of spline patches which globally achieve some order of continuity. Most of these approaches are based on some number of initial "corner cuttings" to regularize the topology, or alternatively place some restrictions on the mesh connectivity [23, 26, 27, 5]. The output of these algorithms is a set of patches of varying, at times rather high, polynomial order and varying shape, typically triangles and quadrilaterals. Once these patches have been generated the surface can be built by subdivision through the de Casteljau algorithm.

2.2 Interpolating Subdivision Schemes

Since the ability to control the resulting surface exactly is very important in many practical applications, a number of modifications of approximating schemes have been developed to force the limit surface to interpolate particular points. Nasri [25] gives modifications to the quadratic scheme of Doo-Sabin to enforce interpolation of vertices and normals by solving a linear system which is global but sparse. Similarly Halstead, Kass and DeRose [20] give an algorithm modifying the cubic scheme of Catmull-Clark to enforce position and normal conditions, again by solving a global and sparse linear system. In both cases there are a number of limitations. For example it is unclear under what conditions the linear system to be solved for the interpolation constraints can become singular. Additionally, the interpolation conditions are only satisfied in the limit. Among the patch based schemes only Peters [26] recently gave one which can incorporate interpolation constraints without requiring the solution of a global linear system. While this could in principle be used to define an interpolating subdivision scheme, we chose a more direct route by considering subdivision schemes which are interpolating by design.

Surprisingly little is known about the construction of such schemes in the topologically irregular setting. Piecewise linear interpolation over arbitrary triangulations certainly belongs to this class and satisfies our simplicity requirement, but for many applications it is not smooth enough. The only other alternative, which satisfies almost all of our requirements, is the Butterfly scheme of Dyn, Gregory and Levin [17] and a later variant [16]. These schemes are interpolating by design, local in that they use a small neighborhood, and simple in terms of the required data structures and algorithms. They are also known to be smooth in the regular setting, where they lead to C^1 limit functions [18, 16]. Topological regularity, however, is a rather severe restriction since all vertices must be of valence six for these results to be applicable. The failure to be smooth for vertices of valence other than six is easily observed in practice as can be seen in the example of Figure 1.

Generalizations of patch based schemes offer simple descriptions of the resulting surfaces in terms of polynomials, except possibly at extraordinary vertices. However, more involved algorithms are required to derive the associated patches. Enforcing interpolation conditions generally leads to the need to solve global but sparse linear systems, although alternatives to this have recently become available [26]. In contrast, interpolating schemes converge to limit surfaces which are harder to characterize and in most cases do not have a closed form solution. This should not be regarded as an obstacle when it comes to evaluating the surface at a point, for which simple and efficient recursive algorithms exist. Even a quantity such as the normal can be computed directly from the mesh. The resulting algorithms are uniform and simple.

3 Interpolating Subdivision Surfaces

Before giving the main results of our analysis we consider what makes a surface smooth. This is the first step in deciding on how to modify a given scheme to create smoother surfaces. In particular it provides us with the motivation to pursue *polynomial reproduction* as a design criterion, i.e., to look for schemes which converge to a polynomial (of suitably low order) if the initial control mesh admits a polynomial parameterization (of the same low order).

3.1 Smoothness

One of the advantages of patch based, *approximating* polynomial schemes is that their *analytic* smoothness properties are well understood and closely correlate with what a human observer would call a smooth surface. Note that this property is only partially covered by such notions as C^k continuity. For example, it is quite possible to have a C^2 function which is nonetheless quite "wiggly," exhibiting surface undulations which are not desirable. On the other hand it is possible to have functions of fractional regularity $\alpha < 2$ whose visual appearance is very smooth nonetheless [9].

These properties of overall smoothness and appearance are often referred to as "fairness." For example, Halstead, Kass and DeRose [20] observed that enforcing interpolation conditions on Catmull-Clark subdivision surfaces resulted in unwanted undulations. They addressed this problem by applying a global optimization pass, minimizing a weighted Sobolev norm over the surface after some number of initial subdivisions.

Ideally one would like to avoid a global minimization step. In the present paper we do not consider the question of optimal fairness. Nonetheless we are attempting to build smooth interpolating subdivision schemes which yield surfaces whose shape is as pleasing as possible. Because no vertex is ever moved once it is computed, any distortion in the early stages of the subdivision will persist. This makes particularly the first few subdivision steps very important, an observation which stands in marked contrast to the asymptotic view.

Considering the local properties of the surface, we find that smoothness corresponds to the existence of a smooth parameterization $(x(s,t), y(s,t), z(s,t))$ of the surface. In our case these functions x, y, and z, can be thought of as defined over a triangulated parameter plane (s,t). In order to construct a smooth surface with our interpolating subdivision scheme we are therefore naturally led to consider the general problem of constructing smooth interpolating subdivision schemes over the real plane. Such a subdivision scheme is then used to build the smooth parameterization of the limit surface. We refer to this as a *planar subdivision*.

3.2 Planar Subdivision

Instead of thinking of surface subdivision in terms of vectors we now consider only one of the coordinate functions. This can be visualized as the graph of some function over the real parameter plane, it is, however, not to be confused with the actual surface. More concretely, instead of thinking about generating new vertices in $\mathbf{R}^3$ by taking local averages of vertices in the control polyhedron, we are now thinking of *scalar values* which are located at vertices in $\mathbf{R}^2$. The planar vertices may be part of a topologically arbitrary triangulation. This triangulation is refined by cutting each triangle into four through the insertion of midpoints in all edges and reconnecting. Now the interpolating subdivision scheme is used to find a new value for the midpoint of some edge by taking suitable averages of nearby values. Note that all new planar vertices will be of valence 6. The neighbors participating in the computation are part of the subdivision *stencil* and their weights characterize the scheme.

Since we assume that all our schemes will be local, i.e., confined to some M-neighborhood, we need to analyze only a small number of possible cases of the relationship between the new site and the topology of its M-neighborhood.

- **Regular:** All vertices in the neighborhood are of valence 6. In this case the Butterfly scheme and its ten point variant are known to produce C^1 limit functions.
- **Isolated K-vertex:** The site in question is immediately adjacent to a $K \neq 6$ valence vertex. In this case the Butterfly stencil looses some of its smoothness properties and a new analysis needs to be

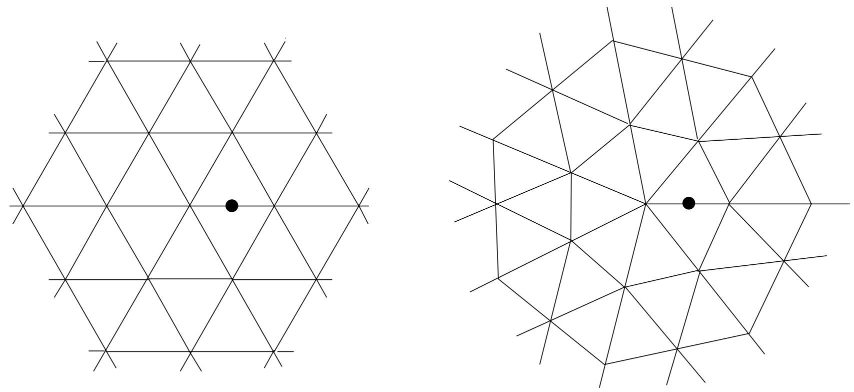

Figure 2: On the left a regular neighborhood in which all vertices have valence 6. On the right an example of a $K = 7$ vertex. In both cases the dot represents a typical midpoint for which we seek to compute a new value. In the regular case the Butterfly stencil is used, while in the $K \neq 6$ a separate analysis is performed leading to a local modification of the weights.

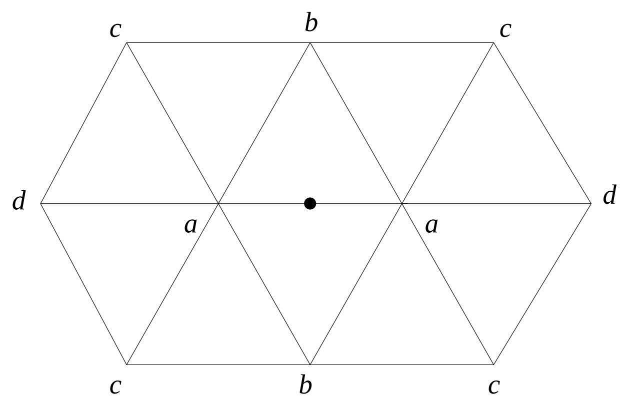

Figure 3: Ten-point stencil.

applied.

It is also possible to have K-vertices which are not isolated, but we postpone that case to a later section.

As mentioned in the introduction, we would like to have smooth limit surfaces. Above we argued that in order to build subdivision surfaces, we need to consider subdivision of functions defined over triangulations of the plane. When we use the stencils from a planar case to build subdivision surfaces, we actually implicitly build a parameterization of the surface, i.e., a local mapping from the surface to the plane. A sufficient condition to get a smooth manifold is that this mapping has a non-zero Jacobian at each point and is smooth itself.

The mapping is defined by three coordinate functions which are functions built by planar subdivision over the parameter space. For each K-vertex in the original mesh, we consider a local mapping to a K-regular triangulation, i.e., a triangulation which has a single K-vertex surrounded by vertices of valence 6. Once we are in that setting, rotational symmetry of the topology allows us to apply a Fourier analysis and reduce the modification of stencils around K-vertices to a simple computation. This derivation is given in the appendix and it amounts to ensuring that the scheme reproduces at least all quadratic polynomials around a K-vertex for $K > 3$. In the case $K = 3$, all quadratic polynomials cannot be reproduced (using only a 1-neighborhood), but good results are achieved for a scheme that reproduces all bilinear functions (Figure 5).

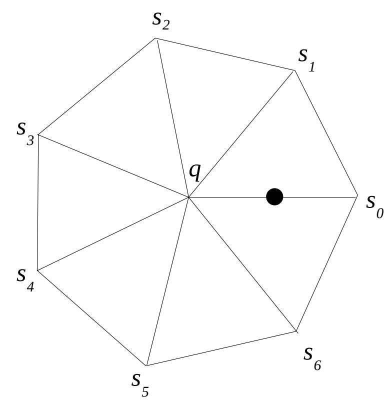

Figure 4: Stencil for a vertex in the 1-neighborhood of an extraordinary vertex

3.3 The Modified Subdivision Scheme

The subdivision scheme computes a new scalar value for each edge midpoint of the triangulation. We distinguish between three situations.

First consider the case where the edge connects two vertices of valence 6. In that case we use the extension of the Butterfly scheme to the ten point stencil [16] (see Figure 3). The dot indicates the midpoint of the edge for which a new scalar value is computed. This is the canonical setting and the weights are given by

$$a = 1/2 - w, \ b = 1/8 + 2w, \ c = -1/16 - w, \ d = w$$

where w can be chosen suitably small [16] (we used $w = 0$).

Second, for an edge connecting a K-vertex ($K \neq 6$) and a 6-vertex, the 1-neighbors of the K-vertex are used in the stencil as indicated in Figure 4. For $K \geqslant 5$ the weights are given by

$$s_j = \frac{1}{K}\left(1/4 + \cos(2\pi j/K) + 1/2\cos(4\pi j/K)\right) \tag{1}$$

with $j = 0, \ldots, K-1$. For $K = 3$ we take $s_0 = 5/12$, $s_{1,2} = -1/12$, and for $K = 4$, $s_0 = 3/8$, $s_2 = -1/8$, $s_{1,3} = 0$.

Finally, in case the edge connects two extraordinary vertices, we take the average of the values computed using the appropriate scheme of the previous paragraph for each endpoint. Since this case can only occur at the topmost level of subdivision the ultimate smoothness of the scheme is not influenced by this choice. However, we have found that the overall *fairness* of the resulting shapes tends to be better with this scheme.

The final algorithm for our interpolating scheme is thus as follows:

1. For all edges with both endpoints of valence 6 compute the value for the midpoint using the coefficients of the ten point scheme.
2. For every edge with one extraordinary vertex, compute the value for the midpoint using coefficients of (1).
3. For every edge that connects two extraordinary vertices, compute values for the midpoint using the coefficients (1) for each vertex and take the average.

4 Results

We built an interactive application supporting general triangular meshes and adaptive subdivision. We use triangular quadtrees and enforce a restriction criterion, which allows easy implementation of such operators as `Neighbor` and `VertexValence`. All subdivision coefficients are precomputed and stored in a table indexed by the valence of a given vertex. We present here some results obtained with this application.

4.1 Bases

We argued earlier that the examination of planar subdivision schemes is fundamental to understanding the surface case. Since the operator which constructs a planar subdivision based on some initial data at all the vertices is linear, we can write the result as a superposition of *basis functions*. A basis function is defined as the subdivision limit of the Dirac sequence, i.e., the value at one vertex is set to 1 while all others are set to zero. Here we consider the case of a K-regular triangulation which consists of isometric triangles where all vertices have valence 6, except the origin, which has valence K. Setting the value at the K-vertex to 1 and all others to zero results in a basis functions whose graph can be visualized over the parameter plane. Figure 5 shows the basis function for a regular triangular tiling (all vertices have valence 6) in the upper left. The row immediately below illustrates the behavior of the unmodified Butterfly scheme for vertices of valence 3, 8 and 13. Note how the smoothness of the bases is lost and the function become very "spiked." In practice this behavior often leads to cusp-like features on the limit surfaces. The top row shows an extreme example, a tetrahedron and its subdivision using the unmodified scheme. Next to it is the result of applying the modified scheme. Now the resulting surface is globally smooth. The basis functions of the modified scheme associated with K-regular settings for $K = 3$, $K = 8$, and $K = 13$ are shown in the third row.

The images also demonstrate the character of the influence of a single vertex on its neighborhood. As it can be seen the support of the basis functions is only slightly larger than their diameter. As the valence of a vertex increases, the basis function has ever more lobes as demonstrated for $K = 13$. A possible way to smooth out the lobes is to go to a scheme which is also modified along edges emanating from K-vertices. The image at the bottom left shows the result of doing this ($K = 13$ here as well) with a scheme we are currently investigating.

4.2 Complex Shapes

Finally we demonstrate the application of these schemes to more complex shapes. In the top half of Figure 6 the geometry of 6 pipes meeting at right angles is shown.[1] On the top left the original control polyhedron. Note that it contains vertices of valence 7 and 4 directly neighboring each other. After applying 6 levels of Butterfly subdivision the shape pictured directly underneath results. It shows various regions where the smoothness of the surface collapsed. In contrast on the top right is the shape resulting from applying the modified scheme to the same original control polyhedron. These images also demonstrate the treatment of boundaries. Edges which are on the boundary are subdivided using the 1-dimensional 4 point scheme ($s_{-1} = -1/16$, $s_0 = 9/16$, $s_1 = 9/16$, $s_1 = 9/16$) [15, 13]. In this case only other edge points participate in the stencil. A consequence of this rule is that two separate geometries, whose boundary is identical, will have a matching boundary curve after subdivision. This is of obvious utility when designing separate parts to be connected later. Edges which are not on the boundary but which have a vertex which is on the boundary are subdivided as before with the proviso that any vertices in the stencil which would be on the other side of the boundary are replaced with "virtual" vertices. These are constructed on the fly by reflecting vertices across the boundary.

Finally we applied our scheme to a dataset of a mannequin head (courtesy University of Washington). The bottom half of Figure 6 shows the original polyhedron and 4 successive levels of subdivision applied to it and a closeup.

5 Summary and Future Work

We have presented a simple interpolating subdivision scheme for meshes with arbitrary topology. Our scheme is based on the Butterfly scheme, with special rules applied in the neighborhood of the extraordinary vertices. The proposed scheme has a number of properties that make it attractive:

- it is interpolating at all levels of subdivision;

[1] This configuration was inspired by a similar configuration of Jens Albrecht, Erlangen University.

- the support of the scheme is minimal;
- it is easy to implement;
- limit surfaces have adequate smoothness;
- subdivision with this scheme can be performed adaptively.

There are explicit formulas for the normals, and due to the simple form of the left eigenvectors of the subdivision matrix one can hope to find geometric conditions for degenerate configurations of points; one such condition is presented in Section A.9. The scheme is especially convenient for multiresolution representation of surfaces and wavelet representation of functions on surfaces.

There are several aspects in which the proposed scheme can be improved:

- Higher degree of smoothness at extraordinary vertices can be achieved by increasing the size of the invariant neighborhood for these vertices and modifying the scheme along the edges of the original mesh.
- Currently we perform only a limited adaptation of the scheme at the boundary. A more detailed analysis of K-vertices on or near the boundary would be desirable.
- By collapsing vertices and edges, this scheme can immediately accommodate mesh tagging approaches such as in [21]. This again relies on the fact that the boundary of a resulting surface only depends on the boundary vertices of the input mesh. Also, the degree of smoothness of the surface can be continuously adjusted by manipulating the remaining degrees of freedom of the scheme. This might allow continuous variations in smoothness in the spirit of tension parameters.
- Little is known about theoretical properties of the scheme. One can hope to analyze its continuity at the extraordinary vertices using methods based on the work of Reif [28].
- We have observed that the "fairness" of the surface is determined by the behavior of the scheme at the first two subdivision steps. A more thorough analysis of the behavior of the scheme and its possible extensions at these subdivision steps may yield new insights on maintaining fairness under the constraint of interpolation.

Another interesting research direction is to explore the properties of the modification of our scheme suggested in Section A.7 which has the unique feature that the same subdivision procedure can be applied to all edges at all levels.

Acknowledgements

This work was supported in part by an equipment grant from Hewlett Packard and funds provided to second author by the Charles Lee Powell Foundation. Additional support was provided by NSF (ASC-89-20219), as part of the NSF/DARPA STC for Computer Graphics and Scientific Visualization. All opinions, findings, conclusions, or recommendations expressed in this document are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

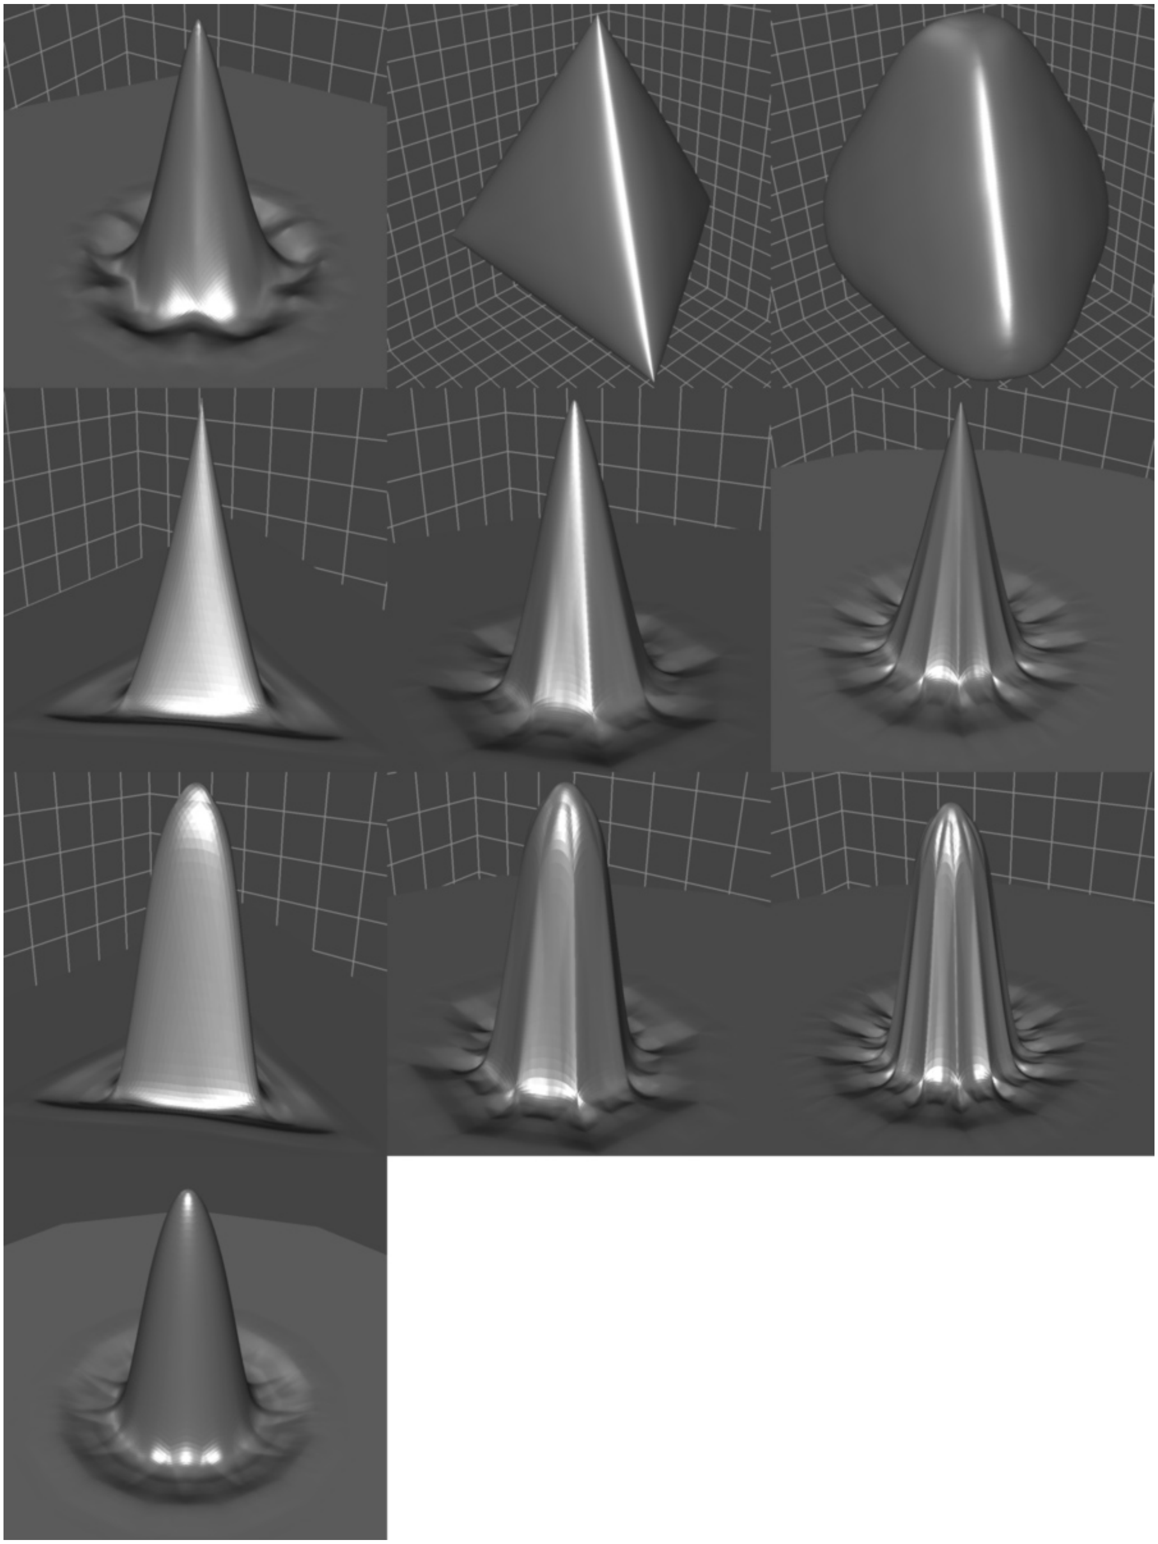

Figure 5: Examples of basis functions. Top left: the basis function of the Butterfly scheme on a regular mesh; top middle: failure of unmodified scheme for tetrahedron; top right: resulting surface for tetrahedron when using the modified scheme; middle row: basis functions for K-regular setting using the unmodified Butterfly scheme for $K = 3$, $K = 8$, and $K = 13$; third row: basis functions of modified scheme in the K-regular setting with $K = 3$, $K = 8$, and $K = 13$; bottom left: a more recent modified scheme for $K = 13$.

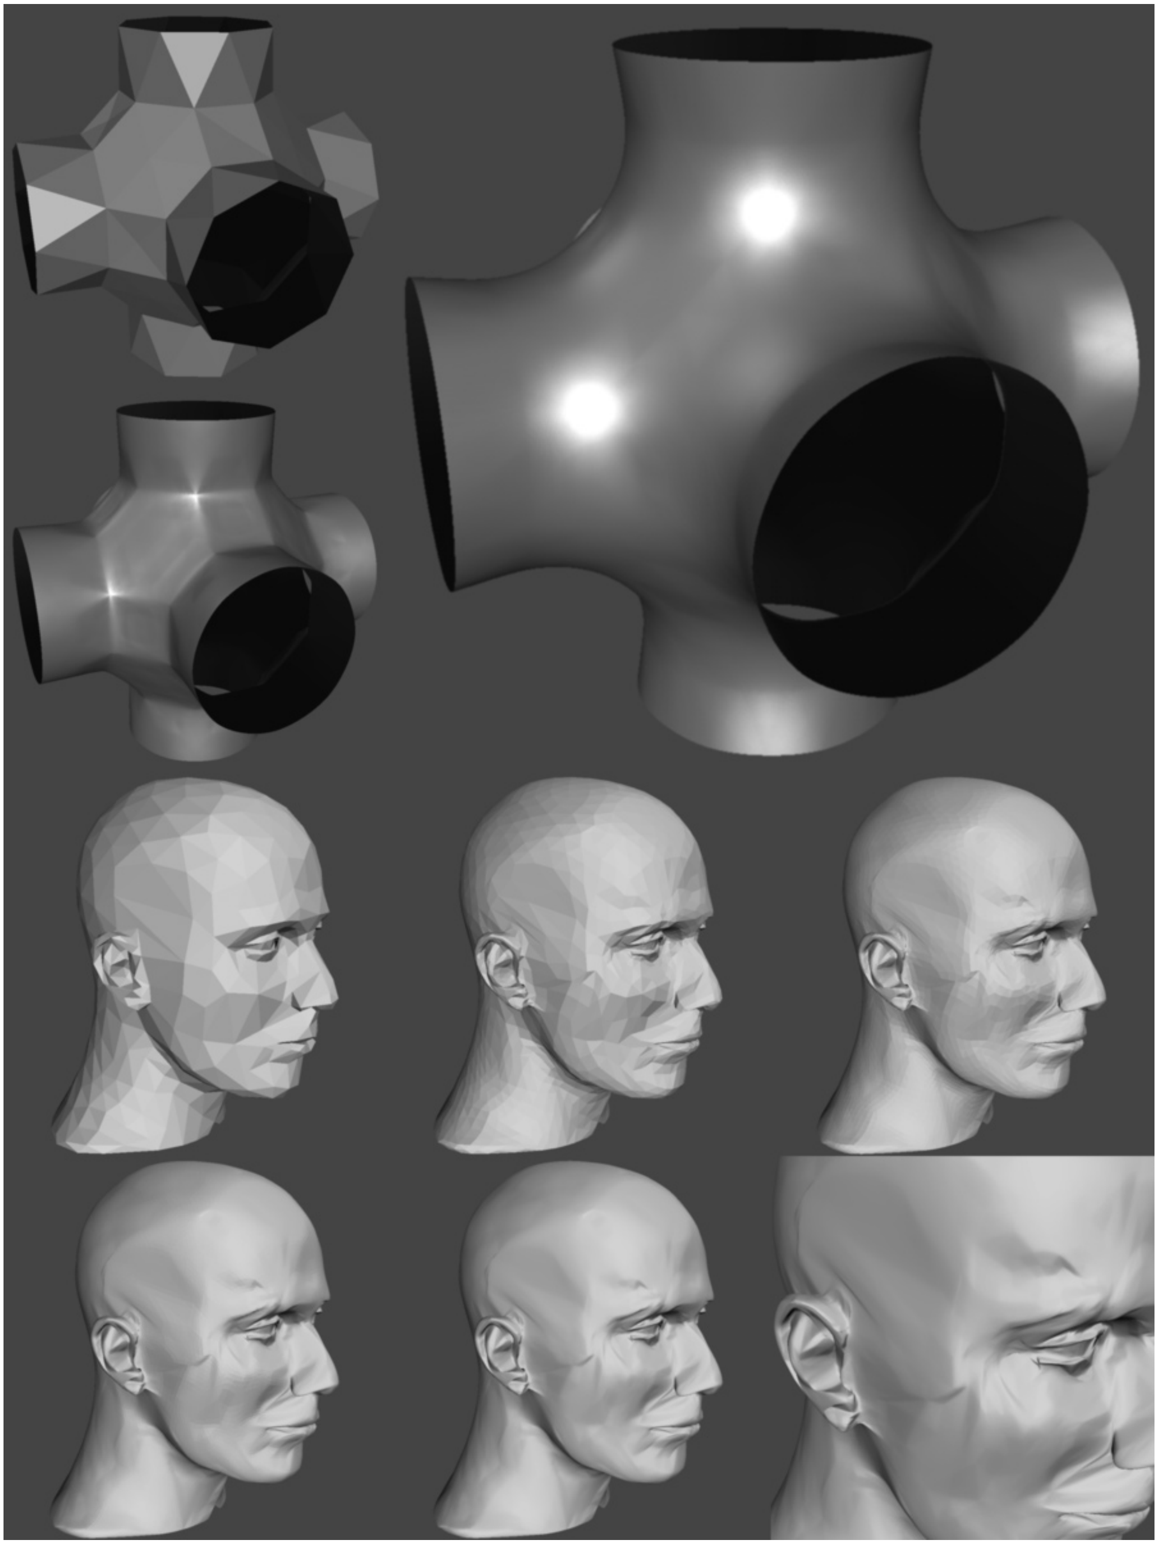

Figure 6: The top images demonstrate the difference between the classical Butterfly scheme and the modified scheme for a complex shape, a set of pipes meeting. On the top left is the original control polyhedron. Below the result of applying the unmodified scheme. Note the pinched regions. Next to it an enlarged image of the modified scheme applied to the same control shape. On the bottom 5 levels of subdivision starting with a simple mesh approximating a mannequin head (courtesy University of Washington).

References

[1] Ball, A. A., and Storry, D. J. T. Conditions for Tangent Plane Continuity over Recursively Generated B-Spline Surfaces. *ACM Transactions on Graphics 7*, 2 (1988), 83–102.

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

[3] Cavaretta, A. S., Dahmen, W., and Micchelli, C. A. Stationary Subdivision. *Memoirs Amer. Math. Soc. 93*, 453 (1991).

[4] Colella, D., and Heil, C. Characterizations of Scaling Functions: Continuous Solutions. *SIAM J. Math. Anal. 15* (1994), 496–518.

[5] Dahmen, W., Micchelli, C. A., and Seidel, H.-P. Blossoming Begets B-Splines Bases Built Better by B-Patches. *Mathematics of Computation 59*, 199 (July 1992), 97–115.

[6] Daubechies, I. *Ten Lectures on Wavelets.* CBMS-NSF Regional Conf. Series in Appl. Math., Vol. 61. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1992.

[7] Daubechies, I., and Lagarias, J. C. Two-scale Difference Equations I. Existence and Global Regularity of Solutions. *SIAM J. Math. Anal. 22*, 5 (1991), 1388–1410.

[8] Deslauriers, G., and Dubuc, S. Interpolation Dyadique. In *Fractals, dimensions non entières et applications.* Masson, Paris, 1987, pp. 44–55.

[9] Deslauriers, G., and Dubuc, S. Symmetric Iterative Interpolation Processes. *Constr. Approx. 5*, 1 (1989), 49–68.

[10] Donoho, D. L. Interpolating Wavelet Transforms. Preprint, Department of Statistics, Stanford University, 1992.

[11] Doo, D. A Subdivision Algorithm for Smoothing Down Irregularly Shaped Polyhedrons. In *Proceedings on Interactive Techniques in Computer Aided Design*, 157–165, 1978.

[12] Doo, D., and Sabin, M. Analysis of the Behaviour of Recursive Division Surfaces near Extraordinary Points. *Computer Aided Design 10*, 6 (1978), 356–360.

[13] Dubuc, S. Interpolation Through an Iterative Scheme. *J. Math. Anal. Appl. 114* (1986), 185–204.

[14] Dyn, N. personal communication, 1995.

[15] Dyn, N., Gregory, J. A., and Levin, D. A Four-Point Interpolatory Subdivision Scheme for Curve Design. *Computer Aided Geometric Design 4* (1987), 257–268.

[16] Dyn, N., Hed, S., and Levin, D. Subdivision Schemes for Surface Interpolation. In *Workshop in Computational Geometry*, A. C. et al., Ed., 97–118, 1993.

[17] Dyn, N., Levin, D., and Gregory, J. A. A Butterfly Subdivision Scheme for Surface Interpolation with Tension Control. *ACM Transactions on Graphics 9*, 2 (April 1990), 160–169.

[18] Dyn, N., Levin, D., and Micchelli, C. A. Using Parameters to Increase Smoothness of Curves and Surfaces Generated by Subdivision. *Computer Aided Geometric Design 7* (1990), 129–140.

[19] Eck, M., DeRose, T., Duchamp, T., Hoppe, H., Lounsbery, M., and Stuetzle, W. Multiresolution Analysis of Arbitrary Meshes. In *Computer Graphics Proceedings*, Annual Conference Series, 173–182, 1995.

[20] Halstead, M., Kass, M., and DeRose, T. Efficient, Fair Interpolation using Catmull-Clark Surfaces. In *Computer Graphics Proceedings*, Annual Conference Series, 35–44, 1993.

[21] Hoppe, H., DeRose, T., Duchamp, T., Halstead, M., Jin, H., McDonald, J., Schweitzer, J., and Stuetzle, W. Piecewise Smooth Surface Reconsruction. In *Computer Graphics Proceedings*, Annual Conference Series, 295–302, 1994.

[22] Loop, C. Smooth Subdivision Surfaces Based on Triangles. Master's thesis, University of Utah, Department of Mathematics, 1987.

[23] Loop, C. Smooth Spline Surfaces over Irregular Meshes. In *Computer Graphics Proceedings*, Annual Conference Series, 303–310, 1994.

[24] Lounsbery, M., DeRose, T. D., and Warren, J. Multiresolution Surfaces of Arbitrary Topological Type. Department of Computer Science and Engineering 93-10-05, University of Washington, October 1993. Updated version available as 93-10-05b, January, 1994.

[25] Nasri, A. H. Surface Interpolation on Irregular Networks with Normal Conditions. *Computer Aided Geometric Design 8* (1991), 89–96.

[26] Peters, J. C^1 Surface Splines. *SIAM J. Numer. Anal. 32*, 2 (1995), 645–666.

[27] Peters, J. Curvature Continuous Spline Surfaces over Irregular Meshes. *Computer Aided Geometric Design* (to appear).

[28] Reif, U. A Unified Approach to Subdivision Algorithms Near Extraordinary Points. *Computer Aided Geometric Design 12* (1995), 153–174.

[29] Sabin, M. *The use of Piecewiwe Forms for the Numerical Representation of Shape.* PhD thesis, Hungarian Academy of Sciences, Budapest, 1976.

[30] Schröder, P., and Sweldens, W. Spherical Wavelets: Efficiently Representing Functions on the Sphere. In *Computer Graphics Proceedings*, Annual Conference Series, 161–172, 1995.

[31] Sweldens, W. The Lifting Scheme: A Construction of Second Generation Wavelets. Tech. Rep. 1995:06, Industrial Mathematics Initiative, Department of Mathematics, University of South Carolina, 1995. ftp://ftp.math.scarolina.edu/pub/imi_95/imi95_6.ps.gz.

[32] Warren, J. Subdivision Methos for Geometric Design. Unpublished manuscript, November 1995.

Appendix A: Planar Subdivision

In this appendix we discuss the construction of planar subdivision schemes that reproduce polynomials. Recall that these schemes will actually be used to build the coordinate functions of the local mapping from the surface to the planar parameter space. We start out by defining the notion of a K-regular triangulation which will form the parameter space for the neighborhood of a vertex with valence K on the original mesh. The Discrete Fourier Transform (DFT) now allows us to exploit the symmetries of the stencil. We use DFT to construct the stencils around K-vertices. Finally, we show how to compute the normals to the limit surface exactly.

A.1 Planar Subdivision Around a K-vertex

Consider an initial triangulation T^0 of the plane with vertices $V^0 = \{v_l^0 \in \mathbf{R}^2 \mid l \in L^0\}$. Here L^0 is an infinite index set. We call a triangulation *K-regular* if one vertex $v_0^0 = (0,0)$ has valence K, all other vertices have valence 6, and all triangles are isometric, see Figure 7 for an example of a 7-regular triangulation. In case $K = 6$ all triangles are equilateral and we call a 6-regular triangulation simply regular. A finer triangulation T^1 with vertices $V^1 = \{v_l^1 \mid l \in L^1\}$ can be obtained from T^0 by adding a new vertex in the middle of each edge of T^0 and connecting the new vertices. All new vertices have valence 6. Repeated application of this procedure yields a sequence of triangulations T^i with $i \geqslant 0$. The vertices of T^i are given by $V^i = \{v_l^i \mid l \in L^i\}$. We have that $V^i \subset V^{i+1}$. We choose the index sets so that $L^i \subset L^{i+1}$ and $v_l^{i+1} = v_l^i$. Note that V^{i+1} consists of "new" vertices $V^{i+1} \setminus V^i$ and "old" vertices V^i. If T^0 is K-regular than all T^i are K-regular as well. In fact, T^i is the dilation of T^{i+1} by a factor of 2. The regular triangulation is in addition invariant under translation.

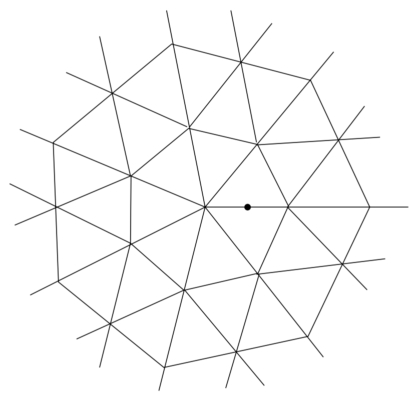

Figure 7: A 7-regular triangulation of the plane. Note how all triangles are isometric.

Consider a set of scalar values $P^0 = \{p_l^0 \mid l \in L^0\}$ assigned to the vertices of the initial triangulation T^0. A *planar subdivision scheme* is a linear operator S that assigns values $P^{i+1} = \{p_l^{i+1} \mid l \in L^{i+1}\}$ to the vertices of the new triangulation T^{i+1} given the values on the vertices of T^i: $P^{i+1} = S\,P^i$. We already mentioned in the section on surface subdivision that the scheme has to be interpolating, i.e., old values do not change, and that the stencil has to be local, i.e., a new value only depends on old values in a finite neighborhood. We add the following two constraints:

- **Affine invariance**: The values computed by the scheme S are invariant under any affine transformation $A = ax + b$. Thus A and S have to commute: $S\,A = A\,S$. By letting $a = 0$ we see that if all old values are equal to a constant, all new values are equal to the same constant. Consequently the coefficients of a stencil sum to one.
- **Symmetry** : The stencil is invariant under isometries that map the triangulation onto itself. A typical example is a rotation with an angle of $2\pi/K$.

We next consider the construction of the stencil in the case when the set of old values is in a 2-neighborhood of the K-vertex. We denote the 2-neighborhood of the K-vertex in T^i with N^i. The set N^i contains $3K+1$

vertices, see Figure 8 for an example in case $K = 7$. We call N^i an *invariant neighborhood* because all values p_l^{j+1} at the vertices of N^{j+1} depend only on the values p_l^j at the vertices of a topologically identical set N^j on a coarser level. As the subdivision is linear, the dependence of the values at the vertices of N^{j+1} on the values at the vertices of N^j can be described by a square $(3K + 1) \times (3K + 1)$ *subdivision matrix* $\mathbf{S}^j$. If the subdivision matrix does not depend on j, we call the subdivision scheme stationary. The smoothness of the limit function at the K-vertex is then completely defined by the subdivision matrix.

A.2 Polynomial Reproduction and Smoothness

As we discussed in the section on surface subdivision, a sufficient condition for the surface to be smooth is that the parameterization is smooth. We thus would like to construct smooth planar subdivision schemes. However there are no necessary and sufficient conditions available for planar subdivision to yield smooth limit functions.

If we consider the one dimensional setting for a moment, the situation is quite different. Partially motivated by wavelet constructions, the smoothness question has been extensively studied and seems to be well understood. We only mention [7, 4, 3, 9] as examples of the literature. A basic result is that in order to achieve a C^k limit function, the subdivision scheme has to reproduce at least all polynomials of degree k. However, very often the resulting smoothness is much less than C^k even though all polynomials of degree $k+1$ are reproduced. Only splines achieve this upper bound. Other families such as the Daubechies' scaling functions [6] or the Deslauriers-Dubuc interpolating functions [9] typically have smoothness proportional to the degree of polynomials reproduced. All these schemes, however, are *constructed* by requiring polynomial reproduction combined with other constraints such as orthogonality or interpolation.

Based on these results and a recent proof given by Warren [32], which shows that polynomial reproduction is also a necessary condition in the planar subdivision setting, we construct our planar subdivision schemes by requiring polynomial reproduction combined with interpolating properties. We start by carefully defining what we mean with polynomial reproduction. Consider a polynomial $P(x)$, $x \in \mathbf{R}^2$ and assume the initial values p_l^0 are samples of this polynomial: $p_l^0 = P(v_l^0)$. We now say that the subdivision scheme reproduces P if all new values are also samples of the same polynomial: $p_l^i = P(v_l^i)$ for $i > 0$ and $l \in L^i$. In the next sections we construct subdivision schemes on K-regular triangulations reproducing all polynomials up to a certain degree. We first simplify the construction using the DFT.

A.3 Fourier Analysis

It is well known that many properties of subdivision schemes are determined by the eigenvalues and eigenvectors of the subdivision matrix [1]. Due to the symmetries of coefficients of a subdivision scheme around a K-vertex a convenient tool for the analysis of such schemes is the Discrete Fourier Transform (DFT) [1].

Consider an invariant 2-neighborhood of the K-vertex in a K-regular triangulation. There are four symmetry types of vertices in this neighborhood, see Figure 8. For any two vertices of the same symmetry type there is a rotation that maps one vertex to the other and maps the triangulation onto itself. Conversely, for two vertices of different type there is no such transformations. We use letters A, B, C to denote values at the vertices of each symmetry type. The center value is denoted with Q. The numbering of vertices of each type is shown in Figure 8.

We arrange the values A_i, B_i, C_i into a vector of length $3K + 1$ in groups corresponding to each type of vertex:

$$P = [Q, A_0 \dots A_{K-1}, B_0 \dots B_{K-1}, C_0 \dots C_{K-1}]^t.$$

The subdivision scheme in the invariant neighborhood can be written as:

$$\begin{aligned}
Q^{j+1} &= Q^j \\
A_i^{j+1} &= s^{AQ} Q^j + \sum_{l=0}^{K-1} (s_{l-i}^{AA} A_l^j + s_{l-i}^{AB} B_l^j + s_{l-i}^{AC} C_l^j) \\
B_i^{j+1} &= s^{BQ} Q^j + \sum_{l=0}^{K-1} (s_{l-i}^{BA} A_l^j + s_{l-i}^{BB} B_l^j + s_{l-i}^{BC} C_l^j) \\
C_i^{j+1} &= A_i^j,
\end{aligned}$$

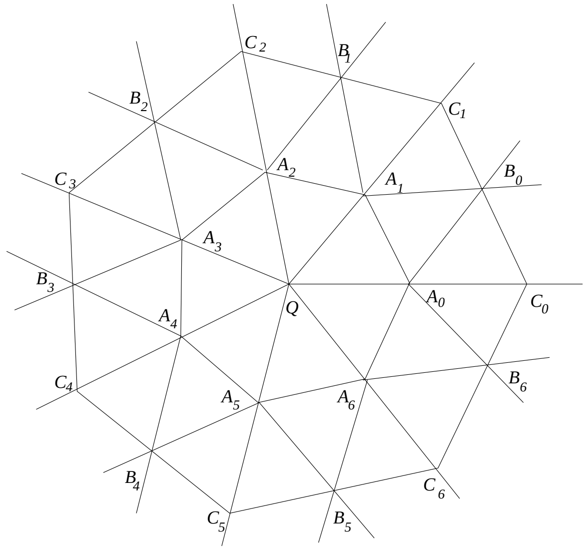

Figure 8: Invariant neighborhood of a 7-point; four symmetry types of vertices are denoted Q, A, B, C.

or as a matrix equation

$$P^{j+1} = \mathbf{S}\, P^j.$$

As the subdivision scheme is interpolating, two types of vertices are inherited from the previous level: Q and C. This means that corresponding rows of the subdivision matrix have one entry equal to 1 and all other entries equal to zero. Old A-values become new C-values on the next subdivision level.

Because constants have to be reproduced each row of the matrix sums up to one. This requirement can always be satisfied by setting the entries in the first column — corresponding to the central vertex — to one minus the sum of the entries of the remaining columns. If we further exploit the assumption of affine invariance we can assume that the value at the central point of the scheme is 0 without loss of generality. Now the central vertex does not contribute to the values at other vertices at all anymore. Therefore, we omit the first column and the first row of the subdivision matrix from now on to simplify exposition. The $3K \times 3K$ subdivision matrix $\mathbf{S}$ can then be separated into nine $K \times K$ blocks:

$$\mathbf{S} = \begin{pmatrix} S^{AA} & S^{AB} & S^{AC} \\ S^{BA} & S^{BB} & S^{BC} \\ \mathbf{I} & \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

The three blocks in the last low are the $K \times K$ unit matrix and zero matrices, reflecting the fact that vertices of type A become vertices of type C on the next subdivision level. Each block is a Toeplitz matrix due to the symmetry of the scheme. It thus simply represents a cyclic convolution. For example, the matrix S^{AA} is given by

$$S^{AA} = \begin{pmatrix} s_0^{AA} & s_1^{AA} & \cdots & s_{K-1}^{AA} \\ s_{K-1}^{AA} & s_0^{AA} & \cdots & s_{K-2}^{AA} \\ \vdots & \vdots & & \vdots \\ s_1^{AA} & s_2^{AA} & \cdots & s_0^{AA} \end{pmatrix}.$$

We next define the Discrete Fourier Transform. Let $W = \exp(-2\pi i/K)$ and let $\mathbf{W}$ be the $K \times K$ DFT matrix:

$$\mathbf{W} = \begin{pmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & W & W^2 & \cdots & W^{(K-1)} \\ 1 & W^2 & W^4 & \cdots & W^{2(K-1)} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & W^{K-1} & W^{2(K-1)} & \cdots & W^{(K-1)^2} \end{pmatrix}.$$

We know that the DFT transforms a Toeplitz matrix into a diagonal matrix. Given that the subdivision matrix has a block Toeplitz structure, we define a corresponding block DFT transform:

$$\mathcal{W} = \begin{pmatrix} \mathbf{W} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{W} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{W} \end{pmatrix}.$$

This is a $3K \times 3K$ matrix with $K \times K$ blocks. We now let the block DFT act on the subdivision matrix:

$$\widehat{\mathbf{S}} = \mathcal{W}\,\mathbf{S}\,\mathcal{W}^{-1} = \begin{pmatrix} \widehat{\mathbf{S}}^{AA} & \widehat{\mathbf{S}}^{AB} & \widehat{\mathbf{S}}^{AC} \\ \widehat{\mathbf{S}}^{BA} & \widehat{\mathbf{S}}^{BB} & \widehat{\mathbf{S}}^{BC} \\ \mathbf{I} & \mathbf{0} & \mathbf{0} \end{pmatrix},$$

where $\widehat{\mathbf{S}}^{AA} = W\,S^{AA}\,W^{-1}$ and similarly for the other blocks. The DFT of a Toeplitz matrix is a diagonal matrix so $\widehat{\mathbf{S}}^{AA} = \mathrm{diag}(\widehat{s}_0^{AA}, \ldots \widehat{s}_{K-1}^{AA})$, and similarly for the other blocks. The vectors $\widehat{s}^{AA} = [\widehat{s}_0^{AA}, \ldots \widehat{s}_{K-1}^{AA}]^t$ and $s^{AA} = [s_0^{AA}, \ldots s_{K-1}^{AA}]^t$ are related by

$$s^{AA} = \frac{1}{K}\mathbf{W}\ \widehat{s}^{AA}.$$

By applying an appropriate permutation, we reduce the subdivision matrix to a block form with 3×3 blocks on the diagonal and zeroes elsewhere:

$$\tilde{\mathbf{S}} = \mathbf{P}\,\widehat{\mathbf{S}}\,\mathbf{P^{-1}} = \begin{pmatrix} \mathbf{B}_0 & 0 & \cdots & 0 \\ 0 & \mathbf{B}_1 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & \mathbf{B}_{K-1} \end{pmatrix},$$

where

$$\mathbf{B}_j = \begin{pmatrix} \widehat{s}_j^{AA} & \widehat{s}_j^{AB} & \widehat{s}_j^{AC} \\ \widehat{s}_j^{BA} & \widehat{s}_j^{BB} & \widehat{s}_j^{BC} \\ 1 & 0 & 0 \end{pmatrix}.$$

The resulting matrix $\tilde{\mathbf{S}}$ has the same eigenvalues as $\mathbf{S}$; moreover, if the eigenvectors of $\tilde{\mathbf{S}}$ are orthogonal, then the eigenvectors of $\mathbf{S}$ are orthogonal, because DFT and permutation matrices are multiples of unitary matrices. If V is an eigenvector of $\tilde{\mathbf{S}}$, then $\mathbf{S}$ has an eigenvector $\mathbf{P}\,\mathbf{W}\,P$. All properties of $\mathbf{S}$ can be described through the properties of a set of 3×3 matrices $\mathbf{B}_j$. It is important to note that the eigenvectors of $\tilde{\mathbf{S}}$ can be chosen to have the form $[0, \ldots, r_0, r_1, r_2, \ldots, 0]^t$ with non-zero entries only in three positions corresponding to the matrix $\mathbf{B}_j$. Here $[r_0, r_1, r_2]^t$ is an eigenvector of $\mathbf{B}_j$. Consequently the eigenvectors of the matrix $\mathbf{S}$ can be chosen to have the form $[r_0\mathbf{e}^{(j)}, r_1\mathbf{e}^{(j)}, r_2\mathbf{e}^{(j)}]^t$, where $\mathbf{e}^{(j)} = [1, W^j, \ldots W^{(K-1)j}]$, the jth row of the $K \times K$ matrix $\mathbf{W}$. We always will take j modulo K in the notation $\mathbf{e}^{(j)}$.

A.4 Construction of Planar Subdivision Schemes

We use the DFT machinery developed in the previous section to construct a polynomial-reproducing scheme at a K-vertex. Polynomial reproduction for the invariant neighborhoods means that vectors of values of monomials at the vertices of the neighborhood should be eigenvectors of the subdivision matrix. As the size of the invariant neighborhood shrinks by 2, the eigenvalues for the monomial eigenvectors are $1/2^m$, where m is the total degree of the monomial.

We represent the coordinates of a point in the plane using complex numbers. Reproducing polynomial functions of the coordinates of points is equivalent to reproducing the polynomial functions of the corresponding complex value $z = x + iy$ and its conjugate $\bar{z} = x - iy$. The vector of values of z at the vertices of the invariant set can be written as

$$[1/2\,\mathbf{e}^{(1)},\ \cos(\pi/K)\,W^{-1/2}\,\mathbf{e}^{(1)},\ \mathbf{e}^{(1)}]^t.$$

The First K components of the vector above $(1/2\,\mathbf{e}^{(1)})$ are the values A_i, the next K $(\cos(\pi/K)\,W^{-1/2}\,\mathbf{e}^{(1)})$ are the values B_i, and last K are the values C_i.

Asking for the scheme to reproduce all polynomials $x^l y^{m-l}$ with $m \leqslant M$ is equivalent to requiring that it reproduce all polynomials $z^l\,\bar{z}^{m-l} = z^{2l-m}$ with $m \leqslant M$. Polynomials of this form lead us to consider

$$P^{(l,m-l)} = [1/2^m\,\mathbf{e}^{(2l-m)},\, \cos^m(\pi/K)\,W^{-(2l-m)/2}\,\mathbf{e}^{(2l-m)},\, \mathbf{e}^{(2l-m)}]^t.$$

Applying the DFT and permutation $\mathbf{P}$ to the vector $P^{(l,m-l)}$ we obtain a vector which has non-zero entries only in three positions starting with $3((2l-m) \bmod K)$. These positions correspond to the eigenvector of the matrix $B_{(2l-m) \bmod K}$.

$$\tilde{P}^{(l,m-l)} = [0,\, \ldots 0,\, 1/2^m,\, \cos^m(\pi/K)\,W^{-(2l-m)/2},\, 1,\, 0,\, \ldots 0]^t.$$

Thus, we arrive at the following simple condition for polynomial reproduction on the subdivision matrix at the extraordinary vertex:

> For polynomial reproduction up to order M on K-regular triangulations, the blocks $B_{(2l-m) \bmod K}$, with $1 \leqslant m \leqslant M$ and $0 \leqslant l \leqslant m$ have to have eigenvectors $\tilde{P}^{(l,m-l)}$ with eigenvalue $1/2^m$.

Note that reproduction of the constant function has already been ensured through the affine invariance of the scheme. Thus the range of m begins at 1 (rather than 0). The total number of distinct polynomials is $2 + \ldots + (M+1) = M(M+3)/2$. Given that we want to fix that many eigenvalues we must have $M(M+3)/2 \leqslant 3K$.

Let us consider this condition in some more detail to see what happens if $K > 3$. Reproduction of linears ($m = 1$) prescribes eigenvalue 1/2 for the blocks $\mathbf{B}_1$ ($l = 1$) and $\mathbf{B}_{K-1}$ ($l = 0$). Reproduction of quadratics ($m = 2$) prescribes the eigenvalue 1/4 for the blocks $\mathbf{B}_0$ ($l = 1$), $\mathbf{B}_2$ ($l = 2$), and $\mathbf{B}_{K-2}$ ($l = 0$). Cubic reproduction ($m = 3$) prescribes eigenvalue 1/8 for the blocks $\mathbf{B}_1$, $\mathbf{B}_3$, $\mathbf{B}_{K-1}$, and $\mathbf{B}_{K-3}$. Obviously some blocks have several eigenvalues prescribed. If for an order M the number of eigenvalues required for a given block is larger than 3, the polynomial reproduction of that order fails. This implies that no matter how large K is we can not hope to do better than reproducing polynomials of order 6 (within a 2-neighborhood), since at that point blocks 0 and $K-1$ have had 3 distinct eigen values assigned to them.

Next take $K = 3$ and $M = 3$. Now $\mathbf{B}_0$ has to have eigenvalues $1/8, 1/4, 1/8$, $\mathbf{B}_1$ has to have eigenvalues $1/2, 1/4, 1/8$, and $\mathbf{B}_2$ has to have eigenvalues $1/2, 1/4, 1/8$. In fact the eigenvectors with eigenvalue 1/8 for $\mathbf{B}_0$ coincide so there is one degree of freedom left.

A.5 The Regular Setting

First we consider regular triangulations. In this case we use the Butterfly scheme first described in [17] and analyzed in [18, 16]. If an interpolating subdivision scheme introduces new points at midpoints of the edges of the original triangulation, the smallest possible stencil consists of two edge endpoints. Clearly, this stencil is not sufficient to reproduce polynomials of degree higher then 1. It turns out that the next larger stencil, consisting of 4 points (Figure 9) is not sufficient for reproducing quadratic polynomials. Adding four more points (see Figure 9) turns out to be sufficient for reproducing quadratic polynomials. The coefficient of the scheme are uniquely defined by the condition of polynomial reproduction: $a = 1/2$, $b = 1/8$, $c = 1/16$. In addition, due to the symmetry of the scheme it turns out that it reproduces cubic polynomials. A natural extension of this scheme in the regular case is a parametric family of schemes on the 10-point stencil as in Figure 10 with coefficients $a = 1/2 - w$, $b = 2w + 1/8$, $c = -1/16 - w$, $d = w$ [16]. Although no scheme from this family reproduces all polynomials of degree 4, the smoothness of the surface can be somewhat improved by choosing a non-zero value for w. All new points introduced by the subdivision are regular, and the simple scheme described above can be used to generate all points of the surface. It was shown in [18] that for regular subdivision of the plane the ten-point schemes produce C^1 limit functions for a range of w.

At the extraordinary vertices the scheme is not always C^1: for vertices of valences 3,4 and greater than 8 it is known not to be C^1 [14]. The following analysis shows that the 10-point scheme in these cases does not always satisfy the necessary conditions for C^1-continuity [28].

The invariant neighborhood for any vertex for a ten-point scheme has size 2, so we can use the results of the previous sections. The blocks $\mathbf{B}_j$ of the subdivision matrix for a vertex of valence K can be written

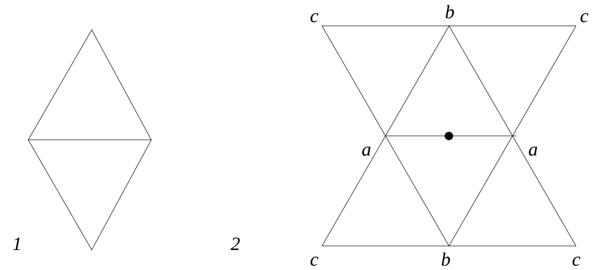

Figure 9: 1. Four-point stencil. 2. Eight-point stencil (Butterfly scheme.)

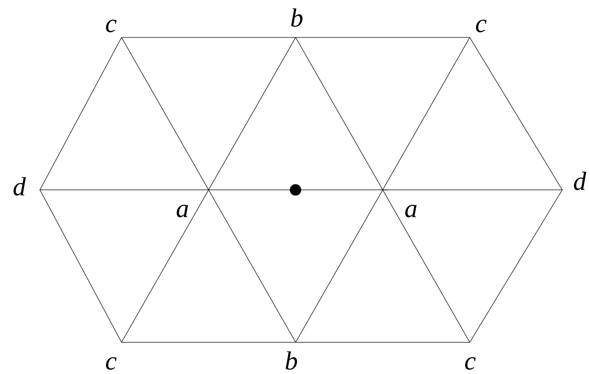

Figure 10: Ten-point stencil. Note that this stencil is the average of two 6 point stencils around the old vertices of the edge containing the new vertex

in the following form ($w = 0$):

$$\mathbf{B}_j = \begin{pmatrix} \frac{1}{2} + \frac{1}{4}\cos(j\omega_K) - \frac{1}{8}\cos(2j\omega_K) & -\frac{1}{16}\left(1 + e^{-ij\omega_K/2}\right) & 0 \\ \frac{1}{2}\left(1 + e^{ij\omega_K/2}\right) - \frac{1}{16}\left(e^{-ij\omega_K/2} + e^{3ij\omega_K/2}\right) & \frac{1}{8} & -\frac{1}{16}\left(1 + e^{ij\omega_K/2}\right) \\ 1 & 0 & 0 \end{pmatrix}$$

where $\omega_K = 2\pi/K$.

For example, for a vertex of valence 3, the largest eigenvalue is $\frac{1}{4}$, and it has multiplicity at least 6. While for large values of k the subdivision matrix may satisfy the necessary conditions for continuity, in general, there tends to be more then one eigenvalue which is close in magnitude to the largest eigenvalue. As a result, the functions generated by the scheme near extraordinary vertices, while they may be C^1 for some valences, informally speaking, do not have enough smoothness to produce adequate interpolating surfaces.

If K approaches infinity, the matrices $\mathbf{B}_j$ and $\mathbf{B}_{j+1}$ for a fixed j go to the same limit. Suppose the largest eigenvalue (1/2) can be found in the block $\mathbf{B}_j$. In almost all cases the eigenvalues of $\mathbf{B}_{j+1}$ converge to the eigenvalues of $\mathbf{B}_j$. As a result, more and more eigenvalues of the subdivision matrix cluster around 1/2. The sequence of decreasing eigenvalues $1/2, 1/2, 1/4, 1/4, 1/4, \ldots$ needed to ensure a C^1 surface thus breaks and the resulting surface is not smooth.

A.6 Extraordinary Vertices

For extraordinary vertices we start out by only considering a 1-neighborhood of the K-vertex to build a subdivision scheme. The stencil then only has $K+1$ coefficients $[q, s_0, s_1, \cdots, s_{K-1}]$, see Figure 11 in the case $K = 7$. After omitting the center vertex, the subdivision matrix $\mathbf{S}$ now is a $K \times K$ Toeplitz matrix. In this case our analysis still applies, but all coefficients except s_j^{AA} have to be set to 0. There is no need to consider a block DFT any more and we can immediately consider $\mathbf{W}\,\mathbf{S}$. The block matrices $\mathbf{B}_j$ simply become real numbers which are eigenvalues of $\mathbf{S}$. The condition for reproducing the polynomials of degree 2 in this case becomes

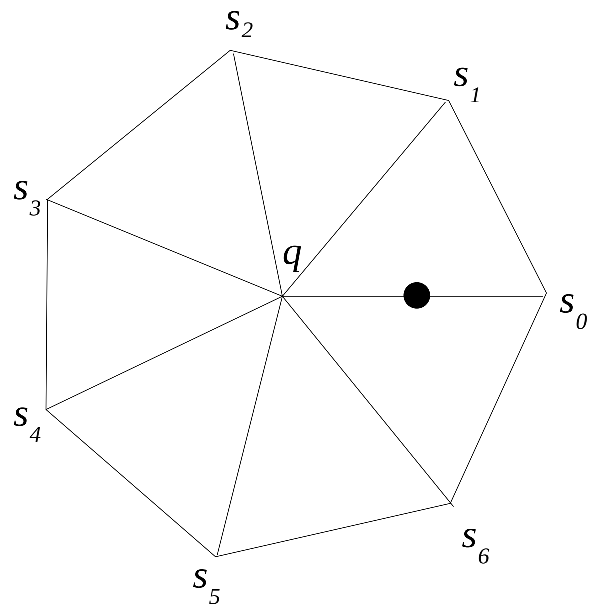

Figure 11: Stencil for a vertex in the 1-neighborhood of an extraordinary vertex.

> The eigenvalues of $\mathbf{S}$ with numbers 1 and $K-1$ should be $1/2$. The eigenvalues with numbers 0, 2, and $K-2$ should be $1/4$.

In other words, this describes 5 eigenvalues: two times $1/2$ and three times $1/4$. Consequently, in the case of $K=3$ and $K=4$ it is impossible to reproduce all quadratics. One then only has either 3 or 4 eigenvalues of $\mathbf{S}$ while quadratic reproduction requires 5 eigenvalues. For $K=3$ we take the eigenvalues $[1/4, 1/2, 1/2]$ which leads to the stencil $q=3/4$, $s_0=5/12$, $s_1=s_{-1}=-1/12$. In case $K=4$ we take the eigenvalues $[1/4, 1/2, 1/4, 1/2]$ which leads to the stencil $q=3/4$, $s_0=3/8$, $s_1=s_3=0$, and $s_2=-1/8$.

For $K \geq 5$ there are $K-5$ eigenvalues of the subdivision matrix that we are free to choose. We found that in general they do not have significant effect on the appearance of the limit surfaces, as long as their absolute value is sufficiently smaller then $1/4$. The simplest choice is to set them all to 0. The sequence of eigenvalues of $\mathbf{S}$ is thus:

$$[1/4, 1/2, 1/4, 0, \ldots, 0, 1/4, 1/2].$$

In case $K \geq 5$, the coefficients $[s_0, s_1, \ldots, s_K]$ are now simply given by the inverse DFT of the vector of eigenvalues. This yields:

$$s_j = \frac{1}{K}\left(\frac{1}{4} + \cos\frac{2\pi j}{K} + \frac{1}{2}\cos\frac{4\pi j}{K}\right). \tag{2}$$

The coefficient q for the central vertex is obtained by subtracting the sum of all other coefficients from 1, resulting in $q = 3/4$ for the central vertex.

A.7 Top-level Subdivision

One question that remains unanswered in the construction above is the choice of the extraordinary vertex for an edge of the original mesh which has extraordinary vertices at both endpoints. In this case we propose to take the average of the positions computed using the appropriate 1-neighbor scheme for each endpoint.

There is an remarkable connection between this procedure and the ten-point schemes. Assume that we would do this averaging also in case both endpoints of an edge have valence 6. Consider thus the regular setting and do a one neighborhood analysis around a 6-vertex. From the previous section we see that the eigenvalues are $[1/4, 1/2, 1/4, -1/24 - 2w, 1/4, 1/2]$, where the 4th eigenvalue has a free parameter w. We thus obtain a one-parametric family of coefficients given by $q = 3/4$, $s_0 = 1/4 - 2w, s_1 = 1/8 + 2w, s_2 = -1/8 - 2w, s_3 = 2w, s_4 = -1/8 - 2w, s_5 = 1/8 + 2w$. If we now average this scheme applied at the left and at the right endpoint we end up *precisely* with the one-parametric 10-point scheme discussed earlier (with the same w parameter).

A.8 Computing the Normals

A disadvantage of the proposed scheme is the lack of the explicit form for the limit surface. However, due to the linear nature of the subdivision process it is possible to compute a number of important functions on the surface explicitly. The neighborhood analysis technique can be used to compute the vertex normals to the limit subdivision surface. In [20] it was shown that the (unnormalized) normal at a vertex can be computed as a vector product of two 1×3 vectors:

$$n = \tau_1 \times \tau_2.$$

Here $\tau_1^t = L_1\mathbf{P}$, $\tau_2^t = L_2\mathbf{P}$, L_1 and L_2 are the left $1 \times 3K$ eigenvectors corresponding to the largest eigenvalue of $\mathbf{S}$, and $\mathbf{P}$ is a $3K \times 3$ matrix containing the coordinate vectors of the surface vertices in an invariant neighborhood. Any subdivision level j can be chosen. The vectors τ_1 and τ_2 are orthogonal tangent vectors to the surface and thus n is the surface normal. Note that there is no immediate guarantee that the vectors τ_1 and τ_2 are non zero.

Normals at extraordinary vertices. In the case of extraordinary vertices we only need to consider the invariant 1-neighborhood. The vectors are thus of length K (as opposed to $3K$). The eigenvectors of the subdivision matrix with eigenvalue $1/2$ correspond to z and $\bar{z}$ which is $1/z$ as we are on the unit circle:

$$\begin{aligned} L_1 &= {P^{(1,0)}}^t = \mathbf{e}^{(1)} = [1, W^{-1}, W^{-2}, \ldots W^{-(K-1)}] \\ L_2 &= {P^{(0,1)}}^t = \mathbf{e}^{(-1)} = [1, W, W^2, \ldots W^{(K-1)}]. \end{aligned}$$

In this case the subdivision matrix is a Toeplitz matrix, and the left and right eigenvectors coincide and are orthogonal. The vectors τ_1 and τ_2 are the second components of the Fourier transform and inverse Fourier transform of P. Although the expressions for τ_1 and τ_2 are complex, real tangent vectors can be obtained by taking the real part of their sum and complex part of their difference.

In this case it is possible to give a simple sufficient (but not necessary) condition for the normal to exist. For any integer j with $0 \leqslant j \leqslant K-1$ we can split the set of indices $\{0, \ldots, K-1\}$ of vertices in the invariant neighborhood into two sets (all indices are taken modulo K): $H_1(j) = \{j, \ldots, j + \lfloor K/2 \rfloor\}$ and its complement $H_2(j) = \{j + \lfloor K/2 \rfloor + 1, \ldots, j + K - 1\}$. For a unit direction in space d consider the orthogonal projection onto d. Take the projections $P_j^d = \mathbf{P}_j\, d^t$ of the points of the invariant neighborhood onto this direction. Then $\tau_1\, d^t = L_1\, \mathbf{P}\, d^t b = L_1\, P^d$. Now consider the length K vector $P^d = \mathbf{P}\, d^t$. If this vector has two zero crossings, its second Fourier component and thus τ_1 is non-zero. In general, if there is a real number C and an integer j, $0 \leqslant j \leqslant K-1$, so that for all $m \in H_1(j)$ $P_m^d < C$ and for $m \in H_2(j)$ $P_m^d > C$, then the normal obtained from Equation 5 is not a zero vector. This condition means that there is a plane such that $\lfloor K/2 \rfloor$ consecutive points of the invariant neighborhood are above this plane and the other $\lceil K/2 \rceil$ are below.

Normals at regular vertices. In this case the invariant neighborhood of a point has size 2. Using the fact that the eigenvectors corresponding to different blocks $\mathbf{B}_j$ are orthogonal it is sufficient to find the left eigenvectors corresponding to the eigenvalue $1/2$ in blocks B_1 and $B_{(K-1)}$. This procedure results in the following formulas for the left eigenvectors:

$$\begin{aligned} L_1 &= \frac{1}{5}[16\,\mathbf{e}^{(1)},\ -\frac{8}{\sqrt{3}}e^{\frac{\pi i}{6}}\mathbf{e}^{(1)},\ \mathbf{e}^{(1)}] \\ L_2 &= \frac{1}{5}[16\,\mathbf{e}^{(-1)},\ -\frac{8}{\sqrt{3}}e^{\frac{-\pi i}{6}}\mathbf{e}^{(-1)},\ \mathbf{e}^{(-1)}]. \end{aligned}$$

In this case it is difficult to find an intuitive geometric criterion for the normal to exist; it is possible however, to write a simple linear condition on the values at the vertices which guarantees that the normal is not zero.