A Unified Framework for Primal/Dual Quadrilateral Subdivision Schemes

Denis Zorin *

New York University

Peter Schröder

Caltech

Abstract

Quadrilateral subdivision schemes come in primal and dual varieties, splitting faces or respectively vertices. The scheme of Catmull-Clark is an example of the former, while the Doo-Sabin scheme exemplifies the latter. In this paper we consider the construction of an increasing sequence of alternating primal/dual quadrilateral subdivision schemes based on a simple averaging approach. Beginning with a vertex split step we successively construct variants of Doo-Sabin and Catmull-Clark schemes followed by novel schemes generalizing B-splines of bidegree up to nine. We prove the schemes to be C^1 at irregular surface points, and analyze the behavior of the schemes as the number of averaging steps increases. We discuss a number of implementation issues common to all quadrilateral schemes. In particular we show how both primal and dual quadrilateral schemes can be implemented in the same code, opening up new possibilities for more flexible geometric modeling applications and p-versions of the *Subdivision Element Method*. Additionally we describe a simple algorithm for adaptive subdivision of dual schemes.

Key words: Subdivision, quadrilateral, primal, dual, approximating, Doo-Sabin, Catmull-Clark, B-splines, averaging

1 Introduction

Historically, subdivision schemes were derived from B-spline knot refinement rules [1] to address the challenge of building smooth surfaces of arbitrary topology [2,3]. Beginning with an arbitrary connectivity control mesh, which forms a topological

* Corresponding author

2-manifold possibly with boundary, the surface is constructed through a limiting process of repeated refinement (Figure 1).

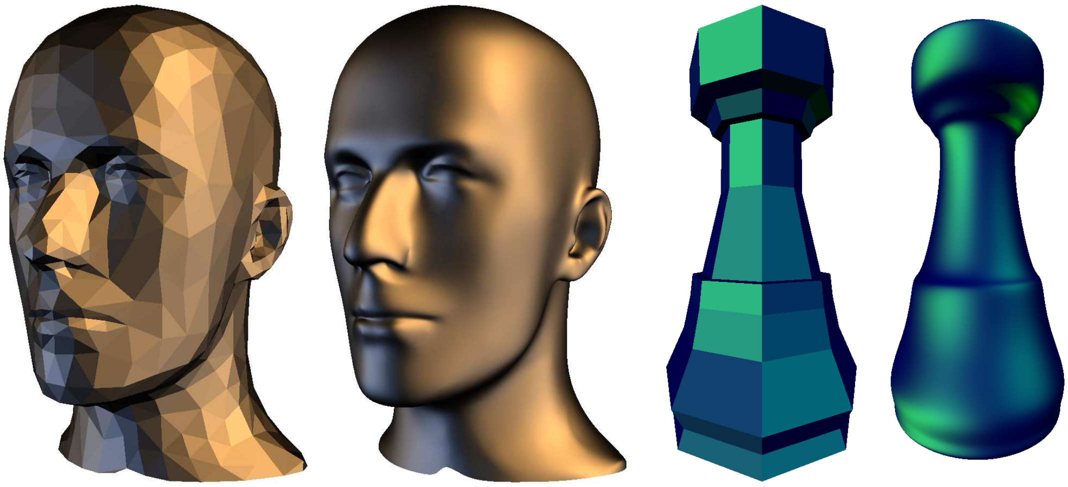

Fig. 1. Examples of subdivision based on triangles and quadrilaterals respectively. On the left Loop's scheme and on the right the scheme of Catmull and Clark.

Subdivision consists of two components, a *topological split* rule describing how the connectivity of the control mesh is refined, and a *geometric rule* which determines the new control point positions from the old positions. Most constructions use smoothing filters of (small) finite support and local definition with constant coefficients depending only on the valence of vertices or faces in the support of the filter. These schemes can be grouped according to a number of basic criteria (Table 1) indicating whether they are

- approximating or interpolating with respect to the original control mesh point positions;
- based on quadrilateral or triangle (hexagon) faces as basic primitive;
- of primal or dual type depending on the split rule (faces or vertices respectively).

Scheme	approx.	interpol.	quad	triangle	primal	dual
Midedge [4,5]	*		*			*
Doo-Sabin [3]	*		*			*
Catmull-Clark [2]	*		*		*	
Loop [6]	*			*	*	
Butterfly [7,8]		*		*	*	
Kobbelt [9]		*	*		*	

Table 1. Classification of major subdivision schemes.

Among approximating subdivision schemes those based on quadrilaterals assume a special role since there exist both primal and dual schemes based on quadrilaterals [3,2,4,5,10]. Typically the topological step of a primal scheme is described as a *face split* while dual schemes employ *vertex splits* (Figure 2).

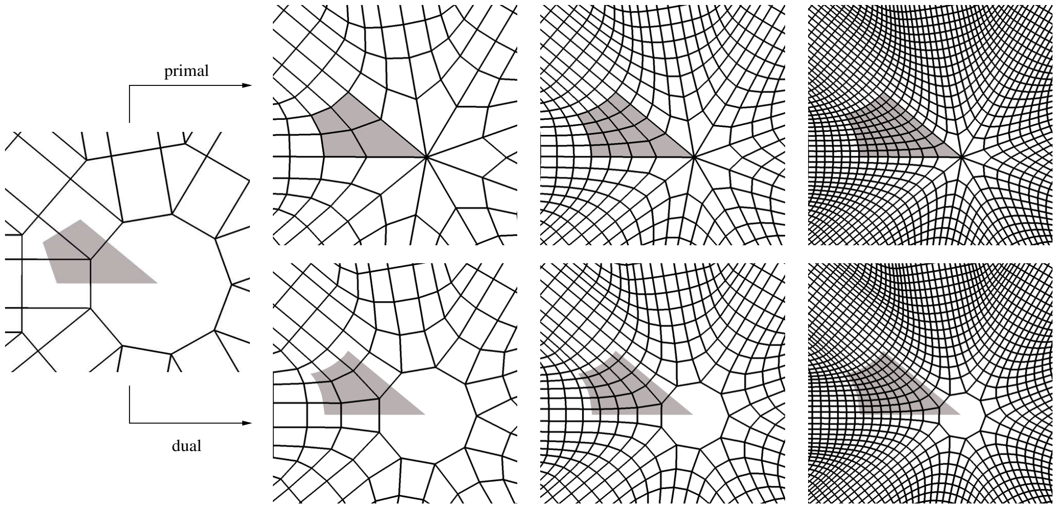

Fig. 2. Primal subdivision schemes are based on face splits while dual subdivision schemes perform vertex splits. Splits are typically four to one, though others are possible [11]. Here we see examples of the Catmull-Clark (top) and Doo-Sabin (bottom) schemes acting on the same input mesh (left). The region corresponding to a limit patch is highlighted.

In this paper we consider the construction of a family of alternating primal/dual quadrilateral subdivision schemes in a unified framework. The approach is based on a straightforward generalization of the repeated averaging construction for B-splines [1] (Section 2). Our construction produces variants of Doo-Sabin and Catmull-Clark as well as a number of novel schemes generalizing B-splines of arbitrary degree. The smoothness of these schemes at irregular surface points up through bidegree nine is proven for all valences (Section 4), and we observe that higher order schemes exhibit problems around high valence faces/vertices which, surprisingly, worsen with increasing order (Figure 19). The schemes are easy to implement and we detail a number of implementation choices which (1) allow the concurrent accommodation of primal and dual schemes within the same datastructure and (2) support adaptive, crack-free tessellations for primal and dual schemes (Section 3). Among all the schemes we single out biquartic subdivision, since it is only slightly more expensive than Doo-Sabin but yields significantly nicer shapes. Even higher order schemes may be of interest in CAGD applications as well as in p-methods for the solution of PDEs based on the Subdivision Element Method [12].

Related Work. Since repeated averaging is a natural approach to the construction of higher order schemes it comes as no surprise that several researchers arrived at the idea nearly simultaneously. For example, a construction of arbitrary order quadrilateral subdivision schemes, identical to our approach, has been proposed by Stam [13]. Stam also considers constructions based on double averaging and an initial classical Catmull-Clark, respectively Doo-Sabin step. A family of repeated averaging schemes similar to ours was described in [14]. However, no analysis was provided and no implementation details were discussed. Similarly, a class of schemes based on 4-8 triangle meshes was derived through generalization of the averaging construction of 4-direction splines by Velho and co-workers [10,15]. In their setting 4-direction splines as well as 2-direction splines of increasing order can be constructed by combining simple bisection refinement steps. Warren [16] considers the construction of bi- and trivariate subdivision schemes based on factored B-splines. Finally, Zheng and Kondo [17] consider adaptive Doo-Sabin subdivision.

Our work focuses on several aspects:

- uniform treatment of dual and primal schemes, based on an elementary barycenter averaging step;
- implementation of adaptive primal and dual subdivision in a common framework;
- theoretical study of the properties of the sequence of schemes, the eigenstructure of the subdivision matrices, and of the characteristic maps.

Basic Assumptions. We assume that the reader is familiar with basic notions of a mesh as a collection of vertices, edges and faces, and with control points assigned to vertices. The input mesh is expected to have 2-manifold topology. For simplicity we further assume that there are no boundaries, a restriction that can be relaxed. We emphasize the important distinction between the *vertices*—which have no geometric positions—and control *points*—which do carry geometric positions. In distinction from the usual setup our meshes carry control points at vertices *and* faces. The former will be referred to as *primal points*, while the latter will be called *dual points*.

2 Quadrilateral Subdivision as Repeated Averaging

Our common approach to primal and dual subdivision schemes is based on the idea of dualization: compute the topological and geometric dual of a mesh by connecting the barycenters of all faces, and repeat the process several times, jumping back and forth. The result is a smoother mesh. For the complete process we need three ingredients: vertex split, control point replication, and barycenter averaging:

- **Vertex Split:** Split each vertex v of valence k into k new vertices one for each face containing v (Figure 3, left to middle). Build the new mesh by connecting these vertices into faces, which are in 1-to-1 correspondence with the old vertices, faces, and edges (Figure 3, right). Recall that each of these faces carries a control point (as yet with undefined value). Since all new vertices have four vertex neighbors the most natural datastructure is that of quadtrees: k trees for each original vertex v of valence k.

- **Control Point Replication:** When **vertex split** is used we initialize the control points of the new vertices through assignment of the control point value of the vertex whose split created them. One may think of this as order zero b-spline control point refinement.
- **Repeated Averaging:** In this step we assign each face control point the barycenter of its incident vertex control points and switch to the (topological) dual mesh. This step is repeated as many times as desired. Note that an even number of such repeated averaging and dualization steps results in a mesh of the same connectivity.

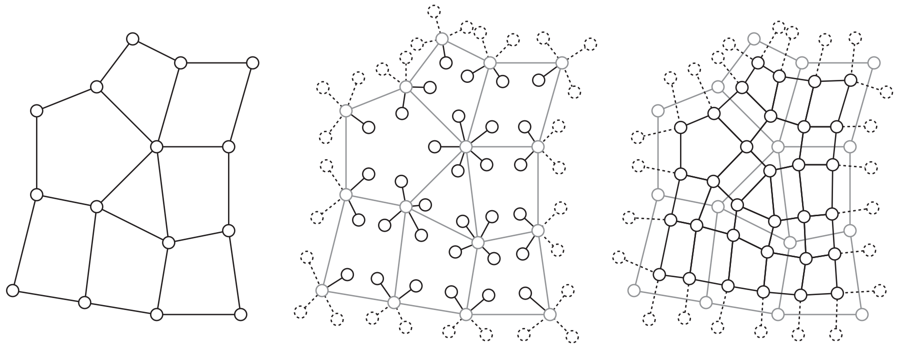

Fig. 3. An input mesh (left) may consist of faces and vertices with arbitrary valence. In a first step each vertex of valence k is split into k, one for each face incident (middle). These are completed into a dual mesh by constructing new faces, one for each original vertex, edge, and face (right).

In practice (Section 3) we construct the dual meshes only implicitly by recourse to the control points associated with faces. With this convention dual and primal meshes exist at the same time, and barycenter averaging just transfers data back and forth between the associated control points. Figure 4 illustrates the averaging process.

2.1 Subdivision by Repeated Averaging

We are now ready to consider the schemes that we obtain by means of repeated averaging. If we only replicate control points, we can regard the resulting process as constant subdivision; all new control points coincide with one of the old ones. Visually, the object does not change. Subsequent averaging steps result in the following sequence of schemes.

First Step: Midpoint Subdivision. The first averaging step computes primal control point positions as averages of dual (initial) control points (Figure 4, left),

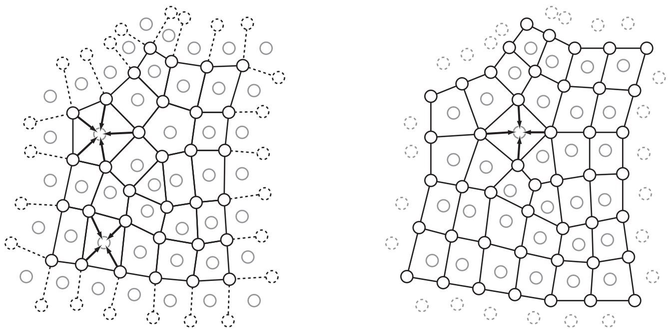

Fig. 4. After the initial vertex split (Figure 3) and point replication a sequence of alternating dual to primal (left) and primal to dual (right) centroid averaging steps builds the ladder of alternating primal dual averaging schemes. Centroid averages are indicated with arrows.

resulting in midpoint, i.e., generalized bilinear, subdivision.

Second Step: Doo-Sabin Subdivision. The second averaging step computes dual control points as averages of neighboring primal control points (Figure 4, right) and produces the Catmull-Clark variant of Doo-Sabin subdivision [2] (Figure 5). For

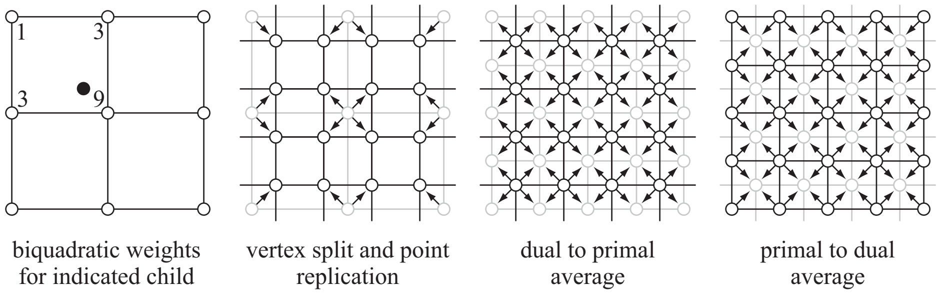

Fig. 5. In the regular setting of biquadratic splines a child is computed with the stencil indicated on the left. The same effect can be achieved through two averaging steps.

irregular faces let p_i^j, $i = 0, \ldots, k-1$ be the points associated with the vertices v_i^j of some face of valence k on subdivision level j. New points are now computed as

$$p_l^{j+1} = \sum_{i=0}^{k-1} w_{i-l} p_i^j$$

with

$$w_{i-l} = \frac{\delta(i-l)}{2} + \frac{\delta(i-l+1)}{8} + \frac{1}{4k}, \tag{1}$$

where δ denotes the Kronecker symbol.

Third Step: Catmull-Clark Subdivision. Continuing with a third averaging step produces a variant of the Catmull-Clark scheme. We do not discuss this case in detail, as it is quite similar to the the previous two cases.

The first new scheme appears after another averaging step and generalizes biquartic splines. It is of particular interest as it requires little more machinery than biquadratic subdivision, but yields significantly better results.

Fourth Step: Biquartic Subdivision. The subdivision stencil resulting from fourfold averaging has support such that the scheme can be implemented with little additional effort beyond the cost of a standard Doo-Sabin scheme. As mentioned before, vertices are best organized as quadtrees in a dual subdivision scheme. It is then natural to compute all four children of a given vertex at the same time. Considering the support size of the stencils for the Doo-Sabin scheme we see that this implies access to all vertices of the faces incident to a given vertex. If these vertices have to be accessed one may as well use non-zero coefficients for all of them for *each* child. Figure 6 (left) shows the support of the resulting stencil for biquartic B-spline subdivision in the regular case. Qu [18] was the first to consider a generalization of the biquartic spline to the arbitrary topology setting. He derived some conditions on the stencils but did not give a concrete set of coefficients. With our simple averaging approach we arrive at a particular set of coefficients. In case that at most one irregular face is incident on a given vertex, [1] stencils shown in Figure 6 result. In these stencils the irregular face contributes *additional* weights as

$$\begin{aligned} nw_i &= \frac{64}{k} + 48w_i + 16w_{i-1} + 16w_{i+1} \\ ne_i &= 32w_i + 16w_{i-1} \\ se_i &= 16w_i, \end{aligned} \qquad (2)$$

where w_i are the weights of Equation 1, $i = 0, \ldots, k-1$ and indices are taken modulo k.

Fifth and Subsequent Steps. Clearly this process can be continued, generating alternating dual and primal schemes. One can ask a number of questions about the behavior of the schemes as the number of averaging steps increases: What degree of continuity do they have at an irregular surface point? What happens to the surface quality? These questions are considered in greater detail in Section 4, where schemes up through order nine are analyzed. A qualitative comparison of basis functions for orders two through nine and valences four and nine is shown in Figures 7 and 8 respectively. Even degree schemes have a square or ninegon in the center with one vertex raised. Odd degree schemes have a valence four or nine

[1] This is always the case after at most two subdivision steps.

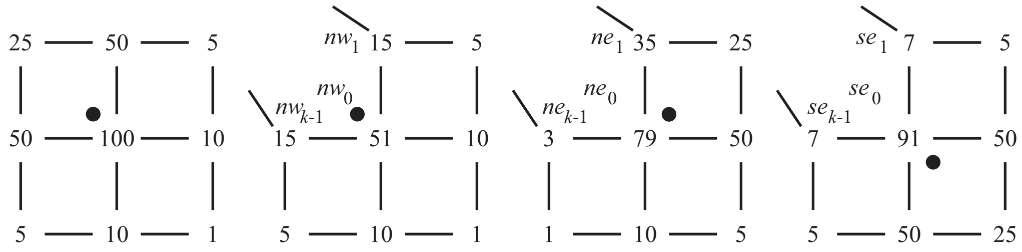

Fig. 6. Biquartic masks in the regular setting (left) and near an irregular face. In the latter case there are three symmetry cases whose weights are given as a sum of fixed coefficients and additional weights (nw, ne, and se) depending on the Doo-Sabin weights (see Equation 2).

vertex in the center which has been similarly raised. The increase in smoothness, decrease in height and increase in support as a function of the number of averaging steps is readily apparent. Figure 9 shows the result of applying degree two, three, four and nine schemes to a more complex shape (control mesh on the left).

3 Implementation

In this section we detail one possible implementation of quadrilateral subdivision schemes which is particularly simple and supports adaptive subdivision. The latter is important in implementations since 4-to-1 splits exhaust memory at much too fast a rate to be practical.

3.1 Mesh Representation

At the coarsest level the input control mesh is represented as a graph. Each input vertex of valence k has k children, each associated with one of the incident faces and carrying the root of a quadtree. These quadtree roots are arranged into a forest in a four connected fashion: each quadtree root having a well defined quadtree root neighbor in each of four directions

```
class QTreeR{
  QTreeR*   n[4];       // four neighbors
  Vector3D  dual;       // holds the initial vertex split
  QTree*    root;       // the actual tree
}
```

A quadtree is given as

```
class QTree{
  QTree*    p;          // parent
```

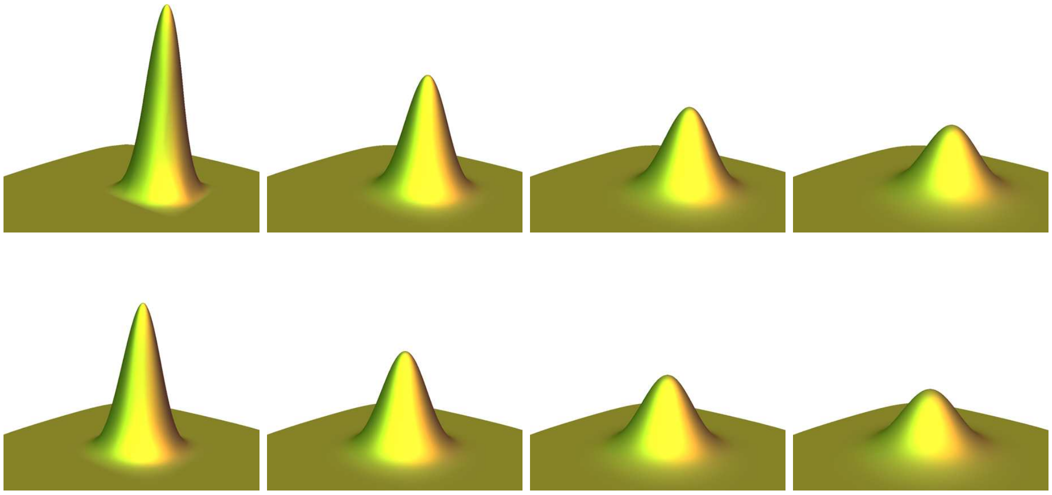

Fig. 7. Basis functions in the regular setting for dual schemes of degree two, four, six, and eight (top) and primal schemes of degree three, five, seven, and nine (bottom).

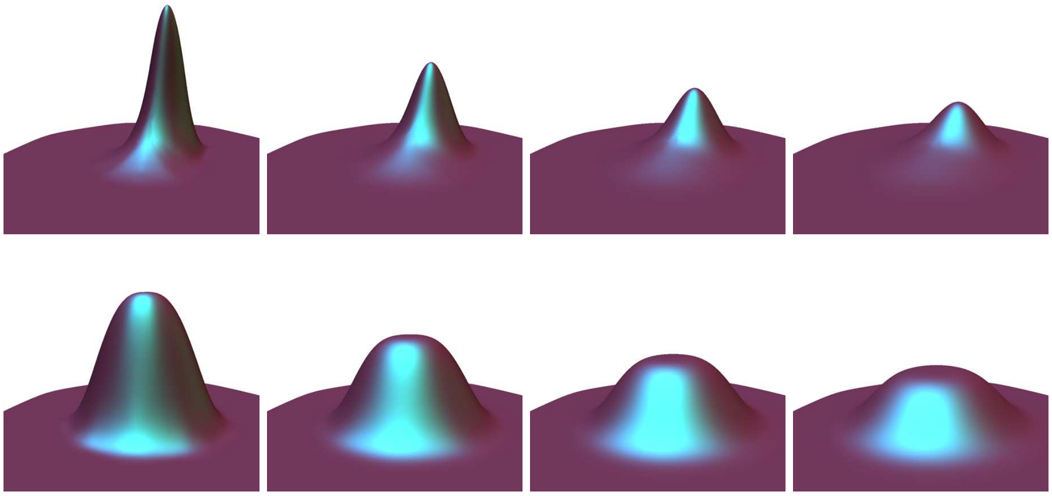

Fig. 8. Basis functions near a nine-gon for dual schemes of degree two, four, six, and eight (top) and at a vertex of valence nine for primal schemes of degree three, five, seven, and nine (bottom).

```
  QTree*     c[4];        // children
  Vector3D   dual;        // dual control point
  Vector3D*  primal[4];   // shared corners
}
```

The organization of quadtrees is depicted in Figure 10. Note that the primal vertices are actually shared among levels, carrying one control point per level reflecting the fact that in primal schemes old vertices receive new point positions at finer levels. This can be achieved by maintaining an appropriately sized stack of point positions at the corners. Both primal and dual subdivision can now be affected by iterating over all faces and repeatedly averaging to achieve the desired order of

Fig. 9. A more complex example showing the initial control mesh on the left followed by the results of applying subdivision of degree two, three, four, and nine.

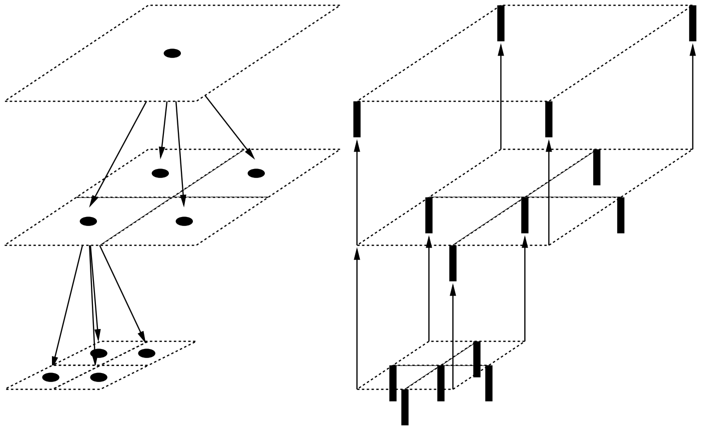

Fig. 10. Quadtrees carry dual control points (left). We may think of every quadtree element as describing a small rectangular piece of the limit surface centered at the associated control point (compare to Figure 2). The corners of those quads correspond to the location of primal control points (right), which are shared among levels.

subdivision [16,10].

For all schemes up through bidegree four one may apply subdivision rules in the more traditional setup by collecting the 1-ring of neighbors of a given control point (primal or dual). Collecting a 1-ring requires only the standard neighbor finding routines for quadtrees [19]. If the neighbor finding routine crosses from one quadtree to another the quadtree root links are used to affect this transition. Nil pointers indicate boundaries. With the 1-ring in hand one may apply stencils directly as indicated in Figure 6, for example. Using 1-rings and explicit subdivision masks can be useful for the proper treatment of boundaries and corners [20].

Boundaries are typically dealt with in primal schemes through special boundary rules. For example, in the case of Catmull-Clark one can ensure that the outermost row of control vertices describes an endpoint interpolating cubic spline (see, e.g., [20]). For dual schemes, for example Doo-Sabin, one can similarly run the associated univariate subdivision scheme on the boundary. Unfortunately these approaches do not in general lead to smooth schemes on the boundary near irregular vertices (respectively faces) [2] . However, in practice these defects are not noticeable and the simplified boundary treatment is the most straightforward to implement. It is the choice we made in our code. Figure 11 shows a pipe shape with boundaries done in this manner for degree two through nine.

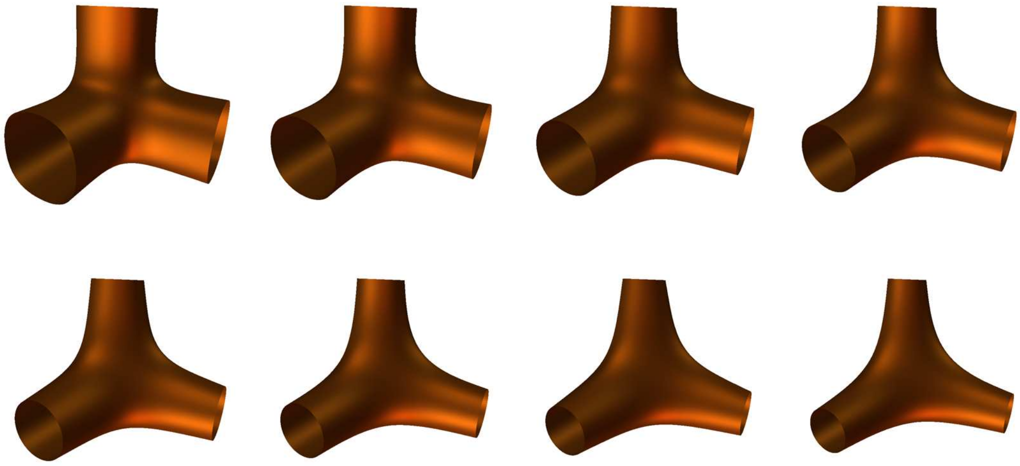

Fig. 11. A three legged pipe showing the effects of boundaries and increasing degrees (two through nine, left to right, top to bottom).

Adaptive Subdivision is mandatory for interactive applications to maintain high frame rates while escaping the exponential growth in the number of polygons with each subdivision step [8]. We first consider adaptive tessellations for primal quad schemes and then show how the same machinery applies to dual quad schemes.

Adaptivity requires two components controlling geometry and connectivity respectively. The former relies on some error measure, while the latter is concerned with managing dependencies between nodes at different levels of the hierarchy as well as their crack-free tessellation during output. We are concerned here only with the connectivity aspects. Data dependencies may be managed in an elegant and purely automatic fashion through lazy evaluation techniques. Crack-free tessellations require a restriction criterion which ensures that edge neighbors are off by no more than one subdivision level [19]. However, a slightly stronger criterion is more useful. Instead of enforcing the restriction criterion only across edges we ask that all neighbors of a given quadtree node be off by no more than one level in the hierarchy. This includes neighbors which share only a common corner, not just edges.

[2] One typical problem arises when the boundary curve tangent vector does not lie in the limit tangent plane as the irregular surface point is approached from the interior of the surface; see [20] for an example.

Such a vertex restriction criterion follows naturally from the traditional subdivision stencils whose support is a 1-ring. In order to create a new point at a finer level, all its parents' 1-ring neighbors must be present. This idea is illustrated in Figure 12 for primal schemes. Once a vertex restricted adaptive quadtree exists one must take

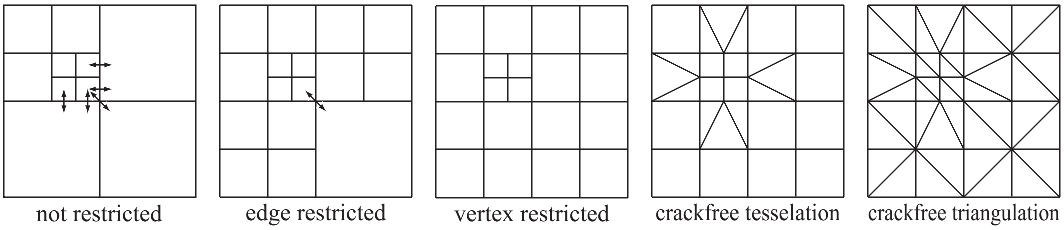

Fig. 12. On the left an unrestricted adaptive primal quadtree. Arrows indicate edge and vertex neighbors off by more than one level. Enforcing a standard edge restriction criterion results in some additional subdivision. A vertex restriction criterion also disallows vertex neighbors off by more than one level. Finally on the right some adaptive tessellations which are crack-free.

care to output quadrilaterals or triangles in such a way that no cracks appear. Since all rendering is done with triangles we focus on the crack-free triangulation case. This requires the insertion of diagonals in all quadrilaterals. One can make this choice randomly, but surfaces have a more visually pleasing appearance if this is done in a regular fashion. Figure 13 illustrates this on the left for a group of four children of a common parent. Here the diagonals are chosen to meet at the center. The resulting triangulation is exactly the basic element of a 4-8 tiling [10]. To deal with cracks we distinguish 16 cases. Given a leaf quadrilateral its edge neighbors may be subdivided once less, as much, or once more. Only the latter case gives rise to potential cracks from the point of view of the leaf quad. The 16 cases are easily distinguished by considering a bit flag for each edge indicating whether the edge neighbor is subdivided once more or not. Figure 13 shows the resulting templates (modulo symmetries). These are easily implemented as a lookup table. For

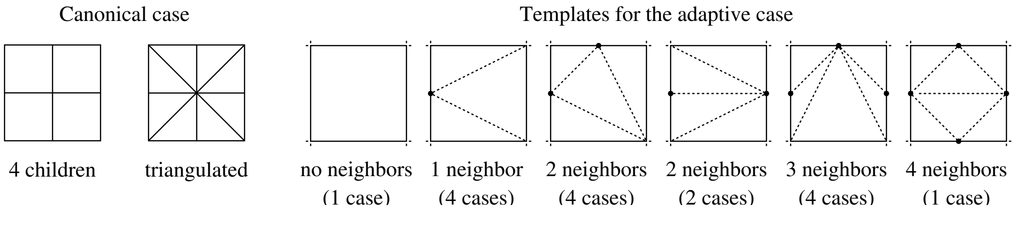

Fig. 13. On the left the canonical triangulation case for a group of four children of a common parent (primal). On the right are the 16 cases for adaptive triangulation of a leaf quadrilateral. Any one of the four edge neighbors may or may not be subdivided one level finer. Using the indicated templates (and their symmetries) one can reliably triangulate an adaptive primal quad tree with a simple lookup table.

dual quadrilateral subdivision schemes crack-free adaptive tessellations are considerably harder to generate. Recall that in a dual quad scheme a quadtree node represents a control point, not a face. It potentially connects to all eight neighbors (see Figure 14, top left). Consequently there are 256 possible tessellations depending on eight neighbor states. To avoid this explosion of cases we instead choose

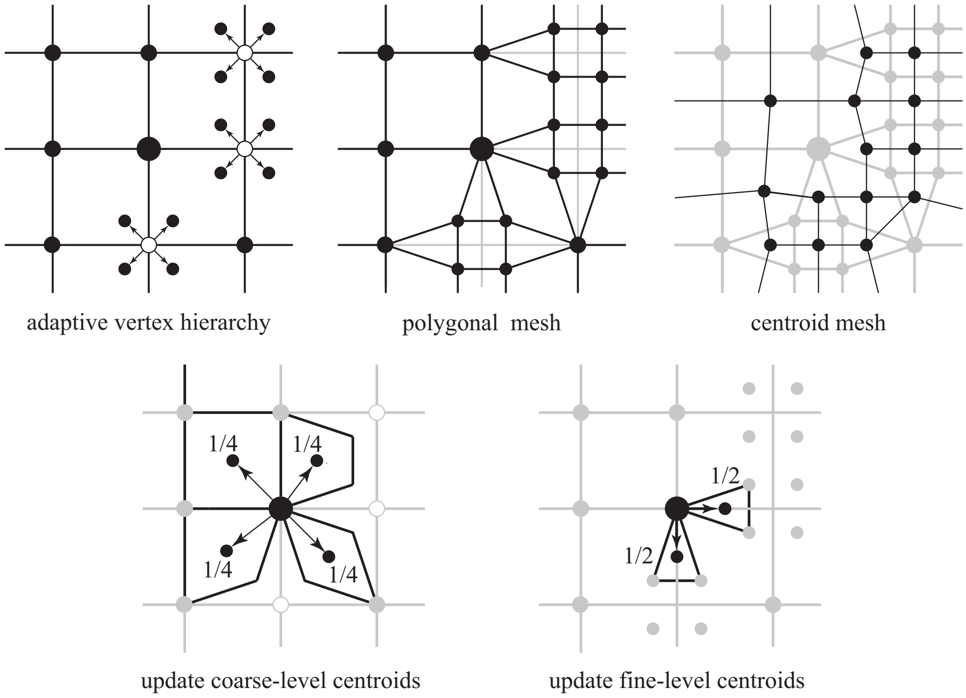

Fig. 14. For dual quad schemes the canonical case is shown on the top left. An adaptive leaf needs to know the state of eight neighbors to affect an adaptive crack-free tessellation (second in top row). Instead we draw a centroid (primal) mesh. It uses control point averages at corners (left bottom) and averages on edges (right bottom). Which is used depends on the neighbor state, i.e., is the neighbor subdivided once more or not.

to output a tessellation of the centroids of the dual control points, i.e., we perform one final averaging step immediately before outputting what is then a primal mesh. Note that this additional averaging step is only performed during output and does not change the overall scheme. Figure 14 shows how to form the additional averages in an adaptive setting. With these drawing averages computed we apply the templates of Figure 13 to render the output mesh. Figure 15 shows an example of such an adaptively rendered mesh.

4 Analysis of Repeated Averaging Schemes

In this section we describe how techniques developed in [21] can be used to establish several facts about averaging subdivision schemes. In particular, we demonstrate that all schemes up to and including nine averaging steps are C^1 but not C^2 continuous. Using our tools we examine the behavior of eigenvalues and eigenvectors and make several important observations about behavior of the schemes near irregular vertices, respectively faces. One of the somewhat surprising conclusions

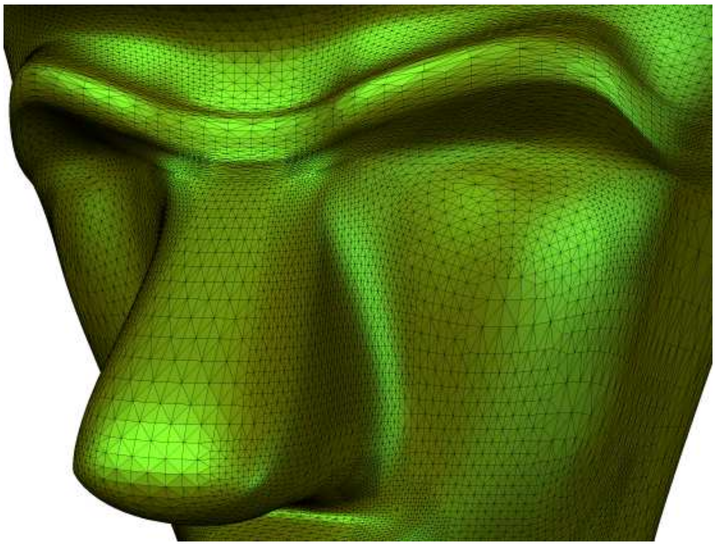

Fig. 15. Detail of an adaptive dual subdivision surface and its crack-free triangulation. Observe the transition between different levels of refinement in highly curved area under the nose and the flat area on the cheek.

is that subdivision schemes of higher order behave in ways inferior to the behavior of low order schemes near irregular vertices, respectively faces, of high valence. We discuss potential reasons for this, and state a conjecture regarding the limit behavior of these schemes as the number of averaging steps approaches infinity.

An important feature of our derivation is an improved way to reduce the subdivision matrix to block-diagonal form. In addition to rotation invariance usually taken into account, we take advantage of mirror symmetries to reduce the diagonal blocks to pure real form, simplifying numerical analysis of the eigenstructure of the matrices.

4.1 Definitions

We start with introducing the basic notation for subdivision matrices. If we consider a primal scheme, we call a mesh k-regular, if it is a quadrangulation of a plane with a single irregular vertex of valence k. For dual schemes a k-regular mesh is the tiling of the plane with a single k-gon and all other faces quadrilateral. For primal schemes, the m-neighborhood of the center consists of m layers of vertices around the irregular vertex, with layers shown in Figure 16. For dual schemes, the m-neighborhood consists of m layers of vertices around the central k-gon, counting the k-gon itself as the first layer.

We call an m-neighborhood *invariant* if all control points in this neighborhood on level $j + 1$ can be computed knowing only the control points on level j. For any subdivision scheme with finitely supported masks there is a *minimal invariant neighborhood*, i.e., an invariant neighborhood of a minimal size m. It is easy to see that this invariant neighborhood has size N for dual schemes using $2N$ and primal schemes using $2N + 1$ averaging steps. Indeed, the first vertex split step doubles the size of the neighborhood to $2m$. Given all control points in the dual l-neighborhood we can compute all points in the primal $l - 1$-neighborhood formed by the centers of faces; one more dualization step allows us to compute all points in a dual $l - 1$-neighborhood. Thus, for each two averaging steps the size of the neighborhood over which points can be computed shrinks by one. After N double averaging steps, the neighborhood shrinks to size $2m - N$. For the neighborhood to be minimally invariant, this size should be exactly N, yielding $m = N$. A similar argument holds for primal schemes.

Each m-neighborhood of a k-regular vertex can be separated into k *sectors*. Each sector contains m^2 vertices for dual k-regular meshes and $m(m + 1)$ vertices for primal k-regular meshes.

Subdivision Matrices. For stationary schemes, the relations between control points on different levels are linear. For an invariant neighborhood we consider a square matrix which maps the vector p^j of control points on level j to the vector of control points p^{j+1} on level $j + 1$. We call such matrices *subdivision matrices*. Two subdivision matrices are of importance to us: the minimal subdivision matrix, acting on the minimal invariant neighborhood of size N, and the subdivision matrix acting on the neighborhood of size $N + 1$. We call the latter the *complete subdivision matrix*, or subdivision matrix for short.

The importance of the complete subdivision matrix arises from its role in the smoothness analysis of subdivision schemes. The latter requires the examination of the characteristic map, initially introduced in [22]. We use a slightly different definition of the characteristic map which can be found in [23]. The minimal subdivision matrix, which is a submatrix of the complete matrix is primarily used to compute subdominant eigenvalues.

Characteristic Map. Assume that the largest eigenvalues of the complete subdivision matrix are λ_0, λ_1, λ_2, and λ_3. Furthermore, assume that the largest eigenvalue $\lambda_0 = 1$, and all these eigenvalues have trivial Jordan blocks. Finally, assume that λ_1 and λ_2 are real, and $1 > |\lambda_1| \geq |\lambda_2| > |\lambda_3|$. The characteristic map in this case is defined using the eigenvectors x_1 and x_2 of eigenvalues λ_1 and λ_2. Note that each eigenvector component corresponds to a vertex in the invariant neighborhood. To each vertex v of the invariant neighborhood assign the pair of the corresponding components of the eigenvectors $(x_1(v), x_2(v))$ as a 2D control point. It is easy to see that given a neighborhood of size $N + 1$ on level 0 we can compute all control points in a $2^j + N$-neighborhood on level j. In the limit we get a dense set of control points for a patch surrounding the center of the k-regular mesh. Thus, the control mesh formed from the components of the eigenvectors can be used to define a map from a regular k-gon into the plane. This map is called the *characteristic map*. A more general definition eliminating limitations on the eigenvalues can be found in [23]. The more limited definition given here is sufficient for the schemes that we consider. In [22] it was shown that for a scheme to be C^1-continuous it is sufficient for the characteristic map to be regular and injective.

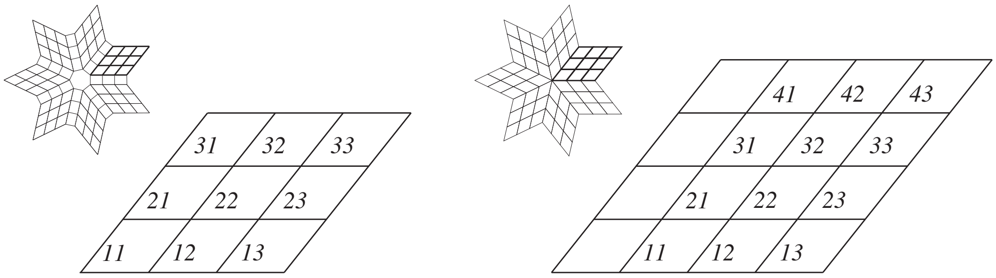

Fig. 16. The indices of the vertices in a sector of the control mesh for the characteristic map. On the left the numbering for dual schemes and on the right the numbering for primal schemes.

4.2 Subdivision Matrices

The subdivision matrices we need to consider are of size $k(N+1)^2$ for dual schemes and $1 + k(N+1)(N+2)$ for primal schemes. To reduce the complexity of the analysis of subdivision matrices, we reduce them to block-diagonal form, using well-known DFT techniques. In addition, we take advantage of the mirror symmetries present in coefficients of schemes invariant with respect to automorphisms of meshes, converting the blocks on the diagonal to purely real matrices, depending on a single parameter $c = \cos(2\pi s/k)$, $s = 0 \ldots k-1$ (s indexes the sector). The transformed matrices are block-diagonal, with k blocks $B(c)$, $c = \cos(2\pi s/k)$, $s = 0 \ldots k-1$ on the diagonal. These blocks have size $(N+1)^2 \times (N+1)^2$ for dual schemes. For primal schemes, the block $B(0)$ has size $(1+(N+1)(N+2)) \times (1+(N+1)(N+2))$ and the rest $(N+1)(N+2) \times (N+1)(N+2)$.

Transforming Subdivision Matrices. To define the transformations of the subdivision matrix, we use the numbering shown in Figure 16. Each control point is indexed by three indices: the first one is the sector number, while the other two define the position of the point inside the sector. In the following we describe our derivation for the dual case; the primal case is similar with changes in numbering required to account for the somewhat different structure of each sector.

Suppose subdivision formulas for control points p^s_{ij}, $s = 0 \ldots k-1$, $i, j = 1 \ldots m$

are given by

$$p'_{sij} = \sum_{t,l,m} a^{tlm}_{sij} p_{tlm}$$

where a^{tlm}_{sij} are the coefficients of subdivision, i.e., the entries in the subdivision matrix, and p'_{sij} are the control points of the subdivided mesh. Rotational symmetry yields relations $a^{tlm}_{sij} = a^{0lm}_{s-t\,ij}$ for any t, s, which makes it possible to reduce the matrix to block-diagonal form through the application of a DFT obtaining

$$p'_{ij}(\omega) = \sum_{l,m} a^{lm}_{ij}(\omega) p_{lm}(\omega)$$

with ω one of the complex k-th roots of unity. Mirror symmetry results in the relations $a^{t\,lm}_{s\,ij} = a^{-t\,ml}_{-s\,ij}$. In the Fourier domain, the mirror symmetry conditions result in $a^{lm}_{ij}(\omega) = a^{ml}_{ji}(\overline{\omega}) = \overline{a^{ml}_{ji}(\omega)}$. From now on, we drop the dependence on ω to simplify the formulas.

We split the control points into three groups: (a) diagonal p_{ii}; (b) upper triangular p_{ij}, $i < j$; and (c) lower triangular p_{ij}, $j < i$. Now the subdivision equations become

$$\begin{aligned}
p'_{ij} &= \sum_{l<m} \left(a^{lm}_{ij} p_{lm} + a^{ml}_{ij} p_{ml} \right) + \sum_{l} a^{ll}_{ij} p_{ll} \\
p'_{ji} &= \sum_{l<m} \left(a^{lm}_{ji} p_{lm} + a^{ml}_{ji} p_{ml} \right) + \sum_{l} a^{ll}_{ij} p_{ll} \\
p'_{ii} &= \sum_{l<m} \left(a^{lm}_{ii} p_{lm} + a^{ml}_{ii} p_{ml} \right) + \sum_{l} a^{ll}_{ii} p_{ll}
\end{aligned}$$

where $i < j$. We introduce new variables: $q_{ij} = (1/2)(p_{ij} + p_{ji})$, $q_{ji} = (1/2I)(p_{ij} - p_{ji})$, for $i < j$, and $q_{ii} = p_{ii}$, where $I = \sqrt{-1}$. We call these simple transformations *symmetrization transforms*. The subdivision equations for the new variables are given as

$$\begin{aligned}
q'_{ij} &= \sum_{l<m} \left(c^{lm}_{ij} q_{lm} + c^{ml}_{ij} q_{ml} \right) + \sum_{l} c^{ll}_{ij} q_{ll} \\
q'_{ji} &= \sum_{l<m} \left(c^{lm}_{ji} p_{lm} + c^{ml}_{ji} q_{ml} \right) + \sum_{l} c^{ll}_{ji} q_{ll} \\
q'_{ii} &= \sum_{l<m} \left(c^{lm}_{ii} q_{lm} + c^{ml}_{ii} q_{ml} \right) + \sum_{l} c^{ll}_{ii} q_{ll}
\end{aligned}$$

with coefficients

$$c^{lm}_{ij} = \frac{1}{2} \left(a^{lm}_{ij} + a^{ml}_{ij} + a^{lm}_{ji} + a^{ml}_{ji} \right)$$

$$
\begin{aligned}
c_{ij}^{ml} &= \frac{I}{2}\left(a_{ij}^{lm} - a_{ij}^{ml} + a_{ji}^{lm} - a_{ji}^{ml}\right)\\
c_{ji}^{lm} &= \frac{-I}{2}\left(a_{ij}^{lm} + a_{ij}^{ml} - a_{ji}^{lm} - a_{ji}^{ml}\right)\\
c_{ji}^{ml} &= \frac{1}{2}\left(a_{ij}^{lm} - a_{ji}^{ml} - a_{ji}^{lm} + a_{ji}^{ml}\right)\\
c_{ij}^{ll} &= \frac{1}{2}\left(a_{ij}^{ll} + a_{ji}^{ll}\right)\\
c_{ji}^{ll} &= \frac{1}{2}\left(a_{ij}^{ll} - a_{ji}^{ll}\right)\\
c_{ii}^{lm} &= \frac{1}{2}\left(a_{ii}^{lm} + a_{ii}^{ml}\right)\\
c_{ii}^{lm} &= \frac{I}{2}\left(a_{ii}^{lm} - a_{ii}^{ml}\right)\\
c_{ii}^{ll} &= a_{ii}^{ll}
\end{aligned}
$$

where $i < j$. Using the mirror symmetry condition $a_{ij}^{lm} = \overline{a_{ji}^{ml}}$ we see that the above coefficients are equal to their conjugates, and thus real.

Observe that the coefficients c_{ij}^{lm} and c_{ji}^{lm} are just the real parts of $a_{ij}^{lm} \pm a_{ij}^{ml}$. Since the c_{ij}^{lm} and c_{ji}^{ml} are Fourier transforms of finite real sequences, their real parts can be regarded as polynomial functions of the variable $c = \cos(2\pi s/k)$. Similarly, the coefficients c_{ji}^{lm} and c_{ji}^{ml} are imaginary parts of the Fourier transforms of real sequences, and can be written in the form $\sin(2\pi s/k)F(\cos(2\pi s/k))$, where F is some polynomial. The same is true for the other coefficients. A final rescaling of q_{ji} by $\sin(2\pi s/k)$ reduces all coefficients to polynomial functions of $c = \cos(2\pi s/k)$.

For example, for the biquartic scheme the resulting transformed matrix is

$$B(c) = \frac{1}{256} \begin{pmatrix} 8\,c^2 + 58\,c + 75 & 20 + 10\,c & 1 & -10 + 10\,c^2 & 0 & 0 & 0 & 0 & 0 \\ 4\,c^2 + 54\,c + 95 & 60 + 26\,c & 5 & -26 + 26\,c^2 & 0 & 0 & 0 & 0 & 0 \\ 18\,c + 99 & 100 + 10\,c & 25 & -10 + 10\,c^2 & 0 & 0 & 0 & 0 & 0 \\ -4\,c - 40 & -24 & 0 & -24\,c + 40 & 0 & 0 & 0 & 0 & 0 \\ 25\,c + 50 & 105 + 50\,c & 10 & 50\,c^2 - 50 & 5\,c + 10 & 1 & 0 & 0 & -5 + 5\,c^2 \\ 5\,c + 50 & 125 + 10\,c & 50 & -10 + 10\,c^2 & c + 10 & 5 & 0 & 0 & c^2 - 1 \\ 25 & 100 & 100 & 0 & 10 & 20 & 1 & 0 & 0 \\ -5 & -10 & 0 & 75 - 10\,c & -1 & 0 & 0 & 5 & 10 - c \\ -25 & -50 & 0 & 95 - 50\,c & -5 & 0 & 0 & 1 & 10 - 5\,c \end{pmatrix}$$

The complete transformed subdivision matrix is block diagonal, with blocks on the diagonal each given as $B(c)$ for $c = \cos(2\pi s/k)$ for $s = 0 \ldots k-1$.

Subdivision Matrices for Averaging Schemes. The above considerations apply to any subdivision scheme with symmetric rules (rules invariant with respect to rotations and mirror reflections of k-regular meshes). Now we consider how the subdivision matrix for averaging schemes can be computed. The crucial observation is that the operation of taking the DFT transform commutes with subdivision; this means that we can iteratively construct complete subdivision rules applying the averaging steps directly to the DFT form of the control points $p_{ij}(\omega)$. More precisely, whenever all four points we average are from the same sector, the Fourier domain averaging takes the form $(1/4)\,(p_{ij}(\omega) + p_{i-1\,j}(\omega) + p_{i\,j+1}(\omega) + p_{i-1\,j}(\omega))$. In the case when we need to use control points from an adjacent sector, we substitute the Fourier transform version. For example, instead of $p_{s-1\,ij}$ we use $\omega p_{ij}(\omega)$. This yields two simple linear matrix operations acting on the two-dimensional arrays of Fourier domain control points $p_{ij}(\omega)$: one for going from dual to primal (Figure 4, left) and one for going from primal to dual (Figure 4, right). Iterative application of these operations to an initial array obtained by replication of all values in symbolic form automatically yields subdivision rules for a scheme of any order. Once the rules are obtained, the subdivision matrix is extracted as the matrix of coefficients of the rules. Finally, symmetrization transforms are applied to obtain real subdivision matrices. We have performed all these operations using the symbolic algebra system Maple V.

The analysis for all averaging schemes is similar to the analysis of the eigenvalues of the subdivision matrix for Kobbelt's quadrilateral scheme presented in [21]. The transformations described in the previous section reduce the analysis of a subdivision matrix to analysis of a family of smaller matrices $B(c)$ parameterized by $c = \cos(2\pi s/k)$. The characteristic polynomial of $B(c)$ can be regarded as a polynomial $P(\lambda, c)$ in two variables. The eigenvalue plots for the first eight nontrivial schemes are shown in Figure 17. From the plot it is clear that the largest eigenvalue

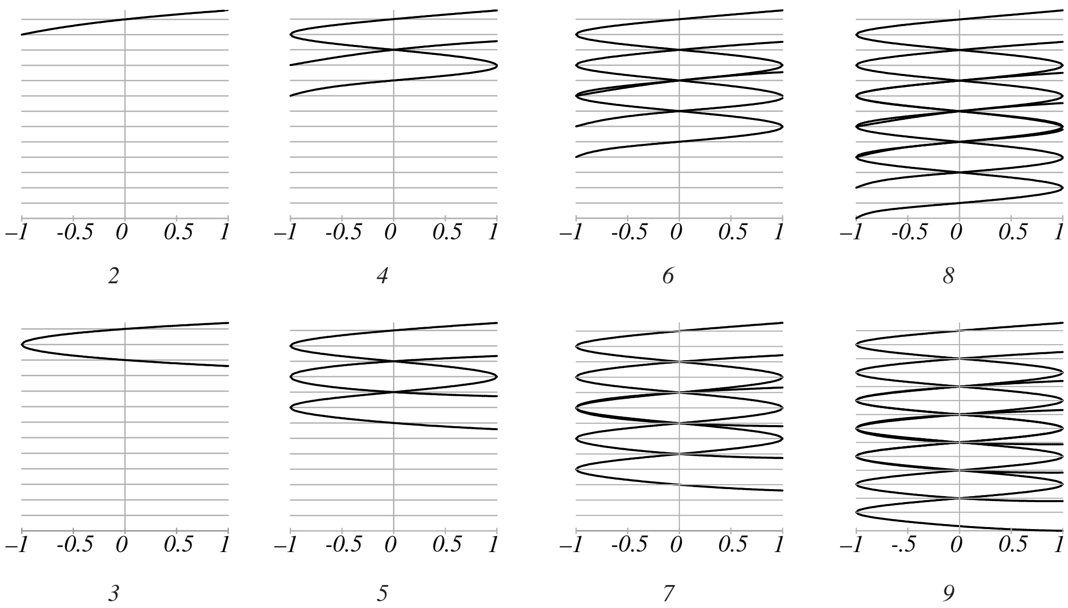

Fig. 17. The magnitudes of the eigenvalues of the diagonal blocks of the transformed subdivision matrices for the averaging schemes of degree two through nine as functions of $c = \cos 2\pi/k$, where k is the valence. Rather than showing discrete points for integer values of k, we show the continuous dependence on c. The plots have logarithmic scale vertically, each horizontal coordinate line corresponds to 2^{-j}, $j = 1 \ldots 14$.

increases as a function of c. Using interval arithmetic methods to get reliable upper and lower bounds for the eigenvalues, we prove the following proposition:

Proposition 1 *For all eight schemes, any valence k, and any $m = 1 \ldots k-1$ the largest eigenvalue is real and unique, and for any subdivision matrix block $B(cos(2m\pi/k))$, $m \neq k-1, 1$ the largest eigenvalue is less than the largest eigenvalue of the blocks $B(2\pi/k)$ and $B(2\pi(k-1)/k)$. The unique largest eigenvalue is the only eigenvalue in the interval $[0.5, 1]$, for $k > 4$.*

The detailed proof with all calculations can be found in a Maple worksheet available from the authors. Here we present an outline of the proof. The proof is performed in several steps:

(1) We show that for $c < 0$, all roots of the characteristic polynomial $P(c, \lambda)$ are less than 0.51.
(2) We show that for any $c \in [0 \ldots 1]$, there is a unique real root $\lambda = \mu$ of $P(\lambda, c)$ in the interval $[0.48 + 0.23c, 0.54 + 0.23c]$ (dual schemes) and $[0.485 + 0.2c, 0.55 + 0.22c]$ (primal schemes), the function $\mu(c)$ is C^1-continuous and increases.
(3) We "deflate" the characteristic polynomial (i.e., divide by $\lambda - \mu$) and verify that all roots of the deflated polynomial are inside the circle of radius 0.5 for $c \in [0, 1]$.

We use the Marden-Jury test [24] to show that the roots of a polynomial are inside a circle of radius 0.51 and 0.5 in the complex plane in steps 1 and 3 respectively. This test requires only a simple algebraic calculation on the coefficients of the polynomial and can be performed easily for symbolic and interval coefficients. See [21] for details regarding the application of this test.

Proposition 1 allows us to compute the subdominant eigenvalue with arbitrary precision for any k and establishes that it always corresponds to the blocks with $m = 1$ and $m = k - 1$. Figure 18 shows comparison of the subdominant eigenvalue behavior for averaging schemes, as the conjectured limit function discussed below. We observe that as the number of averaging steps increases, the subdominant eigenvalues also increase for large valences. This suggests that for large valences convergence of subdivision to the limit near irregular vertices (respectively faces) of large valence will be slow, and the quadrilateral size near irregular vertices will be extremely nonuniform. This is confirmed by experiments (Figure 19).

4.4 Analysis of the Characteristic Maps

Instead of proving that a characteristic map is injective, it is sufficient to verify that it maps only zero to zero, and has index 1. More precisely, the following theorem was proved in [21]:

Theorem 2 *Suppose a characteristic map Φ satisfies the following conditions:*

(1) the pre-image $\Phi^{-1}(0)$ contains only one element, 0;
(2) the characteristic map has a Jacobian of constant sign everywhere on $\mathbb{R}^2$ except zero.

Then the extension of the characteristic map is a surjection and a covering away from 0. In particular, if the winding number with respect to the origin of the image $\Phi(\gamma)$ of a simple curve is 1, the characteristic map is injective and the scheme is C^1-continuous.

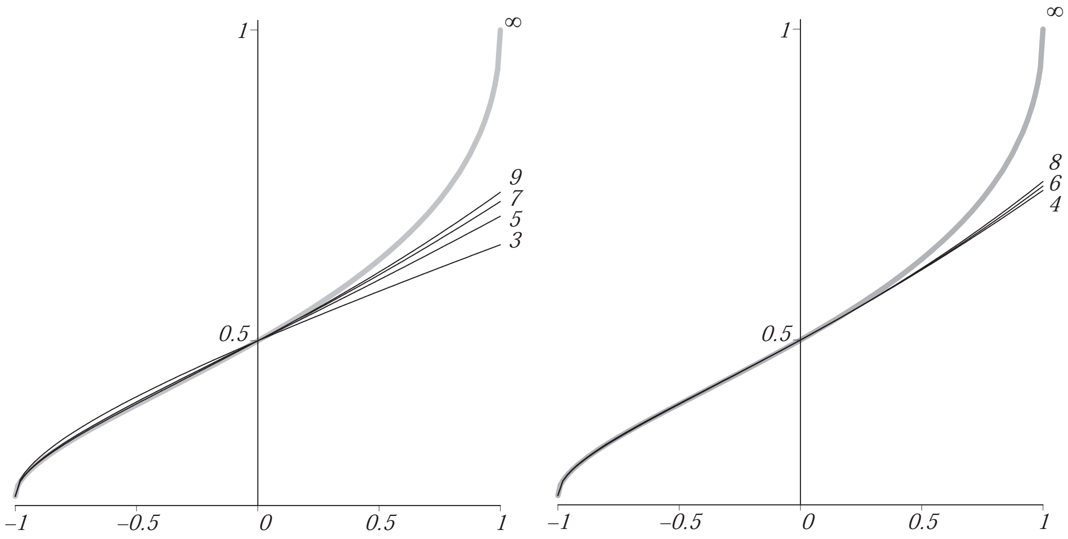

Fig. 18. Subdominant eigenvalues for primal (left) and dual (right) averaging schemes and the conjectured limit curve (gray) for the number of averaging steps increasing to infinity. While the eigenvalues for large valences appear to increase as expected, the data is not sufficient to strongly confirm our hypothesis.

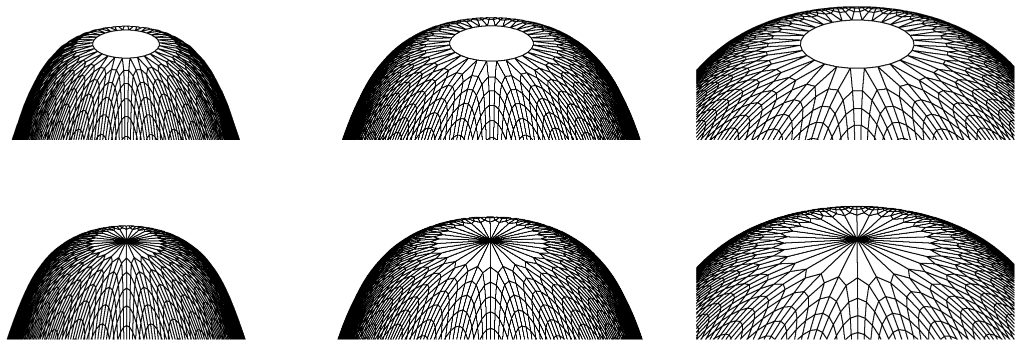

Fig. 19. Behavior of dual schemes of degree 2, 10, and 30 (top row), and primal schemes of degree 3, 11, and 31 (bottom row) near an irregular face (respectively vertex) of valence 30.

For averaging schemes the characteristic map can theoretically be computed in closed form. However, this is relatively difficult even for bicubic schemes and becomes impractical for higher order schemes. Instead we use general tools that evaluate a sufficiently close approximation of the Jacobian of the characteristic map to verify regularity. Similarly, rather than computing the image of a curve enclosing zero under the characteristic map in explicit form, we use a sufficiently close piecewise linear approximation.

The characteristic map Φ satisfies the scaling relation $\Phi(y/2) = \mu\Phi(y)$, where μ is the subdominant eigenvalue and y is a point in the k-gonal planar domain of the characteristic map. To establish regularity of the map, it is sufficient to examine its behavior on a ring around an irregular vertex, such that under the iterative application of the scaling transformation $t \to t/2$, the copies of the ring cover 1-neighborhood of the center. Due to the rotational symmetry of the characteristic map, only one segment needs to be examined; for the biquartic scheme, for example, we end up with the control mesh shown in Figure 20(c), which is defined on a subset of a regular quadrilateral grid. The size of this control mesh does not depend on the valence. For all other schemes, control meshes for several valences are

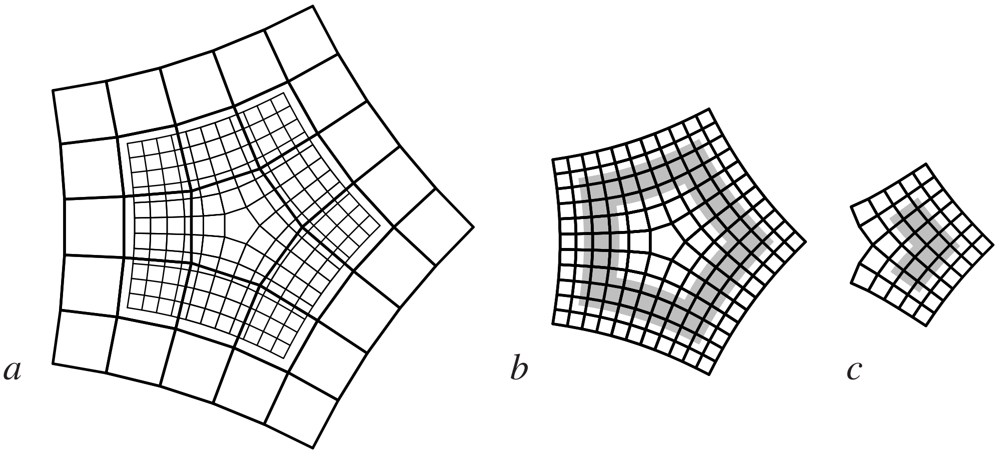

Fig. 20. Examples of characteristic map control meshes for the biquartic scheme. (a) The control mesh for the characteristic map for valence five. The control mesh after two subdivision steps is shown with thin lines. (b) The control mesh for a single ring of the characteristic map (the limit patches are shown in gray). (c) The control mesh for one sector of the ring.

shown in Figure 21.

Let L^∞ be the limit function of subdivision, L^m be its approximation after m steps, and let p^m be the vector of control points of a subdivision surface. Then the following estimate holds [25,21]:

$$\|L^\infty - L^m\|_\infty \le \frac{c}{1-\gamma} D\left(p^m\right), \tag{3}$$

where D is the contraction function, which we choose to be $\|\nabla p\|_\infty$, with ∇ being the vector of directional differences. The constants γ and c can be computed from the coefficients of subdivision.

For a convergent subdivision scheme S, one can always find a *matrix* difference subdivision scheme S' (i.e., a scheme whose coefficients are matrices) acting on the vectors of differences at each vertex, such that $\nabla S p = S' \nabla p$. If the scheme is C^1 on regular grids, then an estimate similar to Equation 3 holds for the difference scheme, with different constants γ^D and c^D used instead of c and γ.

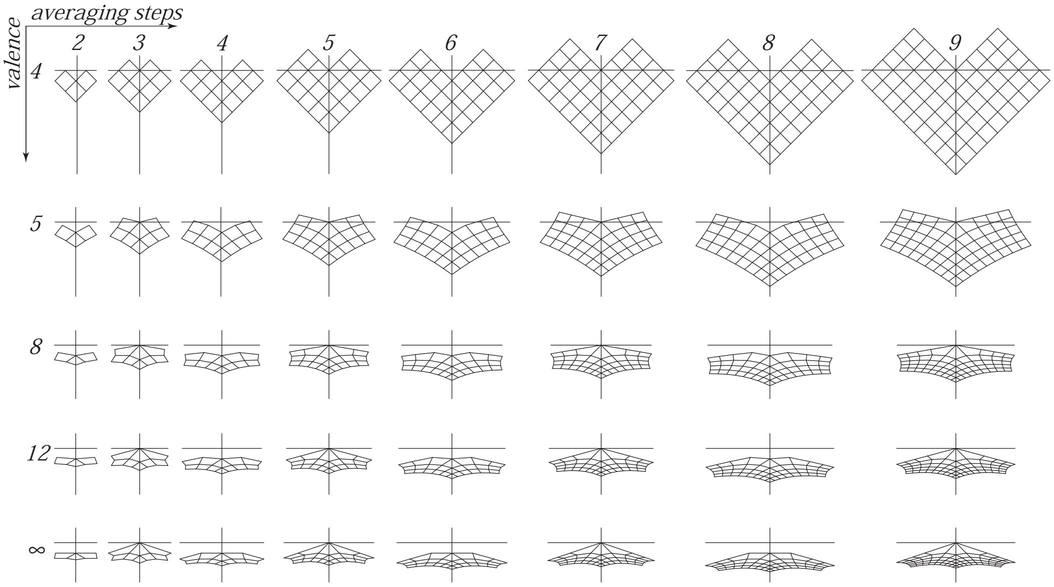

Fig. 21. Control meshes for a single segment of the characteristic map for averaging schemes, shown as a function of valence and numbers of subdivision steps. The scale is chosen so that the sector angle is $\pi/4$ and control points in the 1-neighborhood are at a fixed distance from the center.

These estimates can be used to compute approximations to the characteristic map and its derivatives with guaranteed error bounds; if the lower estimate for the Jacobian of the characteristic map obtained in this way is positive (or upper estimate negative), the map is guaranteed to be regular. For the schemes that we consider $\gamma = \gamma^D = 1/2$, c varies in the range from $1/2$ to $35/32$, and c^D varies in the range from $1/2$ to $75/64$. The behavior of upper and lower bounds for Jacobians is shown in Figure 22.

C^2**-Continuity.** The same analysis that is used to establish C^1-continuity, with small modifications, can be used to show that the schemes are *not* C^2-continuous at the irregular points. Indeed, all eigenvalues of the subdivision matrix are eigenvalues of the blocks $B(c)$; we have identified the largest eigenvalue for every block number m and every valence k such that $c = \cos(2\pi m/k) > 0$, and demonstrated that the corresponding eigenvector produces a nondegenerate limit function. We have also established that $\mu(c)$, the largest eigenvalue of $B(c)$ increases as function of c. Therefore, the third largest eigenvalue is the eigenvalue of block $B(\cos(4\pi/k))$. Clearly, as k increases, the difference between $\mu(\cos(2\pi/k))$ and $\mu(\cos(4\pi/k))$ becomes negligible. Formally, we verify that the third largest eigenvalue is larger than a square of the subdominant eigenvalue using interval arithmetic and our eigenvalue estimates. Analysis identical to C^1-continuity analysis guarantees that the corresponding eigenbasis function is not degenerate. This contradicts necessary conditions for C^2-continuity [23].

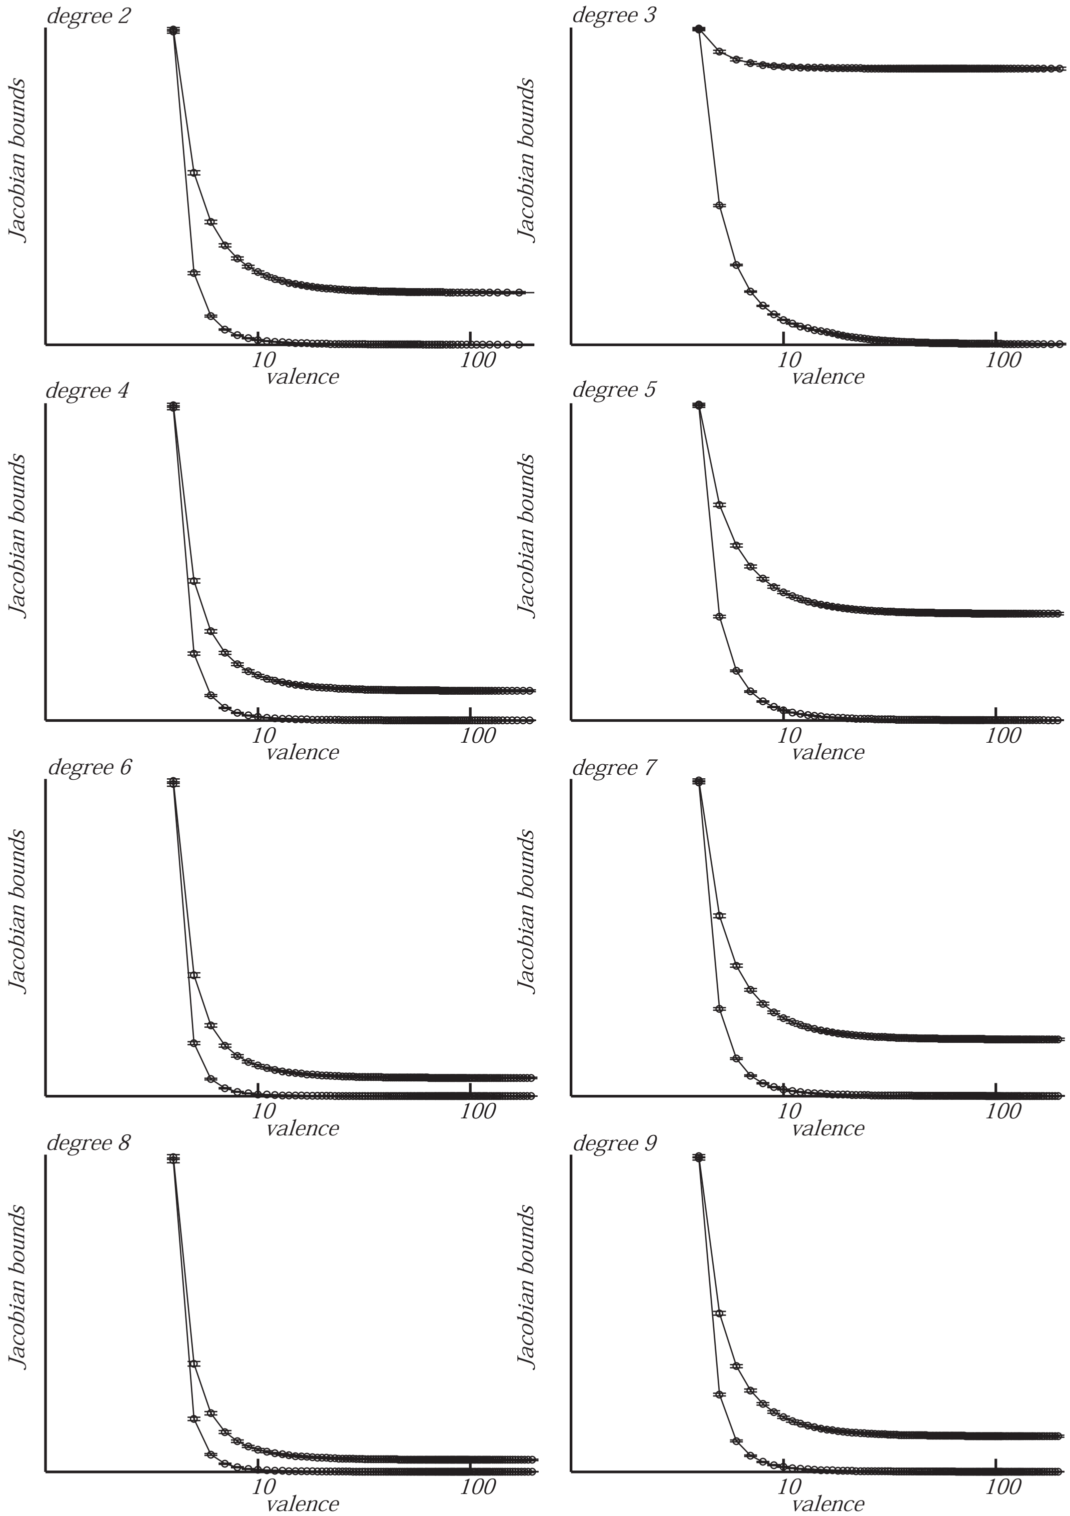

Fig. 22. The upper and lower bounds for the Jacobian (upper and lower curve in each chart) of the characteristic maps as functions of the valence for averaging subdivision schemes of bidegree two through nine (top left to bottom right).

Behavior at Infinity. Our proof of C^1-continuity for high valences is based on the following important fact: the control configuration for a single sector of the characteristic map obtained from the transformed subdivision matrix depends only on $c = \cos(2\pi/k)$, and approaches a non-degenerate limit configuration as $k \to \infty$. Figure 23 shows an example for the biquartic scheme. We are using interval arithmetic to examine approximations to characteristic maps. Suppose the control mesh was computed using the interval $[1-\epsilon, 1]$ for c. If we verify that the Jacobian has constant sign for this control mesh with interval control points, we have verified this fact simultaneously for all control meshes for which $c \in [1-\epsilon, 1]$.

To complete the analysis of the scheme we need to describe the behavior of $\mu(c)$ at infinity. Specifically, to use our algorithm for verification of smoothness for all valences, we need to estimate the corresponding interval value $\mu(c)$ for an interval $c = [1-\epsilon, 1]$, so that the eigenvectors can be computed. As $\mu(c)$ changes slowly, linear approximation is sufficient for our purposes. The upper bound for the derivative μ'_c at $c = 1$ can be computed easily if we regard the characteristic polynomial as a function $F(\mu, c)$ of two variables μ and c, and estimate $\mu'(c)$ using the ratio of the components of the gradient of this function. The upper bound for $\mu'(c)$ in the region of interest is 0.5 for all schemes that we have considered.

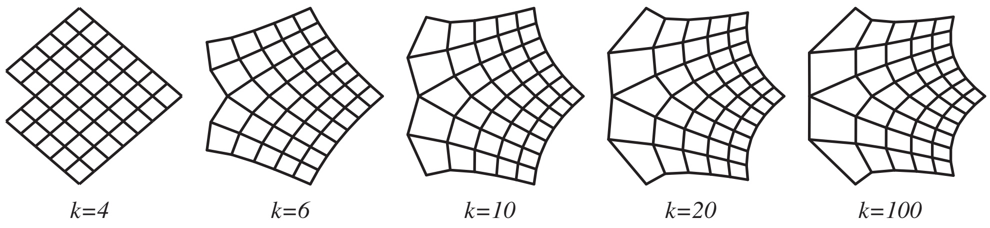

Fig. 23. As $k \to \infty$, $c = \cos(2\pi/k) \to 1$, and rescaled control meshes for a segment of the characteristic map converge to a limit.

Summary of Analysis. All eight schemes that we have analyzed are C^1 at irregular vertices, respectively faces, for all valences. However, these schemes are not C^2 for valences greater than 4. As the number of smoothing steps increases, the eigenvalues for large valences also increase. In addition, the characteristic map become increasingly nonuniform, as can be seen from the behavior of Jacobians (Figure 22) and control meshes (Figure 21).

4.5 Behavior for Large Number of Averaging Steps

While we did not rigorously analyze the behavior of schemes for very large numbers of averaging steps, one can extrapolate from the observed behavior. Using some informal arguments we can state a conjecture regarding the limit of the characteristic maps. These considerations are far from complete.

Consider the sequence of primal schemes obtained by applying double barycenter averaging N times, for $N = 1 \ldots$ acting on a k-regular mesh. If we apply an infinite number of averaging steps at each subdivision step, and at each step the process converges, we get the same configuration up to a scale factor. Denote by S^∞ a single step of the smoothing operator. Suppose we start with an arbitrary initial configuration x, and the sequence $(S^\infty)^j\, x$ converges to x_0. Then $S^\infty x_0 = x_0$, and applying S^∞ an infinite number of times can be interpreted as solving the infinite linear system $(I - S^\infty)x_0 = 0$. Note that for all vertices excluding the center point and its neighbors, and vertices on k rays extending the edges incident on the central vertex, the system has a very simple form and is a discretization of the Laplace equation $\Delta f = 0$ for uniform grids on the plane. Our grid is an affine transformation of the regular grid only in the interior of each sector; across sector boundaries, the smoothing operation is no longer the discretization of the Laplacian.

However, we can reparameterize the k-regular mesh using a parameterization which has the following property: as we move further from zero, our smoothing operator converges to the discretization of the Laplacian. One can demonstrate that this is achieved by a map described in [23], which, once restricted to a single sector and adjusted for quadrilateral schemes, is a composition of an affine transform map h_k mapping a sector with angle $2\pi/k$ to a sector with angle $\pi/2$ and the conformal mapping $z \to z^{4/k}$, mapping it back to the sector with angle $2\pi/k$.

Once the reparameterization is done, the reparameterized limit characteristic map is likely to be a solution of Laplace's equation almost everywhere, excluding a single point. Given the rotational symmetries that we require, the most natural (but not unique) candidate is the identity. Thus, we come to the conclusion that on a single sector the limit characteristic map composed with the reparameterization and the identity is $(\cdot)^{4/k} \circ h_k$. This immediately suggests the limit behavior of eigenvalues: for our schemes, the characteristic map satisfies the scaling relation $\Phi(z/2) = \mu\Phi(z)$ where μ is the subdominant eigenvalue. The map $\Phi^\infty(z) = (h_k(z))^{4/k}$ satisfies the scaling relation with $\mu = 2^{-4/k}$. This assumption appears to be in reasonable agreement with our data (Figure 18).

5 Conclusion

B-splines of arbitrary order can be constructed in the regular setting through a process of upsampling followed by the appropriate number of averaging steps. In the arbitrary topology setting upsampling turns into a topological vertex split operation. Averaging remains well defined by computing dual control points as centroids of their neighboring primal control points and vice versa. Such repeated averaging constructions of subdivision schemes which generalize tensor product B-splines to the arbitrary topology setting are conceptually simple and easy to implement. We have shown in this paper that the resulting schemes are C^1 (but not C^2) at irregular vertices (primal; odd degree) respectively faces (dual; even degree) for up to and including nine averaging steps. All schemes can be accommodated concurrently in the same quadtree datastructures and adaptive subdivision is easily accommodated even in the case of dual schemes.

Potential applications for these schemes can be found in CAGD which at times requires higher order schemes than the most often used bicubics. Additional application areas include the Subdivision Element Method which could benefit from adaptive order basis functions as required by p-methods.

In our current framework boundaries are not well accommodated. It is straightforward to use the appropriate univariate subdivision scheme on the boundary, but it is well known that irregular vertices, respectively faces, on the boundary are problematic without special accommodation in the subdivision stencils. Similar problems appear at corners. In future work it would be interesting to pursue simple local modifications of the atomic averaging steps to build rules that work correctly on the boundaries. On the theory side it would be desirable to find an answer regarding the smoothness of repeated averaging schemes as the number of averaging steps grows without bound.

Acknowledgement

This work was supported in part by NSF (ACI-9721349, ACI-9978147, DMS-9874082, DMS-9872890), Alias|Wavefront, a Packard Fellowship and a Sloan Fellowship. Special thanks to Khrysaundt Koenig for modeling, lighting and texturing and Jianhui Zhang for implementation help.

References

[1] E. Cohen, T. Lyche, R. Riesenfeld, Discrete b-splines and subdivision techniques in computer aided geometric design and computer graphics, Computer Graphics and Image Processing 14 (2) (1980) 87–111.

[2] E. Catmull, J. Clark, Recursively generated B-spline surfaces on arbitrary topological meshes, Computer Aided Design 10 (6) (1978) 350–355.

[3] D. Doo, M. Sabin, Analysis of the behaviour of recursive division surfaces near extraordinary points, Computer Aided Design 10 (6) (1978) 356–360.

[4] A. Habib, J. Warren, Edge and vertex insertion for a class of c^1 subdivision surfaces, CAGD 16 (4) (1999) 223–247, previously available as a TR, Rice University, August 1997.

[5] J. Peters, U. Reif, The simplest subdivision scheme for smoothing polyhedra, ACM Transactions on Graphics 16 (4) (1997) 420–431.

[6] C. Loop, Smooth subdivision surfaces based on triangles, Master's thesis, University of Utah, Department of Mathematics (1987).

[7] N. Dyn, D. Levin, J. A. Gregory, A butterfly subdivision scheme for surface interpolation with tension control, ACM Trans. Gr. 9 (2) (1990) 160–169.

[8] D. Zorin, P. Schröder, W. Sweldens, Interpolating subdivision for meshes with arbitrary topology, Proceedings of SIGGRAPH 96 (1996) 189–192.

[9] L. Kobbelt, Interpolatory subdivision on open quadrilateral nets with arbitrary topology, in: Proceedings of Eurographics 96, Computer Graphics Forum, 1996, pp. 409–420.

[10] L. Velho, Using semi-regular 4-8 meshes for subdivision surfaces, to appear in the Journal of Graphics Tools (2001).

[11] L. Kobbelt, $\sqrt{3}$ subdivision, Proceedings of SIGGRAPH 2000 (2000) 103–112.

[12] F. Cirak, M. Ortiz, P. Schröder, Subdivision surfaces: A new paradigm for thin-shell finite-element analysis, Int. J. Numer. Meth. Engng. 47 (2000) 2039–2072.

[13] J. Stam, On subdivision schemes generalizing uniform b-spline surfaces of arbitrary degree, to appear in CAGD (2000).

[14] H. Prautzsch, Smoothness of subdivision surfaces at extraordinary points, Adv. Comput. Math. 9 (3-4) (1998) 377–389.

[15] L. Velho, D. Zorin, 4–8 subdivision, to appear in CAGD (2000).

[16] J. Warren, H. Weimer, Subdivision for geometric design, to appear (2001).

[17] X. Xu, K. Kondo, Adaptive refinements in subdivision surfaces, in: Eurographics 99 Proceedings, 1999.

[18] R. Qu, Recursive subdivision algorithms for curve and surface design, Ph.D. thesis, Brunel University (1990).

[19] H. Samet, The Design and Analysis of Spatial Data Structures, Addison-Wesley, 1990.

[20] H. Biermann, A. Levin, D. Zorin, Piecewise smooth subdivision surfaces with normal control, Proceedings of SIGGRAPH 2000 (2000) 113–120.

[21] D. Zorin, A method for analysis of c^1-continuity of subdivision surfaces, SIAM Journal of Numerical Analysis 37 (4) (2000) 1677–1708.

[22] U. Reif, A unified approach to subdivision algorithms near extraordinary points, Comput. Aided Geom. Des. 12 (1995) 153–174.

[23] D. Zorin, Smoothness of subdivision on irregular meshes, Constructive Approximation 16 (3) (2000) 359–397.

[24] E. I. Jury, Theory and Applications of the z-Transform Method, Wiley, New York, 1964.

[25] A. S. Cavaretta, W. Dahmen, C. A. Micchelli, Stationary subdivision, Memoirs Amer. Math. Soc. 93 (453).