Curvature-bounded guided subdivision: biquartics vs bicubics

Kęstutis Karčiauskas[a] and Jörg Peters[b]
[a] Naugarduko 24, Vilnius University, LT-2006, Vilnius, Lithuania
[b] Dept CISE, University of Florida, Gainesville FL 32611-6120, USA, jorg@cise.ufl.edu, tel/fax 352.392.1200/1220

Abstract

A sequence of C^2-connected nested subdivision rings of polynomial degree bi-4 can be made to follow a guide surface and completed by a tiny finite cap to serve as a refinable surface representation for design and analysis [1]. This raises the question, both of academic and practical interest, how much and at what cost to surface quality can the efficiency be improved by lowering the number of rings and/or the polynomial degree.

In a systematic exploration, a new bi-4 construction is discovered that requires half the number of surface rings but matches the quality of [1]. For this surface quality, numerous trials indicate that this number of surface rings is minimal and that the degree can not be reduced. Bi-3 constructions have inferior highlight line distributions – although the best of the new bi-3 constructions visibly improve on Catmull-Clark subdivision and its curvature-bounded variants.

Keywords: subdivision surfaces, highlight line distribution, rapid contraction, guided subdivision, curvature-bounded, bi-quartics, bi-cubics

1. Introduction

Recently, progress has been made towards creating subdivision surfaces with good highlight line distributions and low degree. [2] defined a sequence of C^2 bi-5 surface rings of the same formal smoothness as Catmull-Clark but with much better highlight line distribution. [3] replaced the three bi-3 patches per sector of each Catmull-Clark subdivision ring by three 2×2 macro-patches of degree bi-4. This paper shows that the improved surface quality and curvature boundedness can be achieved with just *one* patch and double the speed. This reduction in the number of patches represents an 8-fold gain in efficiency over the earlier constructions. New improved techniques also generate bi-3 subdivision surfaces with bounded curvature at the extraordinary point, good highlight line distribution and twice the contraction of Catmull-Clark. By contrast, curvature bounded subdivision schemes in the literature contract, especially for high valences, much slower than Catmull-Clark.

A second motivation for devising a new generation of subdivision surfaces is to provide surfaces with C^k transitions (as opposed to G^k transitions based on reparameterization, such as [4, 15]). C^k transitions simplify basic geometry processing operations including texturing, path tracing, engineering analysis, localized geometric modification, etc.. For example, introducing localized geometric detail on a G^2 surface is non-trivial whereas C^2 guided subdivision surface rings support nested refinement by knot insertion. Fixing a maximal expected refinement level as suggested in [1] avoids the main drawback of the subdivision approach for engineering analysis: the infinite sequence of subdivision rings that requires careful estimates to compute approximate integrals of derivatives. The suggestion of [1] is to stop subdivision after a few rings and fill the gap with a tiny, but well-behaved G^k cap that is not expected to be refined for the application at hand. Fig. 1 and Fig. 3 show how capped subdivision enables localized design.

Tensor-product bicubic splines are widely used to represent higher-quality surfaces due to their curvature continuity and their ability to model inflections as well as subtle convex shapes. Knot insertion, the origin of subdivision, provides a means to lay down the basic shape and to follow up with modification of detail. Catmull-Clark subdivision [5] was developed to extend these advantages to irregular layout where multiple other than four polynomial surface pieces meet at a point. Due to their simple C^2 transitions almost everywhere, these surfaces lend themselves also for modeling functions on surfaces, e.g. for isogeometric computations [6, 7]. However, it is well-understood that Catmull-Clark surfaces have visible shape deficiencies, e.g. pinched highlight lines near high-valent vertices (see e.g. Fig. 3b).

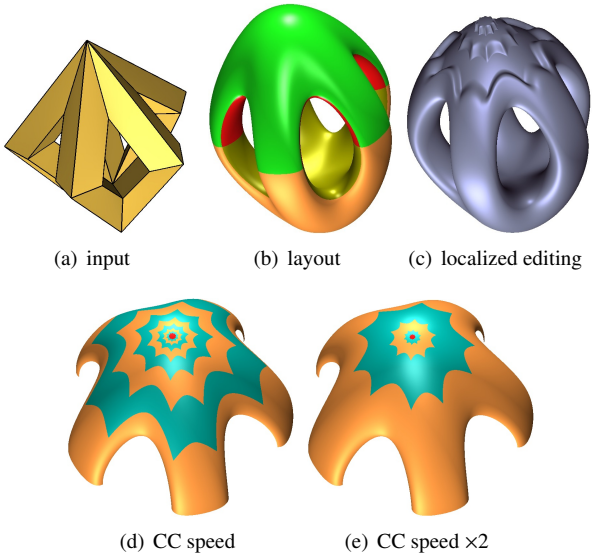

(a) input (b) layout (c) localized editing

(d) CC speed (e) CC speed ×2

Figure 1: Input polyhedron (a) with extraordinary nodes of valence 10. (b) Layout of multi-sided macro-patches. (c) Embossed features taking advantage of refinability. (d) 8 layers forming the green surface in (b) with a contraction speed of Catmull-Clark (CC) subdivision, completed by a tiny red cap. (e) 4 layers at twice the contraction rate.

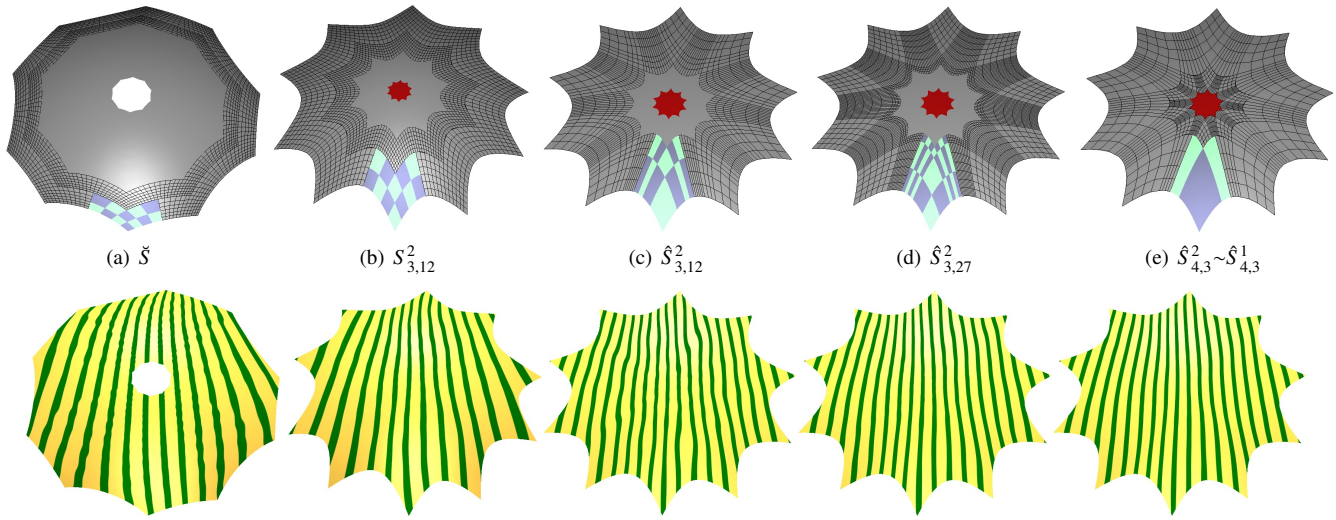

(a) $\check{S}$ (b) $S^2_{3,12}$ (c) $\hat{S}^2_{3,12}$ (d) $\hat{S}^2_{3,27}$ (e) $\hat{S}^2_{4,3}$~$\hat{S}^1_{4,3}$

Figure 2: $S^k_{dp} := C^k$*-sequence of subdivision rings of degree bi-d using p patches per sector.* $\hat{S}$ indicates double the contraction speed. (*top* row) Surface layout: the polynomial pieces of one sector of one subdivision ring are highlighted. (*bottom* row) The corresponding highlight lines near the extraordinary point. (a) Bi-3 curvature bounded scheme $\check{S}$ of [8]. (b) Bi-3 with contraction speed of Catmull-Clark subdivision. (c,d,e) Double contraction speed. (c) Bi-3 with 12 patches per sector. (d) Bi-3 with three 3*x*3 patches per sector, i.e. 27 patches of unequal size (cf. Fig. 4c). (e) Bi-4 with 3 patches per sector ($\hat{S}^2_{4,3}$ and $\hat{S}^1_{4,3}$ are visually alike).

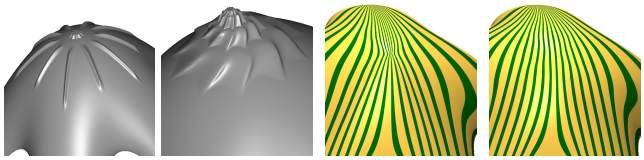

(a) localized editing of $\hat{S}^2_{4,3}$ (b) Catmull-Clark (c) $\hat{S}^2_{4,3} \sim \hat{S}^1_{4,3}$

Figure 3: Valence 10 input from Fig. 1. (a,b) Embossed features: at every refinement, the same coefficients within each sector are (*left*) displaced in the normal direction or (*right*) rotated. (b,c) highlight lines of Catmull-Clark surface vs bi-4 ($\hat{S}^2_{4,3}$ and $\hat{S}^1_{4,3}$ are visually alike).

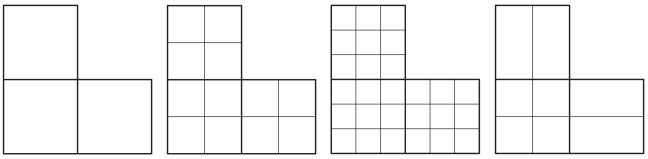

(a) $p = 3$ (b) $p = 12$ (c) $p = 27$ (d) $p = 8$

Figure 4: Layout of subdivision rings for S^k_{dp} and $\hat{S}^k_{dp}$. $\hat{S}$ indicates double the contraction speed of S. (a) Three patches per sector as for Catmull-Clark subdivision surfaces. (b) Three 2×2 macro-patches = 12 patches. (c) Three 3×3 macro-patches = 27 patches. (d) One 2×2 macro-patch and two 2×1 macro-patches = 8 patches.

To highlight the contribution of this paper, Fig. 1a displays new subdivision-plus-cap surfaces for an input mesh with extraordinary nodes of valence 10. This mesh is extreme in that it has no 4×4 submesh suitable for interpretation as a control net of a bicubic spline patch. We chose this high-valence mesh to expose more starkly otherwise subtle differences between various options. For the input mesh in Fig. 1a, Fig. 2a shows the layout (*top*) and highlight line distribution (*bottom*) of Sabin's pioneering bounded curvature subdivision [8]. Note the extremely slow approach to the extraordinary point, a behaviour shared by later improvements in [9, 10]. The new guided subdivision schemes can even double the contraction speed compared the Catmull-Clark subdivision and so reduce the number of rings before applying the final cap. And all but one subdivision in Fig. 2 are C^2 except at extraordinary point. At the extraordinary point they are C^1 and curvature bounded. Guided bi-3 subdivision clearly improves on Catmull-Clark subdivision and its curvature-bounded variants. But even finely-partitioned bi-3 macro-patches yield a less regular highlight line distribution than the bi-4 construction. For example, the 2×2 macro-patches of degree bi-3 and the contraction speed of Catmull-Clark subdivision perform poorer than single bi-4 patches with the more challenging and efficient double contraction speed.

The good highlight line distribution is due to the guided approach that decouples differential smoothness from shape: the subdivision surfaces closely follow a shape of an on-the-fly constructed guide. The guided approach stabilizes the C^2 subdivision sequence to allow for accelerated contraction even for rings consisting of a patchwork of thin and small patches Fig. 2c,d,e.

To summarize, the contributions of this paper are:

- A new class of low-degree, smooth, accelerated guided subdivision algorithms with good highlight line distribution. Acceleration means fewer patches.
- Subdivision-plus-cap hybrid surfaces consisting of several subdivision rings and a tiny cap. The hybrid surfaces are almost everywhere parametrically C^2. They are everywhere smooth.
- A systematic comparison of tuned bi-3, new guided bi-3 and new guided bi-4 guided subdivision schemes.

Overview. Section 2 lays out the setup, operators and technical preliminaries and presents a synopsis of guided subdivision. Section 3 defines the crucial guide sampling mechanism. Section 4 and Section 5 present families of subdivision schemes of degree bi-4 and bi-3 respectively. Section 6 explains the cap construction (not a focus here). Section 7 compares the subdivision surfaces in more detail and Section 8 places the observations in a larger context.

2. Definitions and Setup

2.1. A B-spline-like control net for irregular layout

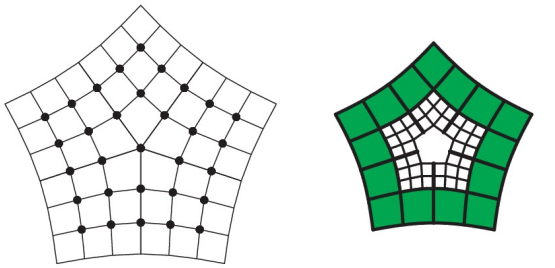

(a) **c**-net extended by 1 layer (b) bi-3 ring + tensor-border

Figure 5: B-spline-like irregular control net and its tensor-border. (a) Extended **c**-net for $n = 5$. (b) Schema of surface ring and its tensor-border (mesh of BB-coefficients) of degree 3 and depth 2. The tensor-border is the input for the rapid surface construction.

We consider as input a network of quadrilateral facets, quads for short. Nodes where four quads meet are regular, else *irregular*. We assume that each irregular node is surrounded by at least one layer of regular nodes. Fig. 5a shows the **c**-net (bullets) of an isolated node of valence $n = 5$. The **c**-net consists of the irregular node plus $6n$ nodes forming two layers of quads surrounding it. Typically a third layer is added for evaluation of local shape (yielding the green surface in Fig. 5b). This allows assessing the highlight line distribution [11] across the transition which is as important as the internal quality of the cap.

Each 4×4 sub-grid of nodes is interpreted as the B-spline control points of a bicubic tensor-product spline surface. Except at the irregular node, well-known formulas can be applied to convert the B-spline form to Bernstein-Bézier form (see e.g. [12, 13]). The tensor-product Bernstein-Bézier (BB) form of bi-degree d is

$$\mathbf{p}(u,v) := \sum_{i=0}^{d}\sum_{j=0}^{d} \mathbf{p}_{ij} B_i^d(u) B_j^d(v) \,,$$

$$(u,v) \in \square := [0..1]^2, \qquad \text{where } B_k^d(t) := \binom{d}{k}(1-t)^{d-k}t^k$$

are the Bernstein polynomials of degree d and $\mathbf{p}_{ij}$ are the BB coefficients. Fig. 5b also shows the C^2 prolongation of this surface ring, i.e. Hermite data represented as a grid (black) of bi-3 BB-coefficients. Specifically, the BB-coefficients $\mathbf{p}_{ij}$, $i = 0,\ldots,3$, $j = 0,\ldots,2$, represent Hermite data of order 2 along one boundary curve $v = 0$. We call these data $\mathbf{t}_{CC}$. More generally, in the remainder of this paper we refer to second-order Hermite data of degree $d \geq 3$ along the loop of boundary curves as **t**.

2.2. Corner jet constructor and related operators

We will construct the tensor-product patches and tensor-borders with the help of *corner jet constructors* in BB-form. A corner jet constructor $[f]^d$ expresses, at a corner of domain square $[0..1]^2$, the expansion of a function f of order 2 in both u and v directions in BB-form of bi-degree d. That is, $[f]^d$ outputs 3×3 BB-coefficients (see Fig. 6a,b). (This geometric interpretation of the BB-form is one of the main reasons for its wide-spread use.) Fig. 6c displays four corner jets $[f]^4$ merged to form a bi-4 patch by averaging the overlapping BB-coefficients. Fig. 6d illustrates the analogous assembly of an L-shaped sector of tensor-border of degree 4 by applying and averaging a jet constructor $[f]^4$ at three corners.

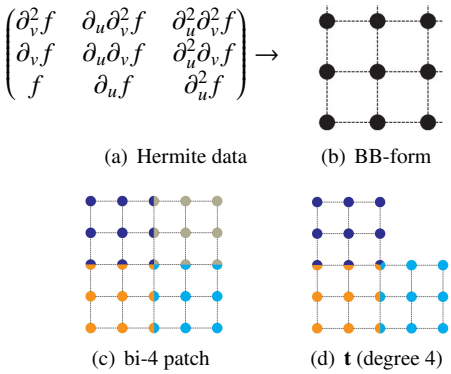

(a) Hermite data (b) BB-form

(c) bi-4 patch (d) **t** (degree 4)

Figure 6: (a) Hermite data as partial derivatives converted to (b) BB-form. (c) A patch of degree bi-4. (d) L-shaped sector of the tensor-border **t** of degree 4.

Fig. 7 illustrates *BB-to-B-extrapolation*: a bi-3 B-spline control point (marked as •) is defined as a weighted average of inner BB-coefficients (marked as circles). The weights $\beta_i := \frac{\Delta_i}{\Delta_{i-1}}$ for a knot sequence t_i in u are the ratios of adjacent knot intervals Δ_{i-1} and Δ_i. The analogous weights for v are $\tilde{\beta}_j$. The extrapolation will be used to construct local B-spline control nets when the inner BB-coefficients do not belong to C^2-connected patches.

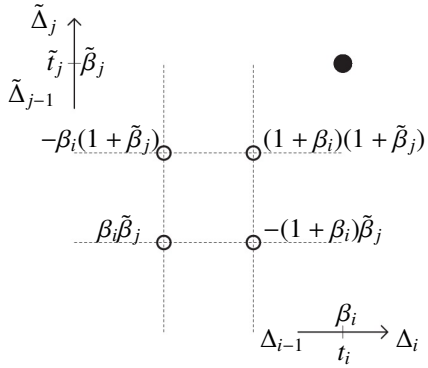

Figure 7: BB-to-B-extrapolation: a B-spline control point • derived from the inner BB-coefficients ∘ of a bicubic Bézier patch.

Several steps of the surface construction use a simple rule. Consider two curve segments in BB-form of the same degree and whose consecutive domain intervals have the ratio $\sigma : 1-\sigma$. The two segments join C^2 at their common end-point (marked as a big bullet in Fig. 8) if and only if the BB-coefficients immediately to its *left* and *right* (small circles indicated by ↓) are defined as a weighted average of the bullets with the displayed weights, where

$$e_0 := \frac{1-\sigma}{2},\; e_1 := \frac{1}{2(1-\sigma)},\; e_2 := 1 - e_0 - e_1, \qquad (1)$$

and $\tilde{e}_i$, $i = 0,1,2$ are obtained from e_{2-i} by replacing σ by $1-\sigma$. We call this averaging the C2-rule.

Fig. 8b displays the weights for the *left* circled BB-coefficient in case of equal lengths, i.e. $\sigma := \frac{1}{2}$ (the *right* BB-coefficient is defined by the mirrored formula). We refer to this special case as *symmetric* C2-rule.

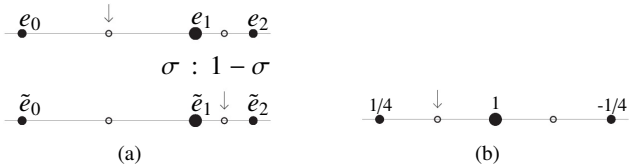

Figure 8: (a) General and (b) symmetric C^2 rule.

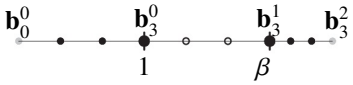

Figure 9: C^2 joining cubics $\mathbf{b}_i^s$, $i = 0,\dots 3$, $s = 0, 1, 2$, defined over the intervals of lengths Δ_s, where $\frac{\Delta_1}{\Delta_0} := 1$, $\frac{\Delta_2}{\Delta_1} := \beta$.

In bi-3 constructions with 3×3 macro-patches we use the arrangement of three C^2-connected cubics displayed in Fig. 9. The BB-coefficients $\mathbf{b}_i^0$, $\mathbf{b}_i^2$, $i = 1, 2$ (small bullets in Fig. 9) uniquely define three C^2 joining cubics $\mathbf{b}^0$, $\mathbf{b}^1$, $\mathbf{b}^2$:

$$\begin{aligned}
\mathbf{b}_3^0 &:= \frac{1}{2\beta(2+\beta)}\big(-\beta(1+\beta)\mathbf{b}_1^0 + \beta(4+3\beta)\mathbf{b}_2^0 + (1+\beta)\mathbf{b}_1^2 - \mathbf{b}_2^2\big);\\
\mathbf{b}_0^1 &:= \mathbf{b}_3^0,\ \mathbf{b}_1^1 := 2\mathbf{b}_3^0 - \mathbf{b}_2^0,\ \mathbf{b}_2^1 := 4\mathbf{b}_3^0 - 4\mathbf{b}_2^0 + \mathbf{b}_1^0; \qquad (2)\\
\mathbf{b}_3^1 &:= \frac{1}{1+\beta}(\beta\mathbf{b}_2^1 + \mathbf{b}_1^2),\ \mathbf{b}_0^2 := \mathbf{b}_3^1\,.
\end{aligned}$$

For $\beta := 1$ we call this assignment *CS2-rule(1)*, for $\beta := \frac{1}{4}$ *CS2-rule($\frac{1}{4}$)*.

2.3. The guide surface $\mathbf{g}$

Since the guided subdivision surfaces follow the shape of a guide surface, its quality is of paramount importance. Among several choices of good quality, we select the G^1 bi-5 guide $\mathbf{g}$ derived in [14]. For completeness and to motivate the reparameterizations and the eigen-structure of the resulting guided subdivision process, we reproduce its main points below.

Maps of total degree d We will use maps $\mathbf{b}^\triangle$ of total degree d whose domain is a regular n-gon D composed from n equal triangles with a common vertex $\mathbf{O}$ at the origin. Fig. 10a shows one such triangle with sides defined by $l_i = 0$, $i = 0, 1, 2$. Each linear barycentric coordinate function l_i is equal to 1 at the vertex opposite to $l_i = 0$. On the triangle we define a map $\mathbf{b}$ of *total degree d* in Bernstein-Bézier form as

$$\mathbf{b} := \sum_{i+j+k=d} \mathbf{b}_{ijk} B_{ijk}^d,\ {\scriptstyle i,j,k\geq 0}, \quad B_{ijk}^d := \binom{d}{ijk} l_0^i l_1^j l_2^k\,. \qquad (3)$$

Fig. 10b labels the BB-coefficients rotationally symmetric. Let $\mathsf{c} := \cos\frac{2\pi}{n}$. Pieces on adjacent sectors $\mathbf{b}^s$ and $\mathbf{b}^{s+1}$ join

$$\begin{aligned}
C^0 \text{ if } \quad & \mathbf{b}_{d-i,i,0}^{s+1} = \mathbf{b}_{d-i,0,i}^s, \quad {\scriptstyle i=0,\dots,d} && (4)\\
C^1 \text{ if } \quad & \mathbf{b}_{d-i,i-1,1}^{s+1} = -\mathbf{b}_{d-i,1,i-1}^s + 2\mathsf{c}\mathbf{b}_{d-i,0,i}^s && (5)\\
& + 2(1-\mathsf{c})\mathbf{b}_{d-i+1,0,i-1}^s, \quad {\scriptstyle i=1,\dots,d}\\
C^2 \text{ if } \quad & \mathbf{b}_{d-i,i-2,2}^{s+1} = \mathbf{b}_{d-i,2,i-2}^s - 4\mathsf{c}\mathbf{b}_{d-i,1,i-1}^s && (6)\\
& + 4\mathsf{c}^2\mathbf{b}_{d-i,0,i}^s - 4(1-\mathsf{c})\mathbf{b}_{d-i+1,1,i-2}^s\\
& + 8\mathsf{c}(1-\mathsf{c})\mathbf{b}_{d-i+1,0,i-1}^s + 4(1-\mathsf{c})^2\mathbf{b}_{d-i+2,0,i-2}^s,\\
& {\scriptstyle i=2,\dots,d}.
\end{aligned}$$

across the *sector boundary* between patch s and patch $s+1$ (modulo n), $s \in \{0, 1, \dots, n-1\}$. The six BB-coefficients $\mathbf{b}_{ijk}^0$ (indicated as red bullets in Fig. 10) define a quadratic expansion $\mathbf{q}$ at the central point $\mathbf{b}_{d00}^0$. This local expansion is propagated to the neighboring sectors by repeatedly enforcing Eq. (4) for $i = 0, 1, 2$, Eq. (5) for $i = 1, 2$ and Eq. (6) for $i = 2$. That is, the $\mathbf{b}_{ijk}^0$ define a unique quadratic expansion of the C^1 map $\mathbf{b}^\triangle$ at $\mathbf{b}_{d00}^0$. The 6 red bullets are unrestricted by C^2 continuity at the central point and $6n$ black bullets are unrestricted by C^1 continuity between sectors. The BB-coefficients forming the vertices of the gold triangles in each sector are unrestricted due to their distances from the sector boundary.

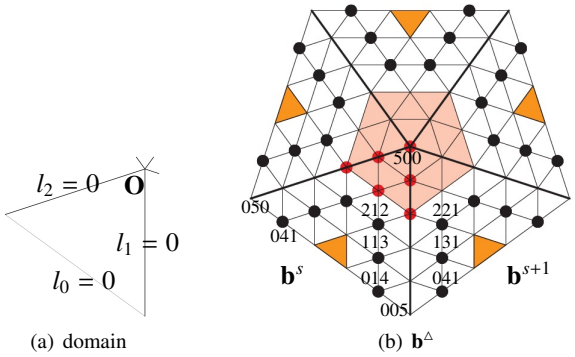

Figure 10: (a) Domain of total degree d map. (b) The $n = 5$ sectors of the C^1 map $\mathbf{b}^\triangle$ of total degree $d = 5$. The 'light red' shaded region indicates the quadratic expansion at the center. The BB-coefficients that remain unrestricted when enforcing C^1 constraints are marked as black and red bullets; the 'gold' underlaid BB-coefficients do not affect C^1 continuity between sectors.

Linear shear map L

We denote by L the linear shear that maps a unit square to the unit parallelogram with opening angle $\frac{2\pi}{n}$ as illustrated Fig. 11. To increase the flexibility of $\mathbf{b}^\triangle$ we set for each sector

$$\mathbf{g} := \mathbf{b}^\triangle \circ L\,.$$

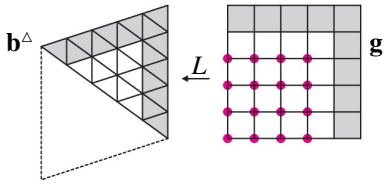

Figure 11: Increasing the flexibility of $\mathbf{b}^\triangle$ via linear shear L.

Then, see Fig. 11, along the sector boundary the gray-underlaid BB-coefficients of the C^1 map $\mathbf{b}^\triangle$ determine the corresponding gray-underlaid BB-coefficients of G^1 map $\mathbf{g}$. But while each sector of $\mathbf{b}^\triangle$ has only three coefficients not influencing C^1 continuity between sectors, each sector of $\mathbf{g}$ has 16 coefficients (magenta bullets) that do not influence G^1 continuity of $\mathbf{g}$. These coefficients provide valuable extra degrees of freedom. Therefore the guide $\mathbf{g}$ has a total of $16n + 6n + 6 = 22n + 6$ degrees of freedom. These are set as the least-squares solution as in [15]. The result is efficiently stored as a matrix expressing the guide's BB-coefficients in terms of the input $\mathbf{c}$-net.

2.4. Synopsis of Guided Subdivision

The idea of guided subdivision [16] is to separate shape finding from the mathematical smoothness constraints of the final

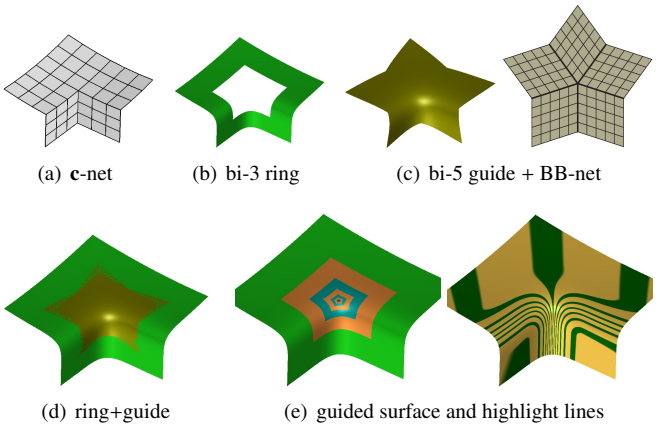

(a) **c**-net (b) bi-3 ring (c) bi-5 guide + BB-net

(d) ring+guide (e) guided surface and highlight lines

Figure 12: Guide and subdivision. (a) The extended **c**-net defines a bi-3 surface ring (b) and the guide (c). The guide does not precisely match the surface ring. (d) Overlay of surfaces. (e) Gold ring derived from (b) and (c), then cyan from gold ring and (c), etc..

output surface. In the guided approach, one first decides the shape. The guide and shape are based on the surrounding surface ring Fig. 12b. The guide surface Fig. 12c need neither be smooth everywhere nor exactly match the surrounding surface Fig. 12d. Guided subdivision then uses the flexibility of the sequence of subdivision rings (see Fig. 13) to smoothly connected to the surrounding surface and between rings while absorbing extra degrees of freedom by closely approximating the guide shape Fig. 12e. [1] focuses on the trade-off between shape and complexity for increased smoothness between the rings at standard contraction rates. Bi-3 surfaces are not considered. Extremely fast contractions, at the cost of surface rings higher polynomial degree, were introduce in [2]. The guided approach has also been used to improve the visible shape of singularly parameterized surfaces [14]. Just like the proposed subdivision surfaces, singularly parameterized surfaces facilitate adaptive refinement and (isoparametric) analysis on surfaces. However, such surfaces have pinched highlight line distributions near the singularity, visible under zoom.

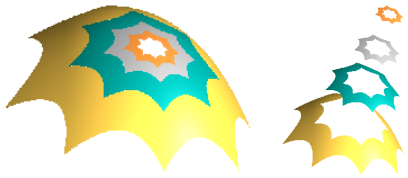

Figure 13: (*left*) Subdivision surface built from (*right*) a sequence of contracting surface rings

3. Characteristic parameterizations for sampling the guide g

The subdominant eigenvalue of bi-3 adjustable speed subdivision [KP09] with 'speed' parameter $\sigma \in (0..1)$ and $\tilde{\sigma} := 1 - \sigma$ is

$$\lambda_\sigma := \frac{\tilde{\sigma}}{2}((1+\mathsf{c})\sigma^2 + 2\tilde{\sigma} + \sigma\sqrt{(1+\mathsf{c})((1+\mathsf{c})\sigma^2 + 4\tilde{\sigma})}), \quad (7)$$

For $\sigma := \frac{1}{2}$ adjustable speed subdivision coincides with Catmull-Clark subdivision.

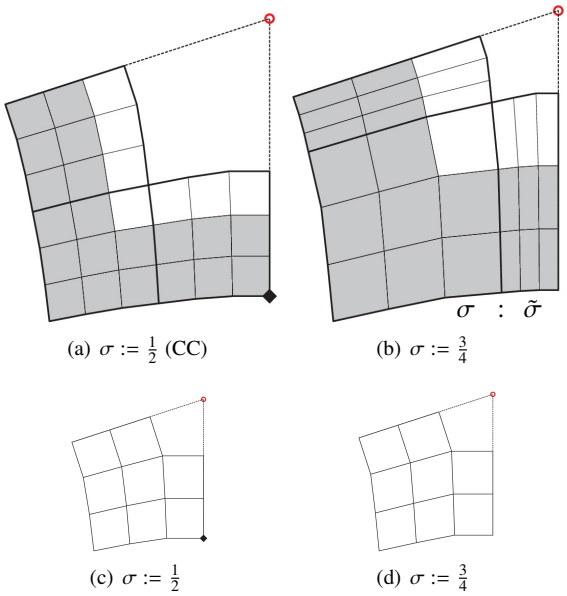

(a) $\sigma := \frac{1}{2}$ (CC) (b) $\sigma := \frac{3}{4}$

(c) $\sigma := \frac{1}{2}$ (d) $\sigma := \frac{3}{4}$

Figure 14: (a,b) χ_σ; (c,d) corresponding $\tilde{\chi}_\sigma$.

Fig. 14a,b display characteristic maps χ_σ for one sector and $n = 5$. Fig. 14c,d display the corresponding characteristic tensor-borders $\tilde{\chi}_\sigma$ of degree 3 and depth 2. In the classical analysis of Catmull-Clark subdivision these tensor-borders are not emphasized since the maps χ_σ already contain the key analytic information. In guided subdivision, however, the tensor-borders $\tilde{\chi}_\sigma$ are in the forefront. The gray underlaid BB-coefficients of χ_σ are the result of splitting $\tilde{\chi}_\sigma$ in the ratio $\sigma : \tilde{\sigma}$ (see Fig. 14b). We consider normalized maps and tensor-borders where the corner BB-coefficients (marked as diamond in Fig. 14a) are at the distance 1 from the red center. While different σ result in different maps χ_σ, their tensor-borders $\tilde{\chi}_\sigma$, although slightly different, look alike. This allows switching to a well-behaved rapidly contracting guided subdivision with $\sigma := \frac{3}{4}$ for input $\mathbf{t}_{CC}$ where $\sigma := \frac{1}{2}$.

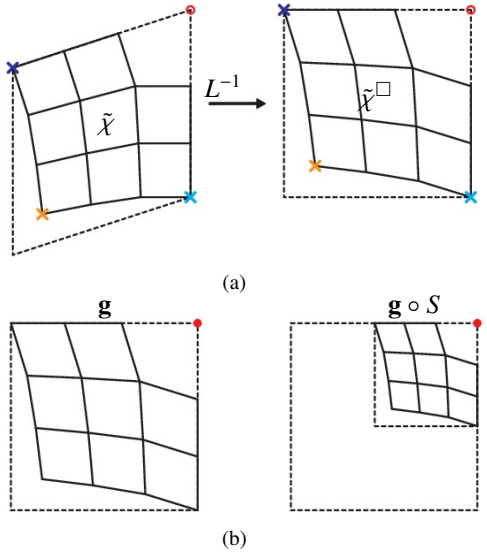

Figure 15: (a) One sector of the transformed characteristic tensor-border; (b) scaled sampling. The origin is marked red.

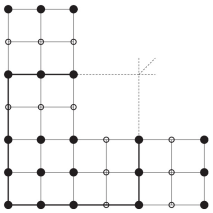

Figure 16: C^2 corrected tensor border $\tilde{\mathbf{t}}$.

We jet-sample the guide $\mathbf{g}$ with $\tilde{\chi}^{\square} := L^{-1} \circ \tilde{\chi}$ in each sector at the locations marked as crosses in Fig. 15a. The contracting guided rings are formed from jets $[\mathbf{g}\ \lambda^r_\sigma\ \tilde{\chi}^{\square}]^d$, $r = 1, \ldots, l$, for fixed σ; l is a maximal anticipated refinement level. Computing $[\mathbf{g} \circ (\lambda\tilde{\chi}^{\square})]^d$ is equivalent to linearly mapping $S : [0..1]^2 \to [0..\lambda]^2$ and sampling $(\mathbf{g} \circ S) \circ \tilde{\chi}^{\square}$; Fig. 15b shows $\tilde{\chi}^{\square}$ and its scaled copy $\lambda\tilde{\chi}^{\square}$. DeCasteljau's algorithm at $u = \lambda = v$ yields the BB-coefficients of $\mathbf{g} \circ S$ as the affine combinations of BB-coefficients of $\mathbf{g}$. These are tabulated as an affine $6^2 \times 6^2$ map (matrix). Jet-sampling symbolic input with $\tilde{\chi}^{\square}$ at the locations marked as crosses in Fig. 15a yields three pre-calculated 9×6^2 affine maps that express the BB-coefficients of 3×3 corner jets as affine combinations of BB-coefficients of a sector of guide $\mathbf{g}$ for an intermediate step of guided subdivision.

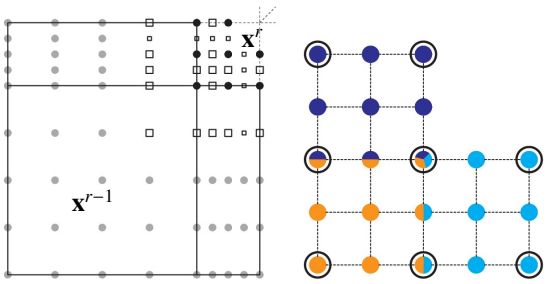

(a) $\mathbf{x}^{r-1} \to \mathbf{x}^r$ (b) $\mathbf{t}$

Figure 17: Construction of C^2 bi-4 surfaces.

4. Biquartic guided surfaces

In this section, we denote by $\mathbf{t}^r$ the tensor-border assembled as in Fig. 6d, from jets $[\mathbf{g} \circ (\lambda^r\tilde{\chi}^{\square})]^4$. Since the sectors of the guide $\mathbf{g}$ are G^1 connected, the sectors of $\mathbf{t}^r$ are C^1 connected. Applying the C2-rule joins the sectors C^2. The C^2 tensor-border is called $\tilde{\mathbf{t}}^r$, Fig. 16 ($\tilde{\mathbf{t}}^r$ is also used for variants when constructing $\hat{S}^2_{4,3}$ below). We choose $\sigma := \frac{3}{4}$ as contraction speed, i.e. twice the speed of Catmull-Clark-subdivision. As introduced in Fig. 2, we abbreviate as S^k_{dp} a C^k-sequence of subdivision rings of degree bi-d using p patches per sector. $\hat{S}$ indicates double the contraction speed.

$\hat{\mathbf{S}}^{\mathbf{2}}_{\mathbf{4,3}}$: For the patch layout of the bi-4 rings displayed in Fig. 2e, the BB-coefficients of ring $\mathbf{x}^r$ marked as gray bullets in Fig. 17a are the split C^2 prolongation of $\mathbf{x}^{r-1}$. If $r = 0$, the gray bullets represent a degree-raised and split input tensor-border $\mathbf{t}_{CC}$. The circled BB-coefficients of $\mathbf{t}^{r+1}$ of Fig. 17b are set to those marked as black bullets in Fig. 17a. Then every second BB-coefficient (marked as big hollow square in Fig. 17a) is determined via the *general* C2-rule from the nearest BB-coefficients marked as black and gray bullets. BB-coefficients marked as small hollow squares are determined with the help of data from adjacent sectors via the *symmetric* C2-rule. This automatically provides the C^2 prolongation for the next ring of $\hat{\mathbf{S}}^{\mathbf{2}}_{\mathbf{4,3}}$.

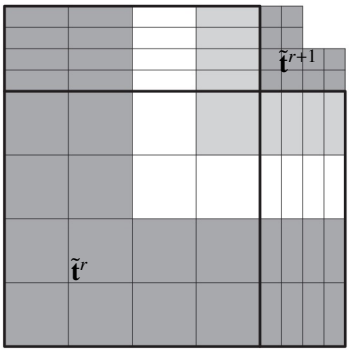

(a) C^1 case

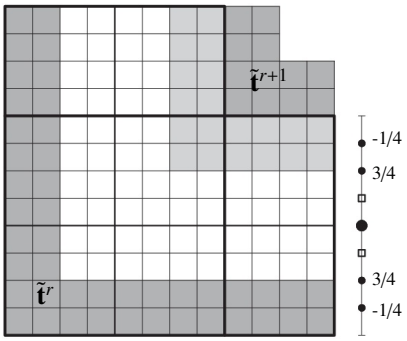

(b) C^2 case

Figure 18: Construction of bi-4 surfaces: (a) C^1 (b) C^2 made uniform by choosing $\sigma := \frac{2}{3}$.

$\hat{\mathbf{S}}^{\mathbf{1}}_{\mathbf{4,3}}$: With the same layout as for $\hat{S}^2_{4,3}$ the dark-gray underlaid BB-coefficients of $\mathbf{x}^r$ (*bottom*,*left* in Fig. 18a) represent the split tensor-border $\tilde{\mathbf{t}}^r$. If $r = 0$, they represent the degree-raised and split input tensor-border $\mathbf{t}_{CC}$. The light-gray underlaid BB-coefficients extend $\tilde{\mathbf{t}}^{r+1}$ backwards so that $\mathbf{x}^r$ joins C^1 with $\tilde{\mathbf{t}}^{r+1}$. We call these surfaces $\hat{S}^1_{4,3}$.

$\hat{\mathbf{S}}^{\mathbf{2}}_{\mathbf{4,8}}$: We can join consecutive rings C^2 by increasing the number of patches. As displayed in Fig. 18b, dark-gray underlaid BB-coefficients of $\mathbf{x}^r$ match the tensor-border $\tilde{\mathbf{t}}^r$ ($\mathbf{t}_{CC}$ for $r = 0$) and the light-gray underlaid BB-coefficients extend $\tilde{\mathbf{t}}^{r+1}$ backwards so that $\mathbf{x}^r$ joins C^2 with $\tilde{\mathbf{t}}^{r+1}$. In the stencil of Fig. 18b, *right*, the smaller bullets determine the boundary BB-coefficient (big black bullet). The symmetric C2-rule then determines the hollow squares. The resulting two once-split patches and the 2×2 corner patch per sector of $\mathbf{x}^r$ are now joined C^2 and each macro-patch is internally C^3.

Implementation.. Simple formulas define the BB-coefficients of the bi-4 rings in terms of the tensor-border $\tilde{\mathbf{t}}$, i.e. splitting $\tilde{\mathbf{t}}$, C^1/C^2 prolongation and applying the stencil of Fig. 18b. The tensor-border $\tilde{\mathbf{t}}$, too, is defined by simple formulas in terms of $\mathbf{t}$: for $\hat{S}^2_{4,3}$ the C2-rules and splitting the previous $\tilde{\mathbf{t}}$; for $\hat{S}^1_{4,3}$ and $\hat{S}^2_{4,8}$ the symmetric C2-rule. Therefore only the part of $\mathbf{t}$ used to define $\tilde{\mathbf{t}}$ is tabulated as a 6×6^2 matrix for $\hat{S}^2_{4,3}$ and a 12×6^2 matrix for $\hat{S}^1_{4,3}$ and $\hat{S}^2_{4,8}$. In summary, we tabulate the guide, de Casteljau's splitting of one sector and the linear sampling process so that the algorithm reduces to matrix operations.

Eigen-structure and continuity.. Starting with the ring $\mathbf{x}^1$, $\hat{S}^1_{4,3}$ and $\hat{S}^2_{4,8}$ inherit the 'eigen-structure' of the guide $\mathbf{g}$ and are curvature bounded C^1 (see the Appendix).

The C^2 bi-4 splines $\hat{S}^2_{4,3}$ are less local in that the black BB-coefficients sampled from $\mathbf{t}_{CC}$ impact all subsequent rings. Since the construction is guided this 'traveling' data does not negatively affect the shape of the surfaces. But it does affect the algebraic structure: the surfaces does not inherit the guide's eigen-structure. Nevertheless, the second part of the Appendix shows that these surfaces are C^1 and have bounded curvature at the extraordinary point.

5. Bicubic guided surfaces

For comparison with the bi-4 surfaces of Section 4, we explore three bicubic constructions that are C^2 except at the extraordinary point where they are C^1 and curvature bounded. Fig. 19 shows the split macro-patches: 2×2 and 3×3, uniform and non-uniform.

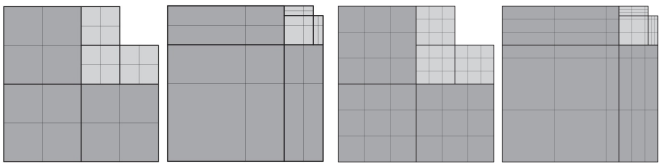

(a) $S^2_{3,12}$ (b) $S^2_{3,27}$

Figure 19: One refinement step of a sector of bi-3 rings with (a) regular speed S and (b) double speed $\hat{S}$. Ring $\mathbf{x}^r$ dark-gray, ring $\mathbf{x}^{r+1}$ light-gray.

For the first two constructions below we use the corner jet constructors $[\mathbf{g} \circ (\lambda^{r+1}\, \tilde{\chi}^{\square})]^3$ restricted to half the domain (i.e. scaled by 1/2), see Fig. 23a. Here λ is the subdominant eigenvalue of the subdivision (with Catmull-Clark speed in Fig. 20a,b and double speed in Fig. 20c,d). Using BB-to-B-extrapolation, averaged at overlapping locations as shown in Fig. 23b yields control points B^r marked as black disks in Fig. 20a,c. (B^r is key to keeping the contracting rings $\mathbf{x}^r$, $r = 0, 1, \ldots$ close to the guide.)

$\mathbf{S}^2_{3,12}$: Choosing $\sigma := \frac{1}{2}$ yields $\tilde{\chi}^{\square}$ of Catmull-Clark-subdivision and $\beta = \tilde{\beta} = 1$ in the BB-to-B-extrapolation of Fig. 7. Fig. 19a shows the refinement. When $r = 0$ the gray bullets A^0 in Fig. 20a represent the twice-refined input $\mathbf{c}$-net. Together with B^0, they define the ring $\mathbf{x}^0$ that joins C^2 with the input data $\mathbf{t}_{CC}$. Refining $A^0 \cup B^0$ yields the gray bullets A^1 in Fig. 20b); together with B^1 they define $\mathbf{x}^1$. Repeating the process, A^r and B^r define $\mathbf{x}^r$; and $\mathbf{x}^r$ and $\mathbf{x}^{r+1}$ are C^2-joined.

$\hat{\mathbf{S}}^2_{3,12}$: Combinatorially the construction is identical to $S^2_{3,12}$. However now λ and $\tilde{\chi}^{\square}$ correspond to subdivision with parameter $\sigma := \frac{3}{4}$. The parameters β and $\tilde{\beta}$ of the BB-B conversion result from the *local* knot sequences displayed in Fig. 20c; due to symmetry with respect to the sector-diagonal the formulas are alike for u and v and due to rotational symmetry for all sectors. Fig. 20d shows the red inserted knots to obtain the refined (light-gray) B-spline control points.

$\hat{\mathbf{S}}^2_{3,27}$: Let $\mathbf{t}^r$ be the tensor-border of degree 5 assembled from jets $[\mathbf{g} \circ (\lambda^r\, \tilde{\chi}^{\square})]^5$ (see Fig. 21a) and $\bar{\mathbf{t}}^r$ its C^2 correction as in Section 4, Scaling to 1/3 at the domain corners and re-represented in bi-3 form, yields the three 3×3 arrangements of BB-coefficients marked as bullets in Fig. 21b. Applying to each layer of BB-coefficients *CS2-rule(1)* yields a piecewise C^2 tensor-border $\bar{\mathbf{t}}^r$.

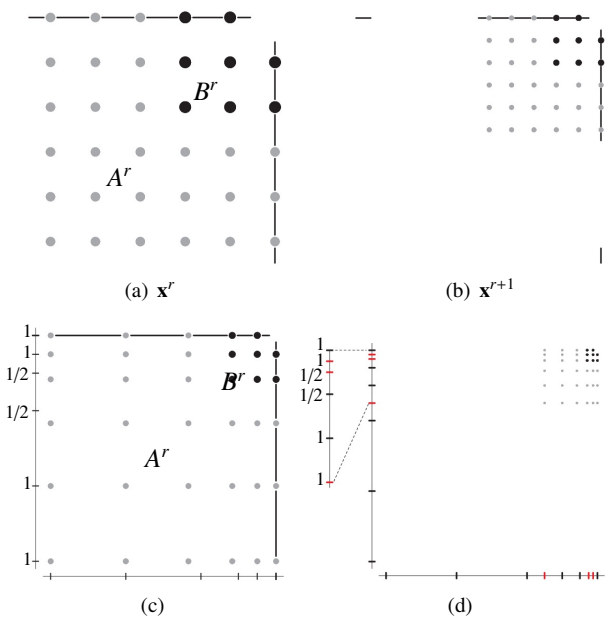

(a) $\mathbf{x}^r$ (b) $\mathbf{x}^{r+1}$ (c) (d)

Figure 20: (a,c): The B-spline control points $A^r \cup B^r$ of $\mathbf{x}^r$ (dark-gray in Fig. 19) are refined to form A^{r+1} of ring $\mathbf{x}^{r+1}$ (light-gray in Fig. 19). If $r = 0$ then the light-gray bullets represent part of the twice-refined input $\mathbf{c}$-net that defines an appropriately split input tensor-border $\mathbf{t}_{CC}$. (b,d): A^{r+1} (light-gray bullets) are determined from $A^r \cup B^r$ by refinement. B^{r+1} (black bullets) drive the refinement based on the guide $\mathbf{g}$.

This $\bar{\mathbf{t}}^r$ is split first at the black tick marks then at the red tick marks in Fig. 21c,d. The different ratios yield $\dot{\mathbf{t}}^r$, respectively $\ddot{\mathbf{t}}^r$. The constructions of $\mathbf{x}^r$ from 3×3 macro-patches in Fig. 22

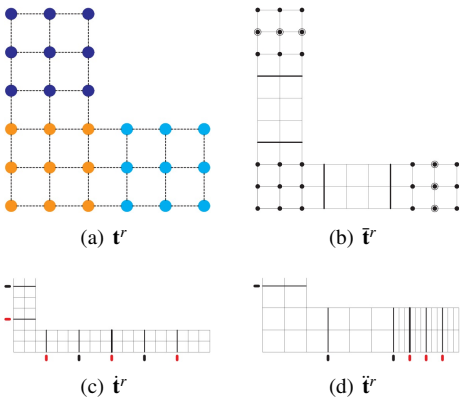

(a) $\mathbf{t}^r$ (b) $\bar{\mathbf{t}}^r$ (c) $\dot{\mathbf{t}}^r$ (d) $\ddot{\mathbf{t}}^r$

Figure 21: (a) Sampled tensor-border $\mathbf{t}^r$ of degree 5 (b,c,d) piecewise C^2 tensor-borders of degree 3.

are similar: the light-gray underlaid part of $\mathbf{x}^r$ is defined by C^2-connection to the tensor-border $\bar{\mathbf{t}}^{r+1}$. When $\sigma := \frac{1}{2}$ the dark-gray part of $\mathbf{x}^r$ is a tensor-border $\dot{\mathbf{t}}^r$ while for $\sigma := \frac{3}{4}$, it is $\ddot{\mathbf{t}}^r$. The macro-patches adjacent to sector separating lines are completed from their boundaries (dark- and light-gray) by $CS2 - rule(1)$ when $\sigma := \frac{1}{2}$ and $CS2 - rule(\frac{1}{4})$ when $\sigma := \frac{3}{4}$. The inner macro-patch is obtained by tensoring the *CS2-rule(1)* when $\sigma := \frac{1}{2}$ and the *CS2-rule*$(\frac{1}{4})$ when $\sigma := \frac{3}{4}$.

5.1. Implementation

The main novelty compared to standard bi-3 subdivision is the computation of the B^r. The corresponding formulas are tabulated

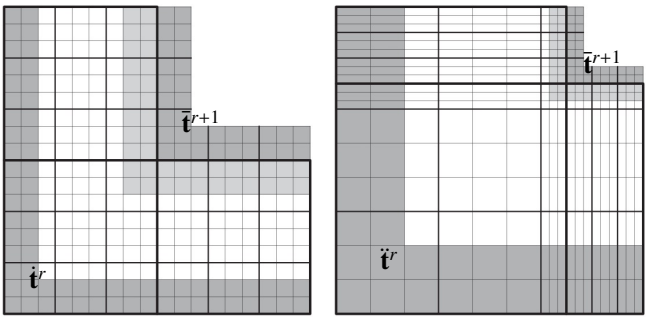

(a) $\sigma := \frac{1}{2}$ (b) $\sigma := \frac{3}{4}$

Figure 22: Construction of the rings $\mathbf{x}^r$.

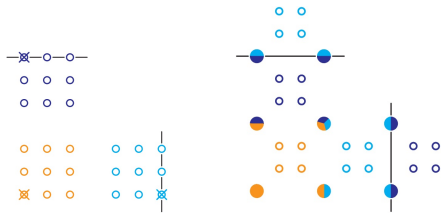

(a) (b)

Figure 23: (a) The sampled bi-3 corner jets after scaling (scaled re-sampling). (b) Averaging B-spline control points.

for *one* sector for each scheme and valence n. For $S^2_{3,12}$ and $\hat{S}^2_{3,12}$ the formulas for the 'inner' B-spline control points orange, orange-cyan, orange-blue and orange-cyan-blue bullets and for two points of each the sector separating line (cyan and blue B-spline control points before averaging) are collected in a 8×6^2 matrix that expresses the 8 points as affine combinations of the BB-coefficients of one sector of the guide $\mathbf{g}$. (4 of the pre-calculated points are subsequently averaged.)

Since the BB-coefficients marked as $\circ$ in Fig. 21b can be obtained by the symmetric C2-rule (using BB-coefficients of adjacent $\bar{\mathbf{t}}$ sectors), for $S^2_{3,27}$ and $\hat{S}^2_{3,27}$ we pre-calculate (the 18 black circles of) $\bar{\mathbf{t}}$ as a matrix of size 18×6^2. In summary, we tabulate the guide, refinement of one sector by de Casteljau's algorithm and the last-mentioned control points so that the algorithm reduces to matrix multiplication.

6. The central caps

Since applications rarely require more than 6 refinement steps, we complete the presented guided subdivision surfaces by a cap that does not compromise the good highlight line distribution of the main surface, even under magnification. We call this cap *tiny* since it fills a hole left after s subdivision steps and s is the maximal refinement needed for a given application. For low valence and double speed the radius of an enclosing sphere is approximately 4^{-s} times the radius of the sphere enclosing the initial hole. The tiny cap also samples the guide via a reparameterization that determines the cap's structure.

Fig. 24 shows the layout of bi-3 and bi-4 caps: each sector is covered by a 2×2 or 3×3 macro-patch, internally C^1 and G^1-connected to the adjacent macro-patches. The caps have a well-defined curvature at the central point. But while the bi-4 cap is C^1-connected to the last guided ring, the bi-3 cap formally is only C^0-connected to the main surface. As [17] showed C^0 continuity does not prevent high surface quality. Indeed, experiments indicate that a slight deviation from C^1 continuity is likely necessary for high quality *finite bicubic* surfaces. Various formally G^1 finite bicubic constructions lead to far worse outcomes.

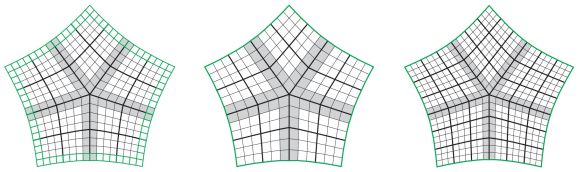

(a) bi-4 cap (b) 2×2 bi-3 cap (c) 3×3 bi-3 cap

Figure 24: The gray underlaid BB-coefficients are involved into G^1 join of adjacent sectors. The green underlaid BB-coefficients are involved into C^0 (bi-3) and C^1 (bi-4) join to main body of a surface.

7. Examples and Comparisons

In this section we compare the new guided subdivision surfaces with classical Catmull-Clark-subdivision [5], as well as its curvature-bounded variant [8]. The comparisons confirm that both Catmull-Clark and [8] subdivision surfaces have problematic highlight line distributions. For Catmull-Clark we observe the typical pinching of highlight lines near irregularities. Tuned curvature-bounded variants such as [9, 18, 19] are aimed at improving the limit shape not the shape of the larger neighborhood. Their characteristic maps hint at slow convergence and artifacts in the larger neighborhood.

Of the published guided subdivision algorithms, we only compare to one with singularly parameterized caps [14] since for [16, 3] the differences in highlight line distributions and curvature are minor. Compared to [16, 3] the advantages of new approach are efficiency and suitability for engineering analysis. Increased stencil complexity is the price all guided subdivision algorithms pay for increase surface quality.

In all examples the input bi-3 ring is green and the quality is measured by the designer-preferred standard of highlight lines [11]. Fig. 25 and Fig. 26 compare bi-cubic constructions to Sabin's scheme for low-valence input $n = 3, 5$. Despite fewer rings, guided subdivision has a better highlight line distribution. Fig. 27 additionally looks at Gauss curvature. Fig. 27g shows the last guided ring and the cap of $\hat{S}^2_{4,3}$. To be able to compare, we generate additional rings Fig. 27e for the slower $S^2_{3,12}$ and many more rings to make the remaining hole sufficiently small to display the surface of [8].

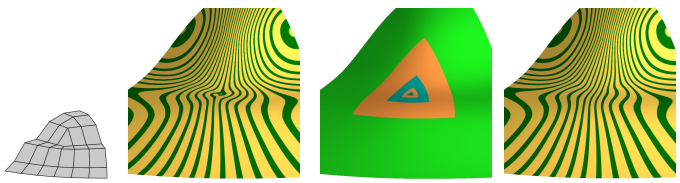

(a) $n = 3$ **c**-net (b) [8] (c) layout (d) $S^2_{3,12}$

Figure 25: (a) Input **c**-net. (b) The highlight lines [8] surface show shape artifacts similar to Catmull-Clark-subdivision. (c,d) Structure and highlight lines of $S^2_{3,12}$: four guided rings + tiny red bi-3 cap.

Fig. 28 shows how bicubic constructions are more sensitive to inter-ring smoothness than biquartics. While the bi-4 surfaces

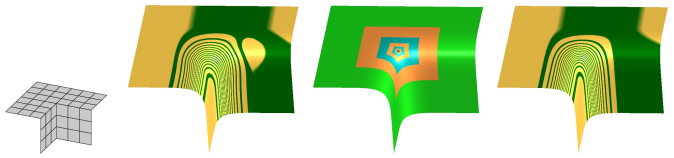

(a) $n = 5$ (b) [8] (c) $S^2_{3,12}$

Figure 26: (a) Input net (classical two-beam corner). (b) While the limit is better than that of Catmull-Clark-surfaces, [8] neglects the transition ring. This results in poor highlight line distribution. (c) Layout and highlight lines of $S^2_{3,12}$ with six guided rings + tiny red bi-3 cap.

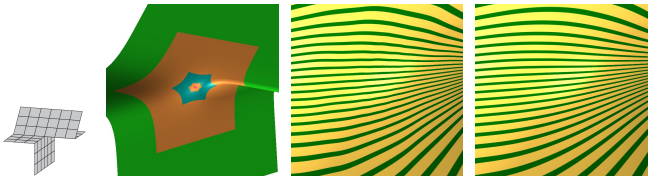

(a) $n = 6$ (b) layout (c) $\hat{S}^2_{3,27}$ (d) $\hat{S}^2_{4,3}$

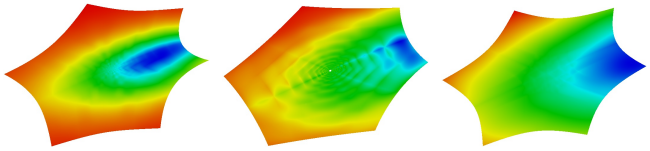

(e) $S^2_{3,12}$ (f) [8] (g) $\hat{S}^2_{4,3}$

Figure 27: (a) Input **c**-net. (b) Structure of $\sigma := \frac{3}{4}$ surfaces: three guided rings + tiny red cap. (c,d) Highlight lines on last two rings + cap. *bottom* row: Gauss curvature.

$\hat{S}^1_{4,3}$ and $\hat{S}^2_{4,3}$ (respectively C^1 and C^2) are visually indistinguishable, a C^1 bi-3 construction (analogous to $S^1_{4,3}$) with 3 patches per sector is clearly inferior to the C^2 construction $\hat{S}^2_{3,27}$ using 27 patches.

Compared to our constructions, the algorithm in [14] (see Fig. 29c,d) requires additional pre-calculation and storage since its last two rings are special. To hide the singular parameterization of the surface cap, the algorithm in [14] generates several bi-3 guided rings contracting at the standard Catmull-Clark-speed Fig. 29d. At the expected maximal refinement, the highlight line distribution appears visually artifact-free Fig. 29c. Only extreme zoom reveals slight oscillations Fig. 29d compared to the new degree 3 variant Fig. 29e. Fig. 30 shows that guided subdivision copes well with complex geometry and high valence.

8. Discussion: C^1 vs C^2, bi-3 vs bi-4

Sampling a locally reparameterized guide of good shape yields high quality refinable surfaces of low degree, namely the bicubic and biquartic C^2 subdivision surfaces whose limit at the extraordinary point is C^1 and curvature bounded. The highlight lines of the new C^1 bi-4 subdivision surfaces are difficult to distinguish from their C^2 counterparts; and completing the subdivision by tiny G^1-C^1 bi-4, respectively or G^1-C^0 bi-3 caps preserves the quality even under zoom. This observation agrees with [20] where it was pointed out that since curvature behaviour is not central to many practical measures of surface quality, the stringent requirements for curvature continuity can be relaxed. The extra freedom can be used to improve practical measures.

(a) $n = 7$ (b) layout (c) Catmull-Clark

(d) C^1 bi-3 (e) $\hat{S}^2_{3,27}$ (f) $\hat{S}^1_{4,3}$

Figure 28: (a) Input **c**-net, $n = 7$. (b) $\sigma := \frac{3}{4}$, four rings + red tiny cap. (*bottom* row) Highlight lines of the last three rings + cap of (c) C^1 bi-3, (3 patches per sector); (d) C^2 bi-3, 3×3 macro-patches; (e) C^1 bi-4, (3 patches per sector).

Compared to pre-existing bi-3 subdivision algorithms the highlight line distribution of the new bi-3 guided constructions is clearly better. This may be due to several reasons. First, tuning conventional subdivision (varying input parameters) is aimed at improving only the limit behaviour without consideration of the neighborhood. Second, guided subdivision focuses on shape without concern of the mathematical properties of the limit – since the analytical properties of the limit are a priori known from subdivision theory [21]. Third, unlike algebraic tuning, guide sampling is geometrically predictable. This includes the desirable control over the contraction speed.

Some applications, for example implementing iso-geometric engineering analysis, favor a uniform degree throughout. For these purposes, the new degree bi-3 surface constructions represent a step forward. However, the experiments indicate that smoothly connected bi-3 B-spline surfaces are more constrained in multi-sided configurations; and the systematic comparison in this paper shows that even 3×3 macro-patches of degree bi-3 fail to match the quality of a single bi-4 patch. This is especially true when the convergence is accelerated to reduce the overall number of patches. Bi-4 surfaces yield both simpler algorithms and better highlight line distributions. C^1 and C^2 guided bi-4 surfaces are remarkably similar: the highlight lines of $\hat{S}^2_{4,3}$ and $\hat{S}^1_{4,3}$, shown in Fig. 2e, are visually indistinguishable. By contrast, in bi-3 subdivision, replacing C^2 (hence 27 patches per sector, cf. Fig. 4c, 2d) with C^1-connected rings (hence 3 patches) drastically decreases surface quality (see Fig. 28).

9. Conclusion

A side-by-side comparison, showed that multi-sided holes can be filled by polynomial rings of degree bi-4 with good highlight line distribution quality even when the number of rings is decreased by increasing the 'contraction speed'. Extensive experiments show double-speed contraction to typically be the maximal speed to retain good shape. Moreover, when reducing the formal mathematical smoothness of the patchwork to C^1, shape and eigenstructure of the guide are preserved. The comparison also indicated that exclusively bi-3 surfaces have notably

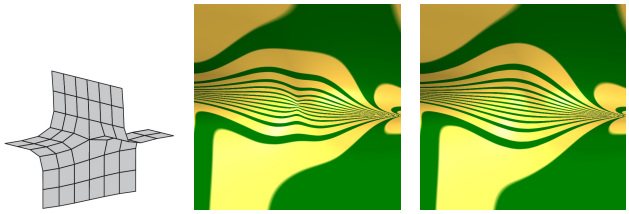

(a) $n = 6$ **c**-net (b) Catmull-Clark (c) $\hat{S}^2_{4,3}$, [14]

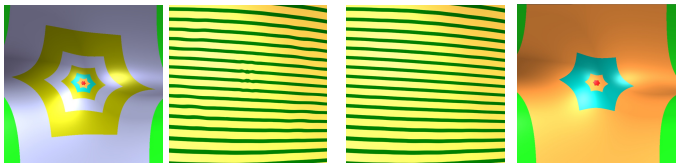

(d) [14], zoom (e) $\hat{S}^2_{3,27}$, zoom

Figure 29: (a) Input **c**-net, $n = 6$. (b) pinched highlight lines characteristic of Catmull-Clark subdivision. (c) At this resolution the highlight line distribution of $\hat{S}^2_{4,3}$ and the construction in (d) look identical. (d) 4 rings of $S^2_{3,27}$ followed by two rings and a C^1 red cap as in [14]. (e) $\hat{S}^2_{3,27}$: 3 rings + C^0 connected internally G^1 3×3 bi-3 red cap. (d,e) zoom to $\hat{S}^2_{3,27}$ + cap.

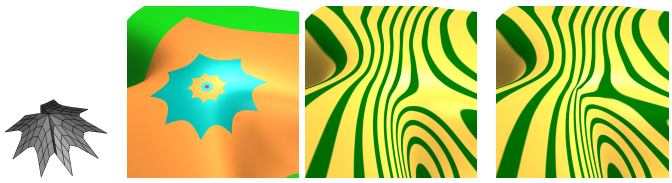

(a) $n = 9$ (b) $\hat{S}^2_{4,3}$ (c) Catmull-Clark

Figure 30: The pinching of highlight lines for Catmull-Clark is absent for $\hat{S}^2_{4,3}$ (4 guided rings + cap).

poorer highlight line distributions. Bi-4 construction is therefore likely the least-degree to provide curvature-bounded subdivision surfaces with consistently visually non-oscillating highlight line distributions. The reasons for the strong quality decrease between bi-3 and bi-4 surfaces are only partly understood. Both bi-4 and bi-3 guided subdivision with adjustable contraction speed have linear and stationary subdivision rules that can be written as large subdivision stencils. The cost of a large stencil over traditional small-stencil subdivision is offset by much improved surface quality.

Acknowledgements This work was supported in part by DARPA HR00111720031 and NIH R01 EB018625.

References

[1] K. Karčiauskas, J. Peters, Refinable bi-quartics for design and analysis, Computer-Aided Design (2018) 1–10.
[2] K. Karčiauskas, J. Peters, Rapidly contracting subdivision yields finite, effectively c^2 surfaces, Computers & Graphics (2018) 1–10.
[3] K. Karčiauskas, J. Peters, A new class of guided C^2 subdivision surfaces combining good shape with nested refinement, Computer Graphics Forum 37 (2018) 84–95.
[4] C. T. Loop, S. Schaefer, G ^{2} tensor product splines over extraordinary vertices, Comput. Graph. Forum 27 (5) (2008) 1373–1382.
[5] E. Catmull, J. Clark, Recursively generated B-spline surfaces on arbitrary topological meshes, Computer-Aided Design 10 (1978) 350–355.
[6] A. Riffnaller-Schiefer, U. H. Augsdörfer, D. W. Fellner, Isogeometric shell analysis with nurbs compatible subdivision surfaces, Applied Mathematics and Computation 272 (2016) 139–147.
[7] K. Bandara, F. Cirak, Isogeometric shape optimisation of shell structures using multiresolution subdivision surfaces, Computer Aided Design 95 (2018) 62–71.
[8] M. Sabin, Cubic Recursive Division With Bounded Curvature, Academic Press, 1991, pp. 411–414.
[9] T. J. Cashman, U. H. Augsdörfer, N. A. Dodgson, M. A. Sabin, Nurbs with extraordinary points: high-degree, non-uniform, rational subdivision schemes, ACM Trans. Graph 28 (3) (2009) 46:1–46:9.
[10] T. J. Cashman, Nurbs-compatible subdivision surfaces, Ph.D. thesis, University of Cambridge, UK, british Library, EThOS (2010).
[11] K.-P. Beier, Y. Chen, Highlight-line algorithm for realtime surface-quality assessment, Computer-Aided Design 26 (4) (1994) 268–277.
[12] G. Farin, Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide, Academic Press, San Diego, 2002.
[13] H. Prautzsch, W. Boehm, M. Paluszny, Bézier and B-spline techniques, Springer Verlag, 2002.
[14] K. Karčiauskas, J. Peters, Improved shape for refinable surfaces with singularly parameterized irregularities, Computer Aided Design 90 (2017) 191–198.
[15] K. Karčiauskas, J. Peters, Improved shape for multi-surface blends, Graphical Models 8 (2015) 87–98.
[16] K. Karčiauskas, J. Peters, Concentric tesselation maps and curvature continuous guided surfaces, Computer Aided Geometric Design 24 (2) (2007) 99–111.
[17] K. Karčiauskas, J. Peters, Can bi-cubic surfaces be class A?, Computer Graphics Forum 34 (5) (2015) 229–238.
[18] T. Ni, A. H. Nasri, Tuned ternary quad subdivision, in: Geometric Modeling and Processing, 2006, pp. 441–450.
[19] U. H. Augsdorfer, N. A. Dodgson, M. A. Sabin, Tuning subdivision by minimising gaussian curvature variation near extraordinary vertices, Computer Graphics Forum 25 (3) (2006) 263–272.
[20] M. A. Sabin, N. A. Dodgson, M. F. Hassan, I. P. Ivrissimtzis, Curvature behaviours at extraordinary points of subdivision surfaces, Computer-Aided Design 35 (11) (2003) 1047–1051.
[21] J. Peters, U. Reif, Subdivision Surfaces, Vol. 3 of Geometry and Computing, Springer-Verlag, New York, 2008.
[22] K. Karčiauskas, J. Peters, Adjustable speed surface subdivision, Computer Aided Geometric Design. 26 (2009) 962–969.

Appendix: Analysis of the Subdivision Scheme

First we observe that the guide **g** is a superposition of homogeneous functions. The Bernstein polynomials $B^d_{0jk} := \binom{d}{0jk} l_1^j l_2^k$ are homogeneous of degree d; for any λ, $B^d_{0jk}(\lambda x) = \lambda^d B^d_{0jk}(x)$. Consequently $B^d_{0jk} \circ L$, of degree $j \times k$, is homogeneous of degree d and for any λ,

$$(B^d_{0jk} \circ L)(\lambda u, \lambda v) = \lambda^d (B^d_{0jk} \circ L)(u, v).$$

Homogeneity is not affected by constraints (4–6). Therefore, if we set one unconstrained BB-coefficient to 1 and the others to 0, the resulting **g** is homogeneous in all sectors. This yields an explicit basis of homogeneous functions of the guide **g**. We count 6 functions of quadratic expansion, $6n$ functions that are non-zero on two adjacent sectors and n groups of 16 functions that are non-zero in only one sector (with the superscript indicating the degree, $(d, ijk) \sim B^d_{ijk} \circ L$: $(4, 022)$, $(5, 032)$, $(5, 023)$, $(6, 042)$, $(6, 033)$, $(6, 024)$, $(7, 052)$, $(7, 043)$, $(7, 034)$, $(7, 025)$, $(8, 053)$, $(8, 044)$, $(8, 035)$, $(9, 054)$, $(9, 045)$, $(10, 055)$). Scaling the arguments of **g** by any λ therefore scales n_ν many basis functions by λ^ν:

$$\begin{array}{lcccccccccccc} \nu = & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 \\ n_\nu = & 1 & 2 & 3 & 2n & 3n & 4n & 4n & 4n & 3n & 2n & n \end{array} \quad (8)$$

The rings $\mathbf{x}^r$ are fully defined by the guide **g** and inherit the guide's homogeneous decomposition. Here $r \geq 2$ of $S^2_{3,12}$ and

$\hat{S}^2_{3,12}$ and $r \geq 1$ for $S^2_{3,27}$, $\hat{S}^1_{4,3}$ and $\hat{S}^2_{4,8}$. We denote by $\mathbf{x}^r_{d,p}$ the rth ring obtained by applying the constructions of Section 5 to the pth homogeneous function $\mathbf{f}_{d,p}$ of degree d. By construction, $\mathbf{x}^r_{d,p} := (\lambda^d)^{r-2}\mathbf{x}^2_{d,p}$.

Since the linear combination of the two homogeneous eigenfunctions of degree 1 reproduce χ and since χ is injective (see [21] for original Catmull-Clark and [22] for general σ), the eigen-spectrum implies that the bi-3 guided subdivision surfaces are C^1 and their curvature at the extraordinary point is bounded.

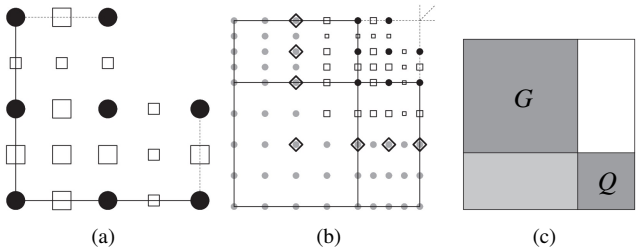

Figure 31: (a,b) Geometric background of the subdivision matrix. (c) The blank sub-matrix is zero.

Properties of $\hat{S}^2_{4,3}$ Two groups of control points define $\hat{S}^2_{4,3}$. Group Q, has $6n$ independent BB-coefficients marked as black diamonds in Fig. 31b. Group G has $22n + 6$ independent BB-coefficients defining the guide $\mathbf{g}$. (Recall that the black bullets in Fig. 31a,b are averaged samples of the guide $\mathbf{g}$. Fig. 31b recalls the construction of $\tilde{S}^2_{4,3}$ and Fig. 31a displays the tensor-border $\tilde{\mathbf{t}}$ constructed in the previous step.) Since some entries of Q depend on G, some entries of the subdivision matrix are dependent. (This dependence between the groups should not be conflated with the linear independence of the degrees of freedom for modeling or engineering analysis that each new ring provides.) Independence of entries in the subdivision matrix is fortunately not a requirement for subdivision theory [21]. We only need stationary refinement rules, i.e. an unchanging subdivision matrix. Indeed the refinement rules are unchanging. (i) de Casteljau's algorithm shrinks the domain of $\mathbf{g}$ so defining a finer sub-matrix G^+. (ii) The Q (diamonds in Fig. 31b) are obtained by splitting the $\tilde{\mathbf{t}}$ constructed in the previous step, see Fig. 31a. (iii) The averaged sampling from $\mathbf{g}$ defining the black bullets and the completion of $\tilde{\mathbf{t}}$ are always the same. That is, Q^+ is always derived from G and Q by unchanging rules, i.e. the subdivision is stationary. (Linearly dependent G and Q are introduced only to prove C^1 continuity at the limit point. Engineering analysis does not use these groups but linearly independent eigen-functions described earlier in this Appendix.)

Fig. 31c shows the block structure of the matrix. The sub-matrix for group G has the eigenstructure of $\mathbf{g}$ and hence the eigenvalues listed in (8). Since the upper right $6n \times (22n + 6)$ sub-matrix is zero, the lower left part of the matrix is irrelevant for calculating the eigenvalues. The eigenvalues of the $6n \times 6n$ sub-matrix corresponding to Q have absolute value less than λ^2. The linear homogeneous functions of the guide determine the eigenfunctions corresponding to the eigenvalue λ and the construction reproduces the characteristic map χ. Together this implies that $\hat{S}^2_{4,3}$ surfaces are C^1 and curvature bounded at the extraordinary point.