

Efficient Shape-Improved Bi-3 Subdivision Surfaces

Kestutis Karčiauskas¹ and Jörg Peters²

¹Vilnius University, Lithuania ²University of Florida, USA

Abstract

Efficient Shape-Improved Bi-cubic Subdivision (ESIS) is a new curvature-bounded bi-cubic subdivision scheme with the same minimal patch layout as Catmull-Clark surfaces, i.e. $3n$ patches per subdivision surface ring, but generating far better shape and using fewer pieces to cover the same n -sided surface region. Fewer pieces means: faster downstream graphics processes and rendering, smaller files for data exchange, and fewer equations for computing on surfaces and simulation. The key technical innovations are: a non-uniform refinement and the relaxing of formal smoothness of subdivision surface rings while preserving good highlight line distributions.

Keywords: subdivision surface, bounded limit curvature, fast convergence

1. Introduction

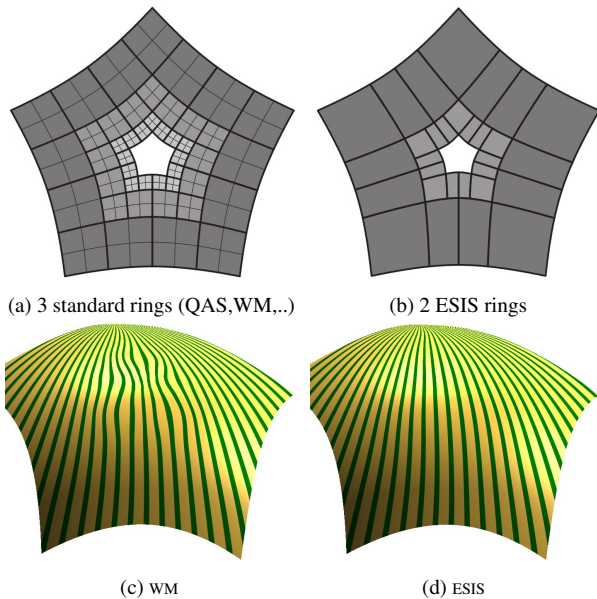


Figure 1: Hole filling by a sequence of nested bi-cubic subdivision surface rings. (a) 3 standard surface rings of current high-end subdivision requiring $12n$ bi-cubics per ring (QAS = [KP24], WM = [WM23]) vs. (b) 2 ESIS rings with $3n$ bi-cubics per ring. Note the non-uniform aspect ratio of the middle patches. The highlight line distribution on equally fast converging (c) WM vs. (d) ESIS.

Subdivision surfaces model the neighborhood of an extraordi-

nary point, where $n \neq 4$ feature merge, as a sequence of contracting surface rings, see Fig. 1 or see the formal exposition by Peters and Reif [PR08]. While a Catmull-Clark Subdivision (CC) ring consists of $3n$ bi-cubic pieces, more recent subdivision algorithms, [KP23, KP24, WM23, WM24], reviewed in Section 2, use $3n$ macro-patches consisting each of 2×2 bi-cubic patches; that quadruples the number of patches. Despite this quadrupling, filling of the hole by these new, higher-quality subdivision surfaces is as slow as for CC. To measure hole filling, we define, for a given subdivision algorithm A as $v_n(A)$ the number of patches to shrink a flat n -sided hole to a prescribed size. We can then compare two algorithms A and B via the coverage ratio $r_n(A, B) := v_n(A)/v_n(B)$. For example, Fig. 1 a,b indicates $r_5(QAS, ESIS) > 6$ since 36 QAS patches per sector [KP24] fill the hole a little less than 6 ESIS patches. Table 3 confirms $r_n(QAS, ESIS) > 6$ for $n = 3, 5, 6, 7, 8, 9, 10$. Faster hole-filling not only decreases the number of refinements needed to visually cover the n -sided hole for more efficient rendering, but reduces data exchange of models and cost of simulation, say when computing integrals up to tolerance for engineering analysis. Manipulating the eigen-spectrum can speed up hole filling and was advertised as improving convergence for engineering analysis on subdivision surface using subdivision elements [WM23]. However, as pointed out, e.g. by Wang and Ma [WM24], the resulting surface shape suffers from this algebraic approach, see the oscillations of the highlight lines in Fig. 1 (c) vs. (d).

Here we propose a different approach, that prioritizes uniform distribution of highlight lines [BC94] over formal, infinitesimal C^2 continuity. This approach is motivated in part by Karciauskas and Peters [KP15] who demonstrated that class A surface quality in the large can be achieved without exact, formal differential smoothness. We instead increase hole-filling contraction by non-uniform refinement and knot spacing. The resulting Efficient

Shape-Improved Bi-cubic Subdivision, short ESIS, has the following properties.

- ESIS generates high-quality surfaces, only slightly worse than the more complicated bi-cubic surface quality-leader QAS [KP23], despite the surface rings being formally less smoothly joined.
- ESIS can be used as a natural companion to QAS by replacing the QAS matrix with the ESIS matrix that maps control nets to finer control nets. This yields a more economical structure after QAS has resolved any highly challenging initial configurations. (By contrast, replacing QAS by Catmull-Clark after the first step(s) yields deficient shape.)
- ESIS yields curvature-bounded surfaces, i.e. formally maximal smoothness at the extraordinary point for bi-3 subdivision.
- ESIS has the same, as-simple-as-possible layout (number of pieces) as Catmull-Clark Subdivision surfaces and can replace Catmull-Clark Subdivision simply by switching out the subdivision matrices, resulting in considerably faster hole filling.

The good shape is due to a guiding surface that is hidden in the derivation and baked into the subdivision matrix. The increased contraction speed is illustrated in Fig. 1 for $n = 5$: three rings of standard (CC-like) contraction fill the hole to a *lesser extent* than two rings of ESIS. Indeed, for $n = 3, 5, \dots, 10$, the hole left in the center after three standard refinement steps is larger by a factor of 1.138 1.128 1.132 1.136 1.139 1.139 1.143 than the hole left after two steps of ESIS. This faster filling by ESIS while retaining good shape is possible due to relaxation of formal smoothness, the remarkable property established in Theorem 1 in Section 5 that the dominant eigen-rings of ESIS join C^2 , and two technical innovations:

- ESIS refinement is derived via a non-uniform characteristic map
- ESIS pioneers refinement based on a mixture of Bézier (not B-spline) control-nets in total degree and tensor-product form.

2. A brief review of recent subdivision algorithms

We focus on bi-cubic subdivision schemes operating on quad meshes. Two possible layout structures of such rings are displayed in Fig. 2 a, for Catmull-Clark Subdivision [CC78] (CC), and in Fig. 2 b, for recent curvature-bounded subdivision such as [MM18]. The latter quadruples the number of patches of the simple(st) layout of Catmull-Clark Subdivision. Catmull-Clark Subdivision is known to suffer from artifacts introduced by the first refinement steps [ADS11]. Mending these artifacts in the limit is the goal of ‘tuning’ by choosing more complex rules in view of the dominating eigen-components of the subdivision matrix [MM18, WM23]. However this focus on the limit easily results in poor surfaces in the neighborhood of the




Figure 2: Subdivision surface rings.

limit extraordinary point, as shown by Karciauskas and Peters [KP22, Fig 1d] and argued in [WM24], where additional operators are applied to smooth out the mesh. Indeed, the fast converging method by Wang and Ma [WM23] results in reduced surface quality.




Figure 3: Guided subdivision separates shape-finding from the mathematical constraints of the final output surface.

Guided Subdivision [KP07] first determines a separate shape, that is sampled to complete to a surface ring the Hermite extension from the surrounding surface, see Fig. 3. Specifically (a) a c-net, cf. Section 3, provides both a surface ring and a guide shape. These pieces of geometry are close but do not join consistently as the ‘pixel-fighting’ in (b) illustrates. Instead Guided Subdivision borrows the smooth extension or prolongation from the existing surface and supplements it with data from the guide to complete (c) an new (bronze) ring that smoothly joins the existing (green) ring and follows the shape of the guide by sampling the guide for data needed to complete the new ring. This approach generically yields excellent highlight lines, see (d). Combining the shape construction and the extension from the previous ring into a single operation, Evolving Subdivision [KP22] generates a new surface ring exactly as for classical subdivision – but of better shape. Upgrading the approach to the quadratic-attraction principle introduced in the shape-quality leader, bi-cubic QAS [KP23], [KP24] achieves uniform contraction speed corresponding to a subdominant eigenvalue $\lambda := \frac{1}{2}$ while preserving good shape.

Yet both tuned bi-cubic subdivision schemes and modern guided subdivision of degree bi-3 quadruple the number of polynomial pieces. Tuning algorithms additionally require that the irregular points are widely separated. This can be achieved by applying Catmull-Clark Subdivision but these initial refinements can already introduce noticeable shape artifacts, such as pinched highlight line distribution [BC94]. ESIS keeps the simple layout of Catmull-Clark Subdivision and accelerates surface covering.

2.1. Extraordinary vertices or faces are not rare

Looking at already subdivided meshes, one may form the impression that the fraction of extraordinary vertices (EVs) or extraordinary faces (EFs) in real life subdivision models is low. However, quite to the contrary, numbers in Table 1 from [TPC*10] and [HZN*18], on practical quad mesh simplification show EVs and

EFs to be quite prominent (see Section 2.1.1) and carefully artist-curated models can even have higher percentage of non-regular sub-nets (see Section 2.1.2).

2.1.1. Subdivision meshes from algorithmic meshing

The exact number EVs/EFs in subdivision meshes generated by algorithmic quad-meshing depends on the topology of the object being meshed, the meshing algorithm and parameters, and error bound(s), target number of facets, and more. Some older algorithms [TPC*10, Table 1] and [HZN*18] report on practical, legacy subdivision meshes. This is excerpted in Table 1. [LWG*25, Table

model	%	model	method	#sing
Moai	56	David	[ACSD*03]	10310
Pensatore	52	Rockerarm	[JTPSH15]	2708
Gargoyle	54	Fandisk	[MK06]	117
Igea	52	Bunny	[TPC*10]	3438
Fertility	53			
Rampart	52			

Table 1: (left) from [TPC*10, Table 1] listing % of not non-regular nodes, (right) from [HZN*18, Table 2] listing the number of singularities.

1] reports a quad-ratio of less than 80% (theirs) and less than 67% otherwise for recent machine-learning approaches, i.e. at least 20%, resp. 33% irregularities. Polycube meshes [THCM04] often have a large percentage of 3 and 5-valent vertices. The meshes generated by Topology Optimization or Generative Design have such a high percentage of extraordinary points, of any valence, that they need to be re-meshed. Recent algorithms [PNA*21, CSH*25] provide, at their coarsest level, roughly 10% EVs (parametrization singularities, cones), see Table 2.




Table 2: Low-poly EVs

model	EVs	Verts
kitten	28	138
fandisk	34	302
indorelax	52	465
dragon	1271	3160

Figure 4: Highlighted EVs and non-quads of Suzanne

2.1.2. Artist-built meshes

The hand-optimized model of ‘Geri’ in Pixar’s short film classic [Pin97] has regions that are almost entirely irregular, both in valence and by inserting triangles. An estimate is that ‘Geri’ is at least 40% irregular. ‘Suzanne’, the Blender model, has 112 irregular vertices among 507 total and 32 of Suzanne’s 50 facets are non-quad, cf. Fig. 4. A well-known ‘Monster Frog’ model has 292 irregular vertices out of 1308. A sample of human quad models with low polygon count for games at ‘turbosquid’ [Men] shows more than 30% irregular nodes.

2.2. Practical relevance of the cost of coverage ratio

Subdivision is now available in industrial design tools such as Rhino3D [M*10]. However, due to imperfect shape of the standard Catmull-Clark subdivision [CC78], see e.g. [KPR04], such models currently serve only as advanced design sketch, and are replaced by NURBS models for downstream use. With ESIS addressing the shape flaws, high-quality capping – by finitely many Bézier patches after a few subdivision steps – provides a transition to CAD tools. However, a high cost of coverage complicates such interchange for all irregular neighborhoods and is very noticeable in the length of the verbose output .step format CAD files. If the geometry is used for engineering analysis as illustrated in Section 7 a factor $r_n(QAS, ESIS) > 6$ means that the most costly part of Galerkin’s Method, computing integral inner products to fill the stiffness matrix, is more than $6\times$ as costly for QAS than for ESIS.

3. Control nets and polynomial pieces

We will refer to three *regular* types of meshes that differ by the interpretation of their nodes/vertices: the B-net, the BB-net and the de Casteljau-net.

- *B-net* vertices are interpreted as control points \mathbf{c}_{ij} of uniform bi-cubic B-splines.

With $N^3(t)$ the uniform degree 3 B-spline [dB78] and $\mathbf{c}_{ij} \in \mathbb{R}^3$ the *B-spline control points*, connecting \mathbf{c}_{ij} to $\mathbf{c}_{i+1,j}$ and $\mathbf{c}_{i,j+1}$ wherever possible yields the *B-net* of the piecewise polynomials

$$\mathbf{f}(u, v) := \sum_{i=0}^3 \sum_{j=0}^3 \mathbf{c}_{ij} N^3(u-i) N^3(v-j).$$


A *c-net* generalizes the B-net by allowing (isolated) points with $n \neq 4$ neighbors. Fig. 5 a displays as solid lines a *c-net* with a central extraordinary node of valence $n = 5$, and an extension, to a net $\check{\mathbf{c}}$, by adding a layer of quadrilateral facets. Interpreting any (overlapping) 4×4 subnet of $\check{\mathbf{c}}$ as a set of (uniform bi-cubic tensor-product) B-spline control points defines a bi-cubic spline ring, green in Fig. 5 b,

• *BB-net* vertices are interpreted as the coefficients \mathbf{b}_{ij} of bi-cubic pieces in tensor-product Bernstein-Bézier form (BB-form);


With $B_k^d(t) := \binom{d}{k} (1-t)^{d-k} t^k$ the Bernstein polynomials of degree d and $\mathbf{b}_{ij} \in \mathbb{R}^3$ the *BB-coefficients*, connecting \mathbf{b}_{ij} to $\mathbf{b}_{i+1,j}$ and $\mathbf{b}_{i,j+1}$ wherever possible yields the *BB-net* [dB87, Far88] of the polynomial

$$\mathbf{f}(u, v) := \sum_{i=0}^3 \sum_{j=0}^3 \mathbf{b}_{ij} B_i^3(u) B_j^3(v), \quad 0 \leq u, v \leq 1.$$

The BB-form and the B-spline form provide alternative polynomial bases for surfaces and the formulas $\mathbf{b}_{3i} := (\mathbf{c}_{i-1} + 4\mathbf{c}_i + \mathbf{c}_{i+1})/6$, $\mathbf{b}_{3i+1} := (2\mathbf{c}_i + \mathbf{c}_{i+1})/3$ and $\mathbf{b}_{3i+2} := (\mathbf{c}_i + 2\mathbf{c}_{i+1})/3$ convert the univariate B-form to the BB-form.



(a) $\mathbf{c}, \check{\mathbf{c}}$



(b) surface ring

Figure 5: B-spline-like control net $\check{\mathbf{c}}$ for valence $n = 5$.




Figure 6

Tensoring (applying in each variable) the univariate rules yields the (bi-variate) *B-to-BB conversion*. The C^1 Hermite extension of the surface ring in Fig. 5 b yields the gray underlaid portion of the partial BB-net in Fig. 6, called *first-order tensor-border*.

- *de Casteljau-net* vertices are interpreted as the coefficients \mathbf{q}_{ij} in total degree Bernstein-Bézier form.

With $B_{ijk}(u, v) := \frac{d!}{i!j!k!}(1-u-v)^k u^i v^j$, where $d = i + j + k$, the Bernstein-Bézier polynomials of total degree d and $\mathbf{q}_{ijk} \in \mathbb{R}^3$ the *deCasteljau-coefficients*, connecting \mathbf{q}_{ijk} to $\mathbf{q}_{i+1, j, k-1}$ and $\mathbf{q}_{i+1, j-1, k}$ wherever possible yields the *deCasteljau-net* of \mathbf{f} of total degree d with associated functions, see [DC63, Far88],

$$\mathbf{f}(u, v) := \sum_{\substack{i+j+k=d \\ i, j, k \geq 0}} \mathbf{q}_{ijk} B_{ijk}(u, v), \quad 0 \leq u + v \leq 1.$$

Fig. 7 shows in red a pentagon consisting of five *de Casteljau-nets* (of depth 2) forming the control data \mathbf{q} at a central vertex and a surrounding reduced BB-net \mathbf{d} consisting of the gray-underlaid points in Fig. 6. The *QAC control structure* consists of pairs (\mathbf{d}, \mathbf{q}) .

\mathbf{d} denotes a net formed by $8n$ points (marked as \bullet), see Fig. 8 a (essentially the *first-order tensor-border* of Fig. 6). Each group of 8 form a sub-net of a bi-3 patch in BB-form, that defines position and first derivative of a circular sequence of patches. Between adjacent patches of this ring transitions are C^1 after setting the shared points \circ as averages of their two neighbors \bullet . Otherwise, the $8n$ BB-coefficients are unconstrained.

\mathbf{q} is a *quadratic expansion* (marked light-red) at the central point \mathbf{p} of a C^1 quartic patch \mathbf{g}^Δ , see Fig. 25 b. \mathbf{p} and its two neighbors \bullet define the central tangent plane. The remaining n \bullet of \mathbf{q} are unconstrained.

4. Derivation of ESIS subdivision

Fig. 7 shows the control data and their refinement. The control data are initialized according to Section 4.1, refined according to Section 4.2 and interpreted as surface rings according to Section 4.3. Section 4.1 explains how to initialize the pair (\mathbf{d}, \mathbf{q}) from a given standard *c-net* of Catmull-Clark Subdivision. Section 4.2 describes

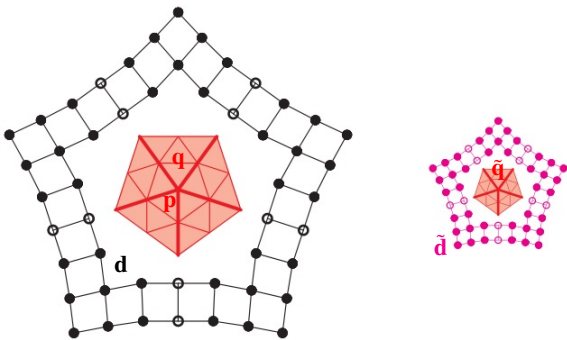


Figure 7: ESIS control data (\mathbf{d}, \mathbf{q}) and its refinement $(\tilde{\mathbf{d}}, \tilde{\mathbf{q}})$. The $8n$ control points \mathbf{d} are marked as \bullet . The BB-coefficients marked \circ are implied by C^1 continuity between sectors.




Figure 8: One bi-3 C^1 ring consisting of $3n$ BB-nets defined by \mathbf{d} and $\tilde{\mathbf{d}}$ (displayed as $8n$ \bullet).

the basic ideas underlying the refinement rules $(\mathbf{d}, \mathbf{q}) \rightarrow (\tilde{\mathbf{d}}, \tilde{\mathbf{q}})$. Since the formulas are linear in (\mathbf{d}, \mathbf{q}) , the relationship can be expressed as

$$(\tilde{\mathbf{d}}, \tilde{\mathbf{q}})^t = M_n(\mathbf{d}, \mathbf{q})^t \quad (1)$$

where n is the valence of the central point and the matrix M_n is of size $m \times m$, $m := n(8 + 5) + 1$ (counting the central point once): per each of the n sectors there are 8 coefficients in \mathbf{d} (that are all independent of each other) and 5 coefficients of \mathbf{q} (that are not independent – but so-listed to reflect symmetry and enable greater flexibility). Section 4.3 presents the construction of C^1 surface rings from \mathbf{d} and its refinement.

4.1. Initialization of \mathbf{d} and \mathbf{q}

B-to-BB conversion of the *c-net* yields in particular the first order tensor-border (gray-underlaid in Fig. 6). Removing the $2n$ BB-coefficients shared by adjacent sectors (implied by the C^1 constraints between sectors) leaves 8 independent BB-coefficients per sector. These $8n$ coefficient form the initial \mathbf{d} . We borrow the initial quadratic expansion $\tilde{\mathbf{q}}$ of [KP23], see the Electronic Supplement, but restate it as \mathbf{q} of total degree 4 by the simple formula (see Fig. 11 for the subscript labeling)

$$\begin{pmatrix} \mathbf{q}_1^5 \\ \mathbf{q}_2^5 \\ \mathbf{q}_3^5 \\ \mathbf{q}_4^5 \\ \mathbf{q}_5^5 \\ \mathbf{q}_6^5 \end{pmatrix} := \frac{1}{12} \begin{pmatrix} 12 & 0 & 0 & 0 & 0 & 0 \\ -3 & 15 & 0 & 0 & 0 & 0 \\ 2 & -10 & 20 & 0 & 0 & 0 \\ -3 & 0 & 0 & 15 & 0 & 0 \\ 2 & -5 & 0 & -5 & 20 & 0 \\ 2 & 0 & 0 & -10 & 0 & 20 \end{pmatrix} \begin{pmatrix} \tilde{\mathbf{q}}_1^5 \\ \tilde{\mathbf{q}}_2^5 \\ \tilde{\mathbf{q}}_3^5 \\ \tilde{\mathbf{q}}_4^5 \\ \tilde{\mathbf{q}}_5^5 \\ \tilde{\mathbf{q}}_6^5 \end{pmatrix}. \quad (2)$$

4.2. Refinement rules

We express the construction process $(\mathbf{d}, \mathbf{q}) \rightarrow (\tilde{\mathbf{d}}, \tilde{\mathbf{q}})$ in symbolic form to arrive at *formulas* in terms of the input nets, see Fig. 7. The formulas are obtained by performing all steps of a geometrically intuitive construction with symbolic control nets and applying a symbolic solver. That is, (\mathbf{d}, \mathbf{q}) are input in symbolic form and, at the end, the formulas of the refined $(\tilde{\mathbf{d}}, \tilde{\mathbf{q}})$ in terms of (\mathbf{d}, \mathbf{q}) form the rows of a matrix M_n . The Electronic Supplement provides the explicit and exact entries of M_n for $n = 3, 5, \dots, 10$.

The construction samples data from a symbolic guiding cap \mathbf{g} that is itself constructed in several steps from the (\mathbf{d}, \mathbf{q}) -data of the previous surface ring. This process is not overly complex but is

lengthly and therefore placed into Appendix A for the interested reader. For example, the BB-net \mathbf{q} enters into the construction of the guide \mathbf{g}^Δ and hence the **Refinement** $(\mathbf{d}, \mathbf{q}) \rightarrow \tilde{\mathbf{d}}$. The formulas are given in Section 1 of the Electronic Supplement.

4.3. C^1 -connected rings of bi-3 patches

The generation of a surface ring in BB-form from two consecutive control nets \mathbf{d} and $\tilde{\mathbf{d}}$ is gratifyingly straightforward by combining the backward C^1 -extension of $\tilde{\mathbf{d}}$ with \mathbf{d} .




Figure 9: Construction of the BB-net of C^1 rings with split ratio $\sigma : 1 - \sigma$. The illustration shows the default $\sigma := \frac{2}{3}$, i.e. a 2:1 split that yields the best results.

Fig. 9 a shows a sector of the input \mathbf{d} -net, slightly extended to its neighbors, and its refinement $\tilde{\mathbf{d}}$. Recall that the sector nets are joined by placing nodes \circ as averages, resulting in C^1 tensor-borders, dark-gray underlaid in Fig. 9 b. As a trade-off between contraction speed and surface quality, we chose $\sigma := \frac{2}{3}$ as the default. Choosing $\sigma < 2/3$ reduces coverage and $\sigma > 2/3$ increases speed at the cost of deteriorating shape. Co-incidentally, for $n = 5, 6$, $\lambda_\sigma \approx 0.4$ and that matches WM, a $12n$ ring layout with poorer shape. The tensor-borders are then, by default, split with ratio 2 : 1 and so match up with the backwards- C^1 -extended tensor-border of $\tilde{\mathbf{d}}$ in the sense that the BB-coefficients satisfy the relation $\diamond := 3 \bullet - 2 \circ$. The constructed bi-3 ring is therefore C^1 -connected to the refined bi-3 ring. Remarkably, the actual smoothness is close to C^2 : Theorem 1 in Section 5 implies that the dominant eigen-component of the rings join C^2 .

5. Limit Analysis

The speed of refinement and hence of hole filling is, up to first order, governed by the contraction of an eigen-surface, called characteristic map. The characteristic map corresponds to the two eigenvectors of the subdominant eigenvalue of M_n , i.e. the next largest to the mandatory largest eigenvalue 1 of M_n that guarantees that the surface neither collapses nor explodes. An essential moment of ESIS is that the symbolic guide \mathbf{g} is sampled by a map χ_σ [KP09] that generalizes χ_{CC} , the characteristic map of Catmull-Clark Subdivision subdivision, by introducing a non-uniform split with ratio $\sigma : 1 - \sigma$ for refinement, i.e. σ determines the ‘speed’ of hole filling. This yields a subdominant eigenvalue

$$\lambda_\sigma := \frac{\bar{\sigma}}{2} \left((1+c)\sigma^2 + 2\bar{\sigma} + \sigma \sqrt{(1+c)((1+c)\sigma^2 + 4\bar{\sigma})} \right), \quad (3)$$

where $\bar{\sigma} := 1 - \sigma$ and $c := \cos \frac{2\pi}{n}$. A larger σ results in a smaller λ_σ , and hence faster contraction. When $\sigma := \frac{1}{2}$ then $\lambda_\sigma = \lambda_{CC}$, the subdominant eigenvalue of CC.




Figure 10: (a) Structure of the subdivision matrix M_n . (b,c,...,h): sector of the ESIS characteristic map χ_E for $n = 3, 5, \dots, 10$.

Balancing between contraction speed and surface quality, we chose $\sigma := \frac{2}{3}$ and this is the value used in Appendix A.

The resulting characteristic maps of ESIS are shown in Fig. 10 b-h. The ESIS subdivision matrix M consists of four sub-matrices, as displayed in Fig. 10 a. One sub-matrix has only zero entries, and $M_d \in \mathbb{R}^{8n \times 8n}$ does not depend on \mathbf{q} . (For comparison, the matrix of QAS is of size $12n \times 12n$.) With the ordering of Fig. 11 , namely

$$\mathbf{p}, \mathbf{q}_2^0, \mathbf{q}_4^0, \mathbf{q}_5^0, \mathbf{q}_5^1, \dots, \mathbf{q}_5^{n-1},$$

the sub-matrix M_q is zero above the main diagonal. The diagonal has entries

$$1, \lambda, \lambda, \lambda^2, \dots, \lambda^{2n}.$$

These entries are the eigenvalues of M_q . Numerical calculation shows that the largest absolute value of eigenvalues of M_d (that do not depend on λ) is less than λ_σ^2 for λ_σ defined by (3). This implies that ESIS generates surfaces with *bounded curvature*.

The first two rows of Table 3 list, for irregular valences up to $n = 10$, the subdominant eigenvalue λ_σ for $\sigma = \frac{1}{2}$ (CC) and for $\sigma = \frac{2}{3}$ (ESIS); the eigenvalues determine the dominant hole-filling contraction. A value greater than 1/2 indicates a slowing down of the hole-filling compared the uniform refinement of the surrounding B-spline surface. The third row of Table 3 measures by how much, after three standard standard refinement steps, the remaining hole is typically larger than the hole after two steps of ESIS.

Remarkably, the geometric derivation of the refinement rules in Appendix A allows reasoning about the characteristic map χ_E without referencing the explicit refinement formulas. The rules leave invariant, except for contraction to the center, a specific input net \mathbf{d} . The surface ring generated from this \mathbf{d} and its refinement according to Section 4.3 can be uniquely identified as defining the characteristic map of ESIS. Theorem 1 not only confirms that ESIS is C^1 in




Figure 11: \mathbf{q}

σ	n=	3	5	6	7	8	9	10
1/2	standard	.41009	.54998	.57968	.59851	.61111	.61993	.62634
2/3	ESIS	.24613	.38400	.41467	.43431	.44754	.45683	.46359
	standard ³ /ESIS ²	1.138	1.128	1.132	1.136	1.139	1.139	1.143

Table 3: Faster hole-filling/ contraction of ESIS vs. standard $\lambda = \lambda_{CC}$ subdivision [CC78, KP23, WM24]. Entries larger than 1 in the last row indicate how much more of a valence n configuration is filled by 2 ESIS rings compared to 3 rings of standard subdivision. Since all numbers are > 1 , the corresponding coverage ratio is $r_n(QAS, ESIS) > 6$ for all n .

the limit but also that ESIS behaves like a subdivision scheme with characteristic map χ_σ that has C^2 -joined surface rings.

Theorem 1 If $\lambda := \lambda_\sigma$ then the characteristic map χ_E of ESIS coincides with χ_σ of [KP09].

The proof requires maps \mathbf{g}^Δ , \mathbf{t}_σ , L and the technical construction of the refinement rules $(\mathbf{d}, \mathbf{q}) \rightarrow (\tilde{\mathbf{d}}, \tilde{\mathbf{q}})$, all defined in Appendix A.

Proof Subdivision theory [PR08] defines χ_E as a planar map that can be chosen rotationally symmetric and symmetric across sector-diagonals. With these restrictions χ_E is unique up to scaling and rotation and has the same symmetries as the map \mathbf{g}^Δ . For the proof, we apply the construction of Appendix A to \mathbf{g}^Δ as the identity map and \mathbf{d} as the planar partial BB-net \mathbf{t}_σ , displayed magenta in Fig. 25 d, left. Then sampling \mathbf{g}^Δ with \mathbf{t}_σ yields back \mathbf{t}_σ ; and sampling the composition of $\mathbf{g} := \mathbf{g}^\Delta \circ L^{-1}$ with $\lambda_\sigma \mathbf{t}_\sigma$ yields $\lambda_\sigma \mathbf{t}_\sigma$ as the refinement of the input $\mathbf{d} = \mathbf{t}_\sigma$. The construction of the next ring according to Section 4.3 yields C^1 -joined sequence of rings of bi-3 patches, as does the λ_σ -scaled sequence of χ_σ (scaled χ_σ join in fact C^2). That is, for the particular choice of \mathbf{g}^Δ and input net, χ_E and χ_σ agree in the first-order Hermite data, at both their inner and outer boundaries. Since the rings are bi-3, this implies $\chi_E = \chi_\sigma$. \square

We can replace the default λ_σ by a much smaller λ such that λ^2 is larger than the largest absolute eigenvalue of $M_{\mathbf{d}}$ – listed for $n = 6, \dots, 10$ in Fig. 12. This preserves curvature boundedness. For comparison with [WM23] we choose $\lambda := 0.4$. Fig. 12 juxtaposes the characteristic map $\chi_{E,0.4}$, red, and the default χ_E , black. Since, numerically, $\partial_u \chi_{E,0.4} \times \partial_v \chi_{E,0.4} > 0$ by a good margin, $\chi_{E,0.4}$ is injective and the limit surface is C^1 .




Figure 12: Characteristic maps of ESIS when $\lambda := 0.4$ for $n = 6, \dots, 10$. Beside the valencies n are listed the corresponding $|\lambda|_{max}$ – the largest absolute eigenvalue of $M_{\mathbf{d}}$. Default χ_E is black, the accelerated $\chi_{E,0.4}$ is red.

6. Comparisons

We compare to the most widely-used CC subdivision [CC78], to the tuned subdivision algorithm WM [WM23] (since WM has

comparable hole-filling speed to ESIS) and to the leading guided bi-cubic subdivision algorithm QAS [KP23]. To demonstrate the shape properties of ESIS listed in the Introduction, we focus on irregular neighborhoods. We use as input the extended net $\tilde{\mathbf{c}}$ with one node of valence n that defines a surrounding bi-3 ring. The layout is color-coded as:

- green: surrounding bi-3 ring.
- red: ESIS rings

The comparison focuses on valences $n=3,5,6$ since in practice – both for generated quad-dominant and hand-crafted non-flat meshes – valence $n > 6$ is rare. For $n = 3$ $n = 5$ and convex or rounded input nets for $n = 6$, ESIS achieves quality on par with QAS.

For sharply changing input nets and $n \geq 6$ (and higher than the already fast default speed) ESIS yields less uniform highlight line distributions than QAS. Due to the structural compatibility, we can insert r QAS rings before switching to ESIS: QAS refinement matrices are available as open source code and the transition from QAS to ESIS is straightforward, see Appendix B. Applying QAS r times before continuing with ESIS yields results on par with QAS, indicating that ESIS is an efficient substitute for QAS once the geometry is sufficiently ‘mollified’ via QAS. Of course this comes at the cost of lower contraction speed and more pieces per ring

For $\sigma := \frac{2}{3}$, i.e. $\lambda := \lambda_\sigma$, we label those surfaces ESIS $_r$ and color QAS rings bronze and bi-3 transition ring yellow. For the analysis of the extreme speed $\lambda := 0.4$, we label the surface as ESIS $_{r,0.4}$. Fig. 13 – Fig. 21 for $n = 3, 5, 6$, demonstrate that typically zero steps of QAS are needed for ESIS to achieve QAS-like surface quality – except, and only, for certain sharp saddle configurations that are not common in large models.

We note that since ESIS converges extremely fast, four levels of subdivision typically suffice for rendering. Nevertheless, we show levels 6,7,8 both to indicate the limit behavior and because other subdivision methods contract much slower.

Valence $n = 3$ permits only a few distinct configurations. Neither full view nor magnification to rings 6,7,8 (f,g,h,m,n,o,p in Fig. 13) reveal major differences in highlight lines, curvature plots or curvature ranges between ESIS, ESIS $_1$ and QAS. Even for the ‘staircase’ example, any difference between ESIS and ESIS $_1$ in Fig. 13 (t) vs (u) is invisible certainly when compared to the corresponding Catmull-Clark surface Fig. 13 (s) where the red arrow points to highlight line oscillation.

The convex $n = 5$ configuration, Fig. 14, does not show major differences between ESIS, ESIS $_1$ and QAS, even in the Gauss curvature (ESIS $_1$ is not shown since indistinguishable from ESIS).

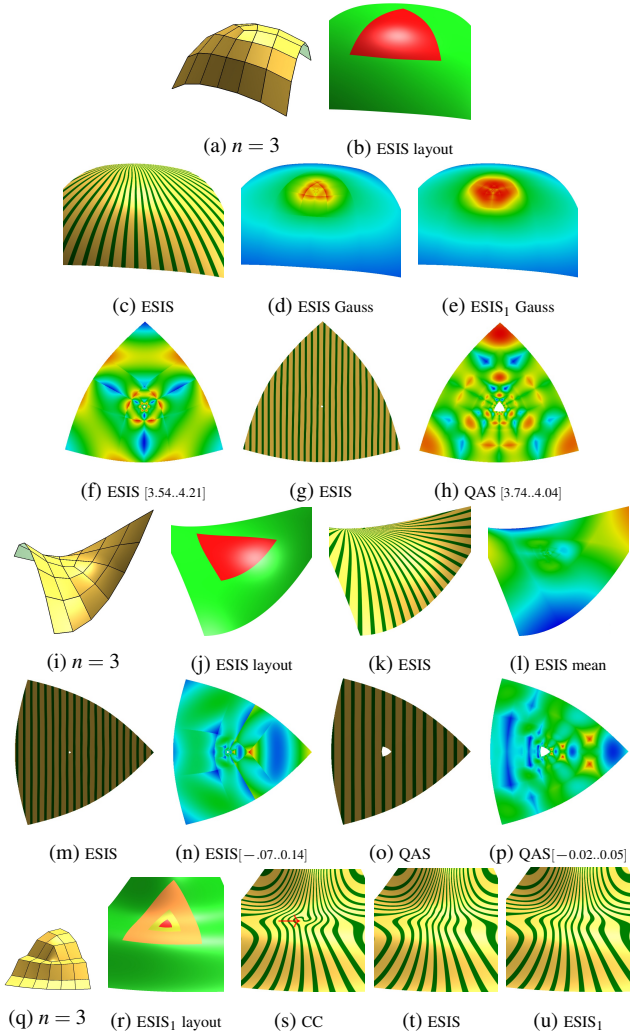



Figure 13: Valence $n = 3$: row 1, 2 highlight lines and Gauss curvature for convex input; (f,g,h) zoom in to rings 6,7,8 of the surface(d) to bring out the fine changes. The interval $[a...b]$ indicates the Gauss curvature range associated with the red-green-blue curvature color bar  in curvature. Rows 4 and 5 show highlight lines and mean curvature for hyperbolic input; (m–p) zoom to rings 6,7,8. Row 6 compares highlight lines on a wavy ‘staircase’ with $n = 3$.

Since ESIS contracts faster than QAS, the remaining hole in (h) is smaller than in (i) and the highlight lines are denser. The major difference in quality is between the undesirable oscillating highlight line distribution of WM, corresponding to dips and raise, versus the desirable uniform distribution of ESIS.

The sharp corner with $n = 5$, Fig. 15, shows in (e) an undesirable oscillation of highlight lines in the first subdivision ring of ESIS. One application of QAS followed by ESIS for the remainder, i.e. $ESIS_1$, removes this oscillation in (e). $ESIS_1$ exhibits good quality throughout and matches the highlight line distribution of QAS both in large – (f) vs (g) – and under extreme zoom where the surface

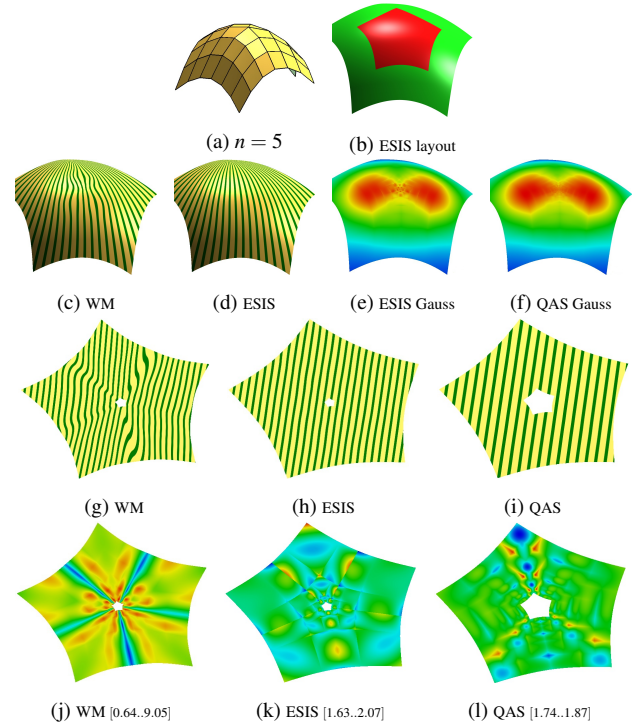


Figure 14: $n = 5$ convex: comparison of highlight lines and Gaussian curvature of WM, QAS and ESIS. last 2 rows: zoom to rings 6,7,8. The interval $[a...b]$ indicates the curvature range.

is entirely ESIS – compare (i) vs (j). By contrast, WM shape is deficient both in large (d) and in vicinity of limit point (h). Degree bi-3 polyhedral-net splines [PLK23] is an open source finite multipatch construction. The sharp transitions in Fig. 15 (c), see \rightarrow , rule this algorithm out as competitor to either QAS and ESIS.

For the convex, $n = 6$ input net, Fig. 16 compares subdivision methods, speeds, preparation of ESIS by QAS. For $n = 6$, ESIS has $\lambda_\sigma = .41467$. Since WM uses $\lambda = 0.4$, we set, without apparent harm to the ESIS surface quality, $\lambda = 0.4$ for ESIS, to ensure a fair comparison. Evidently the highlight line distribution of WM is poor, see Fig. 16 f. On the other hand, the difference between Fig. 16 (c) and (d) is imperceptible, indicating that no preparation of ESIS by QAS is needed for this valence and shape. This observation is confirmed also in the strong zoom limit.

By contrast, the sharply turning two-beam input Fig. 17 benefits from initial ‘mollification’ by QAS, compare (c) to (d). Indeed, using two layers sets up ESIS so that the difference in highlight lines is imperceptible, except for the miniature hole left by slower QAS. The sharp transitions in Fig. 17 (g), see \rightarrow , rule out [PLK23] as competitor to either QAS and ESIS. That is, QAS provides a high-quality initial ring that works well for ESIS.

The transition from QAS to ESIS is straightforward and detailed in Appendix B. The preparation with r steps of QAS, namely $ESIS_r$, reduces efficiency but achieves excellent results. By contrast, as demonstrated in Fig. 18, preparing with Catmull-Clark subdivision in place of QAS leads not just to the expected well-





Figure 15: Two-beam corner $n = 5$: comparison of WM, QAS and ESIS. (g,h,i) zoom to rings 6,7,8, highlight lines and mean curvature in (k,l,m) in the range $[a\dots b]$ mapped to the red-green-blue curvature color bar .

known ‘first step artifacts’ shown in Fig. 18 a, but additionally suffers from a pinching of highlight lines throughout as demonstrated by zooming in on fine rings 6,7,8 in Fig. 18 d. The pinching disappears in the limit when ESIS is applied Fig. 18 e, but is visible in the large Fig. 18 b – in contrast to preparation with QAS when followed up with high-speed $\lambda = 0.4$ ESIS: ESIS_{2,0.4} shows an excellent highlight line distribution in (c). Conversely, we can ask whether QAS pre-processing always yields good shape, also for other algorithms. Fig. 19 illustrates the general finding that QAS

cannot save CC or WM: following QAS by CC or WM yields poor outcomes, whereas QAS followed by ESIS yields excellent shape. The examples of Fig. 20 and Fig. 21 consist entirely of ESIS surfaces without additional surrounding spline surfaces. There is no QAS pre-processing. Spot the extraordinary point in the zoomed-in region of Fig. 20 (b,c). Hint: look at (d) where the center of CC subdivision is easily visible as a pinch point.

High valences $n > 6$ and stress testing: In modeling practice, as opposed for scientific scrutiny, any valence above 6 is unusual for non-flat facets. For extreme configurations, such as twisted higher-order saddle data and $n = 9$ (see the Electronic Supplement), applying ESIS directly to a net yields less uniform highlight line distributions than for the shape-leader QAS – although shape is far better than CC or WM.

7. Sample Applications: Texturing, Engineering Analysis

As a basic application, Fig. 22 illustrates the absence of distortions when a grid-like texture is applied across an extraordinary point of ESIS by using ESIS-subdivided texture coordinates to look up and transfer the checkerboard as, for example, in [DKT23].

Fig. 23 illustrates the use of ESIS both to define the geometry and to compute on the geometry by solving partial differential equations on the surface as the domain. In the spirit of isogeometric analysis [BF84, CSA*02, HCB05], we solve a second-order partial differential equation, the heat equation,

$$\frac{\partial u_\Omega}{\partial t} = \Delta_{\mathbf{x}} u_\Omega + f \quad (4)$$

for the surface temperature u_Ω on the ESIS surface \mathbf{x} , given a heat source f . That is, $\Delta_{\mathbf{x}}$ is the Laplace-Beltrami operator on \mathbf{x} . We implemented the time-stepping simulation up to integer time t in n steps as the equation

$$\left(2M + \frac{n}{t}K\right) \mathbf{u}_n = \left(2M - \frac{n}{t}K\right) \mathbf{u}_h^- \quad (5)$$

where \mathbf{u}_h denotes the numerical solution and \mathbf{u}_h^- the solution at the previous time step, starting with 0 temperature. The function f enters the computation by setting the value of (a collection of) heat sources in \mathbf{u}^- before each time step. The test and trial functions ϕ_j , ψ_i of the weak Galerkin formulation are scalar-valued functions in ESIS-form and the non-zero entries of the mass and the stiffness matrix are, respectively,

$$M_{ij} := \sum_{\alpha} \int_{\square} \phi_i \psi_j |J_{\alpha}| d\square,$$

$$K_{ij} := \sum_{\alpha} \int_{\square} (\nabla \phi_j)^{\top} (J_{\alpha}^{\top} J_{\alpha})^{-1} (\nabla \psi_i) |J_{\alpha}| d\square,$$

with J_{α} the Jacobian and $J_{\alpha}^{\top} J_{\alpha}$ the first fundamental form of patch α and \square the unit square. The integration is up to ring 3 of the ESIS surface. The temperature isolines of u_Ω trace out geodesic distance from the heat source: every tenths isoline is displayed as a white band.

8. Conclusion

ESIS limit surfaces are curvature-bounded, i.e. as smooth as possible for non-regular control nets and bi-cubic degree. Even for

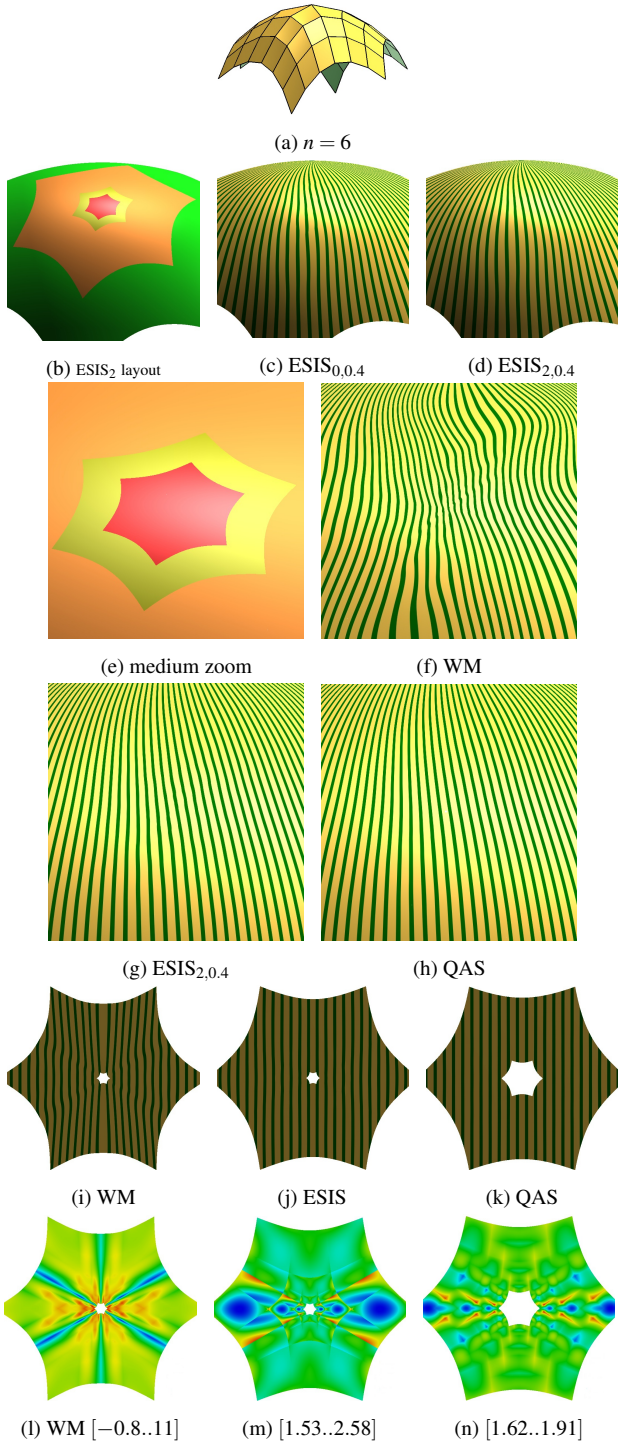


Figure 16: High speed 0.4 ESIS for $n = 6$ convex configuration: (c–h) highlight lines with mild zoom and (i,j,k) high zoom to rings 6,7,8. (l,m,n) The Gauss curvature range indicated as $[a\dots b]$.

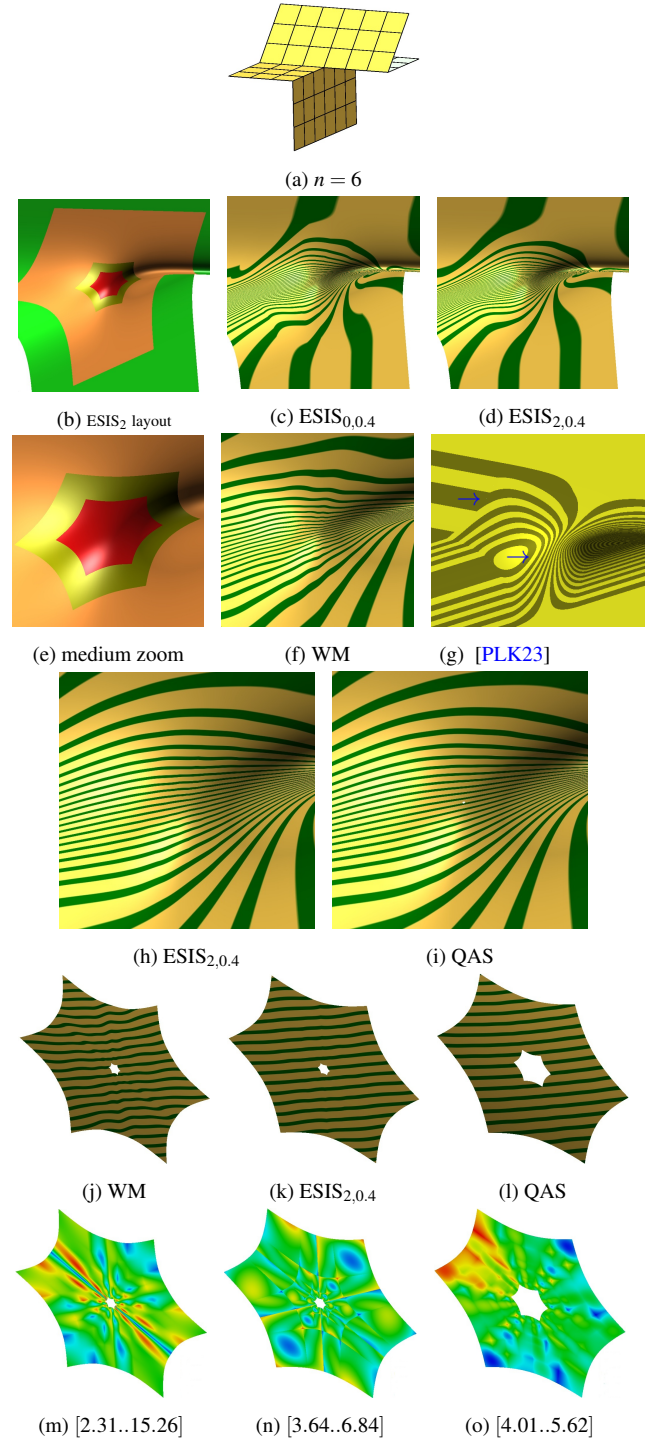


Figure 17: High speed 0.4 ESIS for challenging saddle configuration $n = 6$. (e–i) highlight lines with mild zoom and (j–o) high zoom to rings 6,7,8. (m,n,o) mean curvature in the range $[a\dots b]$.

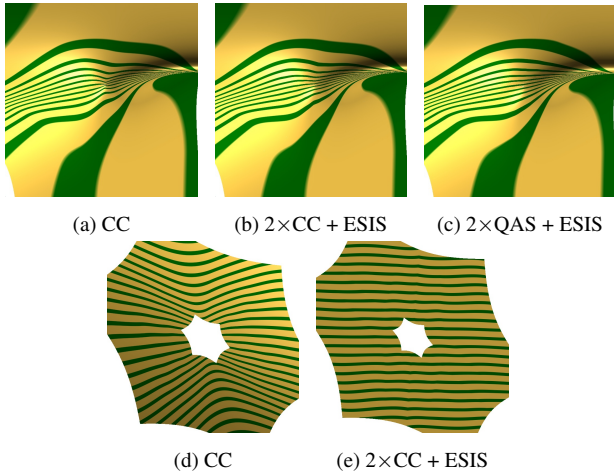


Figure 18: Transition to ESIS using Catmull-Clark Subdivision is poor for the beam join of Fig. 17 a compared to transition via QAS. (d,e) When zooming in to rings 6,7,8, ESIS recovers.

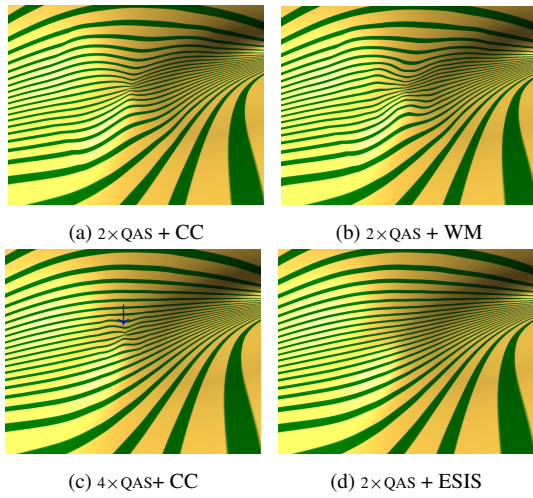


Figure 19: QAS-initialized subdivision algorithms perform poorly (pinched highlight lines) except for ESIS in (d) – on input Fig. 17 (a), viewed as in Fig. 17 (e).




Figure 20: Moebius DNA strip: surface consists entirely of ESIS-rings, $n = 6$. (b,c) ESIS highlight lines and mean curvature in zoom. (d) Catmull-Clark mean curvature (same scaling).




Figure 21: Genus 15 surface consisting entirely of ESIS-rings, $n = 8$; (b) distribution of neighborhoods; (c) mean curvature; (d) highlight lines on a segment with two circular and two half-circular branches.




Figure 22: Textured surfaces for a steering wheel: (a) input net. (b,c) The location of the irregularities are not apparent in the grid-like textures applied using ESIS texture coordinates.




Figure 23: Hot steering wheel with temperature at time t computed by leveraging ESIS both for the surface geometry and for the finite elements on the surface. White strips indicate geodesic distance level sets: The heat spreads on the surface from the constant bright yellow source.

challenging input, when initialized with the structurally compatible QAS, the nominally less smooth transition between subdivision rings yields high-quality surfaces, on par with the shape leader QAS but at 1/4 the cost in the number of pieces.

Relaxation of formal smoothness is the price exacted for achieving bi-cubic curvature-bounded subdivision with only $3n$ pieces per subdivision ring. The minimal $3n$ patch layout, and, to a lesser degree, the fast contraction, mean that ESIS degrades more rapidly with n than QAS. The efficient layout also forces the reduction in formal inter-ring continuity. Remarkably, though, we proved that the sub-dominant eigen-component of the surface rings, the characteristic map χ_E , is C^2 , which helps explain the good shape of ESIS in the limit. ESIS pioneers refinement based on a mixture of Bézier (not B-spline) control-nets in total degree and tensor-product form. Notably, two steps of ESIS cover a multi-sided hole more than do three steps of advanced algorithms in the literature that generate four times as many patches per ring.

Acknowledgements Jörg Peters acknowledges the support of the Fields Institute, Toronto, CA. U.S. Navy Grant number SBIR N251-060

Ethics N/A

Conflict of Interest none

Data Availability see Electronic Supplement.

References

- [ACSD*03] ALLIEZ P., COHEN-STEINER D., DEVILLERS O., LÉVY B., DESBRUN M.: Anisotropic polygonal remeshing. In *ACM SIG-GRAPH 2003 Papers*. 2003, pp. 485–493. [3](#)
- [ADS11] AUGSDÖRFER U. H., DODGSON N. A., SABIN M. A.: Artifact analysis on B-splines, box-splines and other surfaces defined by quadrilateral polyhedra. *Comp Aid Geom Design* 28, 3 (2011), 177–197. [2](#)
- [BC94] BEIER K.-P., CHEN Y.: Highlight-line algorithm for realtime surface-quality assessment. *Comp-Aid Design* 26, 4 (1994), 268–277. [1](#), [2](#)
- [BF84] BRAIBANT V., FLEURY C.: Shape optimal design using B-splines. *CMAME* 44 (1984), 247–267. [8](#)
- [CC78] CATMULL E., CLARK J.: Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design* 10 (Sept. 1978), 350–355. [2](#), [3](#), [6](#)
- [CSA*02] CIRAK F., SCOTT M., ANTONSSON E. K., ORTIZ M., SCHRÖDER P.: Integrated modeling, finite-element analysis, and engineering design for thin-shell structures using subdivision. *Computer-Aided Design* 34, 2 (2002), 137–148. [8](#)
- [CSH*25] CAPOUELLEZ R., SINGH R., HEISTERMANN M., BOMMES D., ZORIN D.: Feature-aligned parametrization in penner coordinates. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–21. [3](#)
- [dB78] DE BOOR C.: *A Practical Guide to Splines*. Springer, 1978. [3](#)
- [dB87] DE BOOR C.: B-form basics. In *Geometric Modeling: Algorithms and New Trends* (1987), Farin G., (Ed.), SIAM, pp. 131–148. [3](#)
- [DC63] DE CASTELJAU P.: Courbes et surfaces à pôles. *André Citroën, Automobiles SA, Paris* 66 (1963). [4](#)
- [DKT23] DEROSE T., KASS M., TRUONG T.: Subdivision surfaces in character animation. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*. 2023, pp. 801–810. [8](#)
- [Far88] FARIN G.: *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide*. Academic Press, 1988. [3](#), [4](#)
- [HCB05] HUGHES T. J. R., COTTRELL J. A., BAZILEVS Y.: Isogeometric analysis: Cad, finite elements, NURBS, exact geometry and mesh refinement. *CMAME* 194, 39–41 (2005), 4135–4195. [8](#)
- [HZN*18] HUANG J., ZHOU Y., NIESSNER M., SHEWCHUK J. R., GUIBAS L. J.: Quadriflow: A scalable and robust method for quadrangulation. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 147–160. [2](#), [3](#)
- [JTPSH15] JAKOB W., TARINI M., PANOZZO D., SORKINE-HORNUNG O.: Instant field-aligned meshes. *ACM transactions on graphics (TOG)* 34, 6 (2015), 1–15. [3](#)
- [KP07] KARČIAUSKAS K., PETERS J.: Concentric tessellation maps and curvature continuous guided surfaces. *Computer Aided Geometric Design* 24, 2 (Feb 2007), 99–111. [2](#)
- [KP09] KARČIAUSKAS K., PETERS J.: Adjustable speed surface subdivision. *Computer Aided Geometric Design*. 26 (2009), 962–969. [5](#), [6](#), [12](#)
- [KP15] KARČIAUSKAS K., PETERS J.: Can bi-cubic surfaces be class A? *Computer Graphics Forum* 34, 5 (August 2015), 229–238. [1](#)
- [KP22] KARČIAUSKAS K., PETERS J.: Point-augmented bi-cubic subdivision surfaces. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 13–23. [2](#)
- [KP23] KARČIAUSKAS K., PETERS J.: Quadratic-attraction subdivision. *Computer Graphics Forum* 42, 5 (2023), e14 900. [1](#), [2](#), [4](#), [6](#), [12](#)
- [KP24] KARČIAUSKAS K., PETERS J.: Quadratic-attraction subdivision with contraction-ratio $\lambda = 1/2$. *Computers & Graphics* 123 (2024), 104001. [1](#), [2](#)
- [KPR04] KARČIAUSKAS K., PETERS J., REIF U.: Shape characterization of subdivision surfaces – case studies. *Computer-Aided Geometric Design* 21, 6 (july 2004), 601–614. [3](#)

- [LWG*25] LIU J., WANG C., GUO S., WENG H., ZHOU Z., LI Z., YU J., ZHU Y., XU J., LEI B., ET AL.: Quadgpt: Native quadrilateral mesh generation with autoregressive models. *arXiv preprint arXiv:2509.21420* (2025). 3
- [M*10] MCNEEL R., ET AL.: Rhinoceros 3d. *Robert McNeel & Associates, Seattle, WA* (2010). 3
- [Men] MENGLOW: Base meshes character starter kit rig. URL: <https://www.turbosquid.com/3d-models/base-meshes-character-starter-kit-rig-1382197>. 3
- [MK06] MARINOV M., KOBELT L.: A robust two-step procedure for quad-dominant remeshing. In *Computer Graphics Forum* (2006), vol. 25, Wiley Online Library, pp. 537–546. 3
- [MM18] MA Y., MA W.: Subdivision schemes with optimal bounded curvature near extraordinary vertices. *Computer Graphics Forum* 37, 7 (Oct 2018), 455–467. 2
- [Pin97] PINKAVA J.: Geri’s Game. Computer animated short film, 1997. Released on November 24, 1997, and shown with the theatrical release of A Bug’s Life. URL: https://pixar.fandom.com/wiki/Geris_Game. 3
- [PLK23] PETERS J., LO K., KARČIAUSKAS K.: Algorithm 1032: Bi-cubic splines for polyhedral control nets. *ACM Transactions on Mathematical Software* 49, 1 (2023), 1–12. 7, 8, 9
- [PNA*21] PIETRONI N., NUVOLE S., ALDERIGHI T., CIGNONI P., TARINI M., ET AL.: Reliable feature-line driven quad-remeshing. *ACM Transactions on Graphics* 40, 4 (2021), 1–17. 3
- [PR08] PETERS J., REIF U.: *Subdivision Surfaces*, vol. 3 of *Geometry and Computing*. Springer-Verlag, New York, 2008. 1, 6
- [THCM04] TARINI M., HORMANN K., CIGNONI P., MONTANI C.: Polycube-maps. *ACM Transactions on Graphics (TOG)* 23, 3 (2004), 853–860. 3
- [TPC*10] TARINI M., PIETRONI N., CIGNONI P., PANOZZO D., PUPPO E.: Practical quad mesh simplification. In *Computer Graphics Forum* (2010), vol. 29, Wiley Online Library, pp. 407–418. 2, 3
- [WM23] WANG X., MA W.: An extended tuned subdivision scheme with optimal convergence for isogeometric analysis. *Computer-Aided Design* 162 (2023), 103544. 1, 2, 6
- [WM24] WANG X., MA W.: A class of new tuned primal subdivision schemes with high-quality limit surface in extraordinary regions. *ACM Transactions on Graphics (TOG)* 43, 6 (2024), 1–17. 1, 2, 6

Appendix A: Derivation of the refinement rules




Figure 24: Assembly of three corner jets (middle) into L-net. The \times marks a corner point of the jet representing f , $\partial_u f$, $\partial_v f$ and $\partial_u \partial_v f$ are evaluated at \times .

For the refinement of (\mathbf{d}, \mathbf{q}) we make *constructive* use of the characteristic map χ_σ – that was introduced in Section 5 for the limit analysis. To extract Hermite data from a guide surface g , we use so-called *tensor-border* \mathbf{t}_σ of χ_σ : \mathbf{t}_σ is a partial BB-net displayed in magenta in Fig. 25 d, left that represents position and derivatives as the union of jets in corners (marked \times in Fig. 24). Tensor-borders naturally combine to form L-shapes, n of which form a ring of BB-coefficients. While χ_σ also offers second-order Hermite data (orange in Fig. 25 d, left) that can be used as in [KP23, KP09], we sample only first order data from a guide surface \mathbf{g} – by extracting parts of the composition $\mathbf{g} \circ \mathbf{t}_\sigma$.

Fig. 25 a displays one sector of the \mathbf{d} -net and two nearest BB-coefficients in each neighbor sector. Given the quadratic expansion \mathbf{q} , illustrated as a light-red filled pentagon in Fig. 25 b (and Fig. 8 a), the bi-4 G^1 guide \mathbf{g} in Fig. 25 e is determined to fit best the input \mathbf{d} -net, following the approach of QAS [KP23] but simpler due to the relaxation from C^2 to C^1 constraints between and within the surface rings.

Constructing \mathbf{g}^Δ We construct a C^1 piecewise total degree 4 pre-cursor \mathbf{g}^Δ , see Fig. 25 b, of the tensor-product bi-4 guide \mathbf{g} by initializing its quadratic expansion at \mathbf{p} by \mathbf{q} . The $4n \bullet$ of \mathbf{g}^Δ will be determined in the next step and the n unmarked points in the outermost layer are not relevant in the refinement derivation. For each $s = 0, \dots, n-1$, the magenta tensor-border coefficients of $\mathbf{g}^\Delta \circ \mathbf{t}_\sigma$ at the corners (marked \circ in Fig. 25 c) are set to the BB-coefficients $\mathbf{d}_{20}^s, \mathbf{d}_{21}^s, \mathbf{d}_{02}^{s+1}, \mathbf{d}_{12}^{s+1}$ provided by \mathbf{d} (see the indices in Fig. 25 a) by solving for the four free \bullet of \mathbf{g}^Δ .

Constructing \mathbf{g} We compose each sector of \mathbf{g}^Δ with the linear shear map L that transforms the unit square into a unit-edge-length parallelogram with opening angle $\frac{2\pi}{n}$ at the common origin \mathbf{O} , as shown in Fig. 25 d. This yields the gray-underlaid BB-coefficients of the piecewise G^1 bi-4 map \mathbf{g} inheriting the G^1 data of \mathbf{g}^Δ . In each sector \mathbf{g}^s , the BB-coefficients (marked as \circ Fig. 25 e) are defined as

$$\mathbf{g}_{2j}^s := \frac{2}{3}(\mathbf{g}_{3j}^s + \mathbf{g}_{1j}^s) - \frac{1}{6}(\mathbf{g}_{4j}^s + \mathbf{g}_{0j}^s), j = 0, 1,$$

and $\mathbf{g}_{02}^s, \mathbf{g}_{12}^s$ are defined by the diagonally symmetric formula. The middle BB-coefficient, marked as \bullet , is

$$\begin{aligned} \mathbf{g}_{22}^s := & \frac{1}{2} \left(\frac{2}{3}(\mathbf{g}_{21}^s + \mathbf{g}_{23}^s) - \frac{1}{6}(\mathbf{g}_{20}^s + \mathbf{g}_{24}^s) \right) + \\ & \frac{1}{2} \left(\frac{2}{3}(\mathbf{g}_{12}^s + \mathbf{g}_{32}^s) - \frac{1}{6}(\mathbf{g}_{02}^s + \mathbf{g}_{42}^s) \right). \end{aligned}$$

Here assignments of the type $\mathbf{b}_2 := \frac{2}{3}(\mathbf{b}_3 + \mathbf{b}_1) - \frac{1}{6}(\mathbf{b}_4 + \mathbf{b}_0)$ reduce

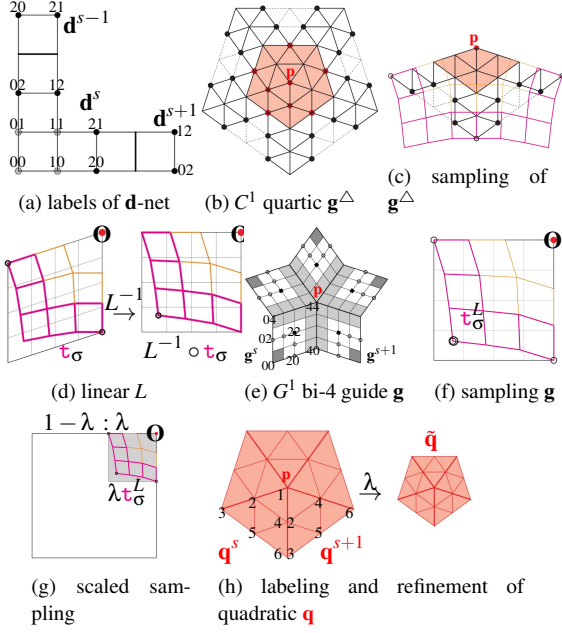


Figure 25: Derivation of refinement rules. *Top row:* (a) Labeling of input and output \mathbf{d} -nets. (b) Piecewise C^1 quartic \mathbf{g}^Δ as intermediate guide. (c) Sampling of \mathbf{g}^Δ using the characteristic tensor-border of \mathbf{t}_σ . *Middle row:* Linear sharing map L and increased flexibility of the intermediate \mathbf{g}^Δ . (d) Transformation of \mathbf{t}_σ . (e) Structure of bi-4 guide \mathbf{g} . (f) Sampling of \mathbf{g} with transformed $\mathbf{t}_\sigma^L := L^{-1} \circ \mathbf{t}_\sigma$. *Bottom row:* (g) Sampling of the final guide \mathbf{g} composed with $\lambda \mathbf{t}_\sigma^L$. (h) Labeling of input and output \mathbf{q} .

the actual degree to 3, a shape-neutral heuristic for shedding extra degrees of freedom. The remaining BB-coefficients \mathbf{g}_{ij}^s , $i = 0, 1$, $j = 0, 1$ (dark-gray underlaid in Fig. 25 e) are set so that the (magenta) corner jet of $\mathbf{g}^s \circ L^{-1} \circ \mathbf{t}_\sigma$ sampled at \circ (Fig. 25 f) equals the corner jet \mathbf{d}_{ij}^s , $i = 0, 1$, $j = 0, 1$.

Refinement $(\mathbf{d}, \mathbf{q}) \rightarrow \tilde{\mathbf{d}}$ (see Fig. 25 g) For each sector, the domain of \mathbf{g} , together with \mathbf{t}_σ^L , is scaled by λ ($\sigma := \frac{2}{3}$) towards the origin \mathbf{O} to sample the corner jets at all three locations marked as \circ . This yields the 8 BB-coefficients of $\tilde{\mathbf{d}}$, see Fig. 25 a. Section 1 of the Electronic Supplement presents explicit formulas for expressing the control points $\tilde{\mathbf{d}}$ as affine combinations of \mathbf{d} and \mathbf{q} .

Refinement $\mathbf{q} \rightarrow \tilde{\mathbf{q}}$. (see Fig. 25 g, h) The domain of \mathbf{g}^Δ is scaled down by λ_σ towards the origin. The restriction of \mathbf{g}^Δ to the scaled domain provides the new $\tilde{\mathbf{q}}$, Fig. 25 h. The explicit form of refinement is

$$\begin{pmatrix} \tilde{\mathbf{q}}_1^s \\ \tilde{\mathbf{q}}_2^s \\ \tilde{\mathbf{q}}_3^s \\ \tilde{\mathbf{q}}_4^s \\ \tilde{\mathbf{q}}_5^s \\ \tilde{\mathbf{q}}_6^s \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1-\lambda & \lambda & 0 & 0 & 0 & 0 \\ (1-\lambda)^2 & 2(1-\lambda)\lambda & \lambda^2 & 0 & 0 & 0 \\ 1-\lambda & 0 & 0 & \lambda & 0 & 0 \\ (1-\lambda)^2 & (1-\lambda)\lambda & 0 & (1-\lambda)\lambda & \lambda^2 & 0 \\ (1-\lambda)^2 & 0 & 0 & 2(1-\lambda)\lambda & 0 & \lambda^2 \end{pmatrix} \begin{pmatrix} \mathbf{q}_1^s \\ \mathbf{q}_2^s \\ \mathbf{q}_3^s \\ \mathbf{q}_4^s \\ \mathbf{q}_5^s \\ \mathbf{q}_6^s \end{pmatrix}. \quad (6)$$

Appendix B: Initialization of ESIS,

When inserting one or more steps/rings of QAS before the more efficient ESIS subdivision, initialization of (\mathbf{d}, \mathbf{q}) is from the last

QAS ring r and its central expansion $\hat{\mathbf{q}}$: Eq. (2) yields \mathbf{q} from this $\hat{\mathbf{q}}$ when $\lambda := \lambda_{CC}^{r+1}$ in Eq. (6).




Figure 26: (dark-gray) QAS ring; (light-gray) transition ring; (white) ESIS rings.

The bi-3 C^1 transition ring for switching from QAS to ESIS, consists of $3n$ patches, Fig. 26. The construction is detailed in Fig. 27 :

- displays unchanged QAS double rings (dark-gray), and a ring to be replaced by the transition ring (lighter gray).
- The outer QAS double ring is C^1 -extended inwards (prolongated) to form 4 uniformly-spaced (meaning 1 : 1 ratio of parameter intervals) tensor-border pieces (dark-gray in Fig. 27 b). The tensor-borders join C^2 . The inner QAS double ring (light gray) is C^1 -extended inwards to yield three light-red underlaid corner jets.
- The three corner jets (2×2 BB-subnets, light-red) are scaled by 2 in both directions to yield the 8 \bullet of the \mathbf{d} -net to start ESIS. Denote the BB-coefficient of one corner \mathbf{b}_{ij} , $i = 0, 1$, $j = 0, 1$, and those of corner jets $\tilde{\mathbf{b}}_{ij}$. Then, with the remaining coefficients marked \bullet determined by symmetry,

$$\begin{aligned} \mathbf{b}_{00} &:= \tilde{\mathbf{b}}_{00}, \quad \mathbf{b}_{10} := 2\tilde{\mathbf{b}}_{10} - \tilde{\mathbf{b}}_{00}, \quad \mathbf{b}_{01} := 2\tilde{\mathbf{b}}_{01} - \tilde{\mathbf{b}}_{00} \\ \mathbf{b}_{11} &:= \tilde{\mathbf{b}}_{00} - 2(\tilde{\mathbf{b}}_{10} + \tilde{\mathbf{b}}_{01}) + 4\tilde{\mathbf{b}}_{11}. \end{aligned}$$

After C^1 -joining the tensor-borders of adjacent sectors, the L -shaped tensor-borders are outwards- C^1 -extended to complete the transition ring with the brown-underlaid BB-coefficients.




Figure 27: Construction of the transition bi-3 ring.