Contents lists available at ScienceDirect

Computers & Graphics

journal homepage: www.elsevier.com/locate/cag

Refinable smooth surfaces for locally quad-dominant meshes with T-gons

Kęstutis Karčiauskas[a], Jörg Peters[b,*]

[a]Naugarduko 24, Vilnius University, LT-2006, Vilnius, Lithuania
[b]Dept CISE, University of Florida, Gainesville FL 32611-6120, USA

ARTICLE INFO

Article history:
Received May 21, 2019

Keywords: T-gon, G1 surface, refinability, geometric tuning, accelerated subdivision, highlight lines

ABSTRACT

Multi-sided facets in polyhedral models and meshes serve to connect regular submeshes (star-configurations) and to start or end quad-strips (T-configurations). Using the polyhedral mesh as control net, recursive subdivision algorithms often yield poor shape for these non-quad configurations. Polynomial surface constructions such as geometrically smooth splines (G-splines) do better, but lack subdivision-like refinability. Such refinability is useful for hierarchical modeling and engineering analysis.

This paper introduces a new class of G-splines that generalizes bi-quadratic C^1 splines to polyhedral control nets with star- and T-configurations and that is refinable.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

A standard approach to interactive shape modeling is to sketch, extrude and manipulate coarse polyhedral control nets and have them algorithmically refined into smooth surfaces. A classical algorithmic refinement is Doo-Sabin subdivision [1]. Doo-Sabin subdivision generalizes bi-2 (bi-quadratic) splines and can therefore be viewed as the most thrifty option for generating smooth surfaces. Unfortunately, when looked at closely, Doo-Sabin surfaces (and to a lesser degree the bi-3 generalizing Catmull-Clark surfaces [2]) have shape deficiencies notably in the irregular neighborhoods for which they were invented (see e.g. Fig. 4b,c), namely star-configurations, where multiple quad-meshes join in a point or polyhedron, Fig. 2.

Classical subdivision also fails to address a second challenge arising in polyhedral shape modeling. Quad-dominant meshes, where most facets are quadrilateral, are preferred by artists to capture the flow of principle directions. Where quad-strips start or end (to accommodate or remove detail, see Fig. 1), Doo-Sabin subdivision creates surfaces with poor highlight line distributions for T-configurations Fig. 3b,c. (Uniformity of high-

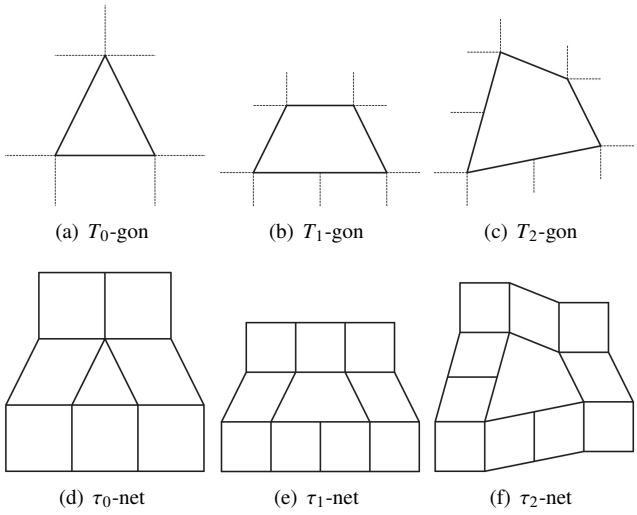

Fig. 1. The main T-configurations

*Corresponding author: jorg.peters@gmail.com

light lines is a measure commonly used by practitioners to evaluate surface quality, see [3])

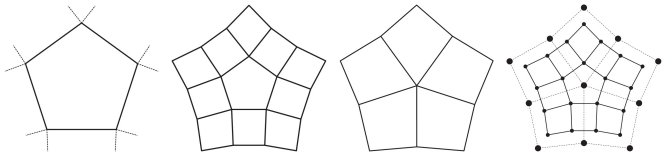

(a) DS-gon (b) DS-net (c) $S\overset{\square}{5}$-net (d) $S^{\square}$-net→DS

Fig. 2. Main star-configurations and conversion by bi-2 spline refinement.

However, subdivision naturally excels at hierarchical refinement for adaptive fine-grained modeling. This paper therefore introduces a new *nestedly refinable* class of G-splines that generalize bi-2 C^1 spline surfaces to include T- and star-configurations and provides much improved highlight line distributions. (Nested refinement means hierarchical refinement, where a coarse generalized spline control net is replaced by a finer one without changing the surface.) It suffices to focus on the configurations shown in Fig. 1 and Fig. 2 since, by local reconnection without adding control points, complex T-junctions in free-form design can be reduced to the three types of T_n-gons, $n = 0, 1, 2$ shown in Fig. 1. Among star-configuration DS-gons, Fig. 2a, are n-sided facet whose vertices are all of valence 4. DS-gons are the hallmark of Doo-Sabin subdivision but star-configurations are often more convenient for design. Since one step of B-spline knot insertion turns a star-net into a DS-net, see Fig. 2d, the constructions can nevertheless restrict attention to just T-gons and DS-gons. T-gons and DS-gons surrounded by quadrilateral facets (quads) provide the data for multi-sided caps and are respectively called τ-nets and DS-nets. Quad-dominant meshes with no other irregularities than DS-nets, star-nets and τ-nets are *locally* quad-dominant.

This paper introduces new constructions for locally quad-dominant meshes that

- are smooth with fewer oscillations and sharp creases in the highlight line distribution than existing bi-2 subdivision algorithms;
- consist of a finite number of patches both for $T-$ and star-configurations;
- are (nestedly) refinable and so enable localized editing (see Fig. 3e, 4e);
- leverage geometric accelerated tuning for improved shape at star-configurations.

To achieve refinability, surface pieces have bi-degree up to 4 near irregularities: bi-2 in regular regions, bi-degree (4,2) for T_0- or T_1-gons, bi-4 for T_2, and bi-3 at star-configurations, optionally completed by a tiny bi-4 cap.

1.1. Literature: Generalizing bi-quadratics

Due to our focus on generalizing least-degree smooth bi-2 splines, we do not review the ample literature on generalizing bi-3 splines. We note though that [5] introduced a T-construction for bi-cubic splines that includes a review

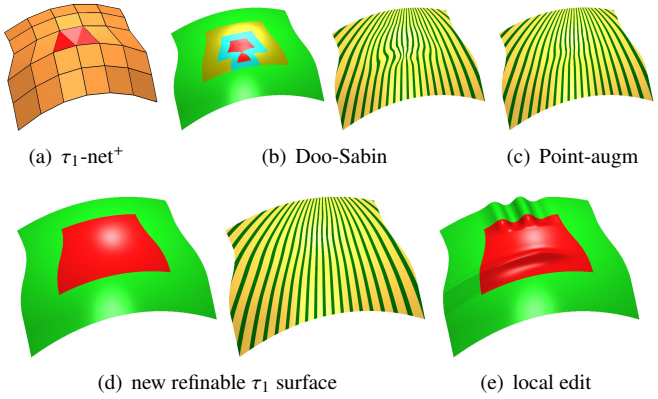

(a) τ_1-net$^+$ (b) Doo-Sabin (c) Point-augm

(d) new refinable τ_1 surface (e) local edit

Fig. 3. Bi-2 subdivision is unsatisfactory for T-junctions. (a) τ_1-net$^+$ = τ_1-net surrounded by one layer of quads. (b) Surface layout: input bi-2 patches, patches generated by first refinement step, by second refinement; subsequent 3-sided and 5-sided areas of Doo-Sabin subdivision. (c) Wavy highlight lines of point-augmented subdivision surface [4]. (d) Layout and highlight lines of the construction in this paper: input bi-2 patches, τ_1-cap of bi-degree $(2, 4)$. (e) Localized edits are used to create vertical grooves on top and a horizontal bulge across the bottom transition from frame to cap.

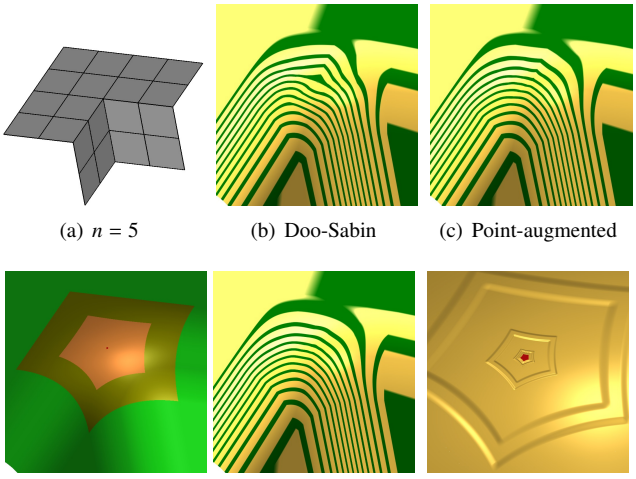

(a) $n = 5$ (b) Doo-Sabin (c) Point-augmented

(d) geometrically tuned (new) (e) local edit

Fig. 4. Bi-2 subdivision is unsatisfactory for star-configurations. (b,c) Highlight lines of Doo-Sabin and point-augmented subdivision surfaces (d) Construction of this paper and (e) localized edits leveraging hierarchical refinement.

and comparison of bi-3 alternatives, including Catmull-Clark-subdvision surfaces (that fail for T-junctions in a manner similar illustrated for Doo-Sabin in Fig. 4). Section 5 and Fig. 22 compare the surface quality of our new construction vs [5] with bi-2 vs bi-3 splines. [5] can be combined with existing constructions for star-configurations, e.g. [6, 7], but cannot as tightly integrated with them as the new construction and does not offer simple refinability.

There are two major generalizations of bi-quadratic splines using subdivision: Doo-Sabin [1] and point-augmented [4]. Even though point-augmented subdivision considerably improves the highlight line distribution over Doo-Sabin-subdivision, treating T-junctions as irregular 3-valent nodes – as subdivision algorithms do – results in poor shape. This is illustrated in Fig. 3b,c. Fundamentally, T-gons have a distinct vertical direction of change (strip density) while star-configurations

are direction-agnostic. (Here and in the following, *vertical* and *horizontal* refer to the standard layout in Fig. 1.) Even for star-configurations, Fig. 4 makes clear that Doo-Sabin [1] and point-augmented [4] subdivision leave ample room for improving the uniformity of highlight lines. A new *geometric, accelerated tuning* improves shape and allows treating the first two or three subdivision rings as the bulk of a cap that consists of finitely many patches.

Both types of irregularities can be covered by geometrically smooth splines (G-splines) that consist of a finite number of polynomial pieces joined along *G-curves*, where derivatives match only after a change of variables. Indeed, state-of-the-art G-spline constructions yield good highlight line distributions, even when the degree is low [8, 5] (see also the purely bi-2 construction [9]). However, these algorithms are not conveniently nestedly refinable. Refinement of G-splines requires that G-curves remain G-curves and therefore forces careful book keeping and special rules to reproduce shape with more control handles. G-curves prevent memory-less uniform hierarchy that comes naturally with subdivision surfaces.

(a) (b) (c)

Fig. 5. Doubling a ribbon's width with T_0-gons. This natural doubling of parameter space is inimical to hierarchical (T-)splines.

Inserting additional local knot lines provides extra degrees of freedom in tensor-product splines (see e.g. [10, 11, 12] for bi-2 splines). However this assumes a disciplined, strictly hierarchical surface design and requires keeping track of knot distances to assure global compatibility. If a prior model or polyhedral design prescribes a distribution of T-junctions a compatible global knot distribution may not exist. For example, Fig. 5 shows an increasing widening of a ribbon that closes back up at twice the initial height. For hierarchical splines, the (total sum of) knot spacings of the initial and the doubly high strip ends have to be equal, and that is not possible unless many knot intervals are zero. Efficient design can therefore combine hierarchical T-spline for refinement of regular sub-meshes, [5] for treatment of irregular nodes and the new construction for T-gons.

Road map: After introducing the input control nets and smoothness between polynomial pieces, Section 3 introduces the construction for T-configurations. The basic idea is to choose a quadratic re-parameterization to relate derivatives across the vertical boundaries of the τ-surface cap. This quadratic accommodates the change of quad-density in the vertical direction while elegantly enabling nested refinement. Section 4 details the construction for star-configurations introducing the concept of *geometric tuning* as opposed to the classical (analytic) eigen-spectrum tuning of subdivision algorithms. The idea is to use extra degrees of freedom in the subdivision rings to improve shape right from the start and not just in the limit. Section 5 explains the choices in developing the new low-degree refinable smooth surface construction and compares the resulting surfaces.

2. Definitions and Setup

Fig. 1 lists the relevant T-configurations. A T-junction is where two quads on one side meet one facet on the other. For $n > 0$, a T_n-*gon* is an $n+4$-gon surrounded by quads and such that n vertices are T-junctions and the remaining have valence 4. A T_1-gon has one T-junction (formally a pentagon Fig. 1c), a T_2-gon has two T-junctions (a hexagon Fig. 1d). For $n = 0$, a configuration that frequently arises in quad-dominant meshing algorithms [13, 14], the triangle has two vertices of valence 4 and one of valence 5 and is surrounded by 7 quads. The T-gon and its surrounding layer of quads is called τ-net.

The main star-configurations and the conversion from a star-net ($S^{\square}$-net) to DS-net are displayed in Fig. 2.

To calibrate what quality of surface can be expected for τ-configurations when the surrounding surface is a C^1 bi-2 spline. the nets are surrounded by one layer of quads to form a net^+.

A GT-spline is a collection of tensor-product patches in Bernstein-Bézier form (BB-form; see e.g. [15]):

$$\mathbf{p}(u,v) := \sum_{i=0}^{d_1}\sum_{j=0}^{d_2} \mathbf{p}_{ij} B_i^{d_1}(u) B_j^{d_2}(v)\,, \quad (u,v) \in [0..1]^2,$$

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. This paper uses $d \in \{2,3,4\}$, i.e. the BB-patches are bi-2 (bi-quadratic), bi-3 (bi-cubic), bi-4 (bi-quartic) or of bi-degree (4,2). Connecting $\mathbf{p}_{ij}$ to $\mathbf{p}_{i+1,j}$ and $\mathbf{p}_{ij+1}$ wherever possible yields the *BB-net* of BB-coefficients. A useful operation on polynomials in BB-form is their splitting into two pieces, say a left half and a right half, by the well-known *de Casteljau algorithm* [15].

For any 3×3 grid in the mesh, the vertices can be interpreted as the control net of a bi-quadratic uniform B-spline. Expressing this spline in bi-2 BB-form is called *B-to-BB conversion*, see Fig. 6a. A partial conversion from a partial mesh yields a sub-net of the BB-net that defines position and first derivatives across an edge, a *tensor-border* of degree 2 and depth 1 denoted by $\mathbf{t}$, see Fig. 6b.

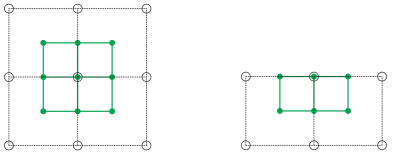

(a) (b)

Fig. 6. Bi-2 B-to-BB conversion. Circles ∘ mark B-spline control points, solid disks • mark BB-coefficients. The rules are: $\bullet_{\text{center}} = \circ_{\text{center}}$, $\bullet_{\text{edge}}$ = the average of $\circ_{\text{center}}$ and a neighbor ∘, and $\bullet_{\text{corner}}$ = the average of the four surrounding ∘.

The T-constructions and the finite cap for star-configurations are based on the concept of *geometric continuity*, i.e. smoothness after change of variables, see e.g. [16, 17]. Patches $\mathbf{p}$ and

$\mathbf{q}$ that share a *G-edge* parameterized by $(u, 0 = v)$ are G^1 connected if they have matching derivatives after change of variables $\rho(u, v) := (u + b(u)v, a(u)v)$:

$$\partial_v \mathbf{q}(u, 0) - a(u)\partial_v \mathbf{p}(u, 0) - b(u)\partial_u \mathbf{p}(u, 0) = 0\,. \qquad (1)$$

2.1. C^1 continuity in the direction along a G-curve

Key to refinability of the τ surfaces is the C^1 continuity of patches $\mathbf{q}^s(u, v)$, $s = 0, 1$ *along* a G-curve (see Fig. 7a). We denote by $\mathbf{jp}$ the first-order expansion of $\mathbf{p}$ with respect to v (effectively across the G-curve) and abbreviate $\mathbf{jp}^0 \sim \mathbf{jp}^1$ if the first order expansions join parametrically C^1 at their respective u-parameters $1 - s$.

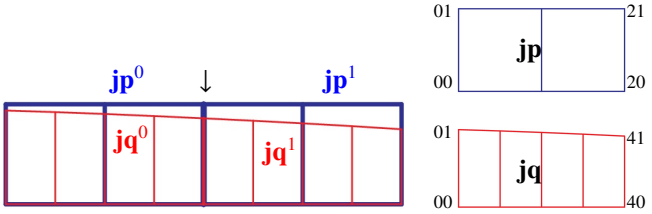

(a) Lemma 1 (b) Lemma 2

Fig. 7. BB-nets of $\mathbf{jp}^s$ (degree 2) and $\mathbf{jq}^s$ (degree 4) illustrating (a) C^1 smoothness of position and of first cross-derivative at the split point where the pieces meet (marked by ↓). (b) Labels of BB-coefficients in Lemma 2.

Lemma 1. *Let*

$$\mathbf{jp}^0 \sim \mathbf{jp}^1, \quad \mathbf{q}^s := \mathbf{j}(\mathbf{p}^s \circ \rho^s), \quad \rho^s(u, v) := (u + b^s(u)v, a^s(u)v).$$

If $a^0 \sim a^1$ and *either*
(i) $\mathbf{p}^s(u, 0)$ *join* C^2 *at* $u = 1 - s$ *and* $b^0 \sim b^1$ *or*
(ii) $b^s(u^s) = 0$
then $\mathbf{jq}^0 \sim \mathbf{jq}^1$.

Proof Differentiating the expansion

$$\mathbf{j}(\mathbf{p}^s \circ \rho^s)(u, 0) = \mathbf{p}^s(u, 0) + (b^s(u)\partial_u \mathbf{p}^s(u, 0) + a^s(u)\partial_v \mathbf{p}^s(u, 0))v.$$

with respect to u yields

$$\partial_u \mathbf{p}^s(u, 0) + (\overline{b^s(u)\partial_{uu}\mathbf{p}^s(u, 0) + \partial_u b^s(u)\partial_u \mathbf{p}^s(u, 0)} + a^s(u)\partial_u\partial_v \mathbf{p}^s(u, 0) + \partial_u a^s(u)\partial_v \mathbf{p}^s(u, 0))v. \qquad (2)$$

If (i) holds then the overlined terms are continuous; if (ii) holds then only the smooth join of a^0 and a^1 matters. |||

Define the quadratic and matrix maps

$$\mathbf{a} := [a_0, a_1, a_2], \quad \mathbf{1} := [1, 1, 1], \quad \mathbf{a}B^2(u) := \sum_{i=0}^{2} a_i B_i^2(u),$$

$$\mathbf{Q}_j := \begin{pmatrix} \mathbf{q}_{0j} \\ \mathbf{q}_{1j} \\ \mathbf{q}_{2j} \\ \mathbf{q}_{3j} \\ \mathbf{q}_{4j} \end{pmatrix}, \; M(\mathbf{a}) := \begin{pmatrix} a_0 & 0 & 0 \\ \frac{a_1}{2} & \frac{a_0}{2} & 0 \\ \frac{a_2}{6} & \frac{2a_1}{3} & \frac{a_0}{6} \\ 0 & \frac{a_2}{2} & \frac{a_1}{2} \\ 0 & 0 & a_2 \end{pmatrix}, \; \mathbf{P}_j := \begin{pmatrix} \mathbf{p}_{0j} \\ \mathbf{p}_{1j} \\ \mathbf{p}_{2j} \end{pmatrix}, \; j = 0, 1.$$

Then the G^1 join characterization, Lemma 2, is easily verified by substitution.

Lemma 2. *Let*

$$(\mathbf{jp})(u, v) := \sum_{i=0}^{2}\sum_{j=0}^{1} \mathbf{p}_{ij} B_i^2(u) B_j^2(v), \quad \rho(u, v) := (u, v\mathbf{a}B^2(u))$$

and $\mathbf{q}_{ij}$, $i = 0, \ldots, 4$, $j = 0, 1$, *the BB-coefficients of* $\mathbf{jq} := \mathbf{j}(\mathbf{p} \circ \rho)$ *in bi-degree* $(4, 2)$ *form. Then*

$$\mathbf{Q}_0 = M(\mathbf{1})\mathbf{P}_0, \quad \mathbf{Q}_1 = M(\mathbf{a})\mathbf{P}_0 + M(\mathbf{1} - \mathbf{a})\mathbf{P}_1. \qquad (3)$$

If $\check{\mathbf{q}}_{ij}$, $i = 0, \ldots, 4$, $j = 0, 1$, *are BB-coefficients of the expansion of* $\mathbf{q}(u, v)$ *represented in bi-degree* $(4, 4)$ *form then* $\check{\mathbf{q}}_{i0} = \mathbf{q}_{i0}$ *and* $\check{\mathbf{q}}_{i1} = \frac{1}{2}(\mathbf{q}_{i0} + \mathbf{q}_{i1})$, $i = 0, \ldots, 4$.

3. Constructions for T-configurations

A T-junction represents a change of quad-strip density when traversing in vertical direction. It is therefore natural to reparameterize along the way and this is achieved by a carefully chosen quadratic re-parameterization of the cross-derivatives along the vertical boundaries of the τ-surface cap. While the quadratic re-parameterization increases the final patch degree to 4, Section 3.1 will show that it enables nested refinement.

Applying the partial B-to-BB conversion to the τ_n-nets, generates a frame of tensor-borders denoted as $\mathbf{t}_i$, $\underline{\mathbf{t}}_j$, $\bar{\mathbf{t}}_k$ with appropriate subscripts displayed in Fig. 8a, Fig. 9a, Fig. 10a. Not all tensor-borders need to be labeled since the τ_0 and τ_1 construction will be symmetric with respect to the central vertical line and τ_2 diagonally symmetric.

The tensor-borders $\underline{\mathbf{t}}$, $\bar{\mathbf{t}}$ are only suitable for inclusion into a C^1 cap once they are split in the ratios shown in Fig. 8b, Fig. 9b, Fig. 10b. After splitting, the frame is not consistent at some corners and so requires reparameterization. While choosing a linear $a(u)$ and a quadratic $b(u)$ that is zero at the split point allows constructing a C^1 bi-3 cap (cf. [5]) this 'zero-b' property cannot be preserved under refinement: due to (2), the refined construction is generally only C^0. The choice $b(u) := 0$,

$$\rho^s :[0..1]^2 \to \mathbb{R}^2, \quad \rho^s(u, v) := (u, a^s(u)v), \quad a^s(u) := \mathbf{a}^s B^2(u), \qquad (4)$$

allows refinement but forces $a(u)$ to be of degree 2 and the final surfaces to be of degree 4 in the u-direction. The tensor-borders $\mathbf{t}^s$ are therefore reparameterized with the $\rho^s(u, v)$ of (4). (In Fig. 8b, Fig. 9b, Fig. 10b, the black arrow heads indicate the u-direction, hollow ones the v-direction.)

τ_0 construction (Fig. 8, 'left-right' symmetric)

1. Reparameterize the tensor-borders $\mathbf{t}^0$, $\mathbf{t}^1$ (Fig. 8a) using

$$\mathbf{a}^0 := [1, 1, \frac{3}{4}], \quad \mathbf{a}^1 := [\frac{3}{4}, \frac{1}{2}, \frac{1}{2}]$$

and represent $\mathbf{j}(\mathbf{t}^s \circ a^s)$ in bi-degree $(4, 2)$ form.

2. Together with the right counterpart and the *top* and *bottom* tensor-borders (split and degree-raised to 4 in the vertical direction to be consistent with the reparameterization) this defines a depth 1 tensor-border frame for the cap Fig. 8c.

3. Set the central vertical layer BB-coefficients, marked ∘ in Fig. 8c, as midpoints of their horizontal neighbors.

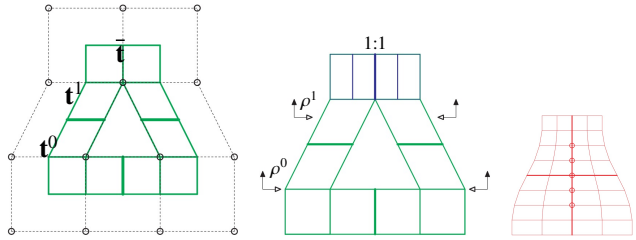

Fig. 8. τ_0-construction. (a) τ_0-net nodes and frame of input tensor-borders. (b) Reparameterization and splitting the tensor-borders. (c) BB-net of the cap of degree 2 in the horizontal and degree 4 in the vertical direction.

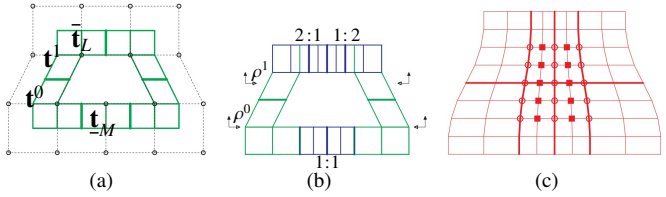

Fig. 9. τ_1-construction. (a) τ_1-net nodes and frame of input tensor-borders. (b) Reparameterization and splitting the tensor-borders. (c) BB-net of the cap of degree 2 in the horizontal and degree 4 in the vertical direction.

Fig. 24 in Appendix lists the complete set of formulas to enable the readers both to confirm their understanding of Lemma 2 and to directly implement the τ_0 construction.

τ_1 construction (Fig. 9, 'left-right' symmetry)
1. and 2. are very similar to the τ_0 construction, but use

$$\mathbf{a}^0 := [1, 1\frac{5}{6}], \quad \mathbf{a}^1 := [\frac{5}{6}, \frac{2}{3}, \frac{2}{3}]. \tag{5}$$

and split top and bottom tensor-borders as shown in Fig. 9b.
3. The BB-coefficients, marked as boxes in Fig. 9c, complete two vertical sequences of BB-coefficients as C^1-connected quadratic curves represented in degree 4 form, i.e. with BB-coefficients $\mathbf{p}_i^0, \mathbf{p}_i^1, i = 0, \dots, 4$ that satisfy

$$\mathbf{p}_0^1 = \mathbf{p}_4^0 := -\frac{1}{2}(\mathbf{p}_0^0 + \mathbf{p}_4^1) + \mathbf{p}_1^0 + \mathbf{p}_3^1,$$
$$\mathbf{p}_2^0 := -\frac{1}{2}\mathbf{p}_0^0 + \frac{4}{3}\mathbf{p}_1^0 + \frac{1}{6}\mathbf{p}_4^1,\ \mathbf{p}_3^0 := -\frac{1}{2}\mathbf{p}_0^0 + \mathbf{p}_1^0 + \frac{1}{2}\mathbf{p}_4^1$$

and $\mathbf{p}_1^1, \mathbf{p}_2^1$ defined by symmetric formulas.
The BB-coefficients, marked as circles, are the 1:1, respectively 2:1 averages of their horizontal neighbors.

τ_2 construction. (Fig. 10, diagonal symmetry)
The ending of quad-strips in both directions creates additional complexity compared to τ_0 and τ_1.
1. Tensor-borders $\mathbf{t}^s$, $s = 0, \dots, 7$ (and their diagonally symmetric counterparts, see Fig. 10b) are reparameterized using a^s defined in (5) and writing the new maps $a(u)$ with a subscript:

$$a_4(u) = a_0(u) := a^0(\frac{2}{3}u),\ a_5(u) = a_1(u) := a^0(\frac{2}{3}(1-u)+u),$$
$$a_6(u) = a_2(u) := a^1(\frac{1}{3}u),\ a_7(u) = a_3(u) := a^1(\frac{1}{3}(1-u)+u),$$

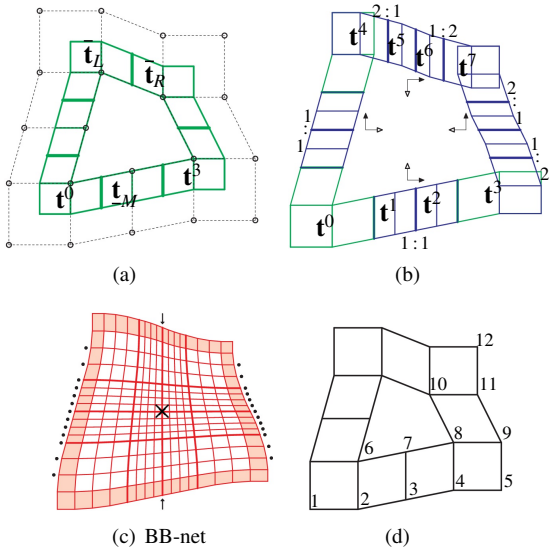

Fig. 10. Construction of the τ_2-cap. (a) τ_2-net nodes and tensor-border frame. (b) split tensor-border; reparameterizations; (c) BB-net of bi-4 cap (consisting of 4×4 pieces); (d) labeling of τ_2-net.

and $\mathbf{j}(\mathbf{t}^s \circ a_s)$ (light-red underlaid in Fig. 10c) is represented in bi-4 form.
2. Completion of the cap, see Fig. 10c.

(i) Construct the central vertical layer marked by the vertical arrows as two C^2-connected cubic curves represented as four pieces of degree 4 (see Appendix).
(ii) Apply the same procedure to complete the horizontal layers of BB-coefficients marked by dots.
(iii) Repeat the Steps (i) and (ii) creating first the central horizontal layer and then the 13 vertical layers.
(iv) To make a construction diagonally symmetric, average the BB-coefficients from Steps (i)+(ii) with those from Step (iii).

The algorithm is executed with the central coefficient, marked as $\times$ in Fig. 10c, undetermined. Then $\times$ is set as the minimizer of the functional $\mathcal{F}_3 f := \int_0^1 \int_0^1 \sum_{i+j=3, i,j\geq 0} \frac{3!}{i!j!} (\partial_s^i \partial_t^j f(s,t))^2 ds dt$, summed over all 16 patches of the cap. The weights w_j of the resulting formula $\times := \sum_{j=1}^{20} w_j \mathbf{c}_j$ in terms of the τ_2-net vertices $\mathbf{c}_j$ labeled in Fig. 10d are (adjusted to have 3 digits and to sum to 1; due to diagonal symmetry only 12 of them are listed):

$w_{1\dots6}$:	0.011	−0.015	−0.02	−0.013	0.007	0.074
$w_{7\dots12}$:	0.215	0.181	−0.022	0.285	−0.019	0.002.

This completes the τ-constructions for T-configurations.

3.1. Nested refinement

A refinement is nested if the finer representation can reproduce any coarser representation. Refinement is useful since it exposes additional degrees of freedom. Since the caps are internally C^1, their refinement amounts to standard knot insertion into a C^1 spline. (For τ_0, τ_1 the spline is of bi-degree $(4, 2)$, for τ_2 of degree bi-4. Since the top and bottom tensor-borders of degree 2 for τ_0, τ_1 are not reparameterized, only the degree

needs to be raised to match that of the cap.) Lemma 3 shows how to refine $\mathbf{jq} := \mathbf{jp} \circ \rho$ for τ_0, τ_1 and τ_2 across the remaining G-curves where $\rho(u, v) := (u, a(u)v)$. The following is then easily verified.

Lemma 3. *Partitioning with breakpoints* $0, h, \ldots$ *in the v-direction, and* d^k *in the u-direction, the pieces along the boundary* $v = 0$ *and between* $u = d^k$ *and* $u = d^{k+1}$ *are defined by maps*

$$\begin{aligned} \mathbf{p}^k(u, v) &:= \mathbf{p}(d^k(1-u) + d^{k+1}u, hv), \\ \mathbf{q}^k(u, v) &:= \mathbf{q}(d^k(1-u) + d^{k+1}u, hv) \\ \dot{\rho}^k(u, v) &:= (u, a(d^k(1-u) + d^{k+1}u)v). \end{aligned} \tag{6}$$

Then $\mathbf{jq}^k = \mathbf{j}(\mathbf{p}^k \circ \dot{\rho}^k)$.

Equations (6) and (3) suggest the following

Refinement Algorithm (see Fig. 11)

- Apply de Casteljau's algorithm to obtain the BB-coefficients of $\mathbf{q}_k$ from $\mathbf{p}$. The new, unconstrained BB-coefficients are marked as red bullets.
- Apply standard bi-2 spline refinement to obtain new, unconstrained control-points marked as green bullets.
- Apply de Casteljau's algorithm to obtain the BB-coefficients of $\dot{\rho}_k$ from ρ.
- Apply B-to-BB-conversion followed by (3) to the new green control points to set the gray-underlaid BB-coefficients.

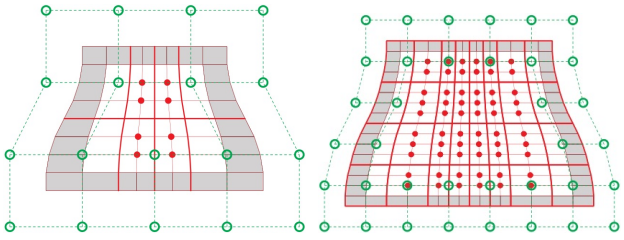

Fig. 11. τ-refinement and degrees of freedom. ○ = regular bi-2 B-spline control points; ● = cap bi-4 BB-coefficients.

Fig. 11 shows the two types of unconstrained handles for the designer: bi-2 spline control points and bi-4 BB-coefficients.

4. Construction for star-configurations

The focus of this section is to improve the shape of existing refinable constructions that generalize bi-2 subdivision to multi-sided configurations. A key new concept is *geometric tuning* of all subdivision rings right from the start, rather than the classical (analytic) eigen-spectrum tuning of subdivision algorithms that has its main effect in the limit. The idea here is to use extra degrees of freedom in bi-3 rings to improve the shape and to accelerate the convergence so that two to three rings cover what otherwise requires four to six rings of Doo-Sabin or augmented

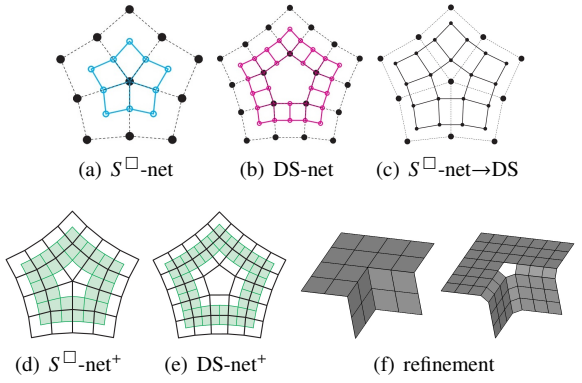

Fig. 12. Star-configurations. (a) $S^{\square}$-net and its degree 2 tensor-border. (b) DS-net and its degree 2 tensor-border. (c) Switching from an $S^{\square}$-net (big bullets) to a DS-net (small bullets). (d,e) $S^{\square}$-net$^+$ and DS-net$^+$ each defining a bi-2 surface ring. (f) geometric embedding of a $S^{\square}$-net$^+$ and its refined net whose innermost quads are the DS-net obtained from an $S^{\square}$-net.

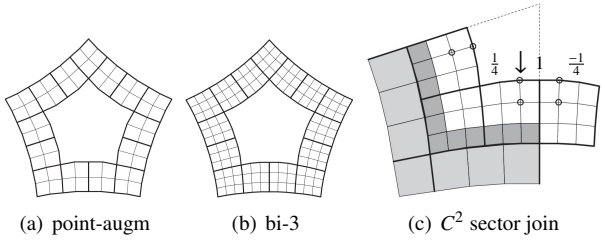

Fig. 13. Geometric tuning. (a) Bi-2 ring of point-augmented subdivision [4] is (b) degree-raised to bi-3 and (c) sectors are joined C^2 by setting the BB-coefficients marked by ○ as weighted average of nearby coefficients. The weights of one BB-coefficient (marked by ↓) are specified.

subdivision. Since a finer subdivision is rarely needed in practice, the two to three rings form a cap, completed by the tiny bi-4 construction in Section 4.1.

Knot insertion, i.e. standard bi-2 spline control net refinement, converts an $S^{\square}$-net to a DS-net as illustrated in Fig. 12c, so that the construction can assume a DS-net as starting point. The construction is based on a 'tuned' (improved) version of point-augmented bi-2 subdivision [4]. For while point-augmented subdivision generates better-shaped surfaces than Doo-Sabin subdivision as demonstrated in Fig. 4b,c, its highlight lines still look too sharp. Unlike conventional *analytic tuning* of subdivision surfaces [18, 19, 20] that focuses on the vicinity of the extraordinary point (by modifying the eigen-spectrum via the refinement rules), the new *geometric tuning* modifies the functions associated with the control net. That is, geometric tuning modifies the generation of surface rings from the refined nets but does not change the refinement rules. Geometric tuning not only improves the highlight lines but stays close to the already reasonable global shape of point-augmented surfaces.

Geometric tuning. See Fig. 13.

- Degree-raise the bi-2 surface rings (Fig. 13a) of augmented-subdivision to bi-3 (Fig. 13b).
- Except for the input ring, join the sectors of each ring C^2 by setting the BB-coefficients marked by ○ in Fig. 13c to

an average of neighbors (weights indicated for one ∘).
- C^1-extend the BB-coefficients of the coarser surface ring (underlaid light-gray in Fig. 13c) and split to form the dark-gray underlaid BB-coefficients of the bi-3 patches.

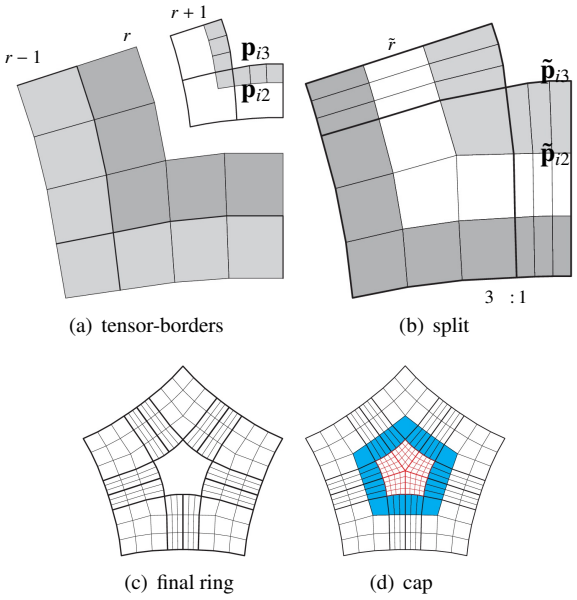

(a) tensor-borders (b) split (c) final ring (d) cap

Fig. 14. Double-speed tuned subdivision. (a) Tensor-borders of subdivision ring. (b) Non-uniform split. (c) Final surface ring. (d) Adjustment and tiny cap.

Geometric tuning additionally supports faster ring contraction by splitting the C^1-extension (dark-gray underlaid BB-net in Fig. 14a) of every second ring (light-gray underlay) in the ratio 3 : 1, see Fig. 14b; the first order Hermite data of the next ring (light-gray underlay in Fig. 14b) is then 'resized' (stretched) to match double-speed contraction: with the labels in Fig. 14a,b: $\tilde{\mathbf{p}}_{i3} := \mathbf{p}_{i3}$, $\tilde{\mathbf{p}}_{i2} := -2\mathbf{p}_{i3} + 3\mathbf{p}_{i3}$. Fig. 14c shows the final ring. Tuned subdivision can be further accelerated by using a third ring and splitting 7 : 1 – but this results in unsatisfactory shape.

Lemma 4. *Geometrically-tuned and double-speed tuned subdivision surface are C^1.*

Proof By construction, tuned surface rings, also when double-speed, are internally C^1 (sectors even join C^2) and join C^1 to one another. The control net refinement rules of tuned subdivision are the same as of [4]. Therefore the subdivision matrix A and its eigen-spectrum is the same, and for double-speed tuned subdivision it is A^2. The necessary conditions for a C^1 limit at the extraordinary point are therefore satisfied as in [4]. Moreover, the characteristic maps are obtained by geometric tuning and speed-doubling of the (sequence of) characteristic map(s) of [4] and are injective. By [21] this completes the proof. |||

Fig. 15 juxtaposes the characteristic maps of [4], tuned and double-speed subdivision for $n = 3, 5, 6$.

4.1. A finite central cap

While standard subdivision (also Catmull-Clark-subdivision) introduce shape artifacts already in the first steps [22], geometrically tuned subdivision preserves the highlight line quality of the regular bi-2 surface. Therefore it makes sense in practice to stop the refinement and smoothly cap the remainder of an otherwise infinite sequence of contracting rings.

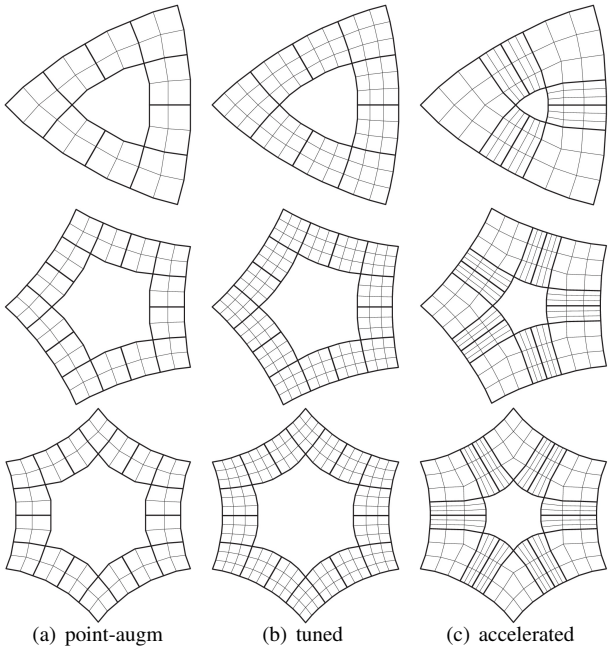

(a) point-augm (b) tuned (c) accelerated

Fig. 15. Characteristic subdivision maps for point-augmented, tuned and accelerated subdivision, respectively.

Skipping, for the last ring, step (2) of geometric tuning (C^2 joining the bi-3 patches), degree-raising and C^1 extending backwards the first-order Hermite data (Fig. 12c) yields the underlaid BB-coefficients in Fig. 14d. This frame admits the surface cap defined in [8]. Since the degree near T-gons is 4, we choose the cap to have single bi-4 patch per sector for $n > 4$ and bi-3 for $n = 3$. According to [8] the bi-4 patches can alternatively be replaced by 2×2 bi-3 patches, and, for quad meshes with valencies restricted to $n = 3, 5$, a special algorithm can additionally yield a single bi-3 patch per sector when $n = 5$.

Central cap algorithm
- As transition, generate one ring of tuned subdivision,
- then three rings of double-speed subdivision, then
- apply the central cap of [8].

Compared to immediate applying to double-speed subdivision, the initial regular-speed transition layer prevents shape artifacts that arise for some challenging configurations.

5. Assessment and Comparison

To calibrate what quality of surface can be expected for τ-nets and star-nets when the surrounding surface is a C^1 bi-2 spline, the nets are surrounded by one layer of quads that defines one input bi-2 outer ring. This extended net is called net$^+$. The central cap is always colored red.

Fig. 16 demonstrates that inclusion of T-gons into quad-dominant mesh does not disrupt the highlight line distribution of such surfaces. Note in Fig. 16b that, as typical for many

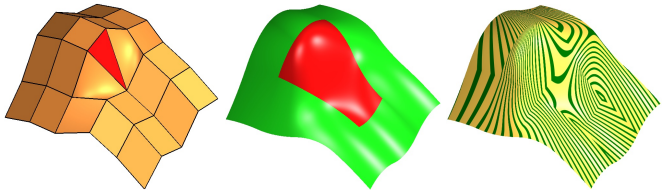

(a) τ_0

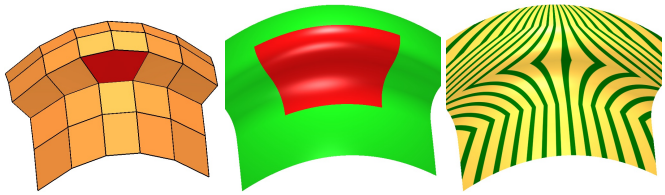

(b) τ_1

(c) τ_2

Fig. 16. From *left* to *right*: T-net$^+$, frame and cap, highlight lines.

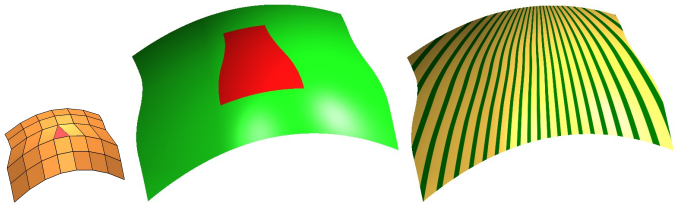

(a) τ_0-surface

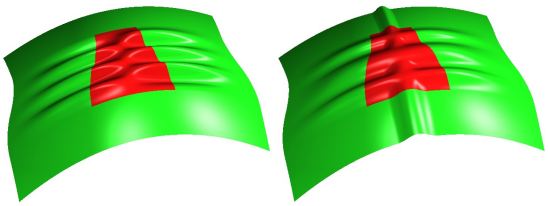

(b) Localized edits

Fig. 17. τ_0-surfaces and edits.

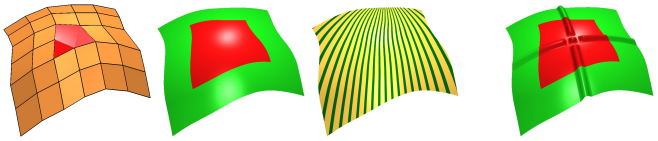

(a) τ_2-surface (b) Localized edit

Fig. 18. τ_2-surface

inputs, the highlight lines of the central cap are smoother than those of the surrounding bi-2 surface.

Fig. 17 for τ_0 and Fig. 18 for τ_2 (τ_1 was shown in Fig. 3) demonstrate the resilience of the algorithms when the input net$^+$ is convex. Such input represents a hard challenge for most constructions as illustrated in Fig. 3b,c.

Fig. 19 shows a two-beam configuration as in Fig. 4. Zooming to the central yellow region reveals sharp highlight line turns for point-augmented subdivision (see Fig. 19b), whereas geometrically tuned subdivision results in smoothed highlight lines displayed in Fig. 19c. As is often the case, smoothing sharp features resulats in slight waviness. Similar sharp vs wavy features reappear in the tiny area marked by □ in Fig. 19a where the last ring reverts back to degree bi-2 to admit the cap of [8]; see the enlarged views of Fig. 19d,e,f.

Fig. 20 tests the algorithm for $n = 3, 6$. As in earlier examples, for $n = 3$ the sharp highlight line turns (kinks) are already part of the given regular bi-2 surface and not of the cap. Note that double-speed is so fast that for all practical purposes two rings suffice before placing the red cap (that is so small that the red is not visible in Fig. 20b). Fig. 20g,h illustrate that complex saddles pose no problem

Fig. 21 shows a design that combines the surfaces developed for T- and star-configurations and features a localized elevation at the center in the horizontal direction.

Comparing generalizations of bi-2 C^1 splines with those of bi-3 C^2 splines makes as much sense as comparing bi-2 C^1 splines to bi-3 C^2 splines. The first two rows of Fig. 22 group together (for the input net of Fig. 3a) the new τ_1-surface and the surface generated by [5] with C^1 bi-2 and C^2 bi-3 surfaces (for the input net re-connected to yield a locally regular net). All four options have good highlight line distributions. By contrast, rows three and four of Fig. 22 accentuate the difference in quality between bi-3 and bi-2 splines and exhibit a corresponding difference between the new τ_1-surface and [5] for the input net of Fig. 16b, *left*.

6. Conclusion

A new class of spline surfaces was presented that generalizes bi-2 C^1 splines to polyhedral control nets with star- and T-configurations. It is based on a rarely-used choice of geometric continuity and a new geometrically-tuned accelerated subdivision. Notably the surfaces are refinable for (iso-geometric) engineering analysis and for localized geometry editing. While applications to engineering analysis are future work, the use of refinability for geometric modeling was demonstrated in numerous examples. Since in practice subdivision is stopped after five or six steps, double-speed contraction delivers in essence a finite refinable spline complex that admits completion by a G^1 cap. Together this considerably increases the flexibility of interactive polyhedral modeling with smooth output surfaces.

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

References

[1] D. Doo, M. Sabin, Behaviour of recursive division surfaces near extraordinary points, Computer-Aided Design 10 (1978) 356–360.

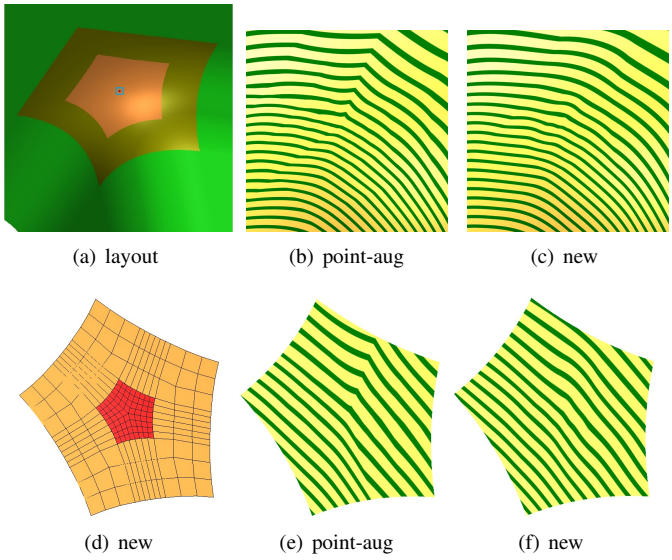

(a) layout (b) point-aug (c) new

(d) new (e) point-aug (f) new

Fig. 19. Input Fig. 4a. (a) Layout from outside inwards: bi-2 input, bi-3 ring, 3 2× speed rings, cap. (b,c) zoom to 2× speed rings using (b) point-augmented subdivision [4] and (c) new construction. (d,e,f) zoom to the area outlined by a □ in (a). (d) BB-net of the last bi-3 ring and bi-4 cap of the 2× speed construction.

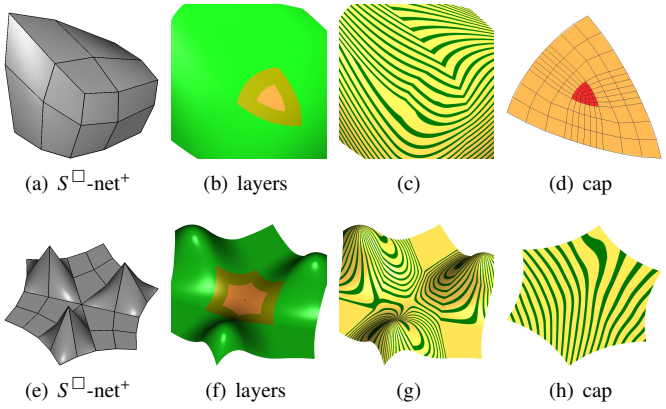

(a) $S^{\square}$-net$^+$ (b) layers (c) (d) cap

(e) $S^{\square}$-net$^+$ (f) layers (g) (h) cap

Fig. 20. *Top* row: valence $n = 3$. (c) The main flaws in the highlight line distribution of the regular bi-2 C^1 surface are extended to the first transition towards the 3-sided cap. (d) BB-net of the last ring and tiny cap (too small to be visible in (b)). *Bottom* row: valence $n = 6$.

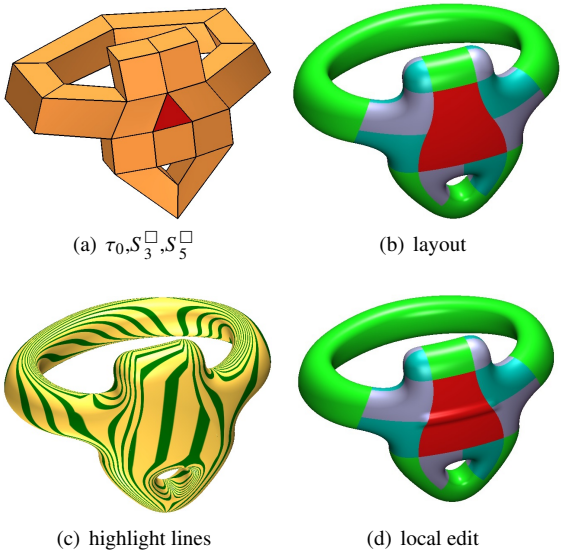

(a) τ_0,$S^{\square}_3$,$S^{\square}_5$ (b) layout

(c) highlight lines (d) local edit

Fig. 21. Design with T- and star-configurations. (a) Mesh with τ_0-gon and irregular nodes of valencies 3 and 5. (b) Surface layout with regular bi-2 context, τ_1-cap and multi-sided caps. (d) Localized edit across T-configuration.

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

[3] K.-P. Beier, Y. Chen, Highlight-line algorithm for realtime surface-quality assessment, Computer-Aided Design 26 (4) (1994) 268–277.

[4] K. Karčiauskas, J. Peters, Point-augmented biquadratic C^1 subdivision surfaces, Graphical Models 77 (2015) 18–26.

[5] K. Karčiauskas, D. Panozzo, J. Peters, T-junctions in spline surfaces, ACM Tr on Graphics, ACM Siggraph 36 (5) (2017) 170:1–9.

[6] C. T. Loop, S. Schaefer, G ^{2} tensor product splines over extraordinary vertices, Comput. Graph. Forum 27 (5) (2008) 1373–1382.

[7] K. Karčiauskas, J. Peters, Improved shape for multi-surface blends, Graphical Models 8 (2015) 87–98.

[8] K. Karčiauskas, J. Peters, Smooth multi-sided blending of biquadratic splines, Computers & Graphics 46 (2015) 172–185.

[9] U. Reif, Biquadratic G-spline surfaces, Computer Aided Geometric Design 12 (2) (1995) 193–205.

[10] J. Deng, F. Chen, L. Jin, Dimensions of biquadratic spline spaces over t-meshes, J. Computational Applied Mathematics 238 (2013) 68–94.

[11] C. Zeng, F. Deng, X. Li, J. Deng, Dimensions of biquadratic and bicubic spline spaces over hierarchical t-meshes, J. Computational Applied Mathematics 287 (2015) 162–178.

[12] C. Zeng, M. Wu, F. Deng, J. Deng, Dimensions of spline spaces over non-rectangular t-meshes, Adv. Comput. Math 42 (6) (2016) 1259–1286.

[13] W. Jakob, M. Tarini, D. Panozzo, O. Sorkine-Hornung, Instant field-aligned meshes, ACM Trans. Graph 34 (6) (2015) 189:1–189:15.

[14] N. Schertler, M. Tarini, W. Jakob, M. Kazhdan, S. Gumhold, D. Panozzo, Field-aligned online surface reconstruction, ACM Trans. Graph 36 (4) (2017) 77:1–77:13.

[15] G. Farin, Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide, Academic Press, 1988.

[16] T. D. DeRose, Necessary and sufficient conditions for tangent plane continuity of Bezier surfaces, Computer Aided Geometric Design 7 (1) (1990) 165–179.

[17] J. Peters, Geometric continuity, in: Handbook of Computer Aided Geometric Design, Elsevier, 2002, pp. 193–229.

[18] L. Barthe, L. Kobbelt, Subdivision scheme tuning around extraordinary vertices, Computer Aided Geometric Design 21 (2004) 561–583.

[19] G. Umlauf, Analysis and tuning of subdivision algorithms, in: Proceedings of the 21st Spring Conference on Computer Graphics, SCCG '05, ACM, New York, NY, USA, 2005, pp. 33–40.

[20] 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.

[21] J. Peters, U. Reif, Subdivision Surfaces, Vol. 3 of Geometry and Computing, Springer-Verlag, New York, 2008.

[22] U. H. Augsdörfer, N. A. Dodgson, M. A. Sabin, Artifact analysis on B-splines, box-splines and other surfaces defined by quadrilateral polyhedra, Computer Aided Geometric Design 28 (3) (2011) 177–197.

Appendix: explicit stencils for implementation

For $i = 0, \dots, 4$, $s = 0, \dots, 3$, consider pieces $\mathbf{p}_i^s$ of degree 4 such that $\mathbf{p}_0^0$, $\mathbf{p}_1^0$, $\mathbf{p}_3^3$, $\mathbf{p}_4^3$ are fixed. Determine position and first derivative data in degree 3 form:

$$\hat{\mathbf{p}}_0^0 := \mathbf{p}_0^0, \quad \hat{\mathbf{p}}_1^0 := -\frac{1}{3}\mathbf{p}_0^0 + \frac{4}{3}\mathbf{p}_1^0, \hat{\mathbf{p}}_3^3 := \mathbf{p}_4^3, \quad \hat{\mathbf{p}}_2^3 := -\frac{1}{3}\mathbf{p}_4^3 + \frac{4}{3}\mathbf{p}_3^3.$$

Then the stencils of Fig. 23 yield two cubic curves, C^2-connected at the point marked as ×, split in the displayed ra-

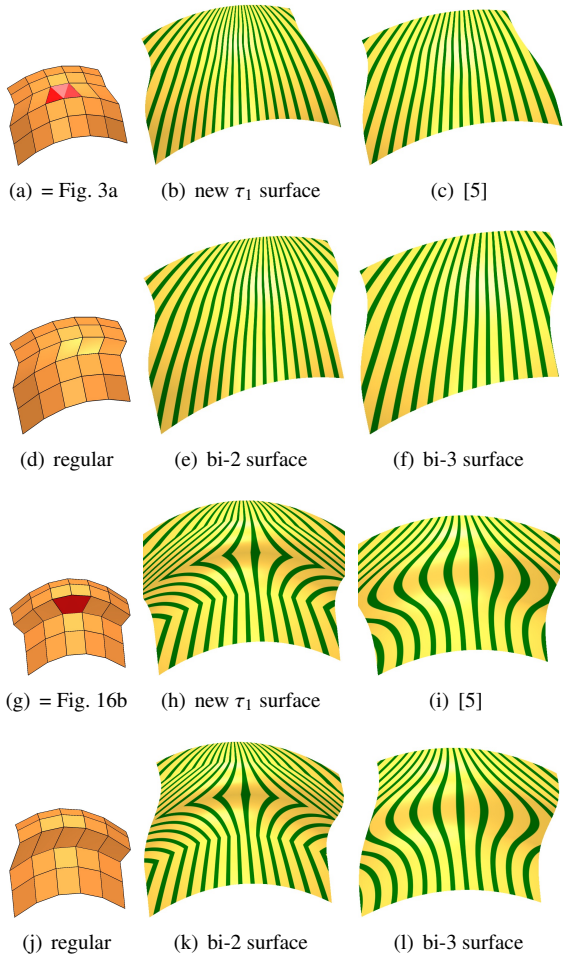

Fig. 22. Comparison of the new τ_1 surface vs the construction of [5]; and the corresponding comparison of bi-2 vs bi-3 surfaces with regular control nets obtained by biased regular re-connections (d) and (j) of the T-meshes (a) and (g) from Fig. 3a and Fig. 16b, *left*. The T-constructions pick up (preserve) the shape of comparable tensor-product splines.

tios (2:1, 1:2) and represented as four pieces $\hat{\mathbf{p}}_i^s$, $i = 0,\dots,3$, $s = 0,\dots,3$. The cubics are degree-raised to quartics.

Fig. 24 lists all the formulas needed to assemble the τ_0-cap constructed in Section 3. Each BB-coefficient of the two left-most vertical layers of Fig. 8c is expressed as a weighted sum of the τ_0-net nodes (lower right inset of Fig. 24) listed as a bracketed group; the 8 numbers in each group represent these weights scaled by 96. The two right vertical layers are obtained by symmetry. The BB-coefficients of the central vertical layer are the midpoints of their left and right neighbors.

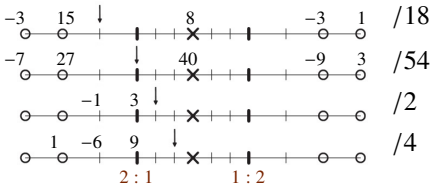

Fig. 23. Four construction stencils (weights sum to 1) of C^2-connected layers that form two C^2-connected cubic curves in four pieces. The curves are defined by two coefficients (marked as the circles) at either end and the central BB-coefficient of the layer (marked as ×).

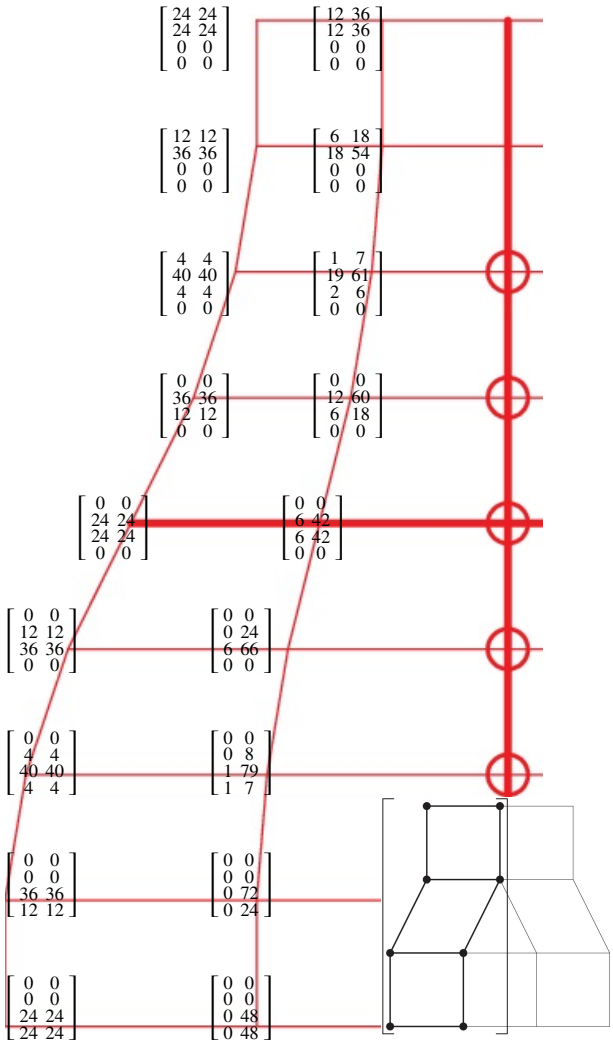

Fig. 24. Complete set of formulas for the τ_0-construction. Cf. Fig. 8c.