Localized G-splines for quad & T-gon meshes

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

[a] *Vilnius University*
[b] *University of Florida*

Abstract

Enriching tensor-product B-spline control nets by allowing T-gons (where strips of quadrilaterals start or end) and irregular nodes (where $n \neq 4$ quadrilaterals meet) reduces the requirements on quad-meshing and increases the flexibility for polyhedral design with associated smooth surfaces. This paper introduces a family of piecewise polynomial, geometrically continuous surface constructions that yield good highlight line distributions also in the presence of *irregular nodes next to a T-gon*. Such tight juxtaposition can further reduce the quad-meshing requirements and increase the space of polyhedral design control structures. The surfaces can be chosen to cover T-gons with G^1 caps of degree bi-4 – or with caps of degree bi-3 that are almost G^1 and preserve the good highlight line distribution of the bi-4 G^1 surfaces.

Keywords: T-junction, smooth surfaces, highlight line distribution, GT-splines, B-splines

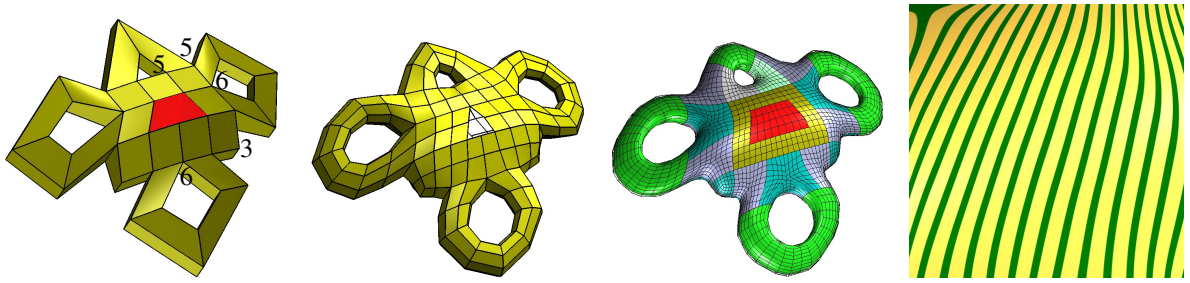

(a) input net (b) refined net (c) BB-net + surface (d) highlight lines

Figure 1: Modeling with T-gons despite nearby irregular nodes. (a) Input control net with red T-gon and neighboring irregular nodes of valences 3, 5, 6. (b) T-refined input net. (c) Surface with Bernstein-Bézier net overlaid: bi-3 regular, bi-3 transition, bi-4 cap, and bi-4 multi-sided: n=3,5,6. (d) Zoom to the transition and cap for the everywhere bi-3 almost-G^1 construction: just like for the bi-4 construction, the highlight lines (see Beier and Chen (1994)) across the surface do not reveal the location of the T-gon.

1. Introduction

Extra freedom to model detail can be added to tensor-product spline surfaces by locally inserting additional knot lines. However this assumes a careful, strictly hierarchical surface design that globally keeps track of knot distances to assure compatibility. If the layout of the control net is given, by a prior model or by a designer using a polyhedral modeler, such a global knot distribution has to be discovered; and it may not exist as simple examples demonstrate, e.g. Fig. 2 of Karčiauskas et al. (2017).

A change in density can be indicated by a control net where a 4-sided facet, short *quad*, has one or more split edges. At the split point two quads on one side meet one facet on the other. This configuration is called a *T-junction*.

*Corresponding author
Email address: jorg.peters@gmail.com (Jörg Peters)

Meshes with T-junctions can arise from remeshing existing polyhedra, or by direct design. An m-gon, $m > 4$, in a quad-dominant mesh will be called *T-gon* if all its vertices are either T-junctions or of valence 4, and if it is surrounded by quads. The three practically relevant cases are: the T_1-gon with one T-junction (a pentagon Fig. 2a), the T_2-gon with two T-junctions (a hexagon Fig. 2b), and the T_3-gon with three T-junctions (a heptagon Fig. 2c). The T-gon and its surrounding layer of quads is called τ-net, or *localized T-net*. As opposed to the more expansive T-net required in Karčiauskas et al. (2017), the τ-net has no restrictions on the nodes that do not belong to the T-gon. That is, all mesh nodes marked as red disks or circles in Fig. 2 can be *irregular*, i.e. have more or fewer than four neighbors, and still admit a smooth surface. If any of the solid red disks are irregular, we call the configuration *tight*.

Allowing tight configurations reduces complex and global constraints of strict quad-meshing (Bommes et al., 2012; Vaxman et al., 2016) to simpler, more lenient ones (Li et al., 2006; Alliez et al., 2003; Kälberer et al., 2007; Jakob et al., 2015; Pietroni et al., 2016).

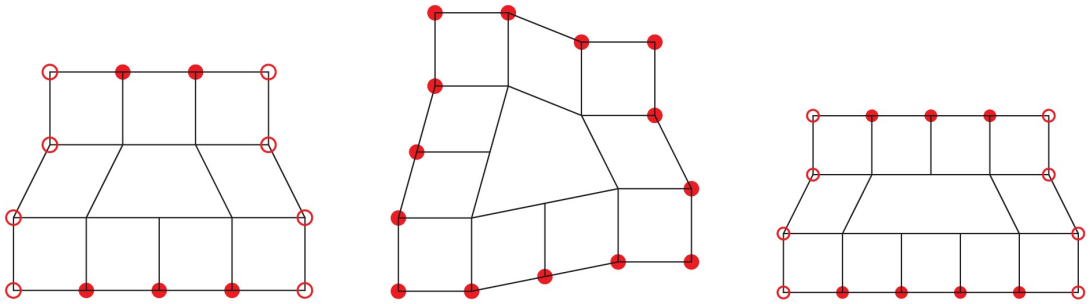

(a) τ_1-net (b) τ_2-net (c) τ_3-net

Figure 2: T_1-gon, T_2-gon T_3-gon with respectively 1,2,3 T-junctions Points marked as red disks may be irregular for a τ-net, but must be regular for GT_{KPP}. Points marked as red circles may be irregular in both cases.

There are several options to merge strips of tensor-product, quadrilateral patches; or to spawn additional strips for introducing more detail. T-splines Sederberg et al. (2003) impose a knot distribution across the surface, and hierarchical constructions like hierarchical B-splines, PHT-splines and LR-splines Kraft (1998); Kang et al. (2015); Dokken et al. (2013) require a global planar parameter space. Alternatively, Catmull and Clark (1978) subdivision can convert a T-gon into a pattern of one m-valent and multiple 3-valent points, one for each T-junction. However, Catmull-Clark subdivision yields visible artifacts. These artifacts rule out the subdivision approach for high-quality surface construction.

Another approach is to introduce rational singularities so that derivatives need not match. This is often called the Gregory-approach (Gregory, 1974; Loop et al., 2009; Hettinga and Kosinka, 2018). Unfortunately, as we explored new singular rational constructions for tight τ_1-net configurations, we found that, akin to other singular constructions, such rational constructions suffer from shape defects (see e.g. Fig. 18).

A third approach reflects the observation that switching density of control points is naturally associated with a change of domain variables, i.e. with geometric continuity. The geometrically smooth modeling of surfaces with T-junctions in Karčiauskas et al. (2017) yield a G^1 (normal-continuous) and strictly local construction, the geometrically-smooth T-, hence GT-spline, whose highlight line distributions are satisfactory over an obstacle course of challenging T-net configurations. We abbreviate this construction as GT_{KPP} in the following. However, GT_{KPP} requires T-junctions to be isolated: all nodes marked by red disks in Fig. 2 must have exactly four neighbors

Our approach improves on the GT_{KPP}-construction, drastically reducing the requirements on the local mesh: the outer nodes of the first ring of quads surrounding the T-gon can all be irregular! This yields more flexibility for the designer, respectively reduced constraints on any quad-dominant meshing algorithm. For example, the new construction allows tight configurations as in Fig. 1a that can not be modeled using GT_{KPP}. The τ-net in Fig. 1a is extreme in that all five nodes adjacent to the T-gon that need to be 4-valent for GT_{KPP} to work (red disks in Fig. 2a) are irregular. The resulting uniform highlight line distribution for this τ-net is remarkable since, when approached in a naive way, already any one of the five adjacent nodes being irregular creates surface quality problems. The complex configuration requires a judicious new localized refinement of the input net in the vicinity of the T-gon Fig. 1b and constructing a

transition Fig. 1c to the central GT-spline cap.

A second contribution of this paper is an alternative construction that provides the option of replacing the G^1 bi-4 cap by a formally G^0 bi-3 cap with essentially identical highlight lines and curvature distribution. Combined with the construction for irregular points from Karčiauskas and Peters (2015a) this allows us to generate everywhere bi-3 free-form surfaces with good highlight line distributions.

A third contribution of this paper is a new approach for T_2-gons. While a construction analogous to the T_1 construction yields good results also for T_3-gons, this type of approach fails for T_2-gons. Instead, we carefully remesh the T_2-net and leverage Karčiauskas and Peters (2015b) to obtain high-quality surfaces. In summary, the contributions of this paper are

- a quality-preserving refinement of the input net in the vicinity of a T-gon to minimize the footprint of the construction;
- a highlight line-preserving transition to the GT-spline cap;
- a quality-preserving remeshing of T_2-gons;
- an everywhere bi-3 construction for free-form surfaces with T-gons and irregular points that yields good highlight line distributions.

Overview. Sections 3, 4 and 5 present, respectively, the constructions of T_1-,T_3-,T_2-surfaces, Section 6 evaluates the constructions, also in comparison to alternative Gregory-type patches.

2. Definitions and Setup

As in GT_{KPP}, we will interpret regular sub-nets, i.e. quads surrounded by quads, two sharing an edge and four a vertex, as bi-cubic uniform single-knot B-spline control mesh. And we represent the polynomial spline pieces of bi-degree d ($d = 3$ or $d = 4$) that make up a GT-spline in Bernstein-Bézier form (short BB-form; see e.g. Farin (1988)):

$$\mathbf{f}(u,v) := \sum_{i=0}^{d}\sum_{j=0}^{d} \mathbf{f}_{ij} B_i^d(u) B_j^d(v)\,, \quad (u,v) \in [0..1]^2. \quad \text{Here } B_k^d(t) := \binom{d}{k}(1-t)^{d-k}t^k$$

are the Bernstein polynomials of degree d and $\mathbf{f}_{ij}$ are the BB-coefficients. The regular grid of BB-coefficients connected whenever their subscripts differ by 1 in exactly one subscript, is called the BB-net (Not to be confused with the much coarser T-net or τ-net defined in the Introduction that are akin to B-spline control nets de Boor (1978).) We refer to the switch from the B-spline control mesh to the BB-net as *B-to-BB conversion*. The algorithms presented in this paper effectively convert τ-nets to BB-form.

We will join adjacent pieces, called patches, with G^1 continuity. Patches join G^k if their cross-derivatives up to order k match along a common boundary after reparameterization. More formally their kth-order jets (one-sided Taylor expansion) match along their common boundary after a change of variables ρ. This characterization is equivalent to formulations of C^k continuity of manifolds in terms of charts, see e.g. Peters (2002). Specifically, two surface pieces $\tilde{\mathbf{f}}$ and $\mathbf{f}$ sharing a boundary curve $\mathbf{e}$ join G^1 if there is a suitably oriented and non-singular reparameterization $\rho : \mathbb{R}^2 \to \mathbb{R}^2$ so that the jets $\partial^k\tilde{\mathbf{f}}$ and $\partial^k(\mathbf{f} \circ \rho)$ agree along $\mathbf{e}$ for $k = 0, 1$. Although ρ is just a change of variables, its choice is crucial for the properties of the resulting surface. Throughout, we will choose $\mathbf{e}$ to correspond to the patch parameters $(u, 0 = v)$. Then the relevant Taylor expansion of the reparameterization ρ with respect to v is $\rho := (u + b(u)v, a(u)v)$ and the chain rule of differentiation yields the G^1 constraints

$$\partial_v\tilde{\mathbf{f}}(u,0) - a(u)\partial_v\mathbf{f}(u,0) - b(u)\partial_u\mathbf{f}(u,0) = 0. \tag{1}$$

3. T_1-surfaces

A T_1-gon is a pentagon with one T-junction, see Fig. 2a. Since we will leverage the center cap of the construction GT_{KPP}, we start by summarizing its main steps.

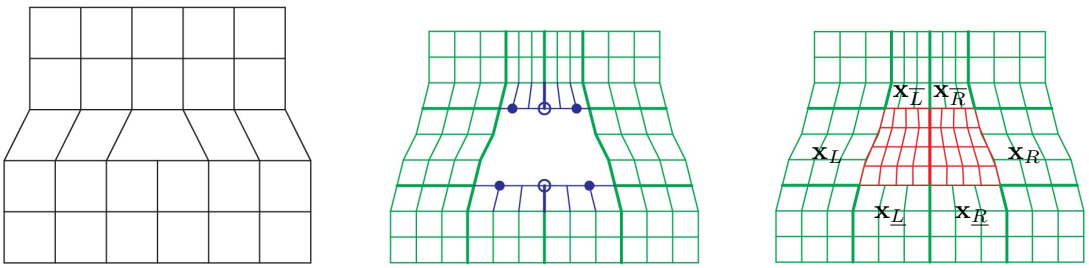

(a) T_1-net (b) structure of bi-3 frame (c) bi-3 ring + bi-4 cap

Figure 3: Construction of T_1 surface using GT_{KPP}.

3.1. *Review of the GT_{KPP} construction for T_1-nets*

Fig. 4a,b and Fig. 3 illustrate the steps of the GT_{KPP} construction. Using a left or right bias, there are two ways to virtually reconnect the T_1-net to the sub-nets sketched in Fig. 4a,b. B-to-BB conversion of these sub-nets yields the patches $\mathbf{x}_L$ and $\mathbf{x}_R$ of a bi-3 *frame* in Fig. 3c. The BB-coefficients marked as $\bullet$ in Fig. 3b C^1-extend the top and bottom curves of $\mathbf{x}_L$ and $\mathbf{x}_R$ (scaled by $\frac{1}{2}$ on *top*). The BB-coefficients marked as $\circ$ are crucial for good shape and are defined by special formulas published with GT_{KPP}. The two remaining BB-coefficients are chosen to C^2-connect the left and right pieces of the top boundary, respectively of the bottom boundary. The resulting frame of bi-3 polynomial patches is filled by a cap consisting of polynomial patches of degree bi-4 whose formulas are listed in Table 1 of GT_{KPP}. Fig. 5b shows the surface together with the BB-net (cf. Fig. 3c). Frame and cap construction honor the distinction between the horizontal and the density-changing vertical direction. By contrast, direction-agnostic multi-sided constructions typically have shape artifacts at T-junctions (as demonstrated for CC in GT$_{KPP}$).

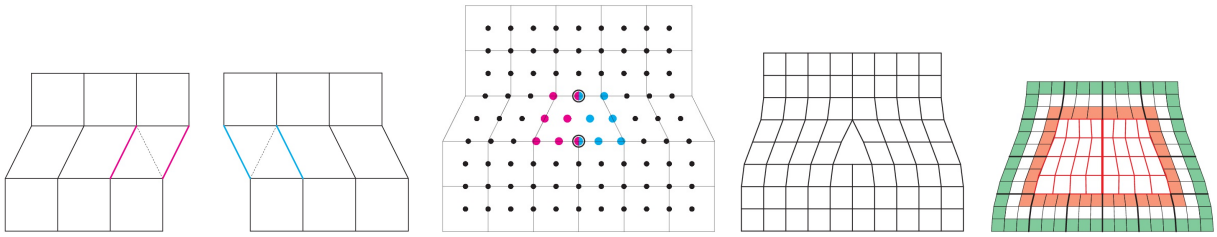

(a) (b) (c) refinement and averaging (d) refined net (e) bicubic transition ring

Figure 4: New localized τ_1 surface construction. (a,b) virtual sub-nets after re-connection. (e) BB-coefficients of transition ring (gold in Fig. 5e).

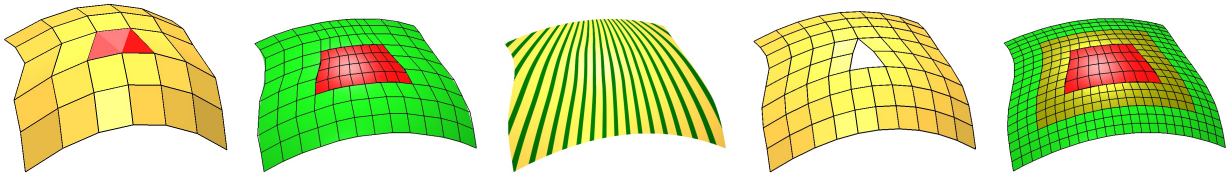

(a) T-net (b) BB-net (c) highlight lines (d) refined net (e) BB-net

Figure 5: (a) Input T-net with red T-gon. (b) Construction of GT_{KPP}: bi-3 frame and a red cap composed of two bi-4 patches. (c) highlight lines of (b) and (e) are alike. New construction: (d) refined net; (e) gold bi-3 transition from the bi-3 input to the bi-4 cap of GT_{KPP}.

3.2. *T_1-refinement and bi-3 transition*

B-spline refinement (subdivision) of the input net of Fig. 3a defines the B-spline control points marked as black bullets in Fig. 4c. The biased sub-net of Fig. 4a additionally defines the magenta and the biased sub-net of Fig. 4b

the cyan coefficients. At the circled overlapping locations in Fig. 4c the coefficients are averaged. This yields the *T-refined net* in Fig. 4d.

The T-refined net defines, by B-to-BB conversion, an *outer frame* (green-underlaid in Fig. 5e) that matches second-order Hermite data along the outer boundary of the original bi-3 frame in Fig. 5b. We focus on the *transition frame* (gold in Fig. 5e). The BB-net of the transition frame and of the central cap are shown in Fig. 4e: the green-gray underlaid part of the transition frame C^1 extends the outer frame; the inner red-gray underlaid part is obtained by splitting the original bi-3 frame of Fig. 5b. This guarantees G^1 continuity with the bi-4 cap of GT_{KPP}, i.e. the τ_1-construction is complete if we accept a bi-4 cap. The following bi-3 cap construction will use the same bi-3 Hermite data as the bi-4 cap, i.e. the same transition frame.

3.3. Constructing a bi-3 cap

Bi-3 splines are widely used to model free-form surfaces. While any given surface could be re-approximated to any desired precision by a quilt of C^0-joined bi-3 patches, the following construction requires only four bi-3 patches and minimally relaxes the G^1 constraints to satisfy the practical prescriptions of Autodesk (2018) in the spirit of (Sabin et al., 2003). To make the degree uniformly bi-3, the center cap in particular, we focus on the first-order Hermite extensions of $\mathbf{x}_L$ and $\mathbf{x}_R$. Fig. 6a labels the BB-coefficients of the $\mathbf{x}_L$ extension. Splitting the extension yields the BB-coefficients marked by circles in Fig. 6b. The C^1 extension, scaled by $\frac{1}{2}$, of $\mathbf{x}_{\underline{L}}$ respectively $\mathbf{x}_{\overline{L}}$ appear as green bullets in Fig. 6b. While the extension of $\mathbf{x}_{\underline{L}}$ agrees with the split extension of $\mathbf{x}_L$, the extension of $\mathbf{x}_{\overline{L}}$ does not. Therefore the extension of $\mathbf{x}_L$ must be reparameterized, see Fig. 6c, by

$$\begin{aligned} \rho_0 &:= u + a_0(u)v\,, \quad a_0 := B_0^2(u) + B_1^2(u) + \frac{3}{4}B_2^2(u)\,; \\ \rho_1 &:= u + a_1(u)v\,, \quad a_1 := \frac{3}{4}B_0^2(u) + \frac{1}{2}B_1^2(u) + \frac{1}{2}B_2^2(u)\,. \end{aligned} \tag{2}$$

Treating the extensions of $\mathbf{x}_R$, $\mathbf{x}_{\underline{R}}$, $\mathbf{x}_{\overline{R}}$ analogously, yields the construction as follows.

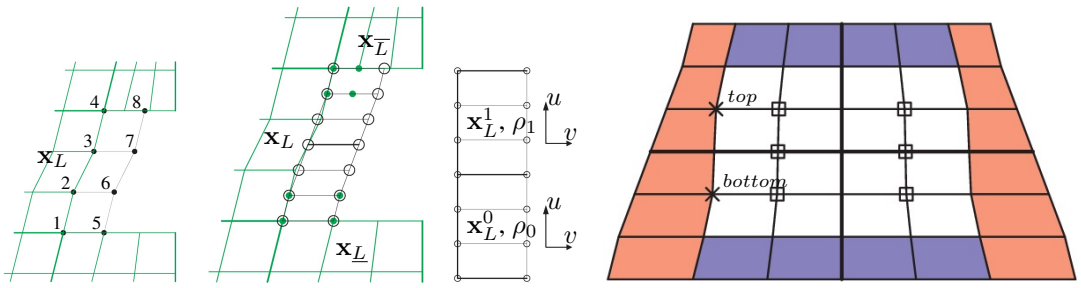

(a) C^1 extension (b) split extension (c) (d) BB-nets of the 2×2 bi-3 cap

Figure 6: τ_1-construction. (c) correspondance of split extensions and reparameterizations.

- The circled BB-coefficients in Fig. 6b define the first-order expansions $\mathbf{x}_L^s(u, v)$, $s = 0, 1$ along the boundary curve $v = 0$.
- Since the first-order expansions of $\mathbf{y}^s := \mathbf{x}_L^s \circ \rho_s$ are of degree 5, we convert $[\mathbf{y}^s, \partial_u \mathbf{y}^s, \partial_v \mathbf{y}^s, \partial_u \partial_v \mathbf{y}^s]$, the Hermite data at the end points, to BB-form of degree bi-3.
- The partial BB-nets are joined to form the left and right layers (underlaid orange-red in Fig. 6d) of the bi-3 cap with boundary curves matching $\mathbf{x}_L$ and $\mathbf{x}_R$.
- Since the layers of $\mathbf{x}_L^0$ and $\mathbf{x}_L^1$ are C^1-connected, it suffices to pre-calculate the BB-coefficients marked $\times$, as affine combinations of the BB-coefficients of the extension of $\mathbf{x}_L$ (resp. $\mathbf{x}_R$). In the layout of Fig. 6a

$$\begin{matrix} 4 & 8 \\ 3 & 7 \\ 2 & 6 \\ 1 & 5 \end{matrix}, \quad \text{the weights for } \times^{top} \text{ in Fig. 6d are: } \begin{bmatrix} 4 & 8 \\ 9 & 15 \\ 6 & 6 \\ 1 & -1 \end{bmatrix} /48 \quad \text{and for } \times^{bottom}\text{: } \begin{bmatrix} -1 & 1 \\ 0 & 12 \\ 3 & 21 \\ 2 & 10 \end{bmatrix} /48. \tag{3}$$

- The 2×2 bi-3 cap is completed by setting the top and bottom blue-gray underlaid BB-coefficients in Fig. 6d by C^1 extension with ratio $\frac{1}{2}$ of the surrounding bi-3 frame and choosing the central curve, from top to bottom, to be a single cubic curve in two segments.

- The BB-coefficients marked as hollow squares are set so that the horizontal layers are C^2 connected.

By construction, the lower and upper halves of the bi-3 cap are C^1 connected and the left and right halves of the bi-3 cap are C^2 connected. The cap is C^1 connected to $\mathbf{x}_{\underline{L}}$, $\mathbf{x}_{\underline{R}}$, $\mathbf{x}_{\overline{L}}$, $\mathbf{x}_{\overline{R}}$. Along the left and right boundary curves the bi-3 cap is formally only C^0-connected to the surrounding bi-3 frame. Testing the construction on a sequence of challenging inputs, we found the mismatch of normals to always be less than $0.1°$. According to the manual Autodesk (2018), a mismatch of less than $0.1°$ is acceptable for smooth surfaces in automobile class A surface design practice. We note that this type of bound is not affected by unitary transformations or uniform scaling. In our experience, meshes whose surfaces exceed the $0.1°$-criterion are artificial in the design context since T-junctions are used where the geometric variation is sufficiently 'tame' to safely decrease the number of quad-strips. For artificial cases like Fig. 7c, splitting the extensions of $\mathbf{x}^s_L$, $s = 0, 1$ and applying the reparameterizations obtained by splitting the original ρ_0, ρ_1 reduces the normal mismatch by an order of magnitude. The resulting eight-piece bi-3 cap is visually alike the 4-piece cap. Since the bi-3 cap uses the same first-order Hermite data along the inner boundary of the bi-3 frame as the bi-4 cap in GT_{KPP}, the bi-3 cap can replace the bi-4 cap both in the T_1 construction in GT_{KPP} and in the new τ_1 construction with transition frame.

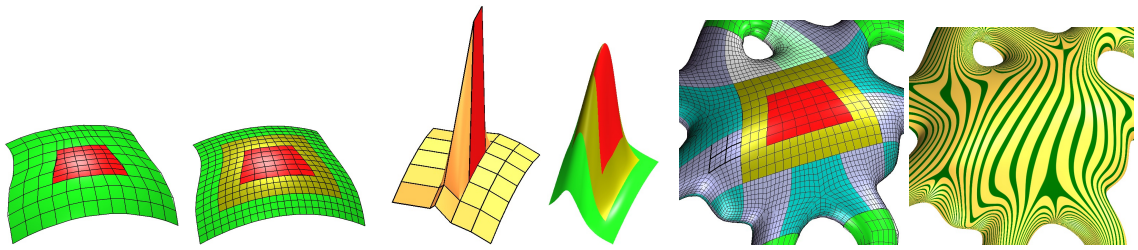

(a) bi-3 of Fig. 5b (b) bi-3 of Fig. 5e (c) $0.1°$ criterion (d) bi-3 of Fig. 1c

Figure 7: Bi-3 τ_3-construction: 2×2 bi-3 patches replace two bi-4 patches of the bi-4 constructions with little loss in highlight line quality. (For multi-sided facets, we use 2×2 bi-3 patches per sector.) (c) Stress-testing for normal mismatch: $0.177°$ before and $0.019°$ after splitting.

4. T_3-surfaces

τ_3-nets can be treated akin to τ_1-nets. Again we leverage GT_{KPP} for the central cap.

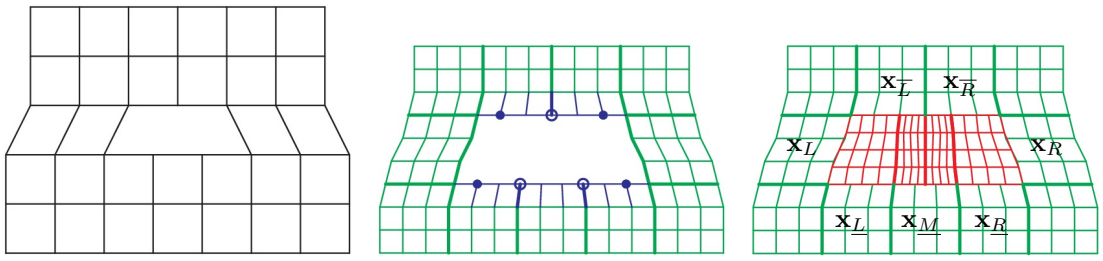

(a) input T_3-net (b) bi-3 T_3-frame (c) frame + central bi-4 T_3-cap

Figure 8: T_3-net and construction GT_{KPP}.

4.1. *Review of the GT_{KPP} construction for T_3-nets*

The bi-3 frame (green in Fig. 8b) is derived from Fig. 8a by B-to-BB conversion of the reconnected subnets Fig. 9a,b. Only the top and bottom curves require special care: the BB-coefficients marked as blue disks in Fig. 8b C^1-extend the left and right neighboring bi-3 patches, $\mathbf{x}_L$ and $\mathbf{x}_R$. The blue circled points are defined via special stencils listed in the GT_{KPP} publication. The remaining coefficients are set so that the two pieces of each curve are C^2 connected. The C^1 extension towards a central cap of bottom-middle patch $\mathbf{x}_M$ is split $1:1$ and the C^1 extensions of top patches $\mathbf{x}_{\overline{L}}$ and $\mathbf{x}_{\overline{R}}$ are split with ratio $2:1$. Then these extensions as well as the extensions of $\mathbf{x}_L$ and $\mathbf{x}_R$ are reparametrized as detailed in Karčiauskas et al. (2017) to become part of the central bi-4 cap.

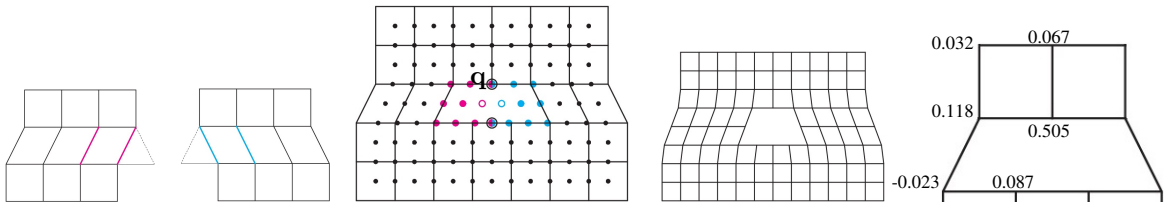

(a) (b) (c) refinement and averaging (d) refined net (e) stencil of $\mathbf{q}$

Figure 9: τ_3 construction. (a,b) the virtual subnets after re-connection. (c) refinement and averaging; (d) refined net; (e) stencil for top-circled point $\mathbf{q}$ (completed by symmetry).

4.2. *T_3-refinement and bi-3 transition*

The refined T_3-net is shown in Fig. 9. The magenta and cyan bullets are obtained from the subnets (a) and (b) and averaged at the overlapping locations marked by black circles (however, unlike the τ_3 case, the hollow magenta and cyan circles are not used in the construction and have been added only for completeness as a by-product of the refinement). Since the choice of the top-circled point $\mathbf{q}$ in Fig. 9c as a straight average of the sub-net BB-coefficients causes shape artifacts, it is re-defined by the stencil Fig. 9e explained below.

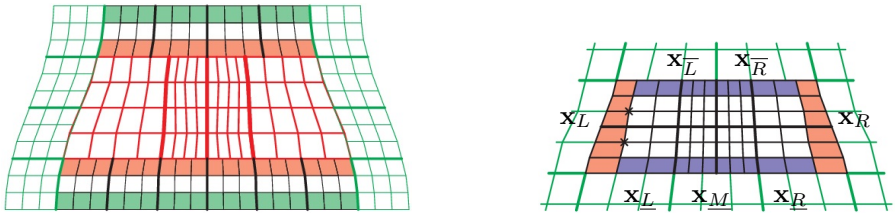

(a) bi-3 transition frame (b) central bi-3 cap

Figure 10: τ_3 construction.

The bicubic transition (see Fig. 10a) is simpler than in the T_1-case: the left and right bi-3 strips stemming from the refined net require no modification. However, slight artifacts (see Fig. 11) motivate a special rule, Fig. 9e,for $\mathbf{q}$ (top, circle in Fig. 9c) by determining

$$\text{argmin} \sum_{p\in\mathcal{P}} \mathcal{F}_2(p) \quad \mathcal{P} := \text{ bi-3 patches affected by } \mathbf{q}\,. \tag{4}$$

$$\mathcal{F}_\kappa f := \int_0^1 \int_0^1 \sum_{i+j=\kappa, i,j\geq 0} \frac{\kappa!}{i!j!} (\partial_s^i \partial_t^j f(s,t))^2 ds dt.$$

The absolute maximal weight of nodes for the argmin outside the τ_3-net is less than 0.01 in the affine expression. We set them to zero. Adjusting the remaining weights so they sum to 1 and have 3 digits yields the localized stencil of Fig. 9e for $\mathbf{q}$. The construction therefore applies even to tight configurations such as Fig. 16a. As for the τ_1-construction, the τ_3-construction can either be completed with the bi-4 cap from GT_{KPP}, or the bi-3 cap specified next.

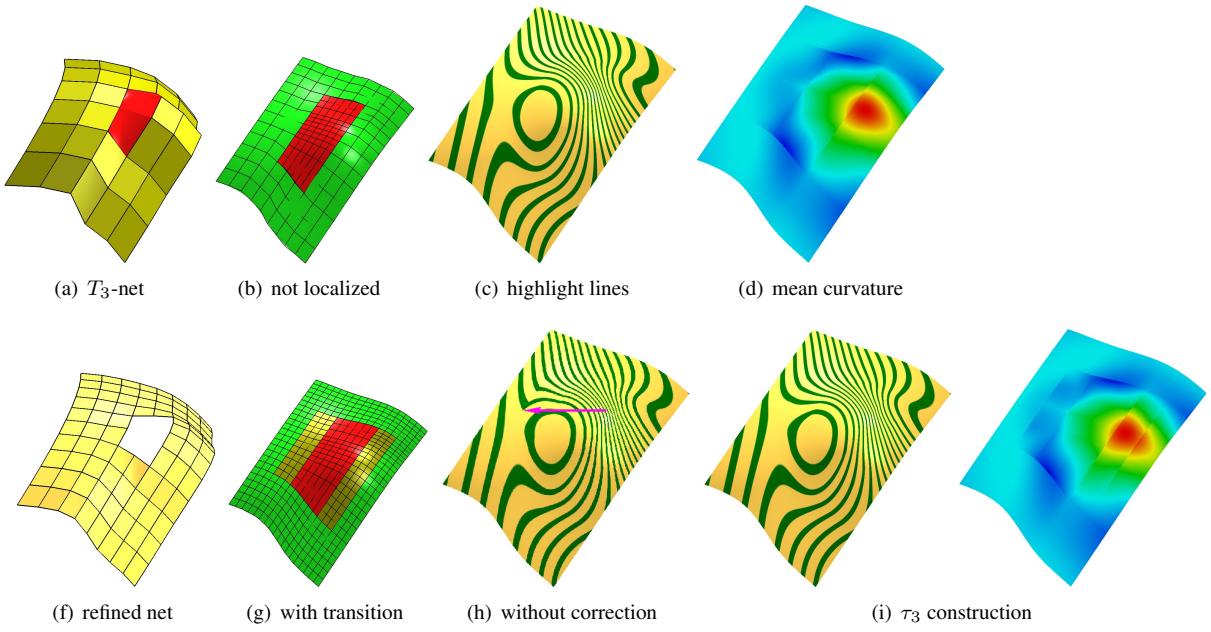

(a) T_3-net (b) not localized (c) highlight lines (d) mean curvature

(f) refined net (g) with transition (h) without correction (i) τ_3 construction

Figure 11: Constructions for T_3-nets. *top*: bi-3 surface without refinement *bottom*: *localized* bi-3 surface with refinement and transition. We note that despite the localization, the highlight line distribution does not suffer and even differences in curvature are slight.

4.3. Constructing a bi-3 cap

To make the degree uniformly bi-3, including the central cap of Fig. 10b, we proceed as in Section 3.3 but use the bi-3 transition ring Fig. 10a described in Section 4.2. The top and bottom transitions are the same as in the bi-4 case. We reparameterize the split extension of $\mathbf{x}_L$ and $\mathbf{x}_R$ with ρ_0 and ρ_1 where

$$a_0 := B_0^2(u) + B_1^2(u) + \frac{5}{6}B_2^2(u), \quad a_1 := \frac{5}{6}B_0^2(u) + \frac{2}{3}B_1^2(u) + \frac{2}{3}B_2^2(u). \tag{5}$$

The reparameterized data determine the red-gray parts of the central cap of Fig. 10b. Analogous to (3), the BB-coefficients marked by $\times$ are determined by the weights

$$\times^{top} : \begin{bmatrix} 4 & 14 \\ 9 & 27 \\ 6 & 12 \\ 1 & -1 \end{bmatrix} /72, \qquad \times^{bottom} : \begin{bmatrix} -1 & 1 \\ 0 & 18 \\ 3 & 33 \\ 2 & 16 \end{bmatrix} /72. \tag{6}$$

The blue-gray part of the BB-net in Fig. 10b is obtained from the top and bottom patches $\mathbf{x}_{\underline{L}}$, $\mathbf{x}_{\overline{R}}$, $\mathbf{x}_{\underline{M}}$, $\mathbf{x}_{\underline{L}}$, $\mathbf{x}_{\underline{R}}$ by C^1 extension and splitting as in the previous section. Except for the red-gray region, the vertical sequences of BB-coefficients each form a single (split) cubic curve. Comparing Fig. 11i vs Fig. 11h illustrates the improvement via the stencil of Fig. 9e.

5. T_2-surfaces

Since the GT_{KPP} construction for T_2-surfaces (see Fig. 13b) involves all nodes of the outermost layer of the T_2-net, it is not surprising that T_2-polygons with adjacent irregular nodes as in Fig. 15a do not offer satisfactory localized GT-splines: removing the outermost nodes from the GT_{KPP} construction results in significantly reduced surface quality over many informed trials. We therefore choose an alternative route that avoids GT-splines altogether.

- Add a new node q to the T_2-net and connect it with three vertices of T_2-polygon, see Fig. 12a,b, to obtain a pair of adjacent irregular nodes of valence 3 and 5.
- Apply one local Catmull-Clark refinement step to separate the irregular nodes.
- Apply the G^1 bi-4 construction of Karčiauskas and Peters (2015b) to the refined net Fig. 12c.

Fig. 12h shows the resulting surface layout.

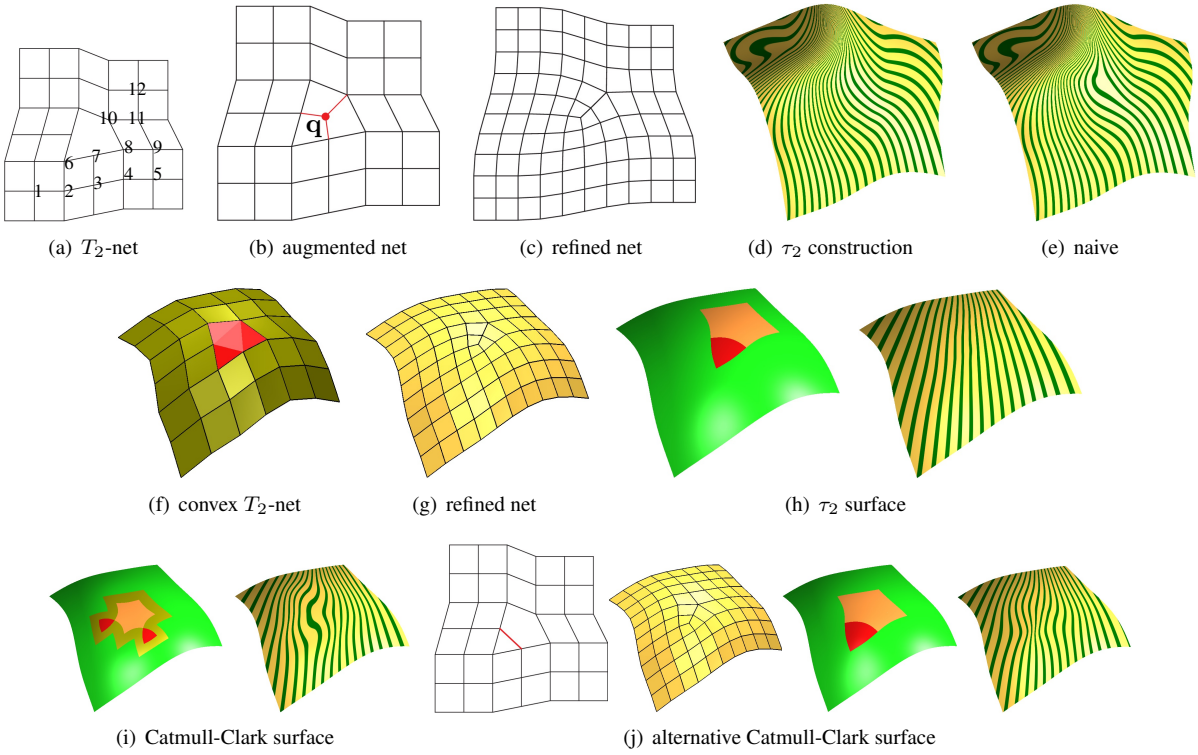

(a) T_2-net (b) augmented net (c) refined net (d) τ_2 construction (e) naive

(f) convex T_2-net (g) refined net (h) τ_2 surface

(i) Catmull-Clark surface (j) alternative Catmull-Clark surface

Figure 12: Construction for a T_2-net. (a) T_2-net (with the labels for definition of new node); (d) highlight lines for surface from T_2-net in Fig. 13a; (e) Result of naive choice of new node as centroid of the T_2-gon.

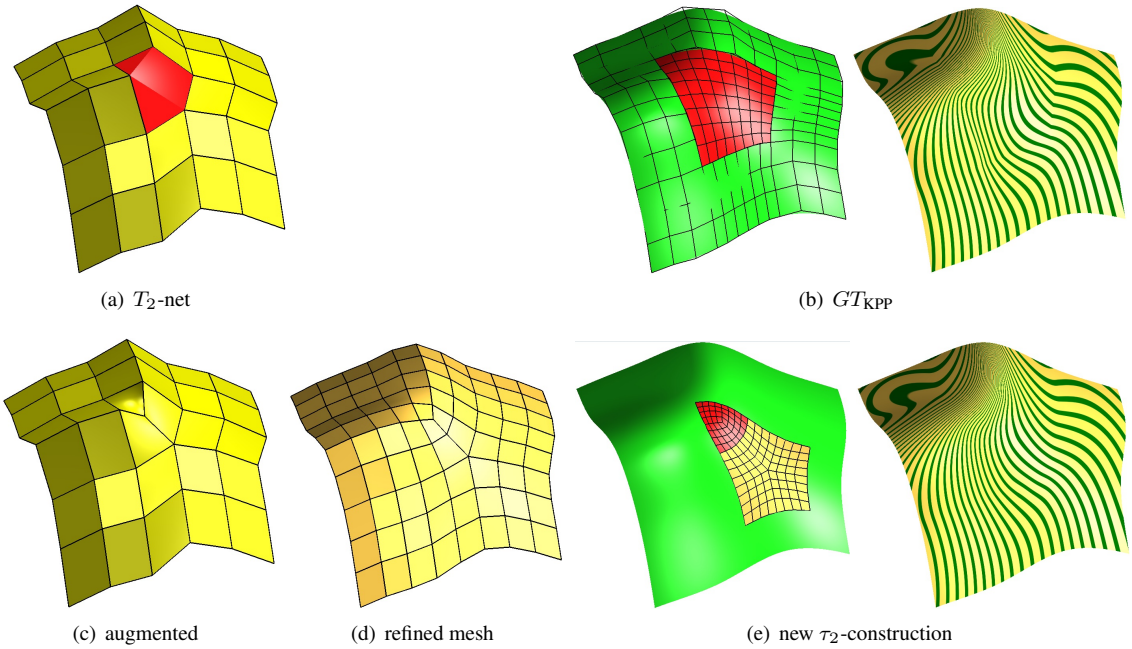

(a) T_2-net (b) GT_{KPP}

(c) augmented (d) refined mesh (e) new τ_2-construction

Figure 13: Comparison of T_1-net construction and T_2-net construction. *Top* row: (a) net with red T_2-gon. (b) GT_{KPP}-construction – green bi-3 ring and four red bi-4 patches; (c) adding a 3-valent node replaces the T-gon by three quads. (d) the refined augmented net (e) the localized τ_2-surface without loss of highlight line quality. Moreover, over many trials, the highlight lines are more uniformly distributed than for the corresponding GT_{KPP} surfaces.

The finite construction yields better results than applying Catmull-Clark repeatedly; and allows placing the shape-critical $\mathbf{q}$ near argmin $\sum_{p\in\mathcal{P}} \mathcal{F}_4(p)$, where $\mathcal{P}$ consists of the patches affected by $\mathbf{q}$ (bi-3 patches where regular, and patches defined by Karčiauskas and Peters (2015b) otherwise: bi-4 patches for 3-valent points and bi-5 patches for $n = 5$-valent points). In the affine expression of the new node, the absolute maximal weight of nodes outside the τ_3-net is less than 0.01. However, in contrast to the τ_3-construction, setting them to zero leads to visible shape artifacts. Deriving a stencil for $\mathbf{q}$ from a localized net requires more care. As indicated by the arrows in Fig. 14, the weights of the outermost nodes (gray bullets) are added to their neighbors (black bullets). Adjusting the resulting weights so they sum to 1 and have 3 digits yields the following stencil for $\mathbf{q}$ in terms of the inner nodes labeled in Fig. 12a (and their diagonally symmetric counterparts):

$$\begin{array}{lcccccc} \{1\dots6\}: & 0.017 & -0.017 & 0.023 & -0.021 & 0.016 & 0.024 \\ \{7\dots12\}: & 0.201 & 0.231 & -0.039 & 0.286 & -0.069 & 0.023 \end{array} \tag{7}$$

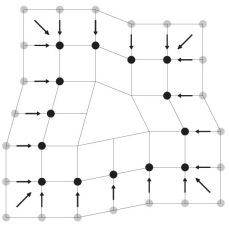

Figure 14: Redistribution of $\mathbf{q}$ weights.

We note that taking the central node naively as the centroid of the T_2-polygon results in distinctly worse highlight line distributions: compare e.g. Fig. 12d to Fig. 12e. The bottom row in Fig. 12 illustrates that both direct application of Catmull-Clark subdivision and more sophisticated re-meshing followed by Catmull-Clark subdivision yields unacceptable highlight lines.

6. Discussion and Examples

6.1. Tight configurations

Just like Fig. 1 demonstrated how the new localized approach enables a close juxtaposition of T-gons and m-valent points for the case of τ_1-nets, so Fig. 15 and Fig. 16 demonstrate this for design with τ_2-nets and τ_3-nets.

Fig. 17 demonstrates for τ_1-gons that the formal decrease in smoothness does not affect the highlight line distribution. In fact, the highlight lines are subtly (not visible in static images) improved since the left and right pieces of bi-3 cap are C^2 connected. And differences between the bi-3 and the bi-4 construction, in the highlight lines across the border of central cap and the transition frame, are also near imperceptible. This indicates that both versions, bi-3 and bi-4, are equally acceptable.

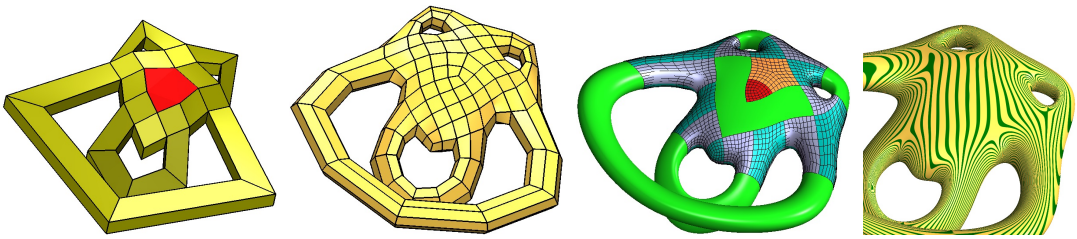

(a) tight configuration (b) refined mesh (c) bi-3 (d) highlight lines

Figure 15: τ_2-construction. (a) mesh with red T_2-gon and irregular nodes of valence 3 and 5; (b) once refined mesh after augmentation; (c) final surface: bi-3 regular patches green, three-sided and five-sided surfaces.

6.2. Localization via inconsistent derivatives: Gregory-type patches

GT-splines scale the C^1 prolongations of the bi-3 frame to equate originally inconsistent derivatives. Alternatively, incompatible inner and boundary BB-coefficients can be made to coexist by introducing a rational patch with different derivatives depending on the approach to the corner, see e.g. (Gregory, 1974; Loop et al., 2009; Hettinga and Kosinka, 2018). Consider Fig. 18a. At the bottom, black bullets mark the BB-coefficients obtained by C^1 extension of $\mathbf{x}_L$ and $\mathbf{x}_{\underline{L}}$ of the bi-3 frame. On top, the black bullets mark only some BB-coefficients of the C^1 extension of

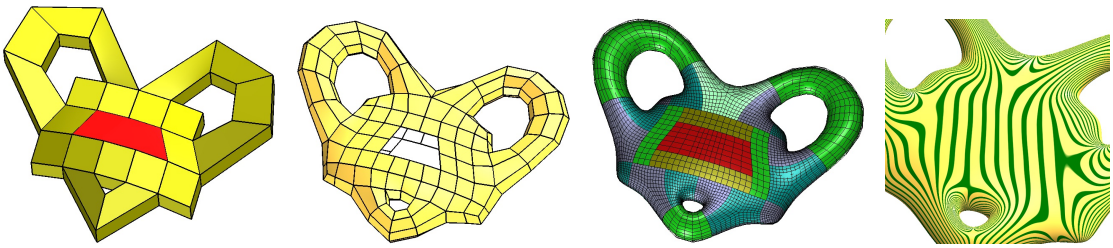

(a) tight configuration (b) refined T-net (c) bi-3 (d) highlight lines

Figure 16: τ_3-construction. (a) Mesh with red T_3-gon and irregular nodes of valencies 3, 5, 6.a (c,d) bi-3 surface and highlight lines.

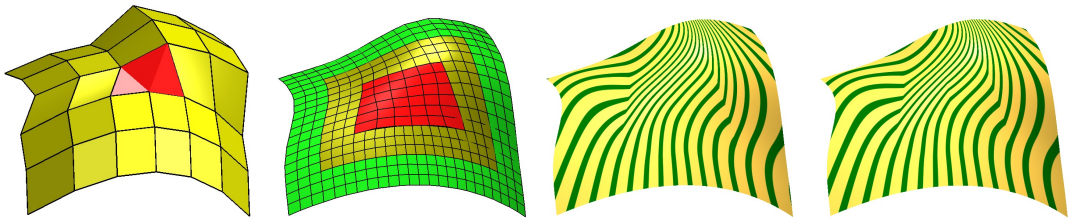

(a) T_1-net (b) bi-3 frame + bi-3 transition frame + bi-3 cap (c) bi-3 frames + bi-4 cap

Figure 17: Comparison of bi-3 and bi-4 central caps.

$\mathbf{x}_{\overline{L}}$, because the BB-coefficients $\boldsymbol{\ell}_{12}, \boldsymbol{\ell}_{13}$ that extend $\mathbf{x}_L$ (gold bullets) do not match the BB-coefficients $\mathbf{t}_{12}$, $\mathbf{t}_{13}$ that extend $\mathbf{x}_{\overline{L}}$. Setting

$$\mathbf{b}_{12} := \frac{u\mathbf{t}_{12} + (1-v)^2\boldsymbol{\ell}_{12}}{u + (1-v)^2}, \quad \mathbf{b}_{13} := \frac{u\mathbf{t}_{13} + (1-v)^2\boldsymbol{\ell}_{13}}{u + (1-v)^2} \tag{8}$$

yields a rational patch of bi-degree 4×5 with a singularity of type $\frac{0}{0}$ at the left-upper corner. This patch joins smoothly with its neighbors. However, like many other singular rational constructions, the resulting surfaces are acceptable but not class A, as may be seen by comparing the surfaces in *bottom* row of Fig. 18. They are acceptable because the singularity does not prevent a graded tessellation: Fig. 18f displays such tessellation with 20×20 quads per patch; and, at first impression, the visual quality is good, see Fig. 18b. However, the rational singularity causes problems when analyzing and interrogating the surfaces. Together with the lack of a 'class A' highlight line distribution, this makes the approach undesirable for industrial applications.

Note, in (8) the point $\mathbf{b}_{12}$ can alternatively be the linear average $\mathbf{b}_{12} := \frac{u\mathbf{t}_{12}+(1-v)\boldsymbol{\ell}_{12}}{u+1-v}$. However, this yields a rational patch of bi-degree 5×6 with slightly worse highlight lines than presented above.

7. Conclusion

T-junctions can reduce the requirements on algorithms for quad-meshing and increase the flexibility of polyhedra as control nets for defining smooth surfaces. The surfaces are affine invariant since their BB-coefficients are linear combinations of the points of the input net. This paper introduces piecewise polynomial geometrically continuous surface constructions for tight configurations, further reducing the quad-meshing requirements and increasing the scope for polyhedral control of design surfaces. At T-configurations the surfaces are G^1 and consist of 2,4,8 patches of degree bi-4 for τ_1, τ_2, τ_3-nets. A key feature, a quality-preserving refinement of the input net in the vicinity of a T-gon, minimizes the footprint of this construction. For τ_2-gons we introduced a quality-preserving remeshing in place of the refinement. Alternatively, we introduced an everywhere bi-3 variant for T-gons that uses twice as many patches as the bi-4 construction and generates surfaces that are formally only C^0. Together with a known construction for irregular points, this variant has the desirable property that the overall degree of the free-form surface is uniformly bi-3, including the center cap. And while not formally smooth, the bi-3 surfaces have good highlight line distributions for T-gons with 1,2, or 3 T-junctions also where vertices in close proximity are irregular (valence other than 4).

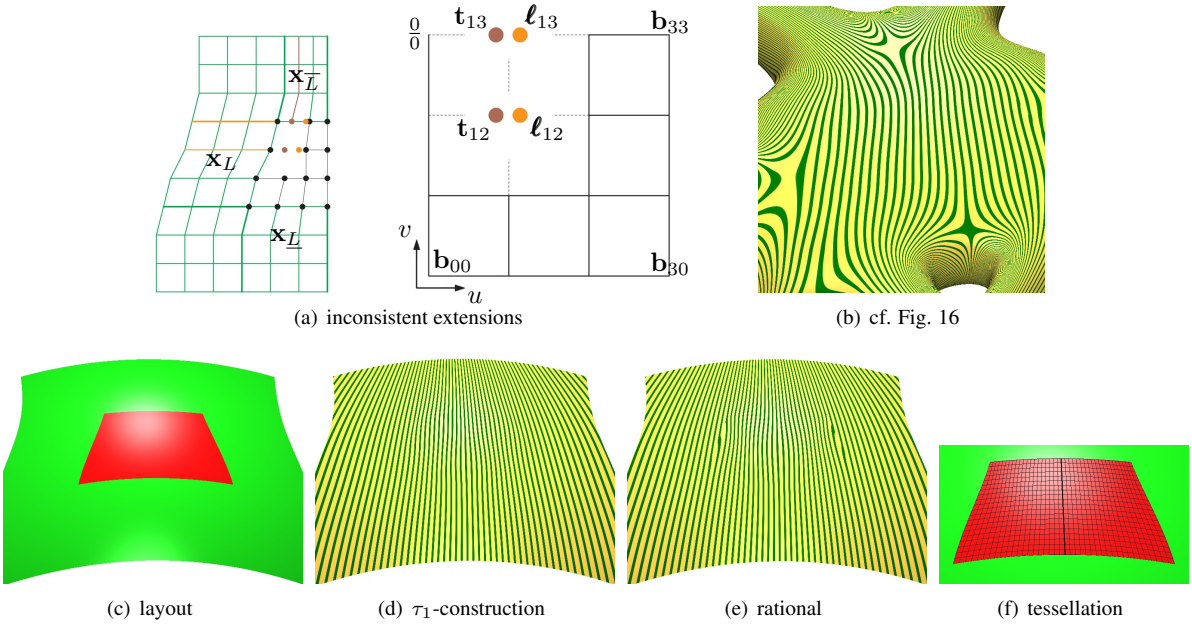

(a) inconsistent extensions (b) cf. Fig. 16

(c) layout (d) τ_1-construction (e) rational (f) tessellation

Figure 18: (a) Input for Gregory-type patch derived from T_1-net and incompatible BB-coefficients for Gregory-type averaging. (b) Gregory-type surface of Fig. 16. (c) Layout of surface from Fig. 5a T_1-net; hl s of (d) GT surface and (e) Gregory-type surface and (f) its tessellation.

Acknowledgements. This work was supported in part by DARPA HR00111720031 and NIH R01 LM011300-01.

References

Alliez, P., Cohen-Steiner, D., Devillers, O., Lévy, B., Desbrun, M., 2003. Anisotropic polygonal remeshing. ACM Trans. Graph. 22, 485–493.
Autodesk, 2018. Tutorial at http://help.autodesk.com/view/alias/2015/enu/?guid=guid-2fce06eb-8ef7-4507-92f7-82a73a0df378. Accessed Nov 14.
Beier, K.P., Chen, Y., 1994. Highlight-line algorithm for realtime surface-quality assessment. Computer-Aided Design 26, 268–277.
Bommes, D., Lévy, B., Pietroni, N., Puppo, E., Silva, C., Tarini, M., Zorin, D., 2012. State of the art in quad meshing, in: Eurographics STARS.
de Boor, C., 1978. A Practical Guide to Splines. Springer.
Catmull, E., Clark, J., 1978. Recursively generated B-spline surfaces on arbitrary topological meshes. Computer-Aided Design 10, 350–355.
Dokken, T., Lyche, T., Pettersen, K.F., 2013. Polynomial splines over locally refined box-partitions. Comp. Aid. Geom. Design 30, 331–356.
Farin, G., 1988. Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide. Academic Press.
Gregory, J.A., 1974. Smooth interpolation without twist constraints. Academic Press. pp. 71–88.
Hettinga, G.J., Kosinka, J., 2018. Multisided generalisations of Gregory patches. Computer Aided Geometric Design 62, 166–180.
Jakob, W., Tarini, M., Panozzo, D., Sorkine-Hornung, O., 2015. Instant field-aligned meshes. ACM Trans. Graph 34, 189.
Kälberer, F., Nieser, M., Polthier, K., 2007. Quadcover - surface parameterization using branched coverings. Comput. Graph. Forum 26, 375–384.
Kang, H., Xu, J., Chen, F., Deng, J., 2015. A new basis for PHT-splines. Graphical Models 82, 149–159.
Karčiauskas, K., Panozzo, D., Peters, J., 2017. T-junctions in spline surfaces. ACM Tr on Graphics, ACM Siggraph 36, 170:1–9.
Karčiauskas, K., Peters, J., 2015a. Can bi-cubic surfaces be class A? Computer Graphics Forum 34, 229–238.
Karčiauskas, K., Peters, J., 2015b. Improved shape for multi-surface blends. Graphical Models 8, 87–98.
Kraft, R., 1998. Adaptive und linear unabhängige Multilevel B-Splines und ihre Anwendungen. Ph.D. thesis. University of Stuttgart.
Li, W.C., Ray, N., Lévy, B., 2006. Automatic and interactive mesh to T-spline conversion.
Loop, C.T., Schaefer, S., Ni, T., Castaño, I., 2009. Approximating subdivision surfaces with Gregory patches for hardware tessellation. ACM Trans. Graph 28, 151:1–151:9.
Peters, J., 2002. Geometric continuity, in: Handbook of Computer Aided Geometric Design, Elsevier. pp. 193–229.
Pietroni, N., Puppo, E., Marcias, G., Scopigno, R., Cignoni, P., 2016. Tracing field-coherent quad layouts. Comp Gr Forum (Pacific Graphics) .
Sabin, M.A., Dodgson, N.A., Hassan, M.F., Ivrissimtzis, I.P., 2003. Curvature behaviours at extraordinary points of subdivision surfaces. Computer-Aided Design 35, 1047–1051.
Sederberg, T.W., Zheng, J., Bakenov, A., Nasri, A., 2003. T-splines and T-NURCCs, in: Hodgins, J., Hart, J.C. (Eds.), Proceedings of ACM SIGGRAPH 2003, ACM Press. pp. 477–484.
Vaxman, A., Campen, M., Diamanti, O., Panozzo, D., Bommes, D., Hildebrandt, K., Ben-Chen, M., 2016. Directional field synthesis, design, and processing. Computer Graphics Forum .