



T-junctions in Spline Surfaces

KĘSTUTIS KARČIAUSKAS

Vilnius University

DANIELE PANOZZO

New York University

and

JÖRG PETERS

University of Florida

T-junctions occur where surface strips start or terminate. This paper develops a new way to create smooth piecewise polynomial free-form spline surfaces from quad-meshes that include T-junctions. All mesh nodes are interpreted as control points of GT-splines, that is, geometrically smoothly joined piecewise polynomials. GT-splines are akin to and compatible with B-splines and cover simple T-junctions by two polynomial pieces of degree bi-4 and more complex ones by four such patches. They complement multi-sided surface constructions in generating free-form surfaces with adaptive layout.

Since GT-splines do not require a global coordination of knot intervals, GT-constructions are easy to deploy and can provide smooth surfaces with T-junctions where T-splines cannot have a smooth parameterization. GT-constructions display a uniform highlight linedistribution on input meshes where alternatives, such as Catmull-Clark subdivision, exhibit oscillations.

Categories and Subject Descriptors: I.3.5 Computing Methodologies [Computer Graphics]: Parametric Curve and Surface Models—*Spline surfaces*

General Terms: T-junctions, Surfaces

Additional Key Words and Phrases: T-junctions, spline, surface, smoothness, highlight line distribution

This work was supported in part by NSF grant CCF-1117695, NSF CAREER award 1652515, NIH grant R01 LM011300-01 and DARPA TRADES HR00111720031.

Authors' addresses: K. Karčiauskas, Vilnius University, LT-2006, Vilnius, Lithuania; New York University, NY, USA; email: Kestutis.Karciauskas@mif.vu.lt; D. Panozzo, Courant Institute of Mathematical Sciences, New York University, 60 5th Ave, 5th floor, New York, NY 10011; email: panozzo@nyu.edu; J. Peters, University of Florida, Gainesville FL 32611-6120, USA; email: jorg@cise.ufl.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2017 ACM 0730-0301/2017/10-ART170 \$15.00

DOI: <http://dx.doi.org/10.1145/3136954>

ACM Reference Format:

Kęstutis Karčiauskas, Daniele Panozzo, and Jörg Peters. 2017. T-junctions in spline surfaces. *ACM Trans. Graph.* 36, 5, Article 170 (October 2017), 9 pages.

DOI: <http://dx.doi.org/10.1145/3136954>

1. INTRODUCTION

Where strips of surface patches are forced together, it is natural to terminate some; and where strips are stretched wide, it is natural to spawn additional strips to keep the size and aspect ratio of the patches within bounds. Stopping or initiating surface strips leads to T-junctions where two finer surface pieces meet one coarser piece. The simplest T-junction-configuration, a \tilde{T} -net (pronounced T1-net), is shown in Figure 1(a): a nominally pentagonal face with exactly one vertex of valence 3 is surrounded by quadrilateral facets. In isolation such transitions are easily modeled by smooth hierarchical splines. But as part of a larger model, their knot-intervals need to be globally coordinated. That coordination is cumbersome. The small quad-mesh in Figure 2 shows that it may even be impossible.

T-junctions allow the introduction of higher geometric detail or to merge two separately-developed spline surfaces. T-junctions also prominently arise when replacing the complex and global constraints of strict quad-meshing [Bommes et al. 2013; Vaxman et al. 2016] by T-meshes, based on triangle meshes [Li et al. 2006; Lai et al. 2008], curvature directions [Alliez et al. 2003; Marinov and Kobbelt 2004], directional fields [Myles et al. 2010, 2014; Pietroni et al. 2016], optimized for planarity [Zadavec et al. 2010; Peng and Wonka 2013], or extracted from local parametrizations [Ray et al. 2006; Jakob et al. 2015].

T-junctions and hierarchical splines. One approach to incorporating T-junctions is hierarchical splines (see, e.g., Kraft [1998], Sederberg et al. [2003], Giannelli et al. [2012], Dokken et al. [2013], and Kang et al. [2015]). Hierarchical splines require that all surface pieces share a single uv -parameterization. For any choice of v , the u -knot intervals must add to the same number, and for any choice of u , the v -intervals must add to one fixed number. This restriction on the knot sums is natural when refining a single patch. But when the input is a given quadrangulation the knot intervals have to be assigned and hence coordinated. Joining many pieces can then become cumbersome since the local knot intervals have to globally add up to matching sums. Where the mesh is not regular, the local construction in Wang et al. [2011] is therefore only C^0 despite generating one order of magnitude more patches than quads. Sederberg et al. [2008] additionally points to shape deficiencies where T-splines with dissimilar knot spacings are merged and proposes special splines of much higher degree. The global construction




Fig. 1. A control net with a single isolated T-junction.




Fig. 2. T-splines require that *the sum of knot intervals on opposing edges of any face must be equal* (Rule 1 of Sederberg et al. [2003]). This forces the width of the horizontal knot intervals of the grey helical strip to be zero, preventing a smooth T-spline parameterization. (cf. Figure 13 for a smooth GT-spline surface.)

of Li et al. [2006] introduces additional T-joints and new extraordinary points that are not motivated by geometry but solely by enforcing knot interval constraints. (Pre-dating [Li et al. 2012], the construction does not guarantee “analysis-suitable” knot-distributions.) Akin to Li et al. [2006, Fig. 5], the example in Figure 2 demonstrates that, without modifying the quad-mesh, global coordination is not always possible without loosing smoothness and even continuity of the parameterization. In the example, each red strip forms a bracelet that is half as wide when it comes back to meet up with the starting edge. Since Rule 1 of any T-spline construction according to Sederberg et al. [2003] mandates that the horizontal knot interval of the red strip be the same where the single-wide edge meets the double-wide edge, the horizontal knot interval of the grey helical strip of patches must be zero. Since, in this example, three consecutive grey horizontal knot intervals are zero, the degree bi-3 spline parameterization is formally C^{-1} . That is, the most basic property of T-splines prevents a smooth parameterization for a class of patch layouts that could be hidden in any large scale quad arrangement. The bracelet implies that joining spline surfaces with a *smooth* T-spline parameterization is not always feasible since this, too, requires making the knot interval sums equal. The slightly larger mesh of Figure 18(a) demonstrates a more complex incompatibility with any assignment of knot-intervals for smooth T-splines. In summary, while *hierarchical splines are naturally suited for introducing T-junctions in quad meshes, they are not naturally suited for generating surfaces from quad meshes with T-junctions*.

T-junctions and Catmull-Clark subdivision. A strictly local construction is provided by Catmull-Clark subdivision [Catmull and Clark 1978]. Here the underlying model is splines with uniform knot spacing and local support. Therefore the parameterizations need not be globally coordinated. However, as Figure 3(a) demonstrates for the convex input \tilde{T} -net of Figure 1(a), the resulting surfaces can be of poor quality: the silhouette dips and rises and the highlight lines oscillate near the T-junction. We note that the oscillations already manifest themselves in the first two subdivision steps and hence rule out Catmull-Clark mesh refinement even as a pre-processor for turning a T-junction into a pair of isolated vertices of valence three and five. (We verified these “first step artifacts” [Augsdörfer et al.

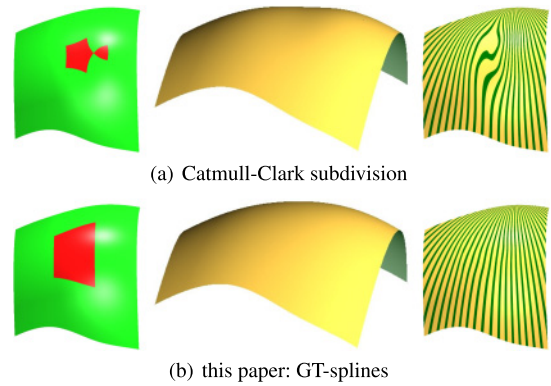


Fig. 3. Surfaces generated from the convex input \tilde{T} -net of Figure 1(b). The red regions in (a) represent an infinite sequence of bi-3 patches covering the five-sided and three-sided extraordinary Catmull-Clark neighborhoods arising from a \tilde{T} -net. The red region in (b) consists of two bi-4 patches by which the GT-construction covers the \tilde{T} -net. Catmull-Clark subdivision produces, in the first two steps, a flattened silhouette and a correspondingly non-uniform highlight linedistribution (right).

2011] by replacing, in a separate computation, the red limit surface in Figure 3(a) by a high-quality surface construction.)

T-junctions and GT-splines. This article develops a new local construction, a geometric approach to T-junctions. This GT-construction does not require global knot interval coordination and yields better shape than Catmull-Clark subdivision (see Figure 3(a) vs. (b)). The GT-construction is based on reparameterization, the natural technique for transitioning between unequal parameterizations on opposite sides of a T-junction. Consequently, the construction leverages the framework of geometric continuity. For \tilde{T} -nets, the resulting surfaces

- consist of a frame of bi-3 (bi-cubic) patches filled by two patches of degree bi-4 (this “cap” is red in Figure 1(b)).
- The bi-4 cap is internally smooth and joins the bi-cubic patches with tangent continuity (G^1)
- The bi-4 cap yields good highlight line distributions on all of a large number of challenging input meshes.
- The Appendix provides simple explicit formulas for all relevant Bernstein Bézier coefficients of the GT-construction in terms of the local \tilde{T} -net.
- GT-splines complement few-piece polynomial constructions such as Karčiauskas et al. [2016] covering extraordinary points to model smooth free-form surfaces of maximal degree bi-4.

We also present two variants of the GT-construction, to generate smooth caps for multiple T-junctions within one facet whose extensions cross or are parallel.

Overview. Section 2 reviews basic concepts of the construction of smooth surfaces. Section 3 defines a bi-3 frame of patches that transitions to the surrounding surface. Section 4 describes the G^1 bi-4 GT-construction of the cap for \tilde{T} -nets. Section 5 and Section 6 develop caps for cases of adjacent T-junctions: for two T-junctions opposite one and for two T-junctions with crossing directions. Section 7 compares the constructions for challenging input data, explains the choices taken along the way and lists the limitations. Section 8 shows how GT-splines collaborate with algorithms in the literature to smoothly cover multi-sided neighborhoods by an overall smooth bi-4 surface.

2. DEFINITIONS AND SETUP

We will construct a T-junction surface by averaging alternative interpretations of the mesh points as regular control points (see Figure 5 and Section 3). This yields a frame of bi-3 polynomial patches to be filled by a cap consisting of polynomial patches of degree bi-4. The following definitions make this approach precise.

The GT-splines are a collection of tensor-product patches in Bernstein-Bézier form (BB-form; see, e.g., Farin [1988]):

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

where $B_k^d(t) := \binom{d}{k}(1-t)^{d-k}t^k$ are the BB polynomials of degree d and \mathbf{f}_{ij} are the BB-coefficients. Adjacent patches join with G^k continuity if their k th-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, for example, Peters [2002]. We use the succinct characterization that *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 $\rho : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ so 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

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

If $\tilde{\mathbf{f}}$ and \mathbf{f} are polynomials, then a and b are rational functions whose degree is bounded in terms of the degree of $\tilde{\mathbf{f}}$ and \mathbf{f} [Peters 1991].

C^1 continuity of the splines over non-uniform knot sequences can be recast as G^1 continuity of patches defined over unit domains (see, e.g., Karčiauskas and Peters [2011]). For example, if \mathbf{f} and $\tilde{\mathbf{f}}$ are consecutive curve segments originally associated with intervals $[-1, 0]$ and $[0, \frac{1}{2}]$ of a C^1 spline with knot sequence $\{\dots, -1, 0, \frac{1}{2}, \dots\}$ then both \mathbf{f} and $\tilde{\mathbf{f}}$ can be newly defined, each on the interval $[0, 1]$, and they then join as $\beta \partial_u \tilde{\mathbf{f}}(1) = \partial_u \mathbf{f}(0)$ with $\beta := 1/2$. When we want to point out that surface patches are, in one variable, related by the identity and, in the other, by C^1 continuity over non-uniform knot sequence, we refer to the transition as: *C^1 with parameter β* .

Our main construction focusses on \tilde{T} -nets that consist, as shown in Figure 1, of quadrilaterals and one nominally five-sided facet. For context and exposition, we can extend the \tilde{T} -net by one layer of quadrilaterals (see Figure 4(a)). This allows applying, away from the five-sided facet, the well-known bi-cubic (bi-3) B-spline to BB-form conversion rules (see, e.g., Farin [1988]). The resulting C^2 bi-3 spline neighborhood is colored brown in Figure 4(b). The smaller \tilde{T} -net (Figure 1) provides a *tensor-border of degree 3 and depth 2*, the C^2 -prolongation of second-order Hermite data shown as a **green** net of BB-coefficients in Figure 4(b). Given the tensor-border, the \tilde{T} -net interior of $4 + 4 + 5 + 5$ control points (see the stencils in Figures 5(d) and (f)) provides all the information for the GT-construction. All formulas (stencils) of the bi-4 cap in terms of the \tilde{T} -net interior are provided in the Appendix.

3. CONSTRUCTION OF A BI-3 FRAME OF PATCHES FOR CAPPING A \tilde{T} -NET

The surface corresponding to a \tilde{T} -net will consist of a frame of bi-3 patches and a central cap consisting of two bi-4 patches. This section

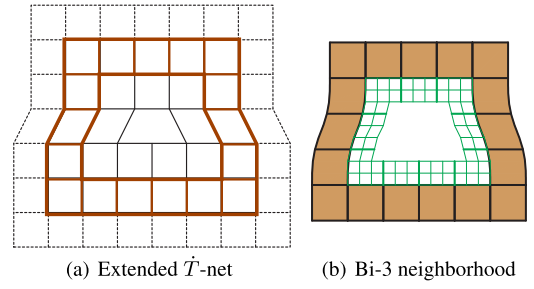


Fig. 4. An isolated T-junction in an extended \tilde{T} -net (a) provides (b) a bi-3 neighborhood (solid, for context only) and a C^2 -prolongation (tensor-border) in BB-form (inner green mesh) that is the only part used for the GT-construction.

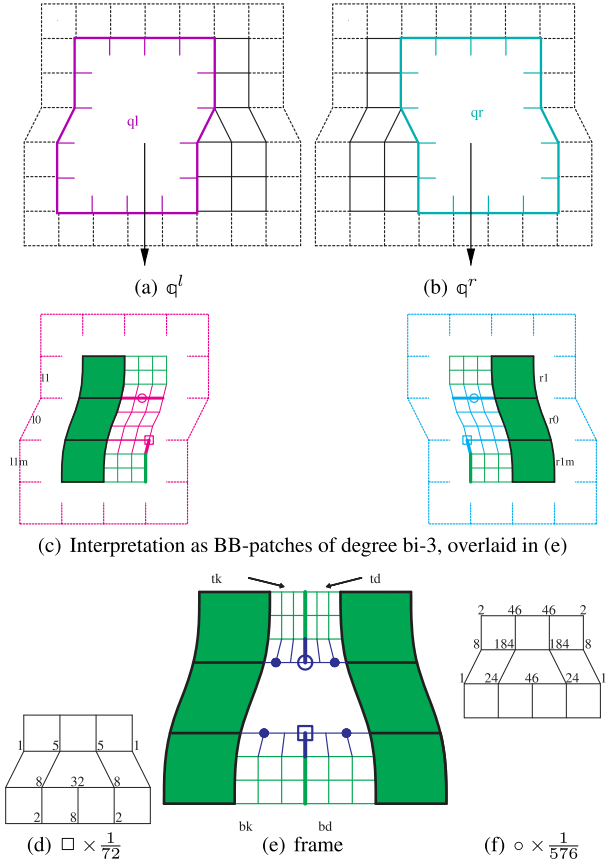


Fig. 5. Constructing the frame. The regular left (a) and right (b) control nets obtained by re-connecting the nodes of the \tilde{T} -net define six bi-3 patches each. (c) Completion of the frame. ((d) and (f)) Stencils of the points $\mathbf{q}_{\square}^{b,-1} = \mathbf{q}_{\square}^{b,1}$ marked \square and $\mathbf{q}_{\square}^{t,-1} = \mathbf{q}_{\square}^{t,1}$ marked by \circ in (e).

builds the frame. For $i = -1, 0, 1$ the frame has left patches $\mathbf{q}^{l,i}$, right patches $\mathbf{q}^{r,i}$ (Figure 5(c)), and, for $j \in \{-1, 1\}$, top patches $\mathbf{q}^{t,j}$ and bottom patches $\mathbf{q}^{b,j}$ (Figure 5(e)). The ribbon is derived by re-connecting the nodes of the \tilde{T} -net to form two regular nets, \mathbf{q}^l from the left (see Figure 5(a)) and \mathbf{q}^r from the right (see Figure 5(b)).

We now interpret \mathbf{q}^l and \mathbf{q}^r as bi-3 B-spline control nets and convert them to BB-form. As illustrated in Figure 5(c) the so-derived patches agree—except for the lower boundary curves of the top patches and the upper boundary of the bottom patches. The bottom

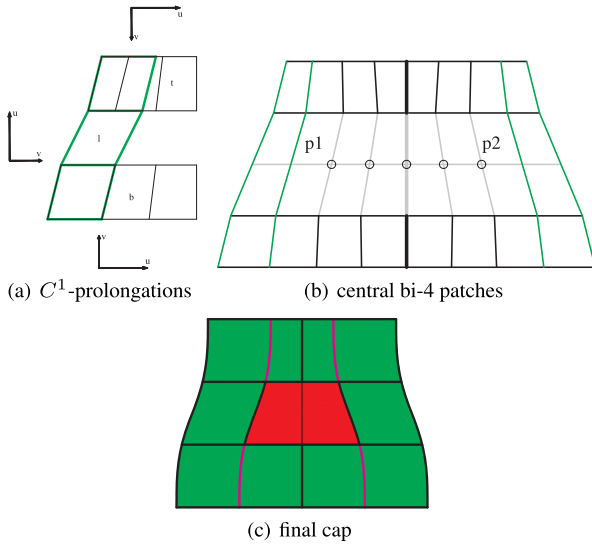


Fig. 6. Construction of the bi-4 cap. (a) The mismatch of the C^1 -prolongations is resolved by reparameterizing them. (b) Interior coefficients minimize the distance to bi-cubics. (c) Final layout: The bi-3 patches are C^1 -connected across the (red) hv-curves.

curves overlap only at their endpoints, marked as a magenta and a cyan box in Figure 5(c). A new common BB-control point, also marked as a blue box in Figure 5(e), is chosen to be the average of these two candidates: Figure 5(d) displays the explicit stencil for the common point $\mathbf{q}_{\square} := (\mathbf{q}_{\square}^{b,-1} + \mathbf{q}_{\square}^{b,1})/2$ obtained by this averaging. The two direct neighbors of \mathbf{q}_{\square} are chosen to make the combined boundary curve C^2 . (The blue disks are defined by the C^1 -prolongation of $\mathbf{q}^{l,-1}$ and $\mathbf{q}^{r,-1}$). The top patches are subdivided at their midpoint and the resulting overlap is then treated like that of the bottom: Figure 5(f) displays the explicit stencil for the common point \mathbf{q}_{\circ} obtained by this averaging. Although this split into $\mathbf{q}^{l,-1}$ and $\mathbf{q}^{r,-1}$ serves only to accommodate the lower boundary of the top patches, extensive experiments show this split to be critical for achieving good shape (see, e.g., Figure 12). The resulting frame of bi-3 patches is C^2 except along the four *hv-curves*, the red curves in Figure 6(c) between the horizontal and the vertical strips of the frame: $\mathbf{q}^{l,-1}$ to $\mathbf{q}^{b,-1}$, $\mathbf{q}^{l,1}$ to $\mathbf{q}^{r,-1}$, $\mathbf{q}^{r,-1}$ to $\mathbf{q}^{b,1}$, $\mathbf{q}^{r,1}$ to $\mathbf{q}^{l,1}$. Across the hv-curves the continuity is C^1 . Since the construction did not change the BB-coefficients derived from \mathbf{q}^l and \mathbf{q}^r that match the tensor-border (Figure 4(b)), the frame joins C^2 with the splines surrounding it.

4. CENTRAL \tilde{T} -NET CAP CONSTRUCTION

To complete the surface, we construct a central cap (red in Figure 6) of degree bi-4 that fills the frame so all transitions are at least G^1 . The bi-4 cap consists of two patches $\mathbf{p}^l, \mathbf{p}^r$; see Figure 6(c). Since the construction of \mathbf{p}^r mirrors that of \mathbf{p}^l , we discuss only \mathbf{p}^l . Figure 6(a) shows the C^1 -prolongations \mathbf{t} of $\mathbf{q}^{l,-1}$ and \mathbf{b} of $\mathbf{q}^{b,-1}$ in black and \mathbf{l} of $\mathbf{q}^{l,0}$ in green. While \mathbf{b} is consistent with \mathbf{l} , the prolongations \mathbf{t} and \mathbf{l} are inconsistent (due to the split of the top patch when constructing the frame). Since we reparameterize \mathbf{l} linearly (to minimize the final patch degree) to match \mathbf{t} , we also need to reparameterize \mathbf{b} after all.

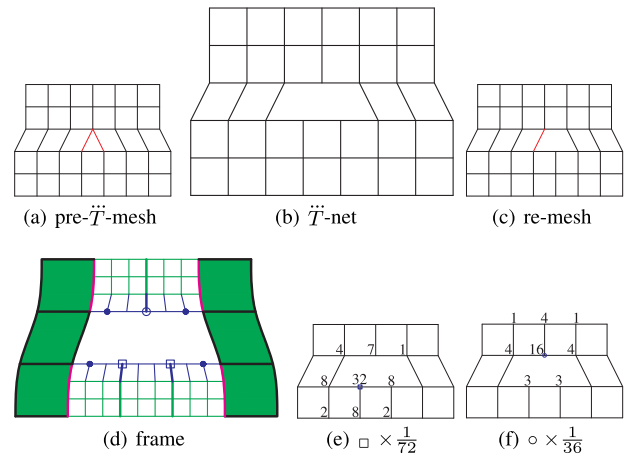


Fig. 7. Construction of the frame for a \tilde{T} -junction.




Fig. 8. BB-coefficients of the bi-4 cap for a \tilde{T} -net.

Together, the choice of parameterizations in Equation (1) are

$$\begin{aligned} \text{left}(\mathbf{l}) : & \quad 1 - \frac{u}{2} & \quad b(u) := & \quad 0 \\ \text{top}(\mathbf{t}) : & \quad 1 & \quad (1-u)u & \\ \text{bottom}(\mathbf{b}) : & \quad 1 & \quad -\frac{1}{2}(1-u)u. & \end{aligned} \quad (2)$$

The interior BB-coefficients (circles in Figure 6(b)) of the bi-4 cap are determined so columns of BB-coefficients form degree-raised curves of true degree 3. By construction, see Figure 6(c), the red bi-4 cap is internally C^1 and joins with G^1 -continuity to the green frame.

5. CAPS FOR PARALLEL T-JUNCTIONS

Configurations with multiple T-junctions can in principle be locally re-meshed to separate them into isolated \tilde{T} -nets. For completeness, and to compare what surface quality can be achieved, we investigate configurations where two T-junctions face another, as shown in Figure 7(b). We call the configuration a \tilde{T} -net (pronounced T3-net). It has one nominally 7-sided face. Such T-junctions can arise, for example, from configuration Figure 7(a) by removing the two red edges of a triangle attached to a point of valence 5. (Asymmetrically removing one yields a mesh with one T-junction as in Figure 7(c)). Applying Catmull-Clark subdivision to \tilde{T} -junctions leads to poor surfaces. As for \tilde{T} -nets, first a bi-3 frame is constructed. Figures 7(e) and (f) provide the stencils for the points marked \square and \circ . Except across the red hv-curves between the horizontal and the vertical strips, the frame is C^2 . In the spirit of the bi-4 \tilde{T} -net construction, a bi-4 cap for the \tilde{T} -net is constructed by subdividing, in the ratio shown in Figure 8(a), the C^1 -prolongation of the frame from the top; and then evenly splitting the middle prolongation of the bottom. Both the top and bottom prolongations are C^1 -connected (in the

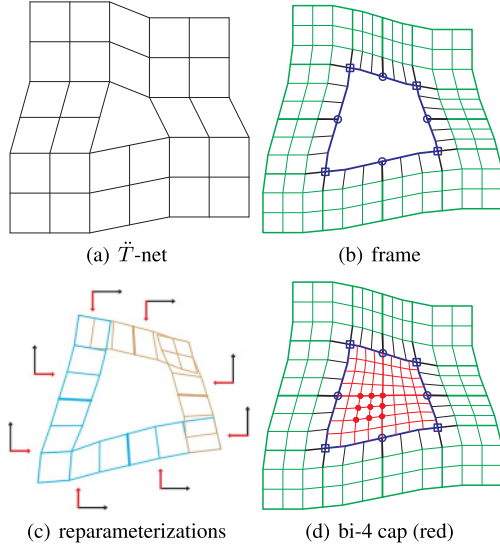


Fig. 9. Bi-4 cap for a \tilde{T} -configuration. The red axes in (c) indicate the v -parameter.

horizontal direction) with the same continuity parameters from left to right: $\beta = \frac{1}{2}$, $\beta = 1$, $\beta = 2$. This implies a continuity parameter of $\beta = \frac{2}{3}$ across the top hv-curves. For the prolongation 1 of $\mathbf{q}^{l,0}$ (see Figure 8(b)) to match the split prolongation of the top, the data are reparameterized according to

$$\begin{array}{lll} & a(u) := & b(u) := \\ \text{left(1)} : & 1 - u/3 & 0 \\ \text{top(left)} : & 1 & \frac{1}{2}(1 - u)u \\ \text{bottom(left)} : & 1 & -\frac{1}{3}(1 - u)u. \end{array} \quad (3)$$

(For the right top and bottom reparameterizations, $b(u)$ is negated.) The remaining (circled) BB-coefficients in Figure 8 are chosen to make their columns have actual degree 3.

6. CAPS FOR TWO T-JUNCTIONS IN CROSSING DIRECTIONS

We also investigate configurations with two T-junctions as shown in Figure 9(a). We call the configuration a \tilde{T} -net (pronounced T2-net). It has one nominally six-sided face. Note that such configurations are explicitly excluded in dyadic T-meshes [Kovacs et al. 2015].

Capping \tilde{T} -nets with good surface quality is similar but more challenging than the earlier constructions. As before, B -spline to Bézier conversion yields the green Bézier control points in Figure 9(b), now with both the top and the right side patches subdivided.

The corner points (blue boxes in Figure 9(b)) and the middle, blue circle coefficients of innermost (blue) boundary curves of the frame are determined only in the last step of the construction. The direct neighbors of the corner points are chosen so adjacent bi-3 patches of the frame connect C^1 (with ratios 1,1 lower left, 1/2,1/2 upper right and 1/2,1 otherwise); the direct neighbors of the middle points are determined so each boundary curve is internally C^2 .

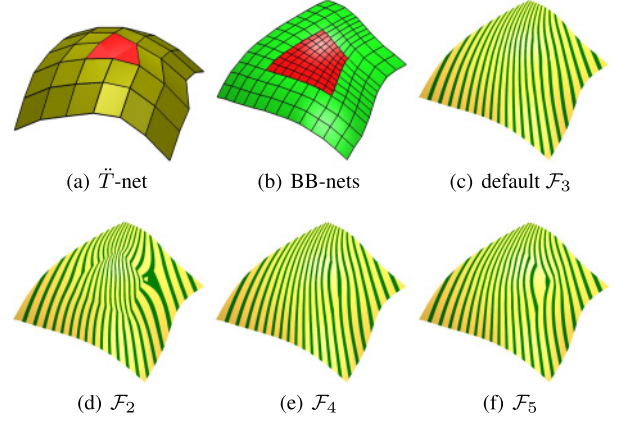


Fig. 10. The surfaces obtained by minimizing functionals. (a) Convex \tilde{T} -net. (b) BB-coefficients of the frame (green) and the central bi-4 cap (red). ((c)–(f)) Highlight lines.




Fig. 11. A \tilde{T} -net (\mathcal{F}_3) surface for a non-convex input mesh.

The C^1 -prolongations are not compatible at the corners (see Figure 9(c)). To make them compatible they are re-parameterized with

$$\begin{array}{lll} & a(u) := & b(u) := \\ \text{top-left} : & 1 - \frac{u}{4} & \frac{1}{2}(1 - u)u \\ \text{top-right} : & \frac{3}{4} - \frac{u}{4} & -\frac{1}{2}(1 - u)u \\ \text{bottom-left} : & 1 - \frac{u}{4} & -\frac{1}{4}(1 - u)u \\ \text{bottom-right} : & \frac{3}{4} - \frac{u}{4} & \frac{1}{4}(1 - u)u. \end{array}$$

This list of the reparameterizations is complete due to the (combinatorially) diagonal symmetry (see the local coordinate systems in Figure 9(c)). The reparameterized tensor-border is of degree 4 and depth 1 and ensures G^1 continuity of the central cap with the frame. Choosing to join C^2 the four 3×3 groups of interior BB-coefficients leaves free one group shown as red disks. Finally, we minimize, over all 11 bi-3 patches of the frame and the 4 bi-4 patches of the central cap, the functional \mathcal{F}_3 , where

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

We minimize the sum with respect to 17 unknown coefficients: 4 corner (blue box), 4 mid-edge (blue circle), and 9 inner ones (red disks in Figure 9(d)). In the implementation, these 17 coefficients enter as affine combinations of \tilde{T} -net points with pre-computed coefficients.

The choice \mathcal{F}_3 is the result of testing a series of input meshes including the challenging elliptic configuration in Figure 10. Figure 11 confirms that this choice also works well for a wave-like input mesh.




Fig. 12. An alternative construction that does not split the patch opposite to the T-junction leads to a poor highlight line distribution.




Fig. 13. The GT-spline surface for the T-mesh of Figure 2 that does not admit a C^1 T-spline. The bi-4 cap in red; the right image shows highlight lines.




Fig. 14. Two regular meshes are merged with a T-junction.

7. DISCUSSION AND COMPARISON

7.1 The Frame Construction: To Split Or Not to Split

Note the flaws in the highlight linedistribution of the surface in Figure 12(b). For the same input mesh Figure 1(b), the highlight linedistribution of Figure 12(b) is much worse than that of the GT-construction Figure 3(b). Despite being formally smooth, surfaces generated by not subdividing the frame patch opposite the T-junction have poor highlight lines. Evidently splitting, though not required by formal smoothness constraints, improves shape quality.

7.2 Comparing GT-spline s to T-splines

In general a comparison to hierarchical splines does not make sense since we proved at the outset that not all meshes with T-junctions admit smooth T-splines. For one such configuration, the bracelet mesh of Figure 2 (that does not admit a C^1 T-spline), Figure 13 demonstrates that applying GT-splines yields a bi-4 surface with an excellent highlight line distribution.

Figure 14 shows the case of two regular meshes of different quad-patch count joined via a T-junction. T-splines require forming a common parameter domain, that, while easy for simple meshes, is impossible for more complex meshes such as Figure 18. The result of directly applying a GT-spline is displayed in Figures 14(b) and (c). To establish an upper bound on the quality of the highlight linedistribution, Figure 15 compares GT-splines to T-splines in the form of standard bi-cubic tensor-product splines. The top row of Figure 15 shows in order (a) the \hat{T} -net input to the GT-construction, (b) the geometrically identical T-mesh, with all knot intervals are 1 except for 0.5 on the thick edges; and (c) the mesh resulting from splitting the cyan knot intervals yielding new (cyan disks) and moved (cyan circles) control points. Since the T-spline surface of input mesh (b) equals the non-uniform C^2 tensor-product

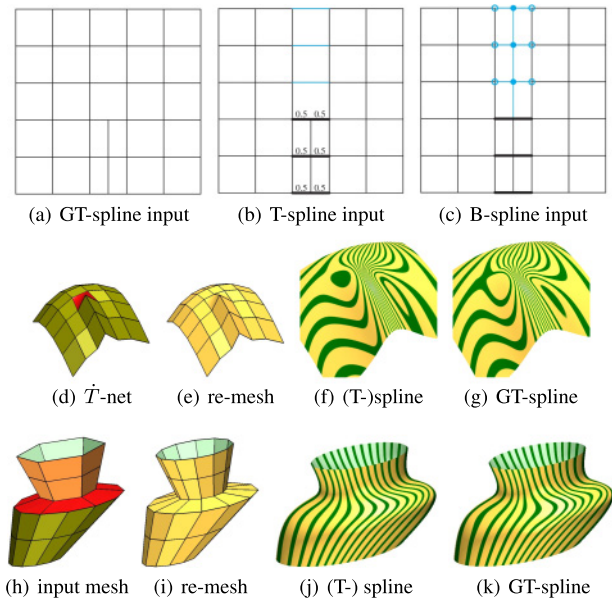


Fig. 15. Comparisons to T-splines and tensor-product splines. Since, in these examples, the T- and B-spline surfaces coincide, we refer to them as (T)-spline surfaces.

B-spline surface of input mesh (c), it suffices to discuss the case (c) in the following. For a concrete comparison in this setting, the mesh in Figure 15(d) is the geometric input both for the GT-spline and for the T-spline, whereas Figure 15(e) is the tensor-product B-spline input mesh. As explained above, the knot-intervals are chosen so the B-spline surface equals the T-spline surface. Figures 15(f) and (g) show the highlight linedistribution on the surfaces.

In Figure 15(h) two regular meshes are connected by a whole ring of T-junctions. This is the (geometric) input mesh both for the GT-spline and for the T-spline, but the knot-intervals differ: for the GT-spline they can all be chosen equal, while the horizontal T-spline knot intervals of bottom mesh are 0.5 when those of the top mesh are 1. Figure 15(i) shows the tensor-product mesh yielding the C^2 surface Figure 15(j) that coincides with T-spline surface from mesh Figure 15(h). Figure 15(k) shows the GT-spline surface obtained directly from Figure 15(h). The highlight lines of the GT-construction are very similar to those of the (T)-spline construction even though the GT-spline is placed at a disadvantage by capping T-junctions while the (T)-spline can take advantage of a tensor-product mesh.

Given the similarity in shape, it is important to recall the essential difference between T-splines and GT-splines. The simple knot intervals used *locally* above can lead to invalid knot-intervals when considering larger meshes: T-splines require a global coordination of knot sequences. Such global coordination is not always possible or may require complex re-meshing. The GT-construction is local, sidestepping the need for global coordination, and producing surfaces of comparable quality.

7.3 Separation and Re-meshing

The GT-constructions in Section 4, 5, and 6 follow a common pattern, of adjusting the bi-3 frame and then forming a central cap of degree bi-4. The minimal submesh required for the GT-constructions is called “net.” Only the outer boundaries of the frame may have

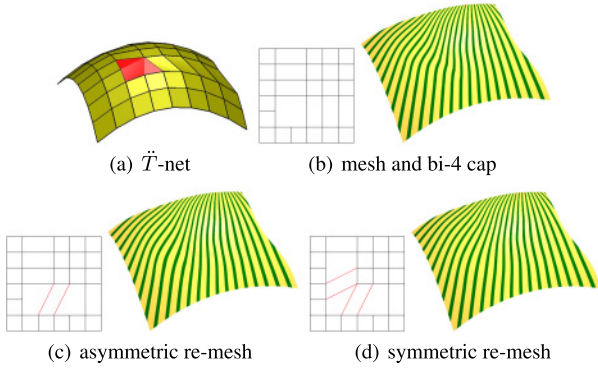


Fig. 16. Effect of re-meshing on the resulting bi-4 cap.

irregular nodes, where $n \neq 4$ quads meet, or T-junctions. Guaranteeing such separation may not be easy in general (see Limitations below) but often local adjustments can be made. While Section 6 demonstrated that \tilde{T} -nets can yield bi-4 surfaces of good quality, better surfaces are often obtained by locally re-connecting the mesh points to isolate the T-junctions as in Figure 16(c). Reconnecting with an even larger footprint for a more symmetric reconnection as in Figure 16(d) can further improve the highlight linedistribution, at the cost of higher construction complexity.

7.4 Limitations

Quadrangulations can contain closely packed T-junctions and irregular points. As our initial example demonstrated, Catmull-Clark refinement is not a good way to separate T-junctions: not only does subdivision increase the number of patches, but, more importantly, it can negatively impact the surface quality. T-mesh subdivision [Kovacs et al. 2015] can also start with irregular points adjacent to T-junctions but does not separate T-junctions from irregular points. Therefore, it too cannot be used for pre-processing.

Although re-meshing can reduce many configurations to the three standard T-nets, GT-constructions are not expected to work with arbitrary T-junction distributions. Here, the quad-meshing algorithm or the designer have to enforce some discipline, already to obtain good shape. Many-sided facets with T-junctions as generated by Alliez et al. [2003] or motorcycle graphs [Eppstein et al. 2008; Myles et al. 2014] are outside the scope of GT-constructions.

In many cases, our approach can allow for tighter packing of T-junctions (for example as in Figure 15, *bottom* row, and Figure 18) and irregularities. However, this article does not attempt to provide a set of recipes for arbitrarily complex T-junctions. On one hand, local re-connection can often reduce the situation to a collection of \tilde{T} -net, $\tilde{\tilde{T}}$ -nets or $\tilde{\tilde{\tilde{T}}}$ -nets, but a principled prescription for such quad-re-meshing is outside the scope. On the other hand, our experiments with complex configurations show that keeping the complexity of capping T-junctions to a minimum results in better shape.

8. COLLABORATION WITH MULTI-SIDED CAPS

Figure 17 demonstrates that sufficiently isolated caps for T-junctions co-exist without problems with irregular vertices where the surface caps are also of degree bi-4 when we apply [Karčiauskas et al. 2016]. Again, tighter configurations are possible, but they may reduce surface quality while increasing the complexity of implementation.

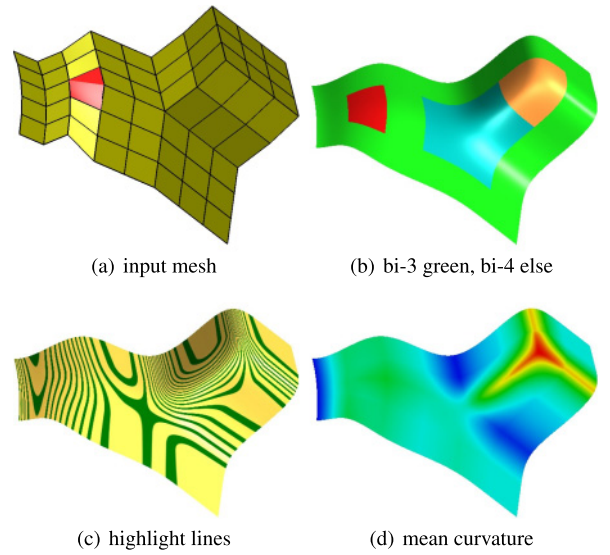


Fig. 17. Mesh and bi-4 surface combining a T-junction with irregular regions of valence 3 and 5 where [Karčiauskas et al. 2016] is applied.

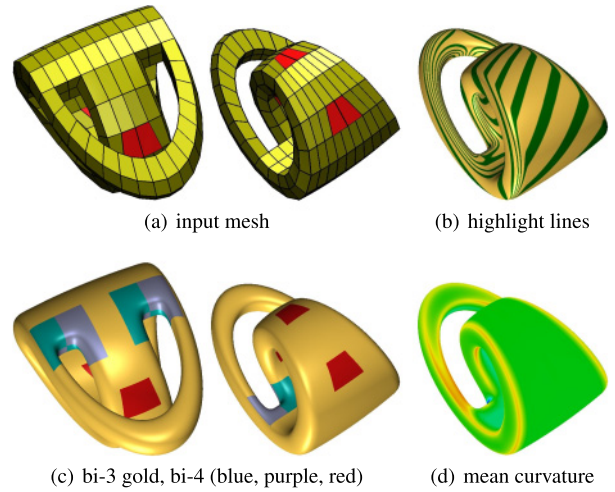


Fig. 18. Mesh and G^1 surface that include horizontally paired \tilde{T} -nets (red) and irregular neighborhoods (blue, purple) treated with [Karčiauskas et al. 2016]. This mesh does not admit a globally consistent (non-zero) knot interval assignment for smooth T-splines.

Figure 18 presents another free-form design that challenges algorithms that require globally consistent knot intervals. As in Figure 2, enforcing T-spline Rule 1 yields zero knot intervals, now also at extraordinary points.

9. CONCLUSION

The article introduced a construction of surface caps for merging and spreading feature lines via T-junctions. In the default \tilde{T} -case, the caps consist of two surface patches of degree bi-4 (see Figure 19), otherwise of four bi-4 patches. The surrounding bi-3 patches are only perturbed where they join the central cap. Since the approach is based on geometric continuity, it does not require




Fig. 19. Distribution of the bi-degree 4 patches (red) for the three basic configurations. The frames (green) are of degree bi-3.

non-local coordination of knot intervals and is not restricted to graphs of functions but applies to general manifolds.

APPENDIX

Due to symmetry and C^2 continuity of the bi-3 frame with the surrounding surface, the \tilde{T} -net-construction is completely defined by the following formulas for one (here the “left”) bi-4 patch. As illustrated for two of the coefficients in Figures 5(d) and (f), each of the 5×5 BB-coefficient is a linear combination of the inner \tilde{T} -net nodes with two rows of four weights and two rows of five weights. In Table I, the stencil for each BB-coefficient, is scaled $\times 144$ and placed into brackets.

Table I. 5×5 Stencils, Scaled by 144 to Convey Relative Size of Entries While Using the Fewest Digits, of the Bi-4 Coefficients of the Construction for \tilde{T} -nets. Each Stencil Lists the Weight of the Control Nodes Surrounding the \tilde{T} -net

[illegible]

REFERENCES

- Pierre Alliez, David Cohen-Steiner, Olivier Devillers, Bruno Lévy, and Mathieu Desbrun. 2003. Anisotropic polygonal remeshing. *ACM Trans. Graph.* 22, 3 (July 2003), 485–493. DOI: <https://doi.org/10.1145/882262.882296>
- Ursula H. Augsdörfer, Neil A. Dodgson, and Malcolm A. Sabin. 2011. Artifact analysis on b-splines, box-splines and other surfaces defined by quadrilateral polyhedra. *Comput. Aided. Geometr. Des.* 28, 3 (2011), 177–197.
- D. Bommers, B. Lvy, N. Pietroni, E. Puppo, C. Silva, M. Tarini, and D. Zorin. 2013. State of the art in quad meshing. *Comput. Graph. Forum* 32, 6 (2013), 51–76.
- E. Catmull and J. Clark. 1978. Automatic and interactive mesh to T-spline conversion. *Computer-Aided Design* 10, 6 (1978), 191–200. <http://www.eg.org/EG/DL/WS/SGP/SGP06191-200.pdf>.
- Tor Dokken, Tom Lyche, and Kjell Fredrik Pettersen. 2013. Polynomial splines over locally refined box-partitions. *Comput. Aided. Geometr. Des.* 30, 3 (2013), 331–356.
- David Eppstein, Michael T. Goodrich, Ethan Kim, and Rasmus Tamstorf. 2008. Motorcycle graphs: Canonical quad mesh partitioning. In *Proceedings of the Symposium on Geometry Processing (SGP'08)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 1477–1486.
- Gerald Farin. 1988. *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide*. Academic Press.
- Carlotta Giannelli, Bert Jüttler, and Hendrik Speleers. 2012. THB-splines: The truncated basis for hierarchical splines. *Comput. Aided. Geometr. Des.* 29, 7 (2012), 485–498.
- Wenzel Jakob, Marco Tarini, Daniele Panozzo, and Olga Sorkine-Hornung. 2015. Instant field-aligned meshes. *ACM Trans. Graph.* 34, 6 (2015), 189. <http://doi.acm.org/10.1145/2816795.2818078>
- Hongmei Kang, Jinlan Xu, Falai Chen, and Jiansong Deng. 2015. A new basis for PHT-splines. *Graph. Models* 82 (2015), 149–159.
- Kęstutis Karčiauskas, Thien Nguyen, and Jörg Peters. 2016. Generalizing bicubic splines for modelling and IGA with irregular layout. *Comput. Aided. Des.* 70 (Jan. 2016), 23–35.
- K. Karčiauskas and J. Peters. 2011. Rational G^2 splines. *Graph. Models* 23, 5 (2011), 286–295. DOI: <https://doi.org/10.1016/j.gmod.2011.05.004>
- Denis Kovacs, Justin Bisceglia, and Denis Zorin. 2015. Dyadic T-mesh subdivision. *ACM Trans. Graph.* 34, 4, Article 143 (July 2015), 12 pages. <https://doi.org/10.1145/2766972>
- R. Kraft. 1998. *Adaptive und Linear Unabhängige Multilevel B-Splines und Ihre Anwendungen*. Ph.D. Thesis. University of Stuttgart.
- Yu-Kun Lai, Leif Kobbelt, and Shi-Min Hu. 2008. An incremental approach to feature aligned quad dominant remeshing. In *Proceedings of the 2008 ACM Symposium on Solid and Physical Modeling (SPM'08)*. ACM, New York, NY, 137–145. DOI: <https://doi.org/10.1145/1364901.1364921>
- Wan-Chiu Li, Nicolas Ray, and Bruno Lévy. 2006. Automatic and interactive mesh to T-spline conversion, Alla Sheffer Konrad Polthier (Ed.). *Proceedings of the Eurographics Symposium on Geometry Processing*.
- Xin Li, Jianmin Zheng, Thomas W. Sederberg, Thomas J. R. Hughes, and Michael A. Scott. 2012. On linear independence of T-spline blending functions. *Comput. Aided. Geometr. Des.* 29, 1 (2012), 63–76.
- Martin Marinov and Leif Kobbelt. 2004. Direct anisotropic quad-dominant remeshing. In *Proceedings of the 12th Pacific Conference on Computer Graphics and Applications (PG'04)*. IEEE Computer Society, Washington, DC, 207–216.
- Ashish Myles, Nico Pietroni, Denis Kovacs, and Denis Zorin. 2010. Feature-aligned T-meshes. *ACM Trans. Graph.* 29, 4, Article 117 (July 2010), 11 pages. DOI: <https://doi.org/10.1145/1778765.1778854>
- Ashish Myles, Nico Pietroni, and Denis Zorin. 2014. Robust field-aligned global parametrization. *ACM Trans. Graph.* 33, 4, Article 135 (July 2014), 14 pages. DOI: <https://doi.org/10.1145/2601097.2601154>
- Chi-Han Peng and Peter Wonka. 2013. Connectivity editing for quad-dominant meshes. *Comput. Graph. Forum* 32, 5 (2013), 43–52. DOI: <https://doi.org/10.1111/cgf.12171>
- J. Peters. 1991. Smooth interpolation of a mesh of curves. *Constructive Approximation* 7, 2 (1991), 221–247.
- J. Peters. 2002. Geometric continuity. In *Handbook of Computer Aided Geometric Design*. Elsevier, 193–229.
- Nico Pietroni, Enrico Puppo, Giorgio Marcias, Roberto Scopigno, and Paolo Cignoni. 2016. Tracing field-coherent quad layouts. *Comput. Graph. Forum* 35, 7 (2016), 485–496.
- Nicolas Ray, Wan Chiu Li, Bruno Lévy, Alla Sheffer, and Pierre Alliez. 2006. Periodic global parameterization. *ACM Trans. Graph.* 25, 4 (Oct. 2006), 1460–1485. DOI: <https://doi.org/10.1145/1183287.1183297>
- Thomas W. Sederberg, G. Thomas Finnigan, Xin Li, Hongwei Lin, and Heather Ipson. 2008. Watertight trimmed NURBS. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 79:1–79:8.
- Thomas W. Sederberg, Jianmin Zheng, Almaz Bakenov, and Ahmad Nasri. 2003. T-splines and T -NURCCs. In *Proceedings of ACM SIGGRAPH 2003 (ACM Transactions on Graphics)*, Jessica Hodgins and John C. Hart (Eds.), Vol. 22(3). ACM Press, 477–484.
- Amir Vaxman, Marcel Campen, Olga Diamanti, Daniele Panozzo, David Bommers, Klaus Hildebrandt, and Mirela Ben-Chen. 2016. Directional field synthesis, design, and processing. *Comput. Graph. Forum* 35, 2 (2016), 545–572.
- Wenyan Wang, Yongjie Zhang, Michael A. Scott, and Thomas J. R. Hughes. 2011. Converting an unstructured quadrilateral mesh to a standard T-spline surface. *Comput. Mech.* 48 (2011), 477–498.
- Mirko Zadravec, Alexander Schiftner, and Johannes Wallner. 2010. Designing quad-dominant meshes with planar faces. *Comput. Graph. Forum* 29, 5 (2010), 1671–1679.

Received November 2016; revised April 2017; accepted July 2017