

Can bi-cubic surfaces be class A?

Kęstutis Karčiauskas^a and Jörg Peters^b

^a Vilnius University ^b University of Florida

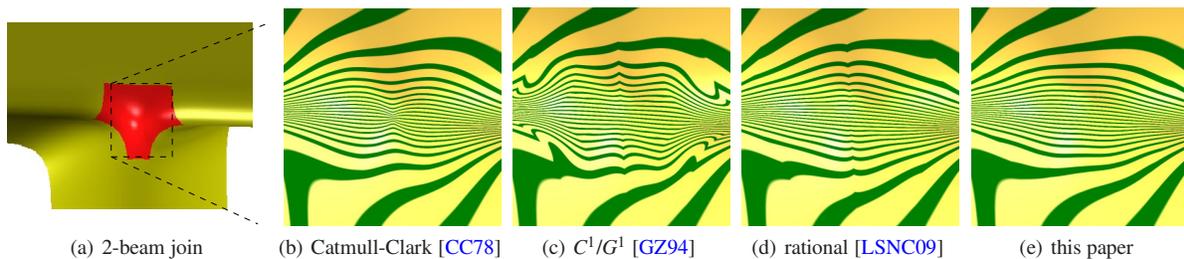


Figure 1: Class A surfaces avoid oscillating highlight lines. (a) Joining two crossing beams, modeled by bi-3 B-spline surfaces, via a multi-sided cap (red). Of the four displayed bi-3 constructions only construction (e) qualifies as class A.

Abstract

‘Class A surface’ is a term in the automotive design industry, describing spline surfaces with aesthetic, non-oscillating highlight lines. Tensor-product B-splines of degree bi-3 (bicubic) are routinely used to generate smooth design surfaces and are often the de facto standard for downstream processing. To bridge the gap, this paper explores and gives a concrete suggestion, how to achieve good highlight line distributions for irregular bi-3 tensor-product patch layout by allowing, along some seams, a slight mismatch of normals below the industry-accepted tolerance of one tenth of a degree. Near the irregularities, the solution can be viewed as transforming a higher-degree, high-quality formally smooth surface into a bi-3 spline surface with few pieces, sacrificing formal smoothness but qualitatively retaining the shape.

Categories and Subject Descriptors (according to ACM CCS): I.x.y [Computer Graphics]: Generation—Surface generation

1. Introduction

Cubic splines allow for inflections and for variation of the second derivative that enables more subtle shapes than quadratic curves. Their tensor-product, with boundaries aligned with feature curves or curvature lines, typically gives the designer the necessary range of freedom to express styling intent. For regular, grid-like layout, bi-cubic (bi-3) tensor-product splines additionally provide smoothness of the normal and continuity of the curvature. Bi-3 tensor-product splines are therefore often chosen as the uniform, standard representation of outer class A surfaces [wik15], i.e. spline surfaces with aesthetic highlight lines [BC94] that satisfy, depending on the application area and contractual agreement, certain hard geometric constraints [And15].

Downstream processing, such as some machine tool controllers, are often restricted to bi-cubic representation. In Computer Graphics, curved surfaces are often introduced by bi-cubic examples, such as Newell’s teapot, or spline patches for ray-traced scenes [PoV15].

Extending the uniform, ubiquitous bi-cubic B-spline representation to class A surfaces with an irregular layout of quadrilateral pieces has motivated early work of John Gregory [Gre74, GZ94] as well as motivated the now-dominant subdivision construction, Catmull-Clark subdivision surfaces [CC78]. However, none of these constructions consistently yields class A surfaces. While a single Catmull-Clark subdivision step is typically beneficial in that it smoothes out sharp transitions near irregular points where $n \neq 4$ patches

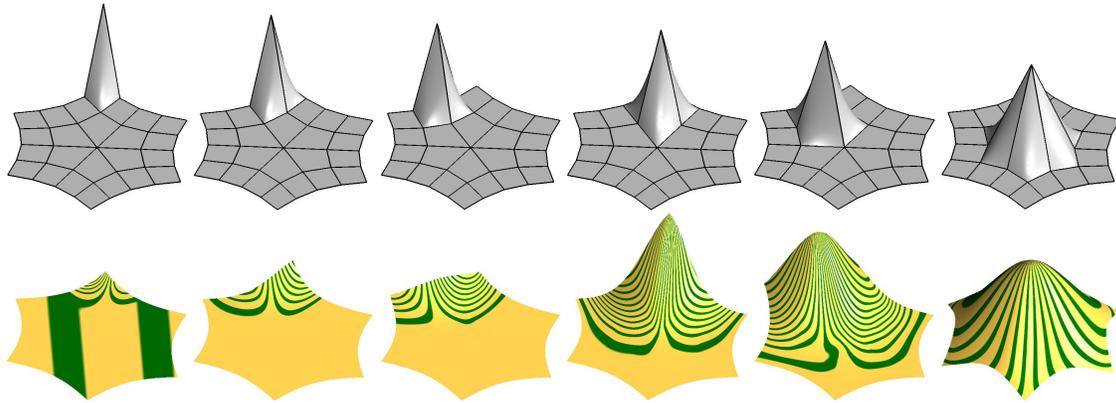


Figure 2: Top row: Control nets (CC-nets) of *basic functions* with layout from the characteristic control net of Catmull-Clark subdivision for $n = 6$. In each a single vertex is pulled up. Bottom row: Basic functions h_1 through h_7 , skipping h_4 which is symmetric to h_2 across the diagonal of the same sector. Applying the construction of Section 2 to the input meshes of the top row yields the (BB-coefficients of the) basic functions h_ℓ in the bottom row. Surface caps are linear combinations of the h_ℓ with quad-net vertices acting as B-spline(-like) control points.

meet, repeated subdivision yields an unnecessary concentration of highlight lines ('bunching up') or to a divergence of highlight lines. Both introduce oscillations of the highlight lines unacceptable for class A surfacing. Moreover, the inherently singular representation and infinite number of polynomial pieces near irregular points of Catmull-Clark surfaces and its variants [Cas12] causes challenges for downstream processing and engineering analysis, including non-uniform size, infinitely many T-joints and a lack of efficient integration rules near irregular points.

Bi-cubic constructions that generate a finite number of polynomial pieces by exploiting the freedom to reparameterize, a.k.a. geometric continuity, have to date failed the test of non-oscillating highlight lines across irregular points. Also constructions with singularities at the vertices, such as Gregory patches [Gre74, LSNC09] fail the class A test. A central question in geometry processing therefore remains: can bi-cubic surfaces be class A?

In the following pages, we answer the question cautiously affirmative: class A surfaces constructed from tensor-product B-splines of degree bi-3 can be extended by bi-3 *surface caps* in the vicinity of irregular points so that the highlight lines are class A everywhere. Our solution requires a slight mismatch of normals where the regular tensor-product arrangement of patches switches to the cap neighborhood of an irregular point. We model each of the n sectors of a n -sided cap by 2×2 polynomial pieces of degree bi-3, akin to the binary refinement of one localized step of Catmull-Clark subdivision. The caps represent a careful transformation of an existing high-quality, but higher-degree surface. The novelty of the approach is that it sacrifices strict formal smoothness to qualitatively retain the class A highlight lines of the guiding surface and lower the degree. The resulting normal

mismatch along the cap boundaries is consistently observed to be below the industry-accepted tolerance of a tenth of a degree [Aut15]. Due to the underlying non-linear relations, deriving formal criteria on the input quad mesh that guarantee the tolerance seems impractical. As a basic check, we analyze, for fixed valences (see Fig. 2 for $n = 6$), the seven basic functions h_ℓ whose linear combinations form the surface caps. (The control nets of the h_ℓ are also some of the more challenging input quad configurations.) We verified, up to $n = 10$ and numerically for higher n that the normal mismatch of the h_ℓ is less than 0.1° (which does, of course, not imply that all their linear combinations obey the bound; hence Section 4 presents many test cases).

Industry accepts such mismatches since they fall below manufacturing tolerances. For example, contracts for CAD in the automotive industry might characterize deliverable class A surfaces as having an angle mismatch of less than $.1^\circ$ [And15]. Correspondingly, 'Fill Surface' and 'Fill Hole' tools in CAD modelers cover multi-sided holes in a tensor-product spline complex by variants of a spline fitting approach [Ead15]: re-sample the geometric data as the Hermite data of a function with respect to one or more datum planes and use one or more spline functions to interpolate the Hermite data. Since the rectangular footprint of the tensor-product spline does not usually match the curved multi-sided hole, extraneous surface pieces are trimmed away. This approach can result in mismatches even of position. Our approach differs from datum plane interpolation, as well as from the notion of approximate smoothness explored in [LM08], in that we leverage an underlying guide surface. Without the guide the degrees of freedom from dropping exact smoothness constraints are very difficult to harness towards class A quality.

Overview. Section 2 explains the ingredients of the 2×2 construction of bi-3 surfaces for arbitrary quad-layout. Section 3 summarizes the formal properties of the resulting surfaces and Section 4 compares the surfaces to the state-of-the-art with numerous, selected examples.

2. Setup, Definitions and Construction

This section formalizes the input and output representation, explains the role of the guide surface and shows how first-order geometric continuity is enforced within the cap.

2.1. Input control net = a quad-net

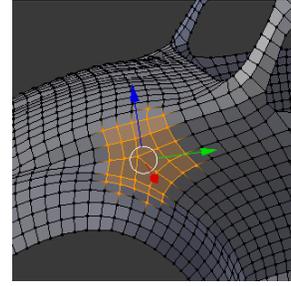
We consider a quilt of quadrilateral facets or *quads* that outline the final surface (see Fig. 3a). The vertices or *nodes* \mathbf{c}_i of this network will act as bi-3 B-spline control points wherever the network forms a tensor-product 4×4 sub-net and as control points (fat points in Fig. 3b) of B-spline-like functions h_ℓ otherwise. That is, as for standard splines, the final surface is the superposition of bi-3 B-splines and B-spline-like functions h_ℓ forming the cap, weighted by the nodes of the input quad-mesh. For each valence n , we need only the contributions to the n -sided cap of seven *basic functions* h_ℓ (see Fig. 2 for $n = 6$) that, by rotational symmetry, define all $6n + 1$ B-spline-like functions associated with two layers of quad facets surrounding an *irregular node* Fig. 3b. An irregular node is a mesh point with $n \neq 4$ neighbors. We assume that each quad has only one irregular node – otherwise we locally apply one step of Catmull-Clark-subdivision. In the following, we also do not worry about surfaces with boundaries but assume an outer layer of regular nodes to which standard spline end conditions can be applied.

Our challenge is to define the functions h_ℓ wherever the 4×4 sub-net does not have tensor-product structure because one of its four interior nodes has valence $n \neq 4$ (see Fig. 3b). We will represent each polynomial piece f of h_ℓ in tensor-product Bernstein-Bézier form (BB form) of degree bi-3 and with BB-coefficients f_{ij} :

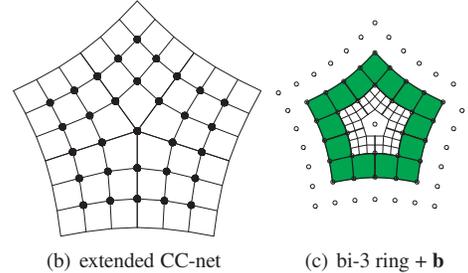
$$f(u, v) := \sum_{i=0}^3 \sum_{j=0}^3 f_{ij} B_i^3(u) B_j^3(v), \quad (u, v) \in \square := [0..1]^2,$$

where $B_k^3(t)$ is the k th Bernstein-Bézier polynomial of degree 3.

Since each tensor-product 4×4 grid of the network is interpreted as the B-spline control net of one bicubic piece of a tensor-product spline surface, the well-known formulas of a B-spline to Bernstein-Bézier (BB)-form conversion [Farin02, Prautzsch02] can be applied to the input quad-net. This also generates position, first and second order derivative data in BB form along the boundary curves of each cap (see Fig. 3c). The Hermite data \mathbf{b} are therefore called the *tensor-border* (of depth 2 and degree 3). The data are fully defined



(a) CC-net in a quad-mesh



(b) extended CC-net

(c) bi-3 ring + \mathbf{b}

Figure 3: The *CC-net*, a sub-net of the input quad-mesh, defines the cap. (a) A *CC-net* highlighted in a larger quad-mesh. (b) The $6n + 1$ B-spline-like control points of a *CC-net* marked by black bullets. (c) The inner net represents the BB-coefficients of the tensor-border \mathbf{b} . The tensor-border is fully defined by the *CC-net*. If all the points of the second ring of the *CC-net* are regular (they need not be) then the *CC-net* plus the next ring define the green ring of bi-3 B-spline patches.

by the $6n + 1$ nodes marked as bullets in Fig. 3b. In our construction, these nodes define the cap as $n \times 2 \times 2$ -patches that each consist of 2×2 polynomial pieces of degree bi-3.

2.2. A guide surface

The second-order Hermite data from the tensor-borders can be used to define a surface of degree bi-5 [KP14] that smoothly joins with the input data and gently propagates curvature from the boundary into the cap. The structure of this surface is shown in Fig. 4a. The BB-coefficients emphasized as red disks define a quadratic polynomial in each sector that, after reparameterization, matches a unique quadratic expansion at the irregular point – so that the guide surface has a well defined curvature at the irregular point. We will use this bi-5 surface \mathbf{g} to guide the shape of our bi-3 cap.

To transfer the information from the high-degree, high-quality guide surface \mathbf{g} to our bi-3 surface, we split its patches \mathbf{g}^k surrounding the irregular point along the parameter lines $u = 1/2, v = 1/2$ into four sub-patches $\mathbf{g}^{k,r}$, $r = 1, \dots, 4$. For each corner of each sub-patch $\mathbf{g}^{k,r}$, the BB-sub-nets defining position and (mixed) first derivative are

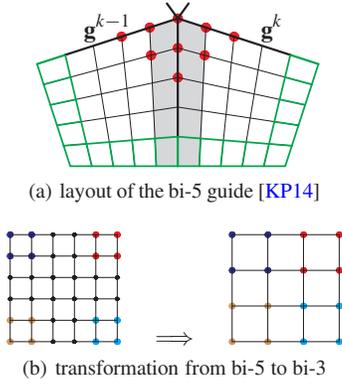


Figure 4: (a) Layout of the *bi-5 guide* surface \mathbf{g} from [KP14]. The choice of BB-coefficients emphasized by red disks implies a well-defined curvature at the irregular point. (b) Each *bi-5* piece (of the subdivided \mathbf{g}) is approximated by a *bi-3* piece by transforming the Hermite data at the four corners of the *bi-5* piece to *bi-3* form.

transformed to *bi-3* form and merged to form patches $\mathbf{a}^{k,r}$ of degree *bi-3* as illustrated in Fig. 4b. By construction the patches $\mathbf{a}^{k,r}$ form an internally C^1 2×2 patch \mathbf{a}^k that joins continuous with the *bi-3* ring Fig. 3c. The transitions between the \mathbf{a}^k are not smooth and this will be addressed next.

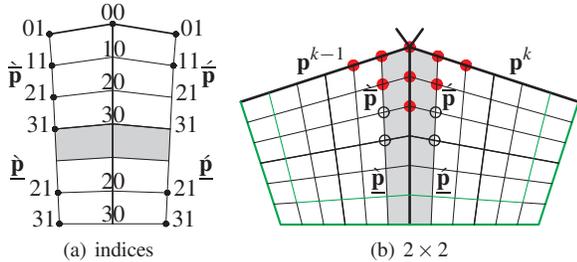


Figure 5: Local indices and layout of two adjacent 2×2 patches of the *bi-3* cap.

2.3. Joining the cap's *bi-3* 2×2 -patches with geometric continuity

As we construct the *bi-3* cap \mathbf{p} , we now focus on a pair of its adjacent 2×2 -patches \mathbf{p}^{k-1} and \mathbf{p}^k , $k = 0, \dots, n-1$ (superscript modulo n) that are initialized to the corresponding auxiliary patches as $\mathbf{p}^k := \mathbf{a}^k$. As illustrated in Fig. 5, the pieces of \mathbf{p}^{k-1} therefore join C^0 with the corresponding pieces of \mathbf{p}^k along the shared boundary curves $\tilde{\mathbf{p}}(u, 0) = \tilde{\mathbf{p}}(u, 0)$ and $\underline{\mathbf{p}}(u, 0) = \underline{\mathbf{p}}(u, 0)$. We preserve the internal C^1 continuity and enforce smoothness between sectors by en-

forcing the G^1 constraints

$$\partial_v \tilde{\mathbf{p}} + \partial_v \underline{\mathbf{p}} = c(2(1-u) + u)\partial_u \tilde{\mathbf{p}}, \quad c := \cos \frac{2\pi}{n},$$

$$\partial_v \tilde{\mathbf{p}} + \partial_v \underline{\mathbf{p}} = c(1-u)\partial_u \underline{\mathbf{p}},$$

or, equivalently, in terms of the BB-coefficients represented as unmarked points in Fig. 5a,

$$\tilde{\mathbf{p}}_{10} := \frac{2(c-1)\tilde{\mathbf{p}}_{00} + \tilde{\mathbf{p}}_{01} + \underline{\mathbf{p}}_{01}}{2c}; \quad (1)$$

$$\tilde{\mathbf{p}}_{20} := \frac{c\tilde{\mathbf{p}}_{00} + 3(c-2)\tilde{\mathbf{p}}_{10} + 3(\tilde{\mathbf{p}}_{11} + \underline{\mathbf{p}}_{11})}{4c}; \quad (2)$$

$$\tilde{\mathbf{p}}_{30} := -\frac{1}{4}\tilde{\mathbf{p}}_{10} + \tilde{\mathbf{p}}_{20} + \frac{1}{4}\underline{\mathbf{p}}_{20}; \quad (3)$$

$$\underline{\mathbf{p}}_{21} := -\frac{2c}{3}\tilde{\mathbf{p}}_{10} + 2\tilde{\mathbf{p}}_{20} + \frac{2c}{3}\tilde{\mathbf{p}}_{30} - \underline{\mathbf{p}}_{21}; \quad (4)$$

$$\underline{\mathbf{p}}_{31} := -c\tilde{\mathbf{p}}_{20} + (2+c)\tilde{\mathbf{p}}_{30} - \underline{\mathbf{p}}_{31}; \quad (5)$$

$$\underline{\mathbf{p}}_{20} := \frac{3(\underline{\mathbf{p}}_{21} + \underline{\mathbf{p}}_{21}) - c\underline{\mathbf{p}}_{30}}{6-c}; \quad (6)$$

$$\underline{\mathbf{p}}_{30} := \frac{1}{2}\underline{\mathbf{p}}_{31} + \frac{1}{2}\underline{\mathbf{p}}_{31}. \quad (7)$$

C^1 continuity is (re-)enforced between $\tilde{\mathbf{p}}$ and $\underline{\mathbf{p}}$ (and between $\tilde{\mathbf{p}}$, $\underline{\mathbf{p}}$; see Fig. 5): the BB-coefficients underlaid in grey are obtained by C^1 -extension of $\tilde{\mathbf{p}}$ (respectively of $\underline{\mathbf{p}}$). This leaves unconstrained only $\tilde{\mathbf{p}}_{21}$ and $\tilde{\mathbf{p}}_{31}$ (circles in Fig. 5b) to be fixed by minimizing

$$\min \sum_{i=2}^3 (\|\tilde{\mathbf{p}}_{i1} - \tilde{\mathbf{a}}_{i1}\|^2 + \|\underline{\mathbf{p}}_{i1} - \underline{\mathbf{a}}_{i1}\|^2).$$

The cap is \mathbf{p} is the union of the $n \times 2$ patches \mathbf{p}^k surrounding the irregular point. Applying the steps of the construction summarized in Fig. 6 to the input CC-nets of Fig. 2, top, yields the basic functions Fig. 2, bottom.

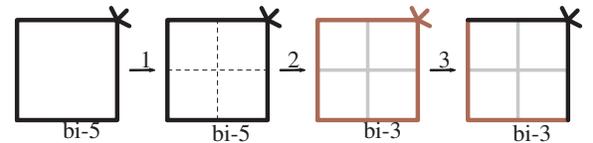


Figure 6: Construction steps of the cap. (1) splitting a *bi-5* patch of [KP14] (see Fig. 4a) into four by DeCasteljau's algorithm. (2) transforming the *bi-5* pieces to degree *bi-3* (see Fig. 4b) and (3) adjusting the internal cap transitions to be G^1 while leaving the cap boundaries C^0 . Black edges: G^1 continuity, gray edges: C^1 continuity, brown edges: C^0 continuity.

3. Properties of the *bi-3* construction

By construction, the 2×2 -patches \mathbf{p}^k are internally C^1 and they are G^1 -connected to their neighbor patches \mathbf{p}^{k+1} that surround the irregular point. Equation (2) implies that

the second-order data at the irregular point (see the BB-coefficients marked by red disks) stem from the guide. Therefore our bi-3 cap \mathbf{p} has a well-defined curvature at the irregular point. By initialization, the cap \mathbf{p} is only C^0 -connected to the tensor-border \mathbf{b} . The guide's C^1 extension across the tensor-border into the interior is not fully captured. Yet, especially after applying one Catmull-Clark refinement step, we consistently observe that the mismatch of normals is below $< 0.1^\circ$, also for challenging models. For example, the input nets of the basic functions h_ℓ Fig. 2 represent hard test data as spiky perturbations of the plane. For all h_ℓ , up to valence 10, we certified that the normal mismatch is indeed less than 0.1° . Not only is such a mismatch visually hard to detect (see e.g. Fig. 2 for an illustration when $n = 6$), but it is acceptable in class A design [Aut15, And15].

Since the mismatch depends non-linearly and non-trivially on the input mesh, characterizing input meshes that guarantee the bound on the normal mismatch does not seem practical. Instead, the Appendix gives the explicit formulas for all BB-coefficients that together with the boundary curve define the normal of the cap. The mismatch can then quickly be evaluated for a particular application and a local design change or one refinement step can be applied in the rare case where the mismatch exceeds the bound. Similarly, rigorously proving that the construction minimizes oscillations no matter what the input, is not realistic. We therefore dedicate the next section to showing, for a representative subset of our tests, that the new 2×2 construction improves on the existing choices and yields class A highlight lines for a battery of challenging input scenarios.

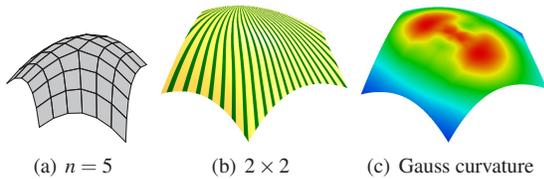


Figure 7: Low valence convex cap. The surface looks curvature continuous (despite a normal mismatch of 0.006°).

4. Discussion

We compare to constructions from the literature and then discuss alternative new constructions that we did not select, since they work less well than the presented 2×2 construction. We conclude with a guided 7-piece construction in Section 4.3.

4.1. Comparisons

Fig. 7 shows that, although formally the new bi-3 surfaces are not smooth, their curvature distribution is often (not only in this example) on par with more complicated higher degree G^2 constructions. The stabilizing effect of the guide is also

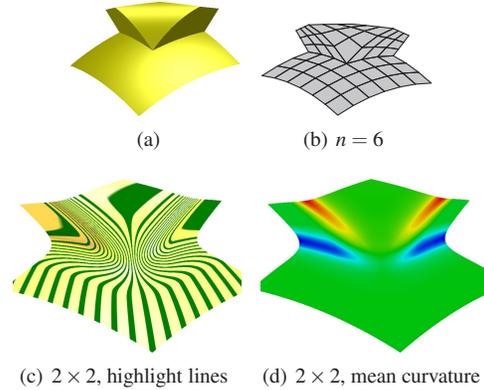


Figure 8: Input design sketch consisting of four quadric pieces (a) converted to a quad mesh (b).

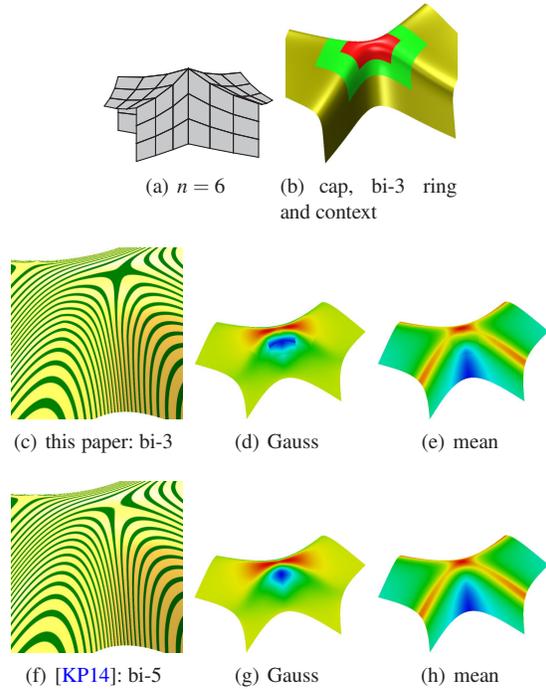


Figure 9: Comparison to the guide surface. The subtle change of the highlight shading reveals a deviation from the guide, but only Gauss curvature shading brings out the difference in smoothness.

visible in Fig. 8. Here a design is sketched out as a juxtaposition of simple quadratic surfaces, approximated by a net. Fig. 9 compares a bi-3 surface to its bi-5 guide [KP14]. The complex input is chosen to reveal the otherwise very small differences. Even so, the mean curvatures are almost identical and only the Gauss curvature shading of the new bi-3

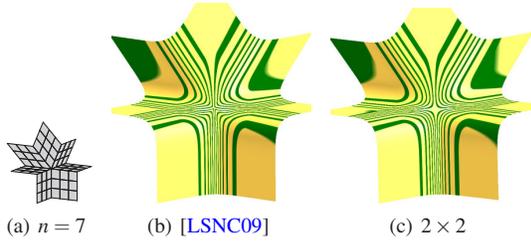


Figure 10: Comparison to a Gregory-patch construction [Gre74, LSNC09]. [LSNC09] suffers from bunched up high-light lines.

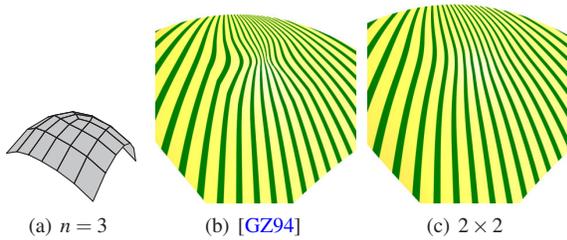


Figure 11: Comparison to [GZ94] when $n = 3$.

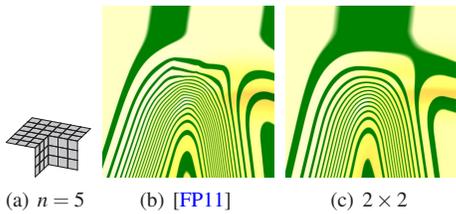


Figure 12: Comparison with a G^1 scheme using a quadratic $b(u)$, $b(u) := 2c(1 - u)^2$, to C^1 -extend the tensor-border. [FP11] fails the oscillation test.

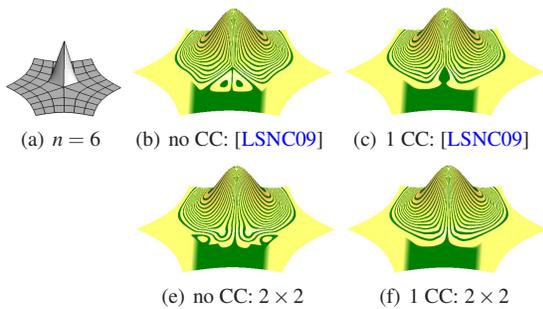


Figure 13: For $n > 4$, one Catmull-Clark-step often improves the shape. (b,e) no Catmull-Clark-step is applied, (c,f) one Catmull-Clark-step is locally applied.

construction reveals an abrupt change across the cap boundary.

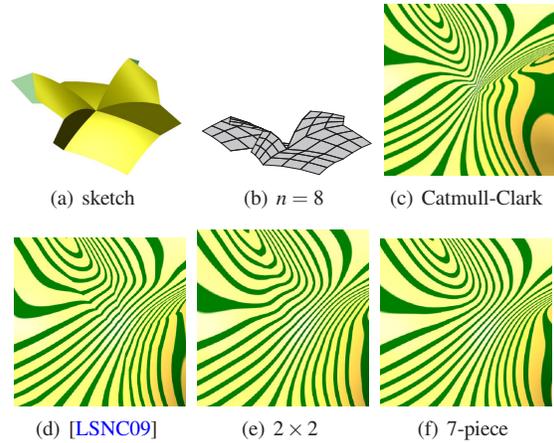


Figure 14: Comparison to Catmull-Clark subdivision and [LSNC09]. Input design sketch consisting of eight quadric surfaces (a) converted into a quad mesh (b). The 7-piece construction (f) is explained in Section 4.3.

Fig. 10 shows that the well-defined curvature of our bi-3 surfaces at the irregular point yields a better highlight line distribution than the Gregory patch [LSNC09]. Fig. 11 and Fig. 12 illustrate an observation that applies to many exact G^1 constructions where the derivatives of the tensor-border are interpolated exactly without reparameterization. For a 2×2 layout, this approach, unlike our bi-3 construction, forces $b(u)$ in the G^1 constraints $\partial_u \vec{p} + \partial_v \vec{p} = b(u) \partial_u \vec{p}$ to be quadratic. A representative example of quadratic $b(u)$ constructions is the classical construction [GZ94], where $b(u) := 2c(1 - u)^2$. As Fig. 11 illustrates, this approach clearly fails to yield class A surfaces. (Our construction shows a slight oscillation of highlight lines motivating the work in Section 4.3). Fig. 12 compares our 2×2 cap to a more sophisticated construction [FP11] where a 3×3 split allows adjacent irregular points and only the middle segments use a quadratic $b(u)$.

Typically (see Fig. 13), but not when $n = 3$, applying one step of Catmull-Clark-refinement improves the distribution of highlight lines. Fig. 14c makes the point that more than one step reduces the quality of subsequent surface constructions: in many configurations the resolution of different highlight lines coming together is late and they bunch up at the irregular point. Although our construction, Fig. 14e, shows more stable highlight lines than [LSNC09] in Fig. 14d, class A for this challenging scenario is only achieved with the 7-piece construction detailed in Section 4.3.

4.2. Alternative constructions

We explored various other patch splits and constructions, using the 7-piece, 9-piece and 10-piece layouts, akin to ap-

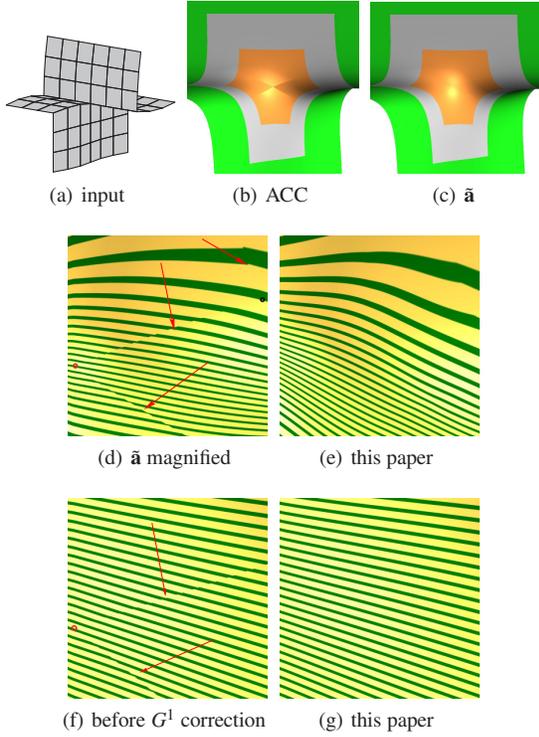


Figure 15: Smoothness and construction steps. One Catmull-Clark-step applied to the input net (a) of Fig. 1 yields the grey surface ring in (b,c). The tensor-border interpolating ACC patches of [LS08] are clearly not smooth. The single patch conversion $\tilde{\mathbf{a}}$ of [KP14], matching the guide's the position and first (mixed) derivatives, appears smooth from afar but fails to be class A under scrutiny. (d) Magnification of (c) with highlight lines: the red circle marks the irregular point, the black circle the midedge of the tensor-border. Red arrows point to lack of smoothness both within the cap and along the tensor-border. (e) The same area of our cap. The transition appears smooth although it is formally only C^0 . (f,g) further magnification focussing on 1/6th of the patch (red circle = irregular point). (f) after step (2) of the pipeline Fig. 6 the patches are not yet G^1 connected.

plying additional Catmull-Clark-subdivision steps. We also tried alternative C^0 caps. Matching the derivatives of the tensor-border exactly, without reparameterization, failed to yield class A constructions even when stabilized with a guide.

Interpolating the full tensor-border and requiring only C^0 continuity at the irregular point yields a cap as in [LS08]. Fig. 15b shows that such a C^0 cap does not yield class A surfaces. We explored applying multiple Catmull-Clark-steps to confine the C^0 spot to a minuscule area. However, not only does ACC fail to meet the bound even after five refinement

steps, but Fig. 1b illustrates that the highlight lines start to oscillate in the larger surrounding area.

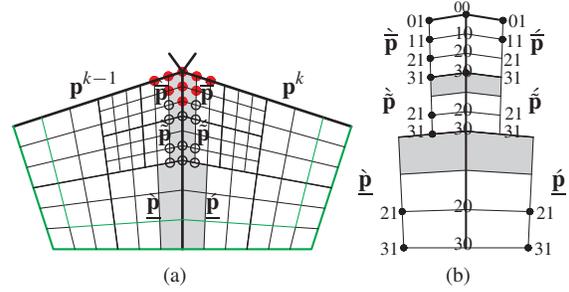


Figure 16: 7-piece layout, akin to applying two Catmull-Clark-subdivision steps.

4.3. Capping with 7-piece patches

Considering the 7-piece split Fig. 16a, an auxiliary cap \mathbf{a} is constructed as in the main 2×2 construction, except that the patch closest to the irregular point is additionally split into four pieces, each transformed into bi-3 sub-patches. Due to the internal T-corners, the 7-piece patches \mathbf{a}^k are not C^1 internally. We enforce the G^1 constraints (see Fig. 16b) between the 7-piece patches as

$$\partial_v \tilde{\mathbf{p}} + \partial_v \hat{\mathbf{p}} = c(2(1-u) + \frac{3}{2}u)\partial_u \hat{\mathbf{p}}, \quad c := \cos \frac{2\pi}{n},$$

$$\partial_v \tilde{\mathbf{p}} + \partial_v \hat{\mathbf{p}} = c(\frac{3}{2}(1-u) + u)\partial_u \hat{\mathbf{p}},$$

$$\partial_v \tilde{\mathbf{p}} + \partial_v \hat{\mathbf{p}} = c(1-u)\partial_u \hat{\mathbf{p}},$$

by setting (the unmarked points in Fig. 16b)

$$\tilde{\mathbf{p}}_{20} := \frac{3c\tilde{\mathbf{p}}_{00} + (5c-12)\tilde{\mathbf{p}}_{10} + 6(\tilde{\mathbf{p}}_{11} + \hat{\mathbf{p}}_{11})}{8c};$$

$$\hat{\mathbf{p}}_{21} := -c\tilde{\mathbf{p}}_{10} + (2 + \frac{c}{3})\tilde{\mathbf{p}}_{20} + \frac{2c}{3}\tilde{\mathbf{p}}_{30} - \tilde{\mathbf{p}}_{21};$$

$$\hat{\mathbf{p}}_{31} := -\frac{3c}{2}\tilde{\mathbf{p}}_{20} + (2 + \frac{3c}{2})\tilde{\mathbf{p}}_{30} - \tilde{\mathbf{p}}_{31};$$

$$\tilde{\mathbf{p}}_{20} := \frac{c\tilde{\mathbf{p}}_{00} + 2(c-3)\tilde{\mathbf{p}}_{10} + 3(\tilde{\mathbf{p}}_{11} + \hat{\mathbf{p}}_{11})}{3c};$$

$$\tilde{\mathbf{p}}_{30} := \frac{4}{9}\tilde{\mathbf{p}}_{20} - \frac{8}{9}\tilde{\mathbf{p}}_{30} + \frac{4}{3}\tilde{\mathbf{p}}_{20} + \frac{1}{9}\tilde{\mathbf{p}}_{20};$$

$$\hat{\mathbf{p}}_{21} := -\frac{2c}{3}\tilde{\mathbf{p}}_{10} + (2 + \frac{c}{6})\tilde{\mathbf{p}}_{20} + \frac{c}{2}\tilde{\mathbf{p}}_{30} - \tilde{\mathbf{p}}_{21};$$

$$\hat{\mathbf{p}}_{31} := -c\tilde{\mathbf{p}}_{20} + (2+c)\tilde{\mathbf{p}}_{30} - \tilde{\mathbf{p}}_{31}.$$

The BB-coefficients $\tilde{\mathbf{p}}_{10}, \tilde{\mathbf{p}}_{20}, \tilde{\mathbf{p}}_{30}$ are defined by (1), (6), (7). (Re-)enforcing internal C^1 continuity (also along the sector boundaries identified by the two grey strips in Fig. 16b), we find unconstrained only the 5 BB-coefficients $\tilde{\mathbf{p}}_{21}, \tilde{\mathbf{p}}_{31}, \tilde{\mathbf{p}}_{30}, \hat{\mathbf{p}}_{21}, \hat{\mathbf{p}}_{31}$ (circles in Fig. 16a). These points are fixed by minimizing the sum of squared distances between the coefficients of \mathbf{p} and \mathbf{a} along the grey strip in Fig. 16a. By construction,

the union of the 7-piece bi-3 patches join G^1 and are C^0 -connected to the input data with a normal mismatch the same as in the 2×2 construction.

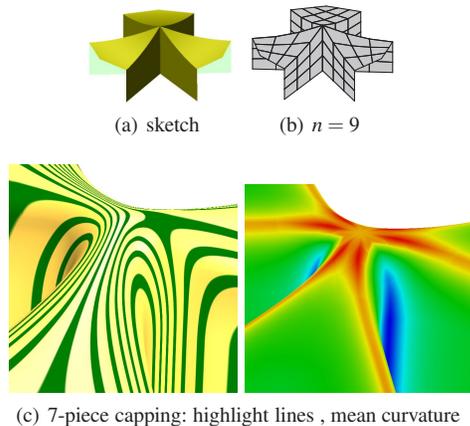


Figure 17: Input design sketch (a) consisting of nine primitive surfaces is (b) approximated by an $n = 9$ CC-net.

Fig. 17 illustrates the amply-tested observation that the 7-piece construction works well also for high valencies and exotic configurations.

5. Conclusions

A new, purely bi-cubic construction generalizing bi-cubic tensor-product B-splines to irregular quad layouts has empirically been shown to satisfy the highlight-line criterion of class A surfacing. Conversely, numerous examples demonstrated that existing bi-3 constructions fail the oscillation test. The key is that, near the irregularities, the new bi-3 surface closely follows a higher-degree, high-quality smooth surface, called guide. The novelty of the construction is a judicious trade-off: sacrificing formal smoothness but qualitatively retaining the good highlight-line distribution of the guide.

G^1 smoothness is exactly enforced between the n sectors of the surface caps since any flaw along those transition lines is conspicuous. Formal smoothness is lost where the regular tensor-product arrangement of patches switches to the neighborhood of an irregular point. We are not able to formally prove conditions so that the normal mismatch along those seams stays always below the industry-accepted bound of 0.1° . However, we verified that the bound on the mismatch holds for all basic functions h_ℓ , up to valence $n = 10$. We note that the control nets of the h_ℓ are difficult input quad-nets and the final surface is a linear combination of B-splines and the basic functions h_ℓ . The bound on the normal mismatch is also verified on many other challenging examples, such as Fig. 18, that are not designed as part of an obstacle course but are based on designed or re-meshed (valence) semi-regular quad meshes [BLP*13].

Another observation is that while Catmull-Clark-subdivision reduces the normal mismatch, subdivision imposes non-class A highlight lines in the process. This suggests that one refinement step is beneficial, but multiple steps are not.

In conclusion, it is now possible to have an everywhere bi-cubic representation of curved surfaces from an irregular network of quadrilateral-faced polyhedra that simplifies the downstream processing of the geometry of class A surfaces.

Acknowledgment The work was supported in part by NSF Grant CCF-1117695. Martin Sarov created a second implementation to verify the correctness of the functions h_ℓ and render the models Wing (authored by James Pence), RockerArm quad remesh (Yu-Kun Lai) and VW hood (grabCad). Tom Grandine pointed out a class of machine tool controllers restricted to bi-3 splines.

References

- [And15] ANDJELIC Z.: What is the definition of a class A surface?, accessed May 2015. <https://grabcad.com/questions/what-is-definition-of-class-a-surface>. 1, 2, 5
- [Aut15] AUTODESK: Tutorial at <http://help.autodesk.com/view/alias/2015/enu/?guid=guid-2fce06eb-8ef7-4507-92f7-82a73a0df378>, 2015. accessed Jan 16. 2, 5
- [BC94] BEIER K.-P., CHEN Y.: Highlight-line algorithm for real-time surface-quality assessment. *Computer-Aided Design* 26, 4 (1994), 268–277. 1
- [BLP*13] BOMMES D., LÉVY B., PIETRONI N., PUPPO E., SILVA C., TARINI M., ZORIN D.: Quad-mesh generation and processing: A survey. *Comput. Graph. Forum* 32, 6 (2013), 51–76. 8
- [Cas12] CASHMAN T. J.: Beyond Catmull-Clark? A survey of advances in subdivision surface methods. *Comput. Graph. Forum* 31, 1 (2012), 42–61. 2
- [CC78] CATMULL E., CLARK J.: Recursively generated B-spline surfaces on arbitrary topological meshes. *Computer-Aided Design* 10 (Sept. 1978), 350–355. 1
- [Ead15] EADIE B.: Solidworks tutorial: Surface fill tips and tricks in solidworks, minute 1:40 accessed May 2015. https://www.youtube.com/watch?v=jgkIJ_ofcqw. 2
- [FP11] FAN J., PETERS J.: Smooth bi-3 spline surfaces with fewest knots. *Computer Aided Design* 43, 2 (Feb 2011), 180–187. *JCAD* 1686. 6
- [Gre74] GREGORY J. A.: *Smooth interpolation without twist constraints*. Academic Press, 1974, pp. 71–88. 1, 2, 6
- [GZ94] GREGORY J. A., ZHOU J.: Filling polygonal holes with bicubic patches. *Computer Aided Geometric Design* 11, 4 (1994), 391–410. 1, 6
- [KP14] KARČIAUSKAS K., PETERS J.: Improved shape for multi-surface blends. *GMOD* (May 2014). Dagstuhl workshop on Geometric Design. 3, 4, 5, 7, 9
- [LM08] LIU Y., MANN S.: Parametric triangular Bézier surface interpolation with approximate continuity. In *Proceedings of the 2008 ACM Symposium on Solid and Physical Modeling, Stony Brook, New York, USA, June 2-4, 2008* (2008), Haines E., McGuire M., (Eds.), ACM, pp. 381–387. 2

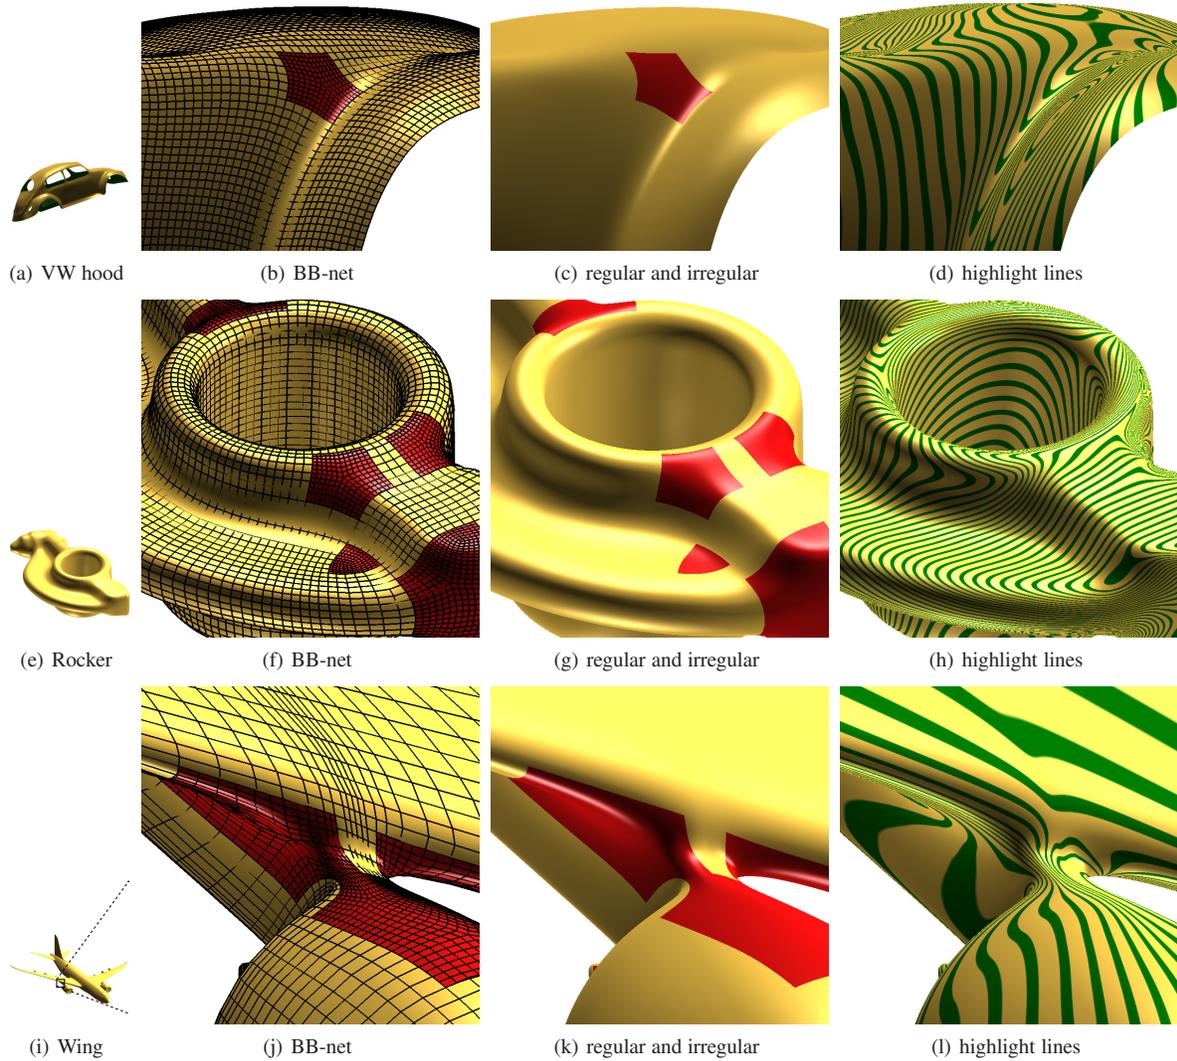


Figure 18: Quad-models with isolated irregular points. Left images show the BB-control net (not the coarser CC-net) superimposed on the surface. Irregular neighborhoods (caps) are rendered red. (The flawed highlight lines on top of the airplane wing (l) are caused by the regular part of the input model and are not part of our cap construction.)

[LS08] LOOP C., SCHAEFER S.: Approximating Catmull–Clark subdivision surfaces with bicubic patches. *ACM Transactions on Graphics* 27, 1 (Mar. 2008), 8:1–8:11. 7

[LSNC09] LOOP C. T., SCHAEFER S., NI T., CASTAÑO I.: Approximating subdivision surfaces with Gregory patches for hardware tessellation. *ACM Trans. Graph* 28, 5 (2009). 1, 2, 6

[PoV15] POVRAY: Bicubic patch object, 2015. <http://povray.org/documentation/view/3.6.1/64/>, accessed Apr 04. 1

[wik15] WIKIPEDIA: Class_a_surfaces, accessed May 2015. http://en.wikipedia.org/wiki/Class_A_surfaces. 1

Appendix: Border Formulas

We denote the BB-coefficients of the input tensor-border \mathbf{b} by \mathbf{b}_{ij} . The index 30 corresponds to the junction between outer boundary curve segments. Recall that the first two boundary layers of Bézier coefficients of the bi-5 guide [KP14] are split and transformed to C^1 -connected pieces $\tilde{\mathbf{b}}^l$ and $\tilde{\mathbf{b}}^r$ of degree bi-3 – that determine the cap normal and hence the normal mismatch along the boundary. Since the cap’s boundary curve is of degree 3 and exactly matched,

only the coefficients the second layer need explicit formulas:

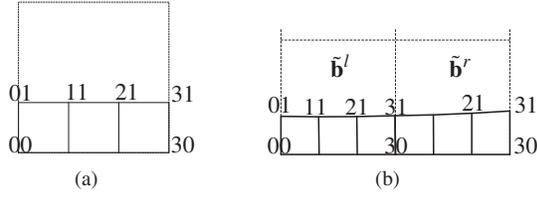


Figure 19: (a) Labels of (a sector of) the tensor-border \mathbf{b} of depth 1 (label 30 corresponds to the junction between sectors); (b) labels for the tensor-border of the 2×2 -patch.

$$\begin{aligned}\tilde{\mathbf{b}}_{11}^l &:= \frac{7+2e(c-2)}{12}\mathbf{b}_{00} + \frac{1-e(c-2)}{12}(\mathbf{b}_{10} + \mathbf{b}_{01}) + \frac{1}{4}\mathbf{b}_{11}, \\ \tilde{\mathbf{b}}_{21}^r &:= \frac{2-e}{4}\mathbf{b}_{20} + \frac{6+e(c-3)}{12}\mathbf{b}_{30} + \frac{e}{4}\mathbf{b}_{21} - \frac{e(c-3)}{12}\mathbf{b}_{31}, \\ \tilde{\mathbf{b}}_{31}^r &:= \frac{2-e}{2}\mathbf{b}_{30} + \frac{e}{2}\mathbf{b}_{31}, \quad \tilde{\mathbf{b}}_{k1}^l := \sum_{j=0}^1 \sum_{i=0}^3 d_{ij}^k \mathbf{b}_{ij},\end{aligned}$$

where $e := 1 + 0.413c + 0.116c^2$,

$$\begin{aligned}d_{i0}^2 &:= \left[\frac{13+e(3c-7)}{48}, \frac{41+e(7c-17)}{96}, \frac{1}{8}, \frac{e(3-c)-3}{96} \right], \\ d_{i1}^2 &:= \left[\frac{4+e(8-3c)}{96}, \frac{3+e(5-2c)}{32}, \frac{2+e(2-c)}{32}, \frac{1-e}{96} \right], \\ d_{i0}^3 &:= \left[\frac{9+e(2c-5)}{64}, \frac{23+e(4c-11)}{64}, \frac{19+e(2c-7)}{64}, \frac{5-e}{64} \right], \\ d_{i1}^3 &:= \left[\frac{1+e(3-c)}{64}, \frac{3+e(9-3c)}{64}, d_{11}^3, d_{01}^3 \right].\end{aligned}$$