Computers & Graphics 33 (2009) 204–210

Contents lists available at ScienceDirect

Computers & Graphics

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

Technical Section

Assembling curvature continuous surfaces from triangular patches ☆

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

[a] Vilnius University, Lithuania

[b] Department of CISE, CSE Bldg, University of Florida, Gainesville, FL 32611-6120, USA

ARTICLE INFO

Article history:
Received 8 December 2008
Received in revised form
27 February 2009
Accepted 3 March 2009

Keywords:
C^2 surface construction
Shape
Guide surface
Triangular patches

ABSTRACT

We assemble triangular patches of total degree at most eight to form a curvature continuous surface. The construction illustrates how separation of *local shape* from *representation and formal continuity* yields an effective construction paradigm in partly underconstrained scenarios. The approach localizes the technical challenges and applies the spline approach, i.e. keeping the degree fixed but increasing the number of pieces, to deal with increased complexity when many patches join at a central point.

1. Introduction

Complex surface blends, for example when capping a C^2 spline surface by n patches, require an increase in either the degree or the number of pieces compared to the surrounding regular spline surface. Typically, the new degrees of freedom do not match the formal continuity constraints and this results in an under-constrained problem. One way to set the extra degrees of freedom is to minimize a geometrically motivated functional, say approximating an integral of the mean, Gauss or total curvature (see e.g. [12,20,4]; these functionals have also been applied to faceted representations—but here we are only concerned with curvature continuous surfaces). Another is to minimize deviation from a space that has too few degrees of freedom, for example by minimizing a quadratic expression of (a convex combination of) higher-order derivatives and hence penalizing higher degree (cf. Definition 1 and Fig. 13). We use a third approach to setting extra degrees of freedom: we first create a surface fragment that captures the shape and then approximate this surface fragment to satisfy the smoothness constraints. This two-stage approach of separating shape from formal smoothness constraints was introduced by Karčiauskas and Peters in [7]; and was already hinted at by the composition with quadratic shapes in [15,19]. We call it *guided surfacing* in the following. Guided surfacing according to [7] has the nice side-effect of localizing, otherwise global, smoothness constraints. This very much simplifies derivation and analysis and avoids the need to invert large matrices during construction.

There is a rich literature on construction of C^2 surfaces based on quadrilateral meshes and patches, e.g. [6,2,15,21,19,3,10,11,13,5,7,8]. Here, we consider *three-sided patches* as in [17], corresponding to control nets with triangular facets. Since box-splines do a good job when those meshes are regular, i.e. all vertices have the same valence 6, the challenge is to complete a C^2 box-spline surface by filling its isolated multi-sided holes (see Fig. 1, right) using n (macro-)patches for an n-sided hole. Our surface construction below is interesting in its own right since the degree of the guided macro-patches output is 8, the same as the lowest total-degree surface constructions in the literature [14,17,18] (a bound also for finite element spaces [9, Ex 5.1]), but has better shape control. Our main aim, however, is to illustrate that separating shape from smoothness constraints allows setting unconstrained parameters in a natural way. We show how a high number n of features such as in high-order saddles, can be made to blend with slowly dissipating curvature differences, not by increasing the degree of the surface, but the number of polynomial pieces in each patch—and by using the underlying guide surface to set unconstrained parameters (such as the $\mathbf{a}_{ij}$ in the local construction on Section 3) in a geometrically intuitive fashion, by sampling.

Concretely, we consider the following setup. We are given a triangulation with *isolated extraordinary vertices*, i.e. vertices of valence $n \neq 6$ such that each direct neighbor is *regular* of valence 6 (Fig. 1, middle). That is, every triangle has at least two regular vertices. Wherever a triangle has three regular vertices, we interpret the vertices of the triangle and their neighbor vertices

☆ This work was supported in part by NSF Grant CCF-0728797.

* Corresponding author. Tel.: +1 352 392 1226; fax: +1 352 392 1220.
E-mail address: jorg@cise.ufl.edu (J. Peters).
URLs: http://www.cise.ufl.edu/~jorg, http://www.cise.ufl.edu/research/SurfLab (J. Peters).

0097-8493/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cag.2009.03.015

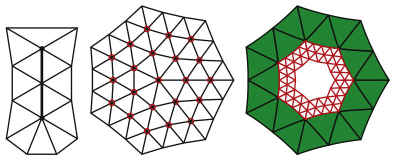

Fig. 1. Input triangulation and boundary data. (left) No two extraordinary vertices are direct neighbors. (middle) Triangulation with a single, isolated central non-6-valent vertex, called extraordinary vertex. We interpret the vertices as box-spline coefficients that define the (green) piecewise degree 4 surface ring right. (This ring is not used in the construction, but will allow us to check the quality of the multi-sided blend to existing data.) The circled vertices (middle) define the *boundary data* **b**: position, first and second derivative, for extending the surface ring. (right) These boundary data are shown as a depth three (red) net of a Bernstein–Bézier (BB) coefficients.

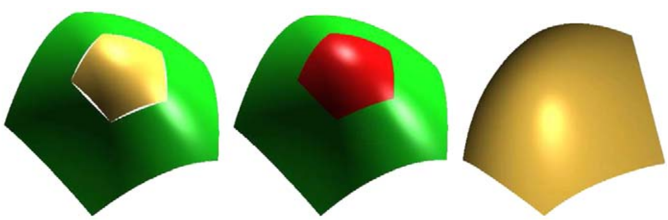

Fig. 2. Guided surfacing. (left) A C^2 guide **g** (yellow cap) is determined with reference to the boundary data **b** defined by a surrounding (green) surface. The guide surface and the surrounding surface are in general not even connected (see Section 2 and the Appendix for the derivation of a guide with the control net structure of Fig. 3, bottom, left). (middle and right) The final piecewise polynomial C^2 surface without gap: the algorithm to be specified constructs the red replacement of the guide that matches the boundary data C^2 after reparameterization.

as (box-spline) control points of a three-sided polynomial patch of total degree 4, defined by the three-direction C^2 box-spline [1] with directions

$$\Xi := \begin{bmatrix} 1 & -1 & 0 & 0 & 1 & -1 \\ 0 & 0 & 1 & -1 & 1 & -1 \end{bmatrix}.$$

Each extraordinary vertex of valence n in the triangular mesh causes an n-sided hole in the regular surface complex. Assuming that such vertices are separated (Fig. 1, left), we want to fill each hole with a *cap* consisting of n macro-patches of degree 8 so that the resulting surface is

- curvature continuous,
- the geometry of the cap incorporates that of the surrounding surface, and
- the cap does not fluctuate in position, normal or curvature.

Construction overview. Our approach is as follows. First, we construct a piecewise C^2 guide surface piece $\mathbf{g} : \mathbb{R}^2 \to \mathbb{R}^3$ (Fig. 2, left). This surface represents the design intent in the sense that the final surface will follow its shape. Even when the guide takes into account the boundary data **b**, it is typically not suitable for a final cap **x** since, as illustrated in Fig. 2, left,[1] it need not even join continuously with the surrounding surface and can have a completely different representation from the one needed for further processing. We define a C^2 reparameterization $\boldsymbol{\rho} : \mathbb{R}^2 \to \mathbb{R}^2$ from the boundary of the domain of **x**, where always six patches meet at a vertex, to the center, where $n \neq 6$ patches meet; and an operator h that takes as input a sufficiently smooth function and constructs surface pieces of degree 8 from the derivatives of its input. Then the cap completing a C^2 surface is defined by $\mathbf{x} := h(\mathbf{g} \circ \boldsymbol{\rho})$.

Paper overview. In Section 2, we construct the C^2 map $\boldsymbol{\rho}$ as well as a prototype guide **g**, whose instantiation, in dependence of the boundary data **b** is described in the Appendix. In Section 3, we construct h piecemeal, via a simple local operator $h^8_\diamond$ acting only at vertices of $\mathbf{g} \circ \boldsymbol{\rho}$. The surface construction then becomes a straightforward localized enforcement of C^2 constraints. In Section 4, we apply the tools of Section 3 to obtain a surface cap that completes a C^2 surface (as in Fig. 2) and has good shape for $n < 9$. In Section 5, we address the cases $n \geqslant 9$ by modeling each of the n segments of the cap by four polynomial pieces. In Section 6, we discuss modifications to obtain good shape both for very high and very low valences; and we illustrate how derivative-based functionals applied in $\mathbb{R}^3$ fail to achieve the same effect as guide surfaces.

2. C^2 functions on a polygonal domain

In this section, we prepare the technical background to be able to focus later only on the high-level construction. We define an n-gon $\Omega \subsetneq \mathbb{R}^2$, a C^2 map $\boldsymbol{\rho}$ that maps n copies of the right angle unit triangle $\varDelta$ to Ω and the guide **g** that maps Ω to $\mathbb{R}^3$. Also the final hole-filling spline cap **x** maps $\varDelta \times \{1, \ldots, n\} \to \mathbb{R}^3$; but it will only be discussed in later sections. Given the purpose of each map, its smoothness and symmetry, the derivations use standard machinery in geometric design. But this does not mean that the derivation is trivial. It reflects important choices of properties and polynomial degree. Until Section 6, we assume

$$n \notin \{3, 4, 6\} \tag{1}$$

since $n = 6$ corresponds to the box-spline construction and $n = 3$ and 4 admit special (simpler) treatment, explained in Section 6.

2.1. The n-sided domain Ω

The n-gon Ω is composed of n triangles

$$\triangle \mathbf{o}\mathbf{v}_i\mathbf{v}_{i+1} := \{\mathbf{o} + u(\mathbf{v}_i - \mathbf{o}) + v(\mathbf{v}_{i+1} - \mathbf{o}), 0 \leqslant u, 0 \leqslant v, u + v \leqslant 1\},$$

where $i = 0, \ldots, n-1$ and, as illustrated in Fig. 3, top left,

$$\Omega := \bigcup \triangle \mathbf{o}\mathbf{v}_i\mathbf{v}_{i+1}, \quad \mathbf{o} := \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad \mathbf{v}_i := \begin{bmatrix} \cos(i\alpha) \\ \sin(i\alpha) \end{bmatrix}, \quad \alpha := 2\pi n. \tag{2}$$

The image of $\triangle \mathbf{o}\mathbf{v}_i\mathbf{v}_{i+1}$ will be the ith *segment* of the map to be constructed. We also define the unit triangle $\varDelta \in \mathbb{R}^2$ and $\angle \mathbf{v}_0\mathbf{v}_1$, a sector of the plane bounded by the rays $\mathbf{o}\mathbf{v}_0$ and $\mathbf{o}\mathbf{v}_1$. With the Bernstein–Bézier (BB) coefficients $\mathbf{p}^i_{jk}$ (see e.g. [16, Chapter 10]) indexed as in Fig. 3, top right, the well-known conditions for two patches $\mathbf{p}^i$ and $\mathbf{p}^{i+1}$ abutting along the boundary defined by $\mathbf{p}^{i+1}_{j0} = \mathbf{p}^i_{0j}$ to join (parametrically) C^1, respectively, C^2 are

$$\mathbf{p}^{i+1}_{j1} = v_0\mathbf{p}^i_{0j} + v_1\mathbf{p}^i_{0,j+1} + v_2\mathbf{p}^i_{1j}, \tag{3}$$

$$\begin{aligned}\mathbf{p}^{i+1}_{j2} = {} & v_0^2\mathbf{p}^i_{0j} + 2v_0v_1\mathbf{p}^i_{0,j+1} + v_1^2\mathbf{p}^i_{0,j+2} \\ & + 2v_0v_2\mathbf{p}^i_{1j} + 2v_1v_2\mathbf{p}^i_{1,j+1} + v_2^2\mathbf{p}^i_{2j},\end{aligned} \tag{4}$$

with scalar weights determined by $\mathbf{v}_{i+1} = v_0\mathbf{o} + v_1\mathbf{v}_i + v_2\mathbf{v}_{i-1}$, as

$$v_0 := 2\bar{c}, \quad v_1 := 2c, \quad v_2 := -1, \quad c := \cos\frac{2\pi}{n}, \quad \bar{c} := 1 - c. \tag{5}$$

[1] This paper is concerned with shape and curvature continuity up close. Therefore it makes little sense to show any large triangulations where shape would be dominated by regular box-splines. The figures in the submission can be enlarged using the pdf zoom capability.