Computers & Graphics 46 (2015) 172–185

Contents lists available at ScienceDirect

Computers & Graphics

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

SMI 2014

Smooth multi-sided blending of biquadratic splines

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

[a] Vilnius University, Lithuania
[b] University of Florida, United States

ARTICLE INFO

Article history:
Received 5 July 2014
Received in revised form
29 August 2014
Accepted 17 September 2014
Available online 14 October 2014

Keywords:
Biquadratic splines
Shape
Multi-sided blends
Reparameterization

ABSTRACT

Biquadratic (bi-2) splines are the simplest choice for converting a regular quad meshes into smooth tensor-product spline surfaces. Existing methods for blending three, five or more such bi-2 spline surfaces using surface caps consisting of pieces of low polynomial degree suffer from artifacts ranging from flatness to oscillations. The new construction, based on reparameterization of the bi-2 spline data, yields well-distributed highlight lines for a range of challenging test data. The construction uses n pieces of degree bi-4 (bi-3 when $n \in \{3,5\}$) and applies both to primal (Catmull–Clark-like) and dual (Doo–Sabin-like) input layouts.

1. Introduction

Many mechanical parts and inner surfaces require only first order smoothness, i.e. continuity of the normal. The simplest tensor-product patch type to match the requirements is the biquadratic (bi-2) spline patch. Keeping the degree lower than bi-3 for the bulk of the surface simplifies downstream use, say for isogeometric computations, and reduces the chance of introducing unwanted visible ripples and oscillations.[1]

A classical challenge, illustrated in Fig. 1, is to smoothly blend, by a surface 'cap', three, five or more bi-2 spline patches. Numerous publications in the late 1980s and early 1990s (see the literature survey in Section 2 below) addressed the algebraic constraints for smoothness, but did not focus on shape. To illustrate this claim, we analyze the canonical representatives ofthe two main alternatives for such blends: subdivision algorithms and Bézier patch constructions. The shape of surfaces generated by the subdivision naturally associated with bi-2 surfaces, the Doo–Sabin algorithm [2], suffers from oscillations near and flatness at the center of the cap (Fig. 2b). Similarly, older (1990s-style) multi-patch Bézier caps of low degree, such as the bi-3 construction of [3] in Fig. 2e, result in poorly-spaced and poorly shaped highlight lines. While our surfaces will not be curvature continuous throughout, as are recent more costly constructions such as [4,5], they are of low polynomial degree, show remarkably good highlight line distribution for a gallery of challenging test data, and use fewer patches than other low-degree CAD-compatible constructions (cf. [6,7]).

Contribution: The focus and contribution of this paper is a simple, yet subtle recipe for constructing *low-degree multi-sided surface caps* that complete a bi-2 C^1 spline complex and do not suffer from the highlight line artifacts of existing methods. The construction applies both to primal (Catmull–Clark-type) and dual (Doo–Sabin-type) input layouts, i.e. to quad meshes including nodes that have $n \neq 4$ neighbors and to faces with $n \neq 4$ edges.

The construction is a recipe in that it expresses all BB- (Bernstein–Bézier) coefficients of the cap directly in terms of the input mesh. No equations have to be solved at runtime: the final surface is a linear combination of precomputed generating functions.

The construction is subtle and different from older constructions in that it (i) leverages carefully chosen *re-parameterizations* β (a bi-variate change of variables) of the Hermite data $\mathbf{b}$ (position and derivative) provided to the cap by the bi-2 splines surrounding the cap and (ii) enforces curvature continuity at the central point. As illustrated in Fig. 3 and explained in more detail in Section 4.1, reparameterizing $\mathbf{b}$ yields better results, visible even without highlight lines. Compare Fig. 3e to f.

Overview: Section 2 reviews the literature and motivates the features of the new algorithm. Section 3 formally introduces the problem to be solved: smoothly filling a multi-sided hole in a bi-2 spline complex. Section 4 introduces the n-sided cap. By default, the cap consists of n patches of degree bi-4, one per sector; when $n \in \{3,5\}$, well-shaped caps of degree bi-3 per sector can be substituted (Section 4.2). (Appendix C demonstrates that if we split each sector, into 2×2 pieces, degree bi-3 work also for $n > 5$, albeit at the cost of slightly lower quality than the main bi-4

* Corresponding author.
E-mail address: jorg@cise.ufl.edu (J. Peters).

[1] In surface interrogation, highlight lines are a commonly-used approximation to a parallel arrangement of tube lights in a car show room [1].

http://dx.doi.org/10.1016/j.cag.2014.09.004

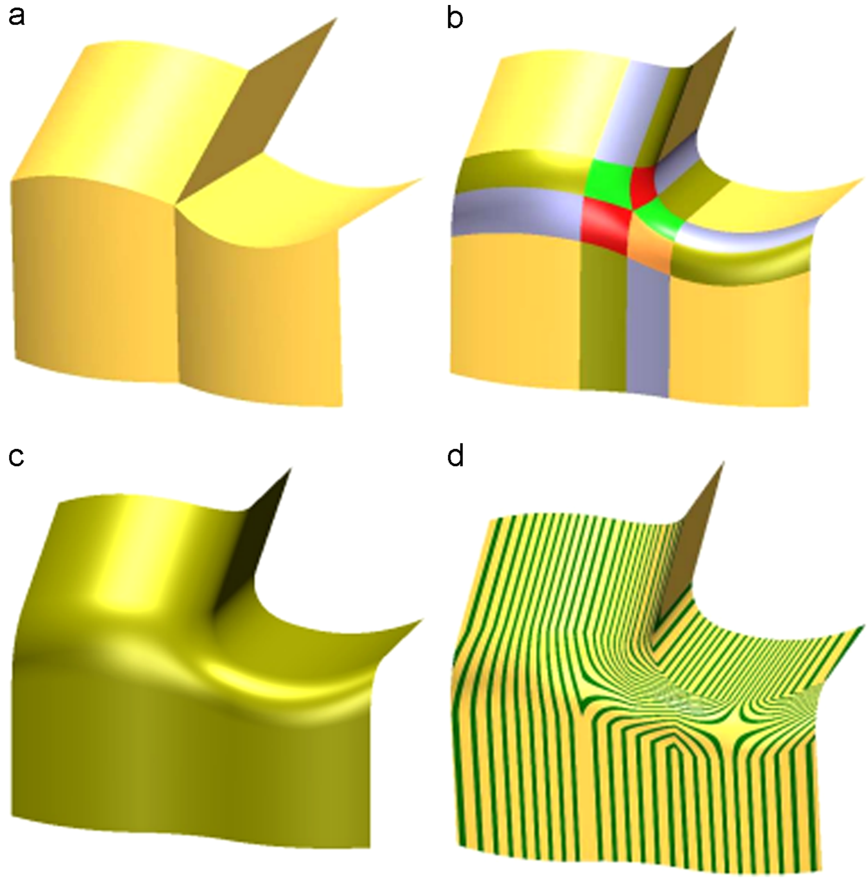

Fig. 1. Pairwise joining and *capping a collection of bi-2 spline surfaces*. (a) Input: one planar and four cylindric pieces. (b) The input is pairwise blended by bi-2 splines and the remaining gap is filled by a G^1 cap consisting of five patches of degree bi-3. (c) Output and (d) surface interrogation.

construction of Section 4.) Section 5 demonstrates the quality of the new G^1 caps when completing bi-2 spline models.

2. Background and motivation

Developed in the late 1980s and early 1990s, the theory of geometric continuity (G^1 continuity [8,9]) provides a foundation for solving the problem of completing a spline surface with tangent plane continuity where n surface pieces meet. The theory spawned a number of innovative algorithms: cap constructions with built-in singularities [10,11]; normalized by averaging, hence rational surface extensions [12]; restrictions of higher-variable constructions [13]; splitting of the domain into three-sided patches [14–16] and global manifold constructions [17,18] employing variants of simplex splines [19], to name just a few. Each class of methods has to overcome the challenge that an object of genus other than one, i.e. topologically different from the torus, does not admit a globally regular, shift-invariant tiling but must have one or more extraordinary points (topological singularities). In fact, when the surface pieces are twice continuously differentiable, all constructions have to solve a vertex-enclosure problem, see e.g. [9]: at even-valence vertices, a matrix of smoothness constraints (in terms of the mixed derivatives that one would naturally associate with the solution) is not invertible.

With the exception of Gregory's singular construction [11], none of these innovative solutions are known to have been adopted by industry. There are multiple reasons for this. First, the methods listed are not compatible with the CAD-industry's NURBS standard. Second, most CAD packages follow a different paradigm to fill multi-sided holes. Solidworks, for example, fits the graph of a function to the data along a boundary and trims off extraneous parts of the resulting many-knot spline surface. The output surfaces may not be watertight and require 'healing' before FEM analysis can be applied. The entertainment industry, led by Pixar, has adopted Catmull–Clark subdivision as a modeling tool. Subdivision surfaces are conceptually simple in that they evoke mesh refinement. Catmull–Clark (CC) [20] subdivision provides sufficient quality for animation but has not entered main stream CAD processing both because it only generalizes uniform polynomial splines and because the resulting surfaces consist of infinitely many pieces. Moreover, the quality of CC subdivision surfaces near the limit point is insufficient [21] (see also Fig. 10a). Doo–Sabin subdivision has had little exposure since it is a 'dual' subdivision method, i.e. shifts the faceted approximation under refinement. The severe artifacts of Doo–Sabin subdivision (e.g. Fig. 2b) do not seem to have been documented before.

There is a growing class of hole-filling constructions that are compatible with tensor-product splines. While constructions of least degree, bi-2 [22], cannot correctly handle higher-order saddles, but generate flat spots, the 1990s constructions of degree bi-3, e.g. [23,3], formally satisfy the smoothness constraints. However, as illustrated in Figs. 2 and 3, the shape of the surfaces generated by these approaches often disappoints. The current paper shines some light on the causes by emphasising the need for boundary data reparameterization.

To improve shape, functionals of the type

$$\mathcal{F}_m f := \int_0^1 \int_0^1 \sum_{i+j=m,i,j\ge 0} \frac{m!}{i!j!}(\partial_s^i f \partial_t^j f)^2,$$

$$\mathcal{F}^*_\kappa f := \int_0^1 \int_0^1 (\partial_s^\kappa f)^2 + (\partial_t^\kappa f)^2,$$