Computer-Aided Design 102 (2018) 104–114

Contents lists available at ScienceDirect

Computer-Aided Design

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

Subdivision surfaces with isogeometric analysis adapted refinement weights

Qiaoling Zhang [a], Malcolm Sabin [a,b], Fehmi Cirak [a,*]

[a] Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK
[b] Numerical Geometry Ltd, 19 John Amner Close, Ely, Cambridgeshire CB6 1DT, UK

ARTICLE INFO

Keywords:
Subdivision surfaces
Finite elements
Thin shells
Isogeometric analysis

ABSTRACT

Subdivision surfaces provide an elegant isogeometric analysis framework for geometric design and analysis of partial differential equations defined on surfaces. They are already a standard in high-end computer animation and graphics and are becoming available in a number of geometric modelling systems for engineering design. The subdivision refinement rules are usually adapted from knot insertion rules for splines. The quadrilateral Catmull–Clark scheme considered in this work is equivalent to cubic B-splines away from extraordinary, or irregular, vertices with other than four adjacent elements. Around extraordinary vertices the surface consists of a nested sequence of smooth spline patches which join C^1 continuously at the point itself. As known from geometric design literature, the subdivision weights can be optimised so that the surface quality is improved by minimising short-wavelength surface oscillations around extraordinary vertices. We use the related techniques to determine weights that minimise finite element discretisation errors as measured in the thin-shell energy norm. The optimisation problem is formulated over a characteristic domain and the errors in approximating cup- and saddle-like quadratic shapes obtained from eigenanalysis of the subdivision matrix are minimised. In finite element analysis the optimised subdivision weights for either cup- or saddle-like shapes are chosen depending on the shape of the solution field around an extraordinary vertex. As our computations confirm, the optimised subdivision weights yield a reduction of 50% and more in discretisation errors in the energy and L_2 norms. Although, as to be expected, the convergence rates are the same as for the classical Catmull–Clark weights, the convergence constants are improved.

1. Introduction

Isogeometric analysis aims to provide a seamless engineering design-analysis workflow by using a single common representation for geometric modelling and analysis. This is usually achieved by representing geometry and discretising analysis models with the same kind of basis functions [1]. The prevailing feature-based CAD modelling systems rely on trimmed NURBS and boundary representations (B-Reps). The resulting non-watertight geometries consisting of several trimmed patches pose unique challenges to finite element analysis. As a generalisation of splines, subdivision surfaces can provide watertight representations for geometries with arbitrary topology. After their early success in computer animation and graphics they are now supported in many CAD systems, including Catia, PTC Creo and Autodesk Fusion 360. Before the advent of isogeometric analysis, it had already been realised that subdivision surfaces provide also ideal basis functions for finite element analysis, in particular, of thin-shells [2–5], see also more recent work [6,7].

Subdivision schemes for generating smooth surfaces were first described in the late 1970s as an extension of low degree B-splines to control meshes with non-tensor-product connectivity [8,9]. In subdivision a geometry is described with a control mesh and a limiting process of repeated refinement. For parts of the mesh containing only regular vertices, with each adjacent to four quadrilateral faces, the refinement rules are adapted from knot insertion rules for B-splines. For the remaining parts with extraordinary vertices the refinement rules are chosen such that they yield in the limit a smooth surface. Subdivision refinement is a linear mapping of coordinates of the coarse control mesh to the coordinates of the refined mesh with a subdivision matrix. Hence, the local limit surface properties can be inferred from the eigenstructure of the subdivision matrix after a discrete Fourier transform [8,10]. The C^1 continuity of the surface and its curvature behaviour at the extraordinary vertex depend on eigenvalues and the ordering, i.e. Fourier indices, of the corresponding eigenvectors. In turn, both

* Corresponding author.
E-mail addresses: zq217@cam.ac.uk (Q. Zhang), malcolm.sabin@btinternet.com (M. Sabin), f.cirak@eng.cam.ac.uk (F. Cirak).

https://doi.org/10.1016/j.cad.2018.04.020

depend on the coefficients of the subdivision matrix that encodes the specific refinement rules applied.

As known, around extraordinary vertices short-wavelength surface oscillations, i.e. ripples, may occur irrespective of C^1 continuity and boundedness of curvature [11,12]. There have been many attempts to improve the fairness of subdivision surfaces, that is, to minimise curvature variations, by carefully tuning the refinement rules, earlier works include [13,14]. More recently, in Augsdörfer et al. [15] the refinement rules for Catmull–Clark and other quadrilateral schemes have been optimised such that the variation of the Gaussian curvature is minimised while ensuring bounded curvatures. Different from the direct search method used in [15], the refinement rules can also be obtained from a nonlinear constrained optimisation problem. Barthe et al. [16] apply such a procedure to triangular Loop and $\sqrt{3}$-subdivision schemes with a multi-objective cost function comprised of terms penalising divergence of curvatures and aiming local quadratic precision. In Ginkel et al. [17] a fairness increasing cost function containing the third derivatives of the surface in combination with C^1 continuity and bounded curvature constraints is optimised.

In the present paper, we optimise the subdivision refinement rules so that their approximation properties are improved when used in finite element analysis of thin-shells. Thin-shells are prevalent in many engineering applications, most prominently in aerospace, automotive and structural engineering, and are equivalent to thin-plates when their unstressed geometry is planar [18]. The thin-shell energy functional, and weak form, depend on the second order derivatives of the stressed surface. Consequently, it is crucial to reduce any short-wavelength oscillations in the subdivision surface. As the included examples demonstrate, meshes with extraordinary vertices usually lead to lower convergence rates than meshes with tensor-product connectivity. For obtaining the improved isogeometric analysis adapted refinement rules we postulate a constrained optimisation problem with a cost function measuring the errors in approximating cup- and saddle-like quadratic shapes. Three of the weights in the Catmull–Clark subdivision scheme around an extraordinary vertex are chosen as degrees of freedom for optimisation. As constraints the C^1 continuity of the surface is strictly enforced and bounded curvatures are enforced as long as non-negative real weights are feasible. The eigenstructure of the subdivision matrix is extensively used in formulating the optimisation problem as usual in previous related work [19, Chapter 4,5] and [20, Chapter 15]. We compute the eigenvalues and eigenvectors numerically after applying a discrete Fourier transform that exploits the local circular symmetry around the extraordinary vertex. The local parameterisation of the subdivision surface required for evaluating the finite element integrals and the cost function is obtained with the algorithm proposed by Stam [21]. Two sets of optimised weights for cup- and saddle-like shapes are obtained. The weights for finite element analysis are chosen depending on the dominant shape of the solution field around an extraordinary vertex.

For completeness, we note that subdivision is not the only approach for creating smooth surfaces on arbitrary connectivity control meshes. Over the years numerous C^k and G^k smooth constructions with $k \geq 1$ have been proposed, too many to name here. The search for sufficiently flexible smooth surface representations, especially with $C^{k\geq2}$ and $G^{k\geq2}$, is still open. It is worth mentioning that none of the existing constructions is widely used in commercial CAD systems. This may well be because their implementation is too complicated. The application of basis functions resulting from smooth constructions for isogeometric analysis is currently a very active area of research. For instance, the utility of G^k constructions with NURBS has recently been explored in [22–24]. Alternatively, C^k constructions relying on manifold-based surface constructions [25–27] and constructions relying on singular parameterisations have also been investigated [28–30]. Some of these schemes are able to provide optimal convergence rates.

The outline of this paper is as follows. In Section 2 the Catmull–Clark subdivision is introduced, with a review of the relevant theory on eigenanalysis of the subdivision matrix. Specifically, the necessary conditions for C^1 smoothness and boundedness of the curvature are motivated, and the local parameterisation of subdivision surfaces using the characteristic map is introduced. These are all classical results and concepts which are mostly unknown in isogeometric analysis. In Section 3 the proposed constrained optimisation problem and its numerical solution are discussed. Two sets of subdivision weights are derived that minimise the thin-plate energy norm errors in approximating locally cup- and saddle-like shapes. Subsequently, it is shown how a finite element solution can be locally decomposed into cup- and saddle-like components. Depending on this decomposition and the following choice of optimal weights, a second more accurate finite element analysis can be performed. In Section 4 the proposed approach is applied to transversally loaded thin-plate problems using meshes with extraordinary vertices and the convergence of the errors in L_2 and energy norms is reported.

2. Catmull–Clark subdivision surfaces

2.1. Refinement weights and the subdivision matrix

Catmull–Clark subdivision is a generalisation of cubic tensor-product B-splines to unstructured meshes [9]. On non-tensor-product meshes the number of faces connected to a vertex, i.e. valence v, can be different from four. The vertices with $v \neq 4$ are referred to as *extraordinary* or *star vertices*. During subdivision refinement each face of the control mesh is split into four faces and the coordinates of the old and new control vertices are computed with the subdivision weights given in Fig. 1. The weights in each of the three diagrams have to be normalised so that they add up to one. The unnormalised weights assigned to the extraordinary vertex (empty circle) are denoted by α, β and γ, respectively. For $v = 4$ and bivariate cubic B-splines the three weights take the values $\alpha = 8$, $\beta = 1$ and $\gamma = 1$. The new vertices introduced by the subdivision process are all regular (with $v = 4$) and the total number of irregular vertices in the mesh remains constant. That is, the irregular vertices are more and more surrounded by regular vertices.

In order to study the smoothness behaviour of subdivision surfaces near an extraordinary vertex, it is sufficient to consider only the vertices in its immediate vicinity. A 1-neighbourhood of a vertex is formed by the union of faces that contain the vertex. The n-neighbourhood is defined recursively as the union of all 1-neighbourhoods of the $(n-1)$-neighbourhood vertices. It is assumed that the considered n-neighbourhood has only one single extraordinary vertex located at its centre. The n-neighbourhood control vertices $\boldsymbol{p}^\ell$ at the refinement level ℓ are mapped to control vertices $\boldsymbol{p}^{\ell+1}$ with the subdivision matrix $\boldsymbol{S}$,

$$\boldsymbol{p}^{\ell+1} = \boldsymbol{S}\boldsymbol{p}^\ell. \quad (1)$$

The square subdivision matrix $\boldsymbol{S}$ can be readily derived from the weights indicated in Fig. 1. The control point coordinates at level ℓ are arranged in this form

$$\boldsymbol{p}^\ell = \begin{bmatrix} p^\ell_{1x} & p^\ell_{1y} & p^\ell_{1z} \\ p^\ell_{2x} & p^\ell_{2y} & p^\ell_{2z} \\ \vdots & \vdots & \vdots \end{bmatrix} \quad (2)$$

with each row containing the coordinates of one control point $\boldsymbol{p}^\ell_j \in \mathbb{R}^3$ with the index j.