Contents lists available at ScienceDirect

Computers & Graphics

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

Special Issue on SMI 2016

Procedural mesh features applied to subdivision surfaces using graph grammars

Wolfgang Thaller [a], Ursula Augsdörfer [a], Dieter W. Fellner [a,b]

[a] Institute of Computer Graphics and Knowledge Visualization, Graz University of Technology, Austria
[b] GRIS, TU Darmstadt & Fraunhofer IGD, Darmstadt, Germany

ARTICLE INFO

Article history:
Received 14 March 2016
Received in revised form
10 May 2016
Accepted 10 May 2016

Keywords:
Procedural features
Feature-based modelling
Free-form surface features
Graph grammars
Catmull–Clark subdivision surfaces
Automating changes in mesh topology

ABSTRACT

A typical industrial design modelling scenario involves defining the overall shape of a product followed by adding detail features. Procedural features are well-established in computer aided design (CAD) involving regular forms, but are less applicable to free-form modelling involving subdivision surfaces. Current approaches do not generate sparse subdivision control meshes as output, which is why free-form features are manually modelled into subdivision control meshes by domain experts. Domain experts change the local topology of the subdivision control mesh to incorporate features into the surface, without increasing the mesh density unnecessarily and carefully avoiding the appearance of artefacts.

In this paper we show how to translate this expert knowledge to grammar rules. The rules may then be invoked in an interactive system to automatically apply features to subdivision surfaces.

1. Introduction

The design of products that have to fulfil engineering requirements as well as aesthetic criteria typically involves free-form shapes. The styling activity of aesthetic product design comprises two steps [1]: First the product's overall shape is defined. This is followed by local refinements where features are added to the overall shape. Often, the same or similar features are repeatedly applied to one design.

A procedural feature consists of a set of parameters along with an algorithm for applying the feature to an underlying model. The designer can then manipulate the feature directly on a high semantic level of abstraction. The underlying shape can be edited while leaving the procedural features in place. Also, the designer is able to control specific aspects of a feature shape, while the overall shape remains fixed. This is referred to as *feature-based modelling*.

For feature-based modelling to work, features must be well-defined in terms of their parameters. Therefore, feature-based modelling was first introduced in the context of solid modelling, where procedural features have become firmly established. A cylindrical hole drilled into an overall shape is the classic example of a procedural feature. The use of free-form surface features, rather than regular-shaped features, is an active area of research [2].

Because feature-based modelling systems need to be compatible with a larger *modelling pipeline*, it is desirable for the system's input and output to be a standard CAD representation.

Non-uniform rational B-splines (NURBS), a patch-based surface representation, are the current standard for free-form modelling in CAD. Adding features to free-form surfaces represented by NURBS frequently involves increasing the resolution of surface patches, which leads to many redundant vertices in the representation.

Subdivision surfaces, already established in the animation industry, have recently gained popularity as an alternative to NURBS in CAD. Subdivision methods are a generalisation of traditional spline patch methods to arbitrary topology; for example, Catmull–Clark [3] generalises bi-cubic patches.

Because subdivision surfaces may also include extraordinary vertices, that is vertices with a valency either more or less than the regular valence, features may be introduced into the surface by locally changing the topology of the control mesh.

However, changes to the topology of the control mesh may give rise to artefacts in the limit surface. To apply free-form design features to a subdivision surface, a CAD expert meticulously adjusts the resolution and the mesh topology of the subdivision control mesh in order to keep changes to the overall shape small and to avoid or hide the visual appearance of artefacts [4]. Once a good topology change is identified, the expert applies the same procedure each time the same or a similar feature has to be applied.

In this paper we describe how this expert knowledge can be formulated as graph grammar rules. We propose to use a graph grammar on top of a scripting language; the basic local modifications are scripted, and the graph grammar allows us to organise the problem.

http://dx.doi.org/10.1016/j.cag.2016.05.014
0097-8493/© 2016 Published by Elsevier Ltd.

2. Related work

Grammars have proved very useful for the procedural modelling of buildings where context-free split grammars have been used [5,6]. There, too, a grammar provides the organising paradigm for something that could theoretically also be hard-coded.

Procedural free-form features have been a lively area of research since the 1990s. For a comprehensive introduction to and overview of feature-based modelling in CAD, we refer the reader to Pernot et al. [2].

2.1. B-spline based methods

In [7], Pernot et al. introduce a method for using three-dimensional curves for defining deformations applied to B-spline surfaces. This method was later extended [8–10] as a method for defining a range of freeform surface features on freeform surfaces. Application of the features is done using a force-based deformation system. The method can operate on NURBS data, but no new control points are added. The input geometry must have sufficient resolution to represent the features.

Chen et al. [11] solve this limitation by using *hierarchical NURBS* [12] to represent the resulting surface. In this approach, higher-resolution patches are used to refine the lower-resolution input patches and add high-frequency detail. The resulting geometry is tied to the use of non-standard hierarchical NURBS as a geometry representation and can therefore not be used in existing modelling pipelines.

Another way to avoid the problem of insufficient resolution in the input control mesh is to generate the feature during rendering instead of attempting to output a control mesh. Displacement mapping [13] is a wide-spread and efficient method for applying high-frequency detail to surfaces when the output is only used for visualisation.

2.2. Subdivision based methods

A popular method for adding a limited repertoire of features to subdivision surfaces has been defined by De Rose et al. [14]. In their approach, edges of a subdivision control mesh may be marked as being exempt from the subdivision process. This introduces sharp and semi-sharp creases into an otherwise smooth limit surface. However, the output surface is not a generalisation of traditional spline patch methods, which typically is a requirement further down the product design pipeline.

Khodakovsky and Schröder [15] describe an algorithm allowing the creation and manipulation of fine scale feature curves on subdivision surfaces. Creation of the features happens during the subdivision process, so there is no sparse control net for the resulting surface.

In the context of Sketch based modelling, Olsen et al. [16] have developed a method for incorporating linear features into subdivision control nets by locally increasing the mesh resolution. The transition between the higher-resolution patched faces and the lower-resolution surroundings is handled by a set of fixed patching template. However, this approach may increase control mesh density more than necessary.

In product design today, topological changes to a subdivision control net to incorporate high frequency features are still modelled manually by the CAD expert.

3. Adding features to subdivision surfaces

The problem of applying free-form features to a subdivision surface can be seen as consisting of three subproblems, namely (1) defining the feature, (2) changing the topology of the subdivision control mesh, and (3) shifting the control points appropriately.

Of these, the second step is the most time consuming and the one where expert knowledge is required.

Typically, a CAD expert with a thorough understanding of the surface representation carefully designs the subdivision control mesh in order to locally increase the resolution around the area where the feature is to be placed. This is not a trivial task.

CAD modelers typically take great care to avoid the appearance of surface artefacts.

In regular regions surface artefacts are known to arise if features are not aligned but run skew to grid lines of the control mesh [4,17,18]. Because of this CAD modelers take great care to align features with the underlying grid.

Locally refining a subdivision control mesh to incorporate the feature typically gives rise to irregular regions: For a subdivision scheme based on quadrilateral meshes these occur around vertices with more or less than four edges, referred to as extraordinary vertices, and non-quad faces, which give rise to extraordinary vertices after one subdivision step.

Subdivision surfaces do not guarantee C^2 continuity at extraordinary vertices [19], and undesirable artefacts are likely to appear [4,20,21].

Typically, when applying features to a subdivision surface, expert mesh modelers meticulously identify changes to the topology of subdivision control meshes such that the visibility of unavoidable artefacts around irregular regions is minimised.

Furthermore, when different features meet in one place on a surface, they have to interact. The interaction of features is domain-specific and cannot be determined without access to domain-specific expert knowledge.

As is clear from the above, incorporating features to a free-form design is time consuming and often requires an expert understanding of the underlying geometric representation in order to integrate features.

Once a method for good integration of a feature to a subdivision control mesh has been identified, the expert CAD designer has to use the same methods repeatedly in order to manually incorporate free-form features of similar type into subdivision surface control meshes. There are currently insufficient intelligent tools to support or automate this laborious task.

Automating some of the repetitive tasks of the feature-based subdivision modelling will accelerate the design process and frees the designer from needing extensive knowledge of the underlying geometry representation.

4. Contribution of the paper

We demonstrate that grammar rules are suitable to translate expert knowledge of how to change the topology of a control mesh to incorporate a feature. Compared to existing methods for automatically applying features to subdivision meshes, the rule-based approach adds significantly fewer control points to the mesh.

Changing the topology of a control mesh unavoidably introduces changes to the limit surface. To keep the alteration of the original surface to a minimum we apply an optimisation step after the changes in mesh topology.

To demonstrate our approach, we have developed a prototype system to automatically change the topology of a Catmull–Clark subdivision control mesh [3] to apply various surface features with only a minor increase in resolution to achieve good visual properties.

Please cite this article as: Thaller W, et al. Procedural mesh features applied to subdivision surfaces using graph grammars. Comput Graph (2016), http://dx.doi.org/10.1016/j.cag.2016.05.014