The Benefits of Subdivision Surfaces for Complex Geometry

Paul SHEPHERD
Research Fellow
Dept. of Arch. & Civil Eng.
University of Bath, UK
p.shepherd@bath.ac.uk

Paul Shepherd has a maths degree from Cambridge, PhD from Sheffield and worked for Buro Happold where he set up their advanced analysis group now known as SMART. He is now Digital Architectonics Research Fellow at the University of Bath.

Summary

In his seminal IASS Journal paper titled "New Challenges for the Structural Morphology Group" in 2010, Andrew Borgart outlined the need for radical new techniques for the design, engineering and construction of complex geometry structures in order to continue to produce innovative and beautiful design solutions in the current economic and environmental climate. He concluded that unorthodox solutions were needed, and that these would only be provided through transfer of technologies from a wide range of disciplines. This paper rises to this challenge, by arguing that the adoption of Subdivision Surfaces as a new modelling framework for engineering design would go a long way towards addressing these problems and would reinvigorate the Shell and Spatial Structure design community.

Keywords: *Digital Architectonics, Subdivision, Surfaces, Complex, Geometry, Mesh.*

1. Introduction

Since computers began to appear in the design office back in the 1960s, engineers have been searching for ways in which they could be used to streamline the design process. However, whilst in the fields of aerospace, automotive and even shipbuilding their computing power has been harnessed to optimise the design itself, in the building industry, computer aided design (CAD) was seen as an electronic version of paper, used for its ease of storage / retrieval more than a tool for analysis. Whilst engineers in other industries were innovating through 3D solid and parametric modelling, building construction industry drawings were being created manually in the same way as had previously been done with a pencil and drawing board.

CAD began in 2D, defining geometry with straight lines and later incorporated Bezier curves and splines. When it moved into 3D surface representation, it took these Splines and arrayed them into grids to make Spline surfaces and NURBS. They required a 2D parameterisation of space in which to array two pseudo-orthogonal sets of splines, and this limitation sometimes led to the need to break down a desired surface into separate patches. Such an approach often introduced problems of discontinuities in tangent (causing creases) or in rate-of-change of tangent (distorting reflections).

Spline surfaces and NURBS can't easily be fabricated for building construction and can be computationally difficult to evaluate. Doubly curved architecturally-driven surfaces are often either post-rationalised by the engineer or contractor into singly-curved or flat panels, or converted into a triangulated mesh, in order to be built.

The use of patches and of post-rationalising the geometry for construction is time-consuming and inevitably leads to a compromise between the surface desired by the architect, the surface representation of the chosen CAD program, and the need for a constructible solution.

2. What are Subdivision Surfaces?

Subdivision surfaces were developed throughout the 1980s for applications in 3D computer graphics and have seen a recent focus on development in the digital entertainment industries for computer animation and gaming. They represent a smooth 3D surface using a polygonal mesh of defined by a set of vertices and an underlying topology. The mesh can be constructed from triangles or quadrilaterals (or a combination of both). For simplicity of explanation this paper will focus on triangular meshes, but the descriptions extend to quadrilateral meshes also.

A mesh is a very simple object to construct and manipulate, but generally has a crease along every edge and therefore cannot describe a smooth surface as such. However, a mesh can be made finer through a process of subdivision. That is to say that each triangular face can be split into four smaller triangles by introducing a new "child" vertex along each edge and joining each of these child vertices to the other two with a new edge, as shown in Fig. 1b. In itself, this purely topological subdivision does not actually change the surface geometry, since each planar face is subdivided into four smaller, but still co-planar, faces. The key to a subdivision surface representation is that the child vertices are not simply placed along the original edge, but have their position carefully calculated as a weighted average of the positions of all the surrounding vertices. In this sense, this averaging out of the coordinates has the effect of reducing the discontinuities, and the resulting mesh has four times as many triangles but is much smoother. There are many different methods for weighting the positions of the surrounding vertices, each known as a Subdivision Scheme. Some only place the child vertices at these weighted positions (interpolating schemes) and others also move the original "parent" vertices (approximating schemes), as shown in Fig. 1.

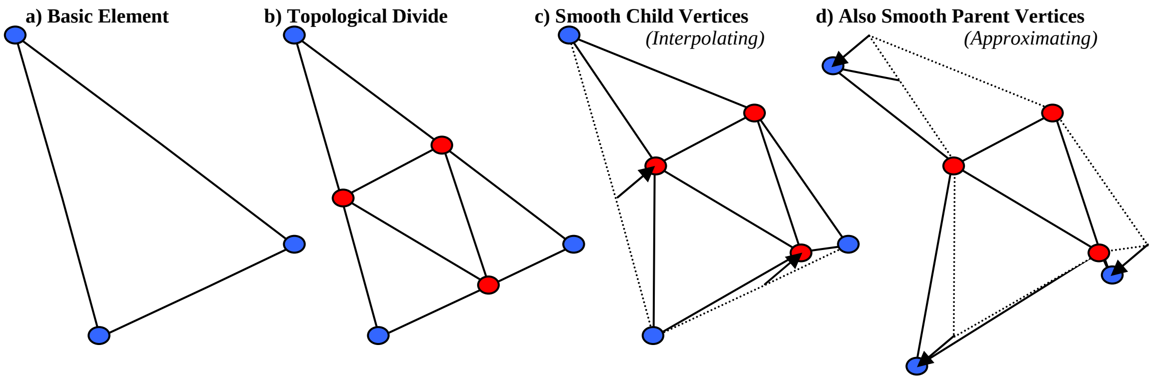

Figure 1 Description of subdivision

Once a mesh has been subdivided, it can be subdivided again. Each "generation" is a finer and finer mesh with four times as many triangles but with a smoother geometry. As more and more subdivisions are carried out, the mesh gets closer and closer to an underlying "limit surface", which is guaranteed not to have any creases and in general has C2 continuity (continuous rate-of-change of tangents).

The advantages provided by this hierarchical series of meshes were capitalised upon by the computer graphics industries, whereby the same object (represented by the limit surface) could be displayed at various levels of detail as required. An object such as a game character far away in the background might only be drawn using a few hundred triangles, whereas the exact same character would be subdivided more when in the foreground and drawn with thousands of triangles. This method was very efficient on computing resources and allowed games to run faster and with more objects, textures and sounds.

3. Why are they so interesting to building designers?

Subdivision surfaces offer many benefits to the building design community.

From an architectural point of view, their guaranteed smoothness creates aesthetically pleasing doubly-curved surfaces. These surfaces can be easily manipulated in real-time using the same techniques as in the computer graphics industries, namely they can be edited, and even stored, at a low level of subdivision and then rendered later at higher levels of subdivision to produce accurate drawings or images.

Engineers often have to convert complex surface geometries into simplified meshes for finite element (FE) or computational fluid dynamics (CFD) analyses. With a subdivision surface representation, the hierarchical level of detail can be used to generate a mesh at exactly the required density for any given analysis application. It is often also the case that a complex surface needs some sort of mesh to represent its support structure. A surface defined using subdivision could be designed using the standard tools such as cutting with sections or draping a grid. But they also come with their own inherent triangulated structural grid, which can also be sampled at various levels of detail to lead to a sensible panel or member size. Fig. 2 shows, for example, what the roof of the British Museum Great Court could have looked like if it had been constructed using a subdivision surface representation.

Fig. 3 British Museum Great Court roof options using various levels of subdivision

Since subdivision surfaces are created using a relatively coarse initial control mesh, they lend themselves very well to optimisation. A complex surface can be defined by a few tens of control vertices. This opens up possibilities of carrying out multi-objective optimisation to assess a proposed structure for any number of structural or environmental performance criteria, and using the results to feed-back and define new positions for the control mesh vertices. Since the control mesh has very few degrees of freedom, any optimisation will be efficient and could provide real-time feedback to a designer on the performance of the current proposal.

In order to be of use in building construction however, subdivision surfaces need to be constrained such that they can be forced to respect a given boundary. In the case of the Great Court described above, with its rectangular shape on-plan, a standard subdivision scheme's goal of smoothing the geometry would result in the corners being rounded off also. Clearly this would not be acceptable in this context and full control is needed to specify where the subdivision can occur and where a given constraint has to be respected – usually at least around the boundary. Subdivision schemes can be adapted to achieve this; however the price paid is the lack of C2 continuity around these constraints. Tangent continuity is preserved however, so no creases appear, and this is therefore viewed as an acceptable compromise.

4. Conclusions

Borgart's [1] call for "radical new techniques for the design, engineering and construction of complex geometry structures in order to continue to produce innovative and beautiful design solutions in the current economic and environmental climate" can be addressed by adopting a subdivision surface representation for complex geometry structures. Subdivision surfaces are aesthetically desirable, and their hierarchy of levels of detail have advantages in terms of providing a wide range of analysis meshes and options for support structure from a single base model.

There are some issues with the one-directionality of subdivision surfaces which will require a shift in the way architects design their buildings. Subdivision surfaces are defined using a coarse control mesh, and are subdivided to find the limit surface on which a design proposal would be based. Currently architects have a clear idea of where they want their surface to be, and would wish to work backwards from this, to discover the coarse control mesh which will result in their desired surface. Whilst software can be developed to back-calculate a control mesh from a given target surface, it is suggested that a more radical change in approach is needed. If a designer only has a coarse control mesh to define, they might think more carefully about where the surface should be and why, and maybe use other feedback mechanisms such as structural or environmental calculations to help make their decision.

The main current limitation on the use of subdivision surfaces is the question of intersections. Currently no robust mechanism has been developed to calculate the line of intersection between two subdivision limit surfaces. Intersection is the basis of all Boolean operations such as union and difference, and such operations will need to be addressed if subdivision surfaces are to be used on real live building projects. An early attempt by the author to develop such functionality is shown below in Fig. 4.

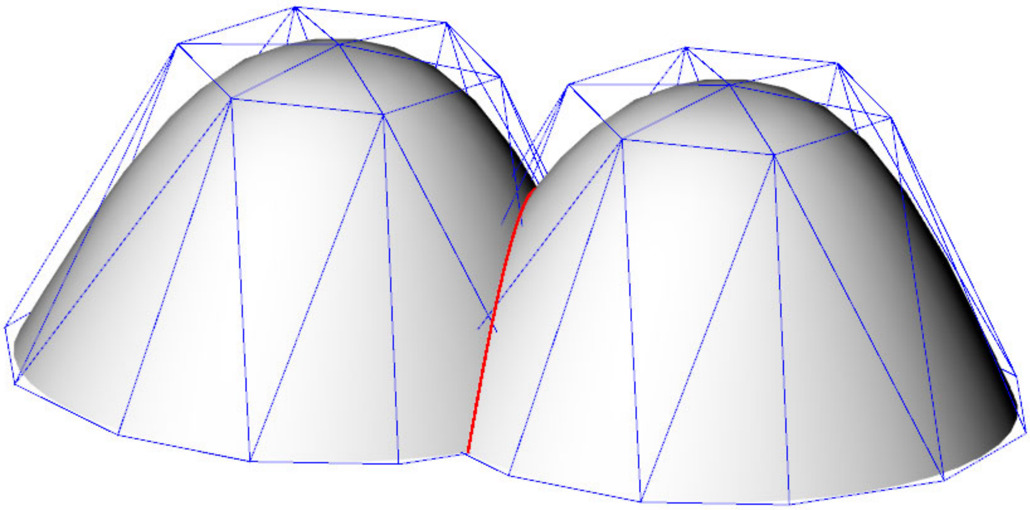

Fig. 4 Coarse control mesh (blue) with limit surface (grey) and intersection (red).

Despite these current limitations, the benefits that subdivision surfaces offer to the building design community are still vast, and building design practitioners should learn to use them or risk losing out.

5. References

[1] Borgart, A., 2010, “New challenges for the structural morphology group”, J. IASS, 51(3), pp183-189.