Title:
C^2-Smooth Subdivision Surface Construction from Cross-Sections

Authors:
Abdulwahed M. Abbas, abbas@balamand.edu.lb, The University of Balamand
Ahmad H. Nasri, anasri.82@gmail.com, Boulder Graphics LLC

Keywords:
Interpolating Surface Construction, Catmull-Clark Subdivision Surfaces, Polygonal Complexes, Extraordinary Vertices

DOI: 10.14733/cadconfP.2019.308-313

Introduction:
This paper presents an alternative to the approach proposed by Gabrielides et al. [3] for the construction of a smooth surface to interpolate data points given on parallel cross sections.

The method presented there focuses on handling the general case of planar cross sections consisting of disjoint smooth Jordan curves, referred to there as contours, whose number may vary from plane to plane (see Fig. 1(a) and Fig. 1(b), borrowed from [3]). Their approach relies on a combination of modeling contexts and a variety of modeling techniques guaranteeing a G1 interpolation surface.

Gabrielides et al. [3] also provide a number of application domains where the solution of this interpolation problem is employed. They also detail a number of methods and tools that were used in the implementation of solutions of this problem.

By comparison, the interpolation problem addressed in this paper concerns the interpolation of the given data points and not the cross sections carrying them.

In this context, one does not fail to see the relationship between this interpolation problem and another problem known as lofting [5] (or skinning [4]), which concerns the generation of a smooth surface interpolating a set of (parallel) curves. Historically, the term lofting is used in the old days when ship bodies used to be designed manually.

Fig. 1: Interpolating surface construction: (a) The contour data set on parallel planes, and (b) One possible interpolating surface.

Compared to [3], this paper advocates the generation of a Catmull-Clark subdivision surface [2] that interpolates the given data points. The benefits here are two folds: first, the simplicity of the approach in that it turns most of the mechanism into a mesh generation process with no particular cases that require special attention and, second, the continuity of the resulting surface will be C2 except at possibly a few extraordinary points were it may drop to C1.

Main Idea:
Our reason for choosing uniform cubic B-spline curves as the representation of curves is motivated by two factors. First, they are very naturally interpolated by Catmull-Clark subdivision surfaces, and second, they can be directly modified to interpolate any point in their vicinity.

Uniform Cubic Refinement of a Polygon
In a refinement step, the control polygon [A_0, B_0, C_0, …] is refined into a polygon [A_0, a_0, b_0, c_0, d_0, …] as follows (see Fig. 2):

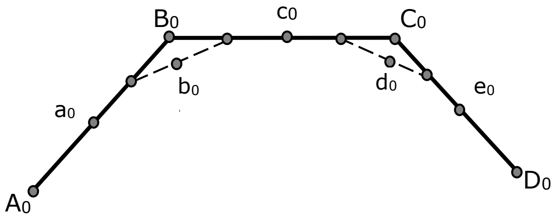

Fig. 2: Uniform cubic refinement of a polygon.

- a_0 (resp. c_0 and e_0) is the midpoint of the edge A_0B_0 (resp. B_0C_0 and C_0D_0), etc.
- b_0 (resp. d_0) is the midpoint of the edge joining the midpoints of a_0B_0 and B_0c_0 (resp. the edge joining the midpoints of c_0C_0, and C_0e_0), etc.

Repeating this refinement process sufficiently often will lead to a cubic B-spline curve. Note that the extremity A_0 of the open polygon will be interpolated under the version of the refinement scheme we are adopting here.

Point Interpolation
Among many equivalent techniques, for any uniform cubic B-spline curve (C) with a control polygon (P) (see Fig. 3(a)), if we replace the point q of (P) with the point:

$$\frac{1}{4}\times\begin{bmatrix}-1 & 6 & -1\end{bmatrix}\times\begin{bmatrix}p & q & r\end{bmatrix}^T \tag{2.1}$$

Then the cubic refinement of the resulting polygon will be a curve that interpolates q itself (see Fig. 3(b)).

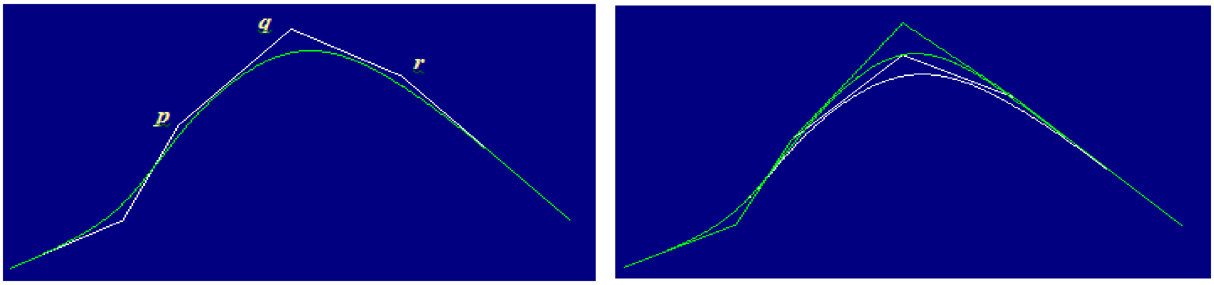

Fig. 3: Interpolating points by curves: (a) Initial control polygon and curve, and (b) Polygon modified to interpolate a given point.

Clearly, this process may be repeated to interpolate any sequence of control points by a single cubic B-spline curve.

Catmull-Clark Subdivision Surfaces

In a Catmull-Clark subdivision step (see Fig. 4(b)), a control mesh M is subdivided [2] into another control mesh M', as follows:

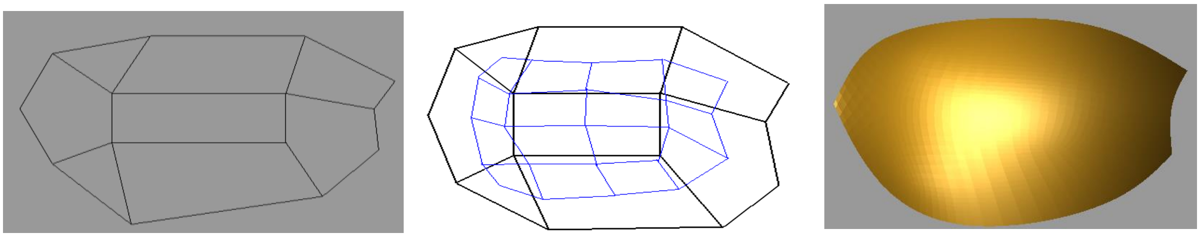

Fig. 4: Catmull-Clark subdivision: (a) Initial control mesh, (b) A single subdivision step, and (c) The limit surface.

Each inner face F of the mesh gives rise to an F-vertex that is the average of the vertices of the face F. Each inner edge E gives rise to an E-vertex that is the average of the vertices of the edge together with the F-vertices of the adjacent faces of E. Each inner vertex v gives rise to a V-vertex given by the following expression:

$$n-2 * v + R + S / n / n \tag{3.1}$$

Where

- n is the number of faces adjacent to v.
- R = $\sum_{n}^{i=1} \mathbf{v}_i$ where vv_i is an edge of the corresponding mesh.
- S = $\sum_{n}^{i=1} \mathbf{v}_{fi}$,

where v_{fi} is an F-vertex of a face f_i of the mesh adjacent to the vertex v.

At the end of this process, each F-vertex is connected to the adjacent E-vertices and each E-vertex is connected to the adjacent V-vertices. The resulting faces will form the next subdivided mesh in the process. Note in this context that repeated application of this subdivision step will converge to a smooth surface (see Fig. 4(c)).

Accordingly, the boundary vertices and edges will not contribute any new vertices. Therefore, the initial boundary vertices will not in general be interpolated by the limit surface.

Continuity of Catmull-Clark Subdivision Surfaces

It is well known that the continuity of Catmull-Clark subdivision surfaces is C2 except at possibly a few extraordinary points were it may drop to C1.

Catmull-Clark Polygonal Complexes and Their Limit Curves

A Catmull-Clark polygonal complex [1] (see Fig. 5) is a 3×n matrix M of points representing three control polygons: top(ti), middle(mi) and bottom(bi), all having the same number n of vertices. These may also be seen as a sequence of pairs of rectangular faces; where each pair of faces of this sequence has a common edge and each two consecutive pairs have common respective edges.

A CC complex is interesting because, under subdivision, it leads to a sequence of thinner and thinner complexes which, at the limit, converges to a smooth curve.

In fact, the limit of a simple CC complex M is a cubic B-spline curve whose control polygon P is given by the following formula (see [4]):

$$1/6 * \begin{bmatrix} 1 & 4 & 1 \end{bmatrix} * M \tag{4.1}$$

Thus, when a complex is embodied within a control mesh, its limit curve will naturally be interpolated by the limit surface of this mesh. More importantly perhaps, if a complex M' is obtained from a CC complex M by substituting the mid-polygon m of M by the polygon (see [4]):

$$m' = 1/4 * \begin{bmatrix} -1 & 6 & -1 \end{bmatrix} * M \tag{4.2}$$

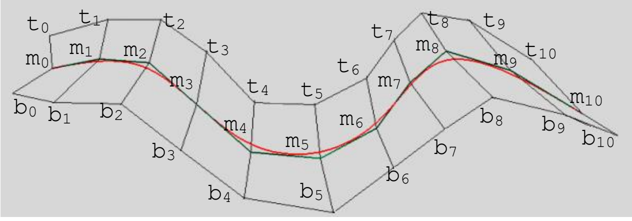

Fig. 5: A Catmull-Clark polygonal complex.

then the limit of M' is a B-spline curve identical to that of m. Thus, given a curve defined by a control polygon (m_i), we can turn it into a polygonal complex M by adding to it two more rows of control points (t_i) and (b_i). In fact, applying Eqn (4.2), we can guarantee that any mesh embodying the complex M' will in fact be interpolating the original curve defined by (m_i).

The Solution of the Interpolation Problem

When solving a problem via subdivision, most of the effort in developing the solution goes into the design of the initial control mesh. After that, the subdivision process will take care of obtaining a limit surface which displays all the desired properties of the solution.

Solution of the Interpolation Problem in a Simplified Case

In order to simplify the exposition somewhat, we start first by showing a solution of the problem where each plane contains a single contour (see Fig. 6).

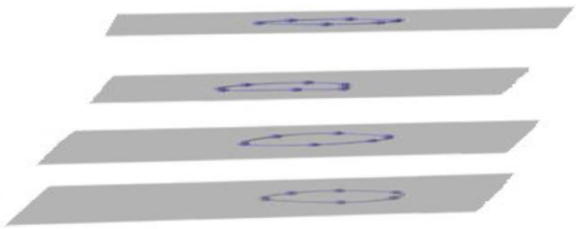

Fig. 6: Interpolating surface construction in a simplified case.

In this case, the solution goes through the following steps:

- Assuming that each contour is a closed planar polygon whose vertices are the initial control points to be interpolated by the final surface, we replace each contour by a control polygon whose corresponding cubic uniform B-spline curve would interpolate the corresponding control points. Thus, from now on, we no longer need to mention these initial control points (see Fig. 7(a)).
- Add on a parallel plane, on either side of each contour, a copy of this contour, in such a way that these new contours are equidistant (see Fig. 7(b)).

Proceedings of CAD'19, Singapore, June 24-26, 2019, 308-313

- For each three consecutive contours in the sequence, add connecting edges in the obvious way, so that the three contours will form a closed polygonal complex in the sense of Section 4. All the faces obtained here will be regular faces.
- For each newly formed polygonal complex, apply the transformation represented by equation (4.2). This way, any surface embodying the complex will end up interpolating the uniform cubic B-spline curves corresponding to the contours and therefore interpolating the initial control points.

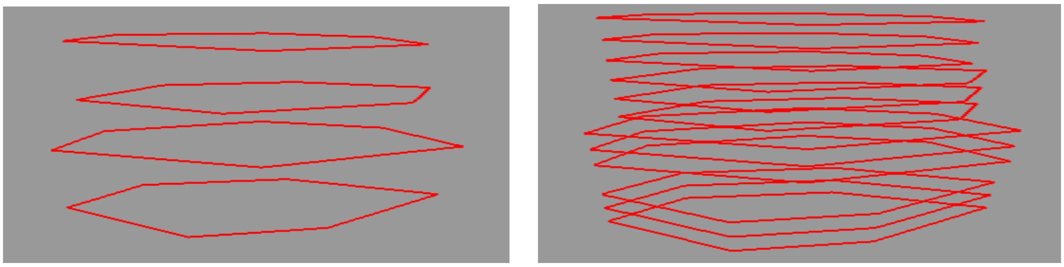

Fig. 7: Control mesh generation: (a) Initial contours, and (b) Augmented contours.

- Apply a heuristic linking procedure that add edges to link neighboring complexes by repeatedly selecting the closest neighboring unconnected pairs of vertices (one of each complex) and join them together by an edge, until no such pair of vertices exist. It is worth mentioning here that, in the case where the two neighboring complexes are different in length, this linking procedure might give rise to extraordinary faces and therefore extraordinary vertices during the subdivision process (see Fig. 8(a)).

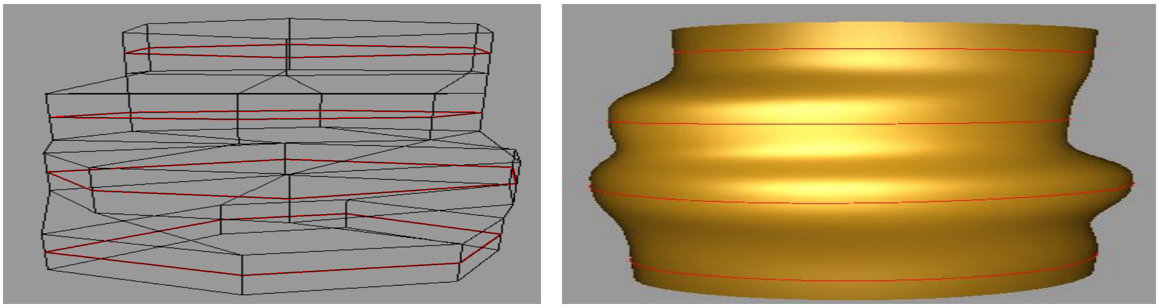

Fig. 8: Control mesh: (a) Before subdivision, and (b) Interpolating surface.

The CC subdivision of the final control mesh obtained this way will give rise to a smooth surface interpolating the initially given control points. This surface will be C2 everywhere, except maybe in the area where the connection procedure gives rise to extraordinary vertices. Here, the continuity of the surface drops to C1 (see Fig. 8(b)).

Solution of the Interpolation Problem in General

In general, each of the parallel planes might contain more than a single contour. Consequently, the interpolation process will have to establish how the interpolating surface has to be extended to involve the interpolation of the given points at either or all the contours at the next parallel plane. See Fig. 1(b) for an idea on how this is tackled in [3]. The technique is referred to there as branching. Gabrielides et al. [3] supplied illustration on how branching is dealt with (see Fig. 11(a) and Fig. 11(b)).

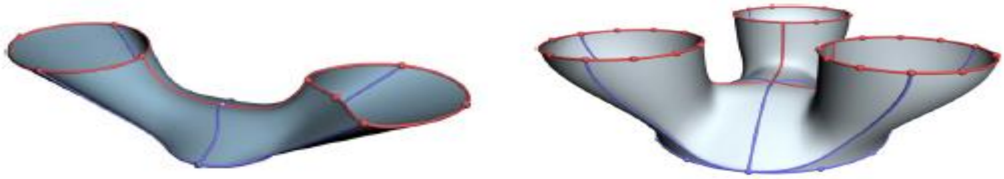

Fig. 11: Branching as tackled by Gabrielides et al. [3]: (a) The final “one-to-two” surface, and (b) The final “one-to-three” branching surface.

It is important to note here (in Fig. 1(b)) that, in the region where two branches meet, the surface will suffer from irregularity, which does not integrate well with B-spline surfaces. This perhaps explains why the technique employed in that paper has to resort to a host of techniques outside the B-spline surface domain causing the continuity of the resulting surface to drop to G1.

In Catmull-Clark subdivision surfaces however, this situation does not arise because the existence of extraordinary vertices in the control mesh will drop continuity there to C1 in the worst situation. In fact, the situations illustrated in Fig. 11 can be integrated in the control mesh illustrating the final surface and then subdivided quite naturally leading to the final target surface.

Conclusions:

The declared aim of this paper is to construct Catmull-Clark subdivision surface that interpolate a set of data points given by non-intersecting cross-contours on parallel planes. The use of subdivision surfaces was needed in order to enhance the continuity of the resulting surface compared to what is obtained by Gabrielides et al. [3].

Fig. 10 shows that the general aim of the paper has largely been achieved. However, the evenness of the resulting surface leaves much to be desired as compared to the results obtained by Gabrielides et al. [3] (see Fig. 1 (b), for example). This is due to the large number of extra points that was required during the process; first, to interpolate the initial data points by Cubic B-spline curves and, later, to obtain the final interpolating surface. This made the general quality of the resulting surface awkward to control, even though the continuity of this surface has been improved.

In this context, Gabrielides et al. [3] mention the use of operators to enhance the quality of the resulting surface. So, as further work, some of such operators could be introduced to our context in order to improve the selection of the extra points needed during the process and thus to obtain similar enhancement of the resulting surface.

Acknowledgements:

The research work reported in this paper is supported by a grant from the Lebanese Council for Scientific Research.

References:

[1] Abbas, A.; Nasri, A.: Lofted Catmull-Clark recursive subdivision surfaces, Proceedings of the International Conference on Geometric Modeling and Processing, Wako, Japan, 83-93, 2002.

[2] Catmull, E.; Clark, J.: Rosalee Wolfe, Recursively generated B-spline surfaces on arbitrary topological meshes, Seminal Graphics, ACM Press, 183-188, 1998.

[3] Gabrielides, N.-C.; Ginnis, A.-I.; Kaklis, P.-D.; Karavelas, M.-I.: G1-smooth branching surface construction from cross sections, Computer-Aided Design 39(8), 2007, 639-651. https://doi.org/10.1016/j.cad.2007.05.004

[4] Nasri, A.; Abbas, A.; Hasbini, I.: Skinning Catmull-Clark subdivision surfaces with incompatible cross-sectional curves, Proceedings of the 11th Pacific Conference on Computer Graphics and Applications, 2003. https://doi.org/10.1109/pccga.2003.1238252

[5] Schaefer, S.; Warren, J.; Zorin, D.: Lofting curve networks using subdivision surfaces, Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing, 103-114, 2004. https://doi.org/10.1145/1057432.1057447