Surface Tiling with Differential Topology

(extended abstract of invited talk)

Herbert Edelsbrunner

Department of Computer Science, Duke University, Durham, and
Raindrop Geomagic, Research Triangle Park, North Carolina, USA.

Abstract

A challenging problem in computer-aided geometric design is the decomposition of a surface into four-sided regions that are then represented by NURBS patches. There are various approaches published in the literature and implemented as commercially available software, but all fall short in either automation or quality of the result. At Raindrop Geomagic, we have recently taken a fresh approach based on concepts from Morse theory. This by itself is not a new idea, but we have some novel ingredients that make this work, one being a rational notion of hierarchy that guides the construction of a simplified decomposition sensitive to only the major critical points.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Boundary representations, Hierarchy and geometric transformations

1. Introduction

This extended abstract is a brief description of the approach to *surface tiling* recently taken at Raindrop Geomagic. We define this problem in terms of input and output and an informal discussion of constraints and properties.

Input: a piecewise smooth embedding of a surface in three-dimensional Euclidean space.

Output: a face-to-face tiling of the surface into four-sided NURBS patches.

The problem is made difficult by a number of often unspoken assumptions, some of which are in the eye of the beholder. One set of assumptions on the output arises from intended, down-stream uses of the tiling. Another such set arises from technical difficulties in generating a light-weight tiling (few patches, small number of control points) that forms a good approximation of the input surface and of its derivatives [Far97]. There are also assumptions on the input, and ours include that the surface be given by a piecewise linear approximation; the actual surface is not known. We also assume that the input surface is the boundary of a member of the class of *mechanical shapes*. It is better to leave it at that, both for the reader and the writer, but it is important to keep in mind that there are substantial gaps in the above problem specification. Figure 1 illustrates the problem by showing an example of the desired output, the boundary of an oil-pump

represented by a collection of NURBS patches. The solution to the surface tiling problem developed at Raindrop Geomagic combines old and new ideas from various disciplines, some of the more important ones from Morse theory. We focus here on the geometry and topology of the construction, ignoring a substantial amount of interesting algorithmic and numerical work. It should be clear that the author of this extended abstract is deeply indebted to the people who did the actual development at Raindrop Geomagic, both in the Research Triangle Park, North Carolina, and in Budapest, Hungary.


Figure 1: *A NURBS representation of the boundary of an oil-pump.*

2. Morse Theory Background

The central objects studied in Morse theory are a manifold $\mathbb{M}$ (without boundary) and a smooth function $f : \mathbb{M} \to \mathbb{R}$; see [Mil63]. We simplify notation by assuming $\mathbb{M}$ is a 2-manifold smoothly embedded in $\mathbb{R}^3$ and distance is measured as the Euclidean length of shortest paths. We may therefore think of the *gradient* at a point $x \in \mathbb{M}$ as a vector $\nabla f(x)$ in $\mathbb{R}^3$. It lies in the tangent plane of x and points in the direction of locally steepest increase in function value. The point is *regular* if $\nabla f(x) \neq 0$ and *critical* if $\nabla f(x) = 0$. The typical picture one finds for critical points assumes the function measures height (distance from a base plane), as below. This leads to a convenient graphical representation but is misleading since the more useful functions are not height functions. A critical point x is *non-degenerate* if the Hessian

minimum   saddle   maximum

(the square matrix of second partial derivatives) is invertible. There are only three types of non-degenerate critical points on a 2-manifold, namely *minima*, *saddles*, and *maxima*, as illustrated above. A *Morse function* is a smooth function $f : \mathbb{M} \to \mathbb{R}$ for which

I. all critical points are non-degenerate;
II. critical points have pairwise different function values.

An important insight is that the Morse functions are dense among the smooth functions on $\mathbb{M}$. We can therefore think of them as the class of generic smooth functions, in the sense that they have generic local structure and every other smooth function can be perturbed to a Morse function.

An *integral line* is a curve $\gamma : \mathbb{R} \to \mathbb{M}$ whose velocity vectors agree with the gradient, $\dot{\gamma}(s) = \nabla f(\gamma(s))$ for all $s \in \mathbb{R}$. Its *image* is $\operatorname{im}\gamma = \{\gamma(s) \mid s \in \mathbb{R}\}$. The curve starts at a critical point, $\operatorname{org}\gamma = \lim_{s \to -\infty}\gamma(s)$, and ends at another, $\operatorname{dest}\gamma = \lim_{s \to \infty}\gamma(s)$. Note, however, that the curve is open and contains neither endpoint. Two integral lines are either disjoint or the same. This suggests we define the *descending manifold* of a critical point x as the set of points on integral lines with destination x, together with the critical point itself,

$$D(x) = \bigcup_{\operatorname{dest}\gamma = x} \operatorname{im}\gamma \ \cup \ \{x\}.$$

If x is a minimum then $D(x) = \{x\}$. If x is a saddle then $D(x)$ is an open curve consisting of x and the images of two integral lines on the two sides. If x is a maximum then $D(x)$ is an open disk consisting of x and a circle of images of integral lines. By construction, the descending manifolds are pairwise disjoint and together they cover $\mathbb{M}$. They also form a complex, in the sense that the boundary of each descending manifold is the union of (lower-dimensional) descending manifolds. We call the collection of descending manifolds the *Morse complex* of f. If we collect integral lines with common origin we get the dual collection of *ascending manifolds*, the Morse complex of $-f$.

3. Work-flow

We divide the construction of the tiling into three stages:

Stage 1. construct a curve network to decompose the surface into the regions of a simplified Morse complex;
Stage 2. extract the feature skeleton, a cleaned-up and locally widened version of the curve network;
Stage 3. refine the decomposition into quadrangles, eventually replacing each by a four-sided NURBS patch.

Stage 1 requires a function on the surface. We experiment with a suite of *basis functions*, which includes approximations of different notions of curvature for piecewise linear surfaces. We get best results by defining f as a weighted sum of the basis functions and allow the user to create this mix interactively, driving the selection by giving positive and negative examples (points on and off the desired feature skeleton). Once obtained, this weighted sum can be used for a family of similar surfaces.

Stage 1: Simplified Morse Complex. Given $f : \mathbb{M} \to \mathbb{R}$, we construct a piecewise linear version of its Morse complex. To describe what this means, we note that $\mathbb{M}$ is represented by a triangulation and f is specified at the vertices and extended by linear interpolation. In this setting, we use lower links of vertices to distinguish between critical and regular points. Most of the subtleties in the construction arise from a


Figure 2: *A simplified Morse complex decomposition with curves running through strips obtained by widening the descending 1-manifolds.*

proper interpretation of Conditions I and II of a Morse function and the occasional simulation of these properties when they are violated. The Morse complex has a region for every maximum, even if it is barely noticeable and perhaps just an artifact of the piecewise linear representation of the function. We therefore simplify the fine complex by cancelling critical points in pairs. This can be done in the order of increasing

persistence [ELZ02], but other measures of importance may also be used. To illustrate this construction, Figure 2 shows a simplified Morse complex obtained for a curvature approximating function.

Stage 2: Feature Skeleton. By construction, the curves (descending 1-manifolds) of the Morse complex follow ridges defined by f. They coincide with what we would intuitively refer to as ridges of the surface only if f is chosen accordingly. These ridges are round, somewhere on the scale between sharp and flat. A sharp ridge is well represented by a curve, possibly after using numerical methods to improve its layout. A round ridge, however, is more appropriately repre-


Figure 4: *Mixing basis functions to place longitudinal boundaries of curve-blends. The dots represent positive and negative examples placed by the user to learn an appropriate weighted sum of basis functions.*

4. Discussion

The author believes that the Morse theory approach works well for the surface tiling problem because it provides a good match for the challenges posed by the subtlety and ambiguity of surface shapes. In most cases, no one particular tiling is obviously the right one. To find a good solution it is important to use a mechanism that switches with ease between alternative solutions following accumulated preference without sacrificing internal consistency.

Figure 3: *A network of curve- and vertex-blends computed from the descending 1-manifolds of the simplified Morse complex.*

sented by a *curve-blend*, covering a strip obtained by widening the curve on both sides, as shown in Figure 3. Widening the curves cannot be done without growing the vertices into little patches, which we refer to as *vertex-blends*. Indeed, we need a variety of different types to accommodate connections to sharp curves as well and curve-blends with different scales of roundness [VaHo98].

Stage 3: Patch Layout. In this final stage, we decide on the exact layout of the quadrangles and we fit NURBS patches to the resulting surface pieces. We start with the longitudinal boundaries delimiting the curve-blends. This layout can be decided with the help of a mix of basis functions, as illustrated in Figure 4. Using the widths of the incident curve-blends, we can now decide on the types and boundaries of the vertex-blends. Once all boundaries of curve- and vertex-blends are determined, we can decompose them into quadrangles. Next, we decompose each region of the simplified Morse complex into quadrangles, making sure that neighboring quadrangles share entire edges. Finally, we decide on the number of control points to be used for each quadrangle and we fit NURBS patches to approximate the surface.

Acknowledgement

References

[ELZ02] H. EDELSBRUNNER, D. LETSCHER AND A. ZOMORODIAN. Topological persistence and simplification. *Discrete Comput. Geom.* **28** (2002), 511–533. 2

[Far97] G. FARIN. *Curves and Surfaces for CAGD. A Practical Guide.* Fourth edition, Academic Press, San Diego, California, 1997. 1

[Mil63] J. MILNOR. *Morse Theory.* Princeton Univ. Press, New Jersey, 1963. 2

[VaHo98] T. VÁRADY AND C. M. HOFFMANN. Vertex blending: problems and solutions. In *Mathematical Methods for Curves and Surfaces II*, 501–527, eds. M. Daehlen, T. Lyche and L. L. Schumaker, Vanderbilt Univ. Press, 1998. 3