Minimum Curvature Variation Curves, Networks, and Surfaces for Fair Free-Form Shape Design

by

Henry Packard Moreton

B.S. (University of New Hampshire) 1979
M.S. (University of New Hampshire) 1983

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Computer Science

in the

GRADUATE DIVISION

of the

UNIVERSITY of CALIFORNIA at BERKELEY

Committee in charge:

Professor Carlo H. Séquin
Professor Beresford Parlett
Professor Forest Baskett
Professor Lawrence A. Rowe

The dissertation of Henry Packard Moreton is approved:

Chair Date

Date

Date

Date

University of California at Berkeley

1992

Minimum Curvature Variation Curves, Networks, and Surfaces for Fair Free-Form Shape Design

by

Henry Packard Moreton

Abstract

Minimum Curvature Variation Curves, Networks, and Surfaces for Fair Free-Form Shape Design

by

Henry Packard Moreton

Doctor of Philosophy in Computer Science

University of California at Berkeley

Professor Carlo H. Séquin, Chair

Traditionally methods for the design of free-form curves and surfaces focus on achieving a specific level of inter-element continuity. These methods use a combination of heuristics and constructions to achieve an ultimate shape. Though shapes constructed using these methods are technically continuous, they have been shown to lack fairness, possessing undesirable blemishes such as bulges and wrinkles. Fairness is closely related to the smooth and minimal variation of curvature.

In this work we present a new technique for curve and surface design that combines a geometrically based specification with constrained optimization (minimization) of a fairness functional. The difficult problem of achieving inter-element continuity is solved simply by incorporating it into the minimization via appropriate penalty functions. Where traditional fairness measures are based on strain energy, we have developed a better measure of fairness; the variation of curvature. In addition to producing objects of clearly superior quality, minimizing the variation of curvature makes it trivial to model regular

shapes such as, circles and cyclides, a class of surface including: spheres, cylinders, cones, and tori.

In this thesis we introduce: curvature variation as a fairness metric, the minimum variation curve (MVC), the minimum variation network (MVN), and the minimum variation surface (MVS). MVC minimize the arc length integral of the square of the arc length derivative of curvature while interpolating a set of geometric constraints consisting of position, and optionally tangent direction and curvature. MVN minimize the same functional while interpolating a network of geometric constraints consisting of surface position, tangent plane, and surface curvatures. Finally, MVS are obtained by spanning the openings of the MVN while minimizing a surface functional that measures the variation of surface curvature.

We present the details of the techniques outlined above and describe the trade-offs between some alternative approaches. Solutions to difficult interpolation problems and comparisons with traditional methods are provided. Both demonstrate the superiority of curvature variation as a fairness metric and efficacy of optimization as a tool in shape design, albeit at significant computational cost.

Table of Contents

List of Figures

List of Symbols

$\hat{e}_1$	principle direction, maximum curvature
$\hat{e}_2$	principle direction, minimum curvature
κ_1	principle curvature, maximum
κ_2	principle curvature, minimum
κ_n	normal curvature
$\vec{\kappa}$	curvature vector
$\hat{l}$	unit vector pointing toward light source
$\hat{n}$	unit normal vector
$\hat{b}$	unit binormal vector
$\hat{t}$	unit tangent vector
ds	arc length differential
dA	area differential
$\int \frac{d\vec{\kappa}}{ds}^2 ds$	MVC functional
$\int \frac{d\kappa_1}{d\hat{e}_1}^2 + \frac{d\kappa_2}{d\hat{e}_2}^2 dA$	MVS fucntional
$\vec{S}_u(u, v)$	partial derivative of $\vec{S}(u, v)$ with respect to u
$\vec{S}_u$	partial derivative of $\vec{S}(u, v)$ with respect to u
$f'(t), f''(t), f'''(t), f^{(4)}(t)$	the first four derivatives of $f(t)$
$\vec{a} \times \vec{b}$	cross product of $\vec{a}$ and $\vec{b}$

G^n	*n*th order geometric continuity
C^n	*n*th order parametric continuity
$f(t)$	a scalar valued function
$\vec{C}(s)$	an arc length parameterized, vector valued function
$\vec{C}(u)$	a vector valued function of arbitrary regular parameterization
$\vec{S}(u, v)$	a bivariate vector valued function

Acknowledgments

The work presented in this thesis was supported and contributed to by many people in many ways, often unknowingly. I would like to recognize some of these people here. My thesis advisor, Carlo Séquin has provided much of the inspiration for this work through his infectious fascination with geometry. His always thorough review of my writings have been a great aid. I would also like to thank the other readers of my thesis. They have contributed greatly to both its english and its technical content. Beresford Parlett caught one major error, noting in red ink, “...this is why people ask for proofs!” Both Forest Baskett and Larry Rowe impressed me with their ability to catch errors many preceding eyes had missed. In his role as my manager, Forest has also been very generous with his support of my project. This support was initiated, at least in part, by several forward thinking people at Silicon Graphics Inc., Jim Clark, Ed McCracken, and Glen Mueller.

In addition to my readers, I had the advice of several valuable consultants. Jim Demmel provided advice concerning the nature of finite precision arithmetic. John Canny patiently helped me work through some of the problems I encountered in differential geometry. Jim Winget was generous with his time and his extensive knowledge of finite element analysis. Finally, Tony DeRose’s interest in my work helped me through that “nobody cares” period. His students, Stephen Mann and Mike Lounsbery continue to be valuable colleagues.

My friends and colleagues have been a great source of support and fun throughout the years. Unfortunately, it is impossible to name each and every person. Ziv Gigus played Mutt to my Jeff, or was it the other way ‘round? Always interested in listening over coffee at Roma. Garth Gibson, first squash partner, then Mystery host, and now CMU professor has been a good friend. Nina Amenta has been a great dinner partner, and even greater party host. Kathryn Crabtree, CS Division resident archeologist and pleasant lunch conversationalist, has been a lot of help with “the process”. Terry Lessard-Smith, Bob Miller and Liza Gabato have ready smiles and have always been willing to help with whatever problem was at hand. At Silicon Graphics, Melissa Anderson has helped keep me in contact with folks there and simply taken good care of me, finding office space, machines, and helping me deal with a company grown 6-fold in my absence.

The Sydeman clan has shown interest and caring: Bill Sydeman and Catherine Madonia, Jay Sydeman, Hope Millholland, Michelle Sydeman, and Ann Sydeman. It is Ann in particular, that I would like to thank for supporting me and staying by me on the roller coaster ride that is a thesis in the making.

Finally, my parents have given me everything, the foundation to work from and unwavering support.

1
Introduction

The field of computer aided geometric design (CAGD) has formed out of the need to design and model curved forms. Shapes are typically represented in a piecewise fashion, composed of primitive elements smoothly joined together to form a larger, more complex whole. Traditionally methods for the design of free-form curves and surfaces focus on achieving a specific level of inter-element continuity. These methods use a combination of heuristics and constructions to achieve an ultimate shape. Though shapes constructed using these methodologies are technically continuous, they have been shown to be of poor quality (lacking fairness), possessing undesirable blemishes such as bulges and undulations.

In this work we present a new technique for curve and surface design that combines a geometrically based specification with constrained optimization (minimization) of a fairness functional. The difficult problem of achieving inter-element continuity is solved simply by incorporating it into the minimization via appropriate penalty functions. Where traditional fairness measures are based on strain energy, we have developed an alternative measure of fairness: the variation of curvature. In addition to producing objects of superior quality, minimizing the variation of curvature makes it trivial to model regular shapes such as, circles and cyclides, a class of surface including: spheres, cylinders, cones, and tori.

In this thesis we introduce: curvature variation as a fairness metric, the minimum variation curve (MVC), the minimum variation network (MVN), and the minimum variation surface (MVS). These three forms are computed to satisfy a set of geometric interpolation conditions while minimizing a fairness functional that measures the variation of curvature.

While interpolating an ordered set of curve specifications, MVC minimize the arc length integral of the square of the arc length derivative of curvature:

$$\int \frac{d\vec{\kappa}}{ds}^2 ds. \qquad \textit{the MVC functional} \qquad \textbf{(1.1)}$$

MVC specifications consist of curve position, and optionally, curve tangent and curvature. Similarly, MVN interpolate a network of surface specifications while minimizing the MVC functional along the arcs of the network. These surface specifications consist of position, and optionally, surface tangent and principle curvatures. Finally, MVS interpolate the same network of specifications while minimizing the area integral of the sum of the square of the derivatives of the principle curvatures:

$$\int \frac{d\kappa_1}{d\hat{e}_1}^2 + \frac{d\kappa_2}{d\hat{e}_2}^2 dA. \qquad \textit{the MVS functional} \qquad \textbf{(1.2)}$$

The resulting models accurately reflect their specifications and are free of unwanted wrinkles, bulges, and ripples. When the given constraints indicate and/or permit, the resulting surfaces take on the desirable shapes of spheres, cylinders, cones, and tori. Specification of a desired shape is straightforward, allowing simple or complex shapes to be described easily and compactly. For example, a “suitcase corner,” the blend of three quarter cylinders of differing radii, is formed by specifying just six sets of constraints (Fig. 1.1) plus three sets of floating vertices (∘).

This work presents in detail new techniques to compute minimum variation curves, networks, and surfaces. Numerous examples demonstrating the efficacy of curvature variation as a fairness metric are provided.

1.1 Overview

In this thesis, we begin with a brief outline of the techniques used to compute the various minimum variation forms. In Chapter 2, we define the terms and properties used in discussing the CAGD of curves and surfaces. In Chapter 3, we present an overview of previous work on curve, network, and surface design and computation. We describe the bulk of our work in Chapters 4, 5, and 6. These chapters are devoted to the details of the computation of minimum variation curves, networks, and surfaces respectively. Though these chapters may be read individually, each chapter builds on its predecessors. The techniques used in the computation of MVC are reapplied in the computation of MVN, and MVN themselves are used in the computation of MVS. In each of these chapters we evaluate minimum variation shapes and compare them with their contemporary

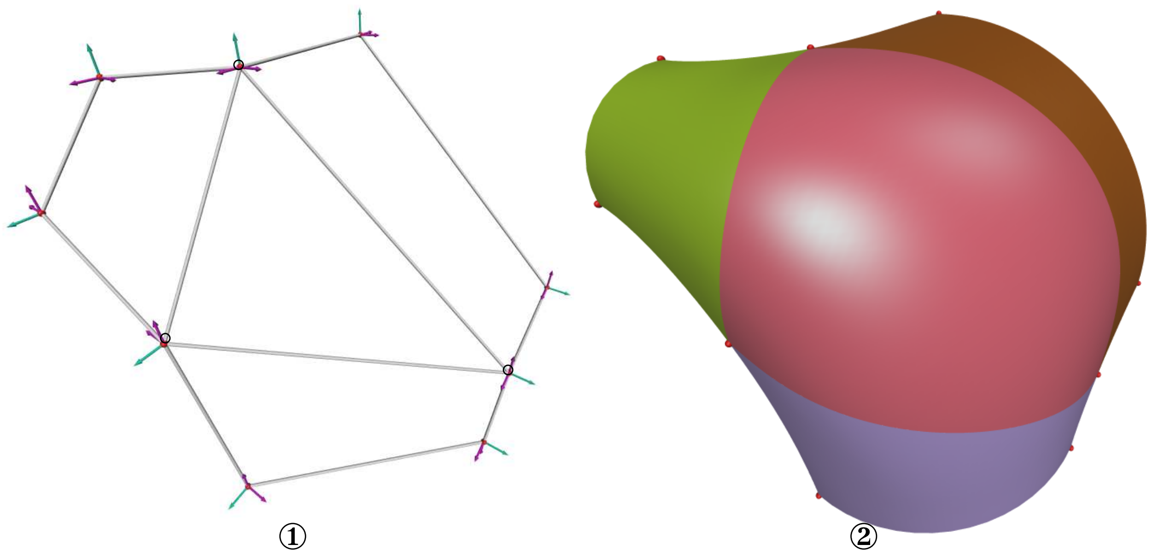

Figure 1.1. A suitcase corner.
① specification with surface normal and curvature constraints. ② the resulting blend.

counterparts. We conclude the thesis with comments about the utility of minimization techniques, the value of curvature variation as a fairness metric, and directions for further research.

1.2 Minimum Variation Curves

In CAD applications, curve design has a potentially conflicting set of requirements. Some applications demand free-form curves, along with regular curves such as circular arcs. In some instances curves must meet a set of exact positional, tangent, and/or curvature constraints. Also in general, a high degree of "fairness" is demanded of all curves. The concept of fairness is typically associated with the curvature characteristics of a curve; a fair curve has smoothly varying curvature, with as few inflections and curvature extrema as possible. The minimum variation curve (MVC) satisfies all these needs.

We cast the problem of computing the MVC as a nonlinear optimization / finite element problem. The curve is broken into a series of quintic polynomial elements constructed to satisfy the given geometric constraints and join with G^2 continuity. The MVC functional is minimized using a gradient descent optimization procedure. A heuristically chosen starting curve greatly accelerates convergence towards minimum variation.

1.3 Minimum Variation Networks

Many surface modeling and data interpolation schemes use a mesh or network of curves as a key component in the construction of a smooth surface [126]. The minimum variation network (MVN) is a G^2 network composed of fair curves (MVCs) that provides an excellent frame on which to build a smoothly curved surface. In fact the MVN is used in the computation of a minimum variation surface (MVS). The MVN acts either as a fixed framework on which to build an MVS or as part of the initialization step in the construction of an MVS. The MVN is computed using the MVC functional (1.1) and MVC optimization techniques with the G^2 curve continuity constructions replaced by G^2 surface continuity constructions. Also, the curve-based specifications are replaced by a network of surface specifications. Heuristics based on the geometry of the constraint network are used to establish a starting point for the optimization.

1.4 Minimum Variation Surfaces

In the computer aided design of curved surfaces there is a wide range of requirements. While it is necessary to model regular shapes such as cylinders, cones, tori, and spheres, it is also important that free-form shapes can be modeled with ease. Often, it is also necessary that surfaces meet a set of exact positional, tangent, and/or curvature constraints. In all cases, surface fairness is of great importance. Like the fairness of curves, surface fairness is related to the variation of curvature across a surface; a fair surface has smoothly varying curvature. These requirements are met by MVS.

As with the computation of the MVC, we compute MVS using optimization techniques to minimize the MVS functional while using constructive methods to satisfy a set of geometric interpolatory constraints. Unlike the MVC computation, MVS inter-element continuity is imposed via penalty functions. We have developed penalty functions for imposing both G^1 and G^2 continuity.

We treat the problem of creating an MVS interpolating a collection of geometric constraints as one of scattered data interpolation. The interpolation problem is broken into three steps (Fig. 1.2); 1) connectivity definition, 2) curve network computation, 3) patch blending. In accordance with the topological type of the desired surface, the geometric constraints are first connected into a network of straight edges. Next, an MVN is calculated for the network of constraints. Finally, an MVS is computed, interpolating the MVN with at least G^1 continuity. In a first approach, the boundaries of the MVS patches are fixed, interpolating the previously constructed curve network. Alternatively, the surface calculation may use the MVN as a starting point and modify its geometry during surface calculation. The latter approach yields even smoother surfaces, but at a

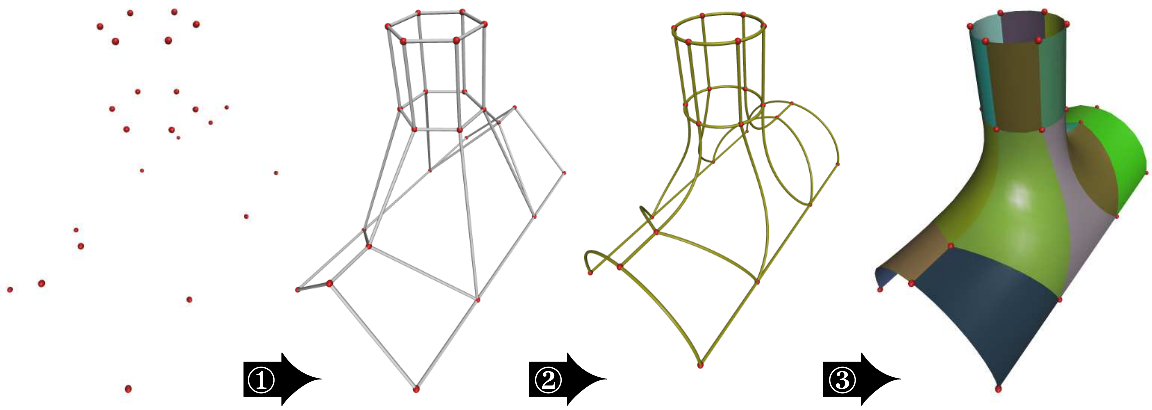

Figure 1.2. The blend of two pipes.

Pipes are blended in three steps: ① The connectivity of the constraints is established. ② Smooth curves are fit to the constraints. ③ Surface patches are fit to the curve network.

substantially higher computational expense. The higher quality surfaces result because the curves of an MVN resulting from a given constraint set do not always lie in the MVS resulting from the same set of constraints.

During the modeling process, the connectivity of the geometrical constraints is typically established as a natural outgrowth of the design process. The techniques described here are also amenable to true scattered data interpolation, in which case connectivity must first be derived with some other method, possibly based on some minimal triangulation on the data points.

Our system is based on triangular and quadrilateral patches. All constraints are located at corners of these patches. Additional vertices and edges may be added to a network of constraints so that it has only three- and four-sided openings. These additional vertices are not constraints and are appropriately positioned by the curve network computation and patch blending phases of the construction.

Before proceeding with a detailed discussion of these techniques, we must establish a common set of properties and terms, and review previous attempts at solving these difficult design problems.