EUROGRAPHICS '93 / R. J. Hubbold and R. Juan (Guest Editors), Blackwell Publishers
© Eurographics Association, 1993

Volume 12, (1993), number 3

Scale-Invariant Minimum-Cost Curves: Fair and Robust Design Implements

Henry P. Moreton and Carlo H. Séquin*

Computer Science Division,

Department of Electrical Engineering and Computer Science,

University of California, Berkeley, CA 94720,

USA

Abstract

Four functionals for the computation of minimum cost curves are compared. Minimization of these functionals result in the widely studied Minimum Energy Curve (MEC), the recently introduced Minimum Variation Curve (MVC), and their scale-invariant counterparts, (SI-MEC, SI-MVC). We compare the stability and fairness of these curves using a variety of simple interpolation problems. Previously, we have shown MVC to possess superior fairness. In this paper we show that while MVC have fairness and stability superior to MEC they are still not stable in all configurations. We introduce the SI-MVC as a stable alternative to the MVC. Like the MVC, circular and helical arcs are optimal shapes for the SI-MVC.

Additionally, the application of scale invariance to functional design allows us to investigate locally optimal curves whose shapes are dictated solely by their topology, free of any external interpolation or arc length constraints.

1. Introduction

The search for the "best" curve satisfying a given set of geometric constraints is open-ended, strongly dependent on the requirements of a particular application area, and typically complicated by poor, subjective, and application dependent definitions of "best". However, the quality or *fairness* of curves independent of application has been studied extensively and has been shown to be closely related to how little and how smoothly a curve bends. Despite this general definition of quality, when the need for perfection **is** coupled with other requirements, such as robustness, and minimum acceptable speed, a plethora of possible methods result.

With the introduction and practicality of global optimization methods, the search narrows to choosing a quantitative measure of curve quality. In terms of optimization this amounts to choosing a particular functional to minimize. The curvature at a point on a curve is a measure of how much the curve is bending at that point. In order to minimize the amount that a curve bends, we therefore minimize the arc length integral of the square of curvature. Bernoulli [4] studied the curves resulting from this minimization, naming them *elastica—an* idealized thin beam passing through a series of frictionless supports. Because the functional is proportional to strain energy, the curves have become known as Minimum Energy Curves (MEC). Such curves and the many approximations of the functional by more efficient, and typically linearized representations are the basis for many spline curves used in computer-aided design (CAD) systems.

---

*Current address: Silicon Graphics Inc., MS 7U-005, 2011 N. Shoreline Blvd., Mountain View, CA 94043.

The MEC has several shortcomings: it is not always stable; it does not readily form regular shapes; and it only permits the specification of position and tangent constraints. In response to these factors, we developed the minimum variation curve (MVC) [5]. Instead of minimizing curvature, the MVC minimizes the variation of curvature. Quantitatively curvature variation is equal to the arc length integral of the square of the arc length derivative of curvature. This new functional produces fairer and more pleasing shapes than the MEC. Whereas MEC only support tangent constraints, MVC permit the specification of curvature constraints. MVC readily form circles or circular arcs since these shapes exhibit a total cost of zero. This optimal shape is important because of the circle's prominent role in manufacturing.

Minimum variation curves are not only fairer than minimum energy curves, they also exhibit greater stability. It is not uncommon for portions of an MEC defined by an inadequate set of positional constraints to run away to infinity while reducing the functional. However, as we will show in this paper, there are situations where even an MVC will escape. Additionally, we wish to study generically minimal shapes, which are defined solely by topological constraints and not by any positional or dimensional constraints. The MVC is unstable in most of the unconstrained cases. It is only in those cases that end in a circular shape that the standard MVC remains finite. The reason for the instability of MEC and MVC is that, while the shape of a curve is independent of uniform scaling of the defining constraints, the value of thefunctional is reduced when the scale is increased.

To counteract the dependence on scaling we developed scale-invariant functionals leading to the scale-invariant minimum energy curve (SI-MEC) and the scale-invariant minimum variation curve (SI-MVC). Scale-invariant curves are obtained by multiplying the cost integral with suitable powers of the curve's arc length so as to counteract the decrease in the curvature or its derivative when the shape is scaled up uniformly. As a consequence, when uniform scaling is applied to constraints, the functional does not change value. The curves resulting from these functionals are considerably more stable.

A through review of previous work on minimum energy curves, minimum variation curves, and optimization for curve design may be found in [5,7].

2. Scale-Invariant Cost Functionals

Curve fairness and smoothness are related to the geometric character of a curve. Geometrically, smoothness is defined as the continuous differentiability of a curve. Fairness refers to the distribution of curvature along the curve. Classically, the arc length integral of the square of curvature is minimized to maximize fairness. Since this integral is proportional to strain energy, minimum energy curves (MEC) result.

$$\int \vec{\kappa}^2 ds$$

Recently, Moreton and Séquin [5] introduced a curve minimizing the variation of curvature. Specifically, these minimum variation curves (MVC) minimize the arc length integral of the square of the arc length derivative of curvature.

$$\int \left(\frac{d\vec{\kappa}}{ds}\right)^2 ds$$

MVC have fairness superior to that of MEC, with a very smooth and lean appearance. Additionally, MVC have greater continuity through interpolation points, allow for curvature specifications, form circular and helical arcs when constraints permit, and are more stable than MEC. A simple example of an MEC's instability is the curve exhaustively studied by Horn [2], defined by parallel tangential constraints at its two ends. If the two tangential constraints are opened even infinitesimally, the MEC will expand forming an infinitely large loop, while the MVC still assumes a circular arc (Fig. 1).

However, in some situations even MVC become unstable and run away. This happens when the curve has an inflection point between constraints and turns through too large an angle. The MVC also turns out to be

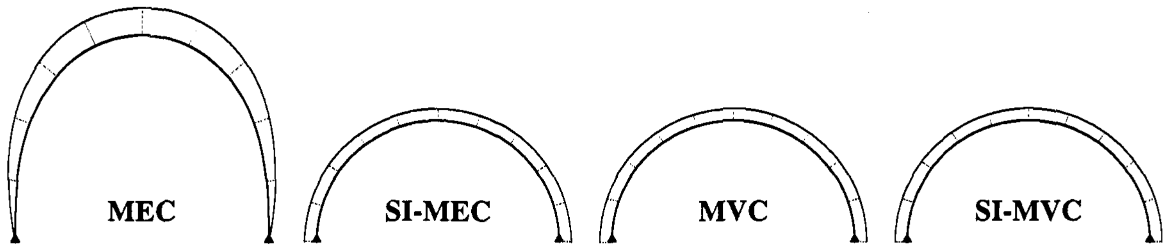

Figure 1. *Horn's test case interpolates two point-tangent constraints, all but the MEC form semi-circles. The width of the fin attached to the curve is proportional to curvature (see also section 4).*

unstable for totally unconstrained curves in the topology of a figure-8. In all cases, the force that drives the curve towards unbounded size originates from the reduction in the value of the cost functional as the scale of the curve is increased. The arc length of the curve grows linearly with scale, and its curvature decreases reciprocally. The MEC functional integrates the square of curvature, and thus its value decreases in direct proportion to the scale factor. In the MVC functional, the derivative of curvature with respect to arc length falls as the square of the scale, and it appears squared in the functional. Thus, overall, the functional decreases in proportion to the cube of the scale factor.

To offset these tendencies to run-away and to obtain scale-invariant functionals, we multiply the MEC integral by a factor that increases with the scale factor, the arc length. We multiply the MVC integral by the cube of the arc length to counteract its growth tendency.

This leads to two scale-invariant functionals which are the subject of this paper. The scale-invariant MEC is found by globally minimizing the expression:

$$\left(\int ds\right)\left(\int \vec{\kappa}^2 ds\right) \qquad \textbf{SI-MEC}$$

subject to the given external constraints; and the scale-invariant MVC is the result of minimizing the product:

$$\left(\int ds\right)^3 \left(\int \left(\frac{d\vec{\kappa}}{ds}\right)^2 ds\right). \qquad \textbf{SI-MVC}$$

Though both SI-MEC and SI-MVC are stable, we feel SI-MVC are superior because they inherit the positive qualities of the MVC: higher order continuity, greater fairness, and simple capture of circular and helical arcs. We have included the SI-MEC in this study to further understand the effect of scale invariance on the stability and fairness of curves resulting from a particular functional.

3. Computation of SI-MEC and SI-MVC*

The curves resulting from any of the four functionals do not have known closed form representations. We cast their computation as a nonlinear optimization / finite element problem whose solution is an accurate approximation. Curve representation plays an important role in making these curves useful to applications. The curve is broken into a series of parametric polynomial elements that satisfy the given geometric constraints, and join

*Portions of this section are taken from [7].

with G^2 continuity and G^1 continuity for MVC and MEC, respectively. Among polynomials, quintic elements are most suitable because they have sufficient descriptive power to simultaneously satisfy arbitrary constraints on position, tangent direction, and curvature.

3.1 Curve Representation

There are several possible representations for a quintic polynomial element (e.g. B-spline, Bezier). We chose the Hermite form because of the ease with which the geometric specifications can be mapped to the defining parameters of the Hermite segments. Also, this form is easily converted into other polynomial representations that are typically used in geometric modeling systems. Quintic Hermite curves are specified by the position of the endpoints and by the first two parametric derivatives at these locations. In vector notation this can be expressed as:

$$\vec{C}(u) = \begin{bmatrix} \vec{C}(0) \\ \vec{C}'(0) \\ \vec{C}''(0) \\ \vec{C}''(1) \\ \vec{C}'(1) \\ \vec{C}(1) \end{bmatrix}^T \cdot \begin{bmatrix} H_0(u) \\ H_1(u) \\ H_2(u) \\ H_3(u) \\ H_4(u) \\ H_5(u) \end{bmatrix} = \vec{C}(0)H_0(u) + \vec{C}'(0)H_1(u) + \ldots + \vec{C}(1)H_5(u)$$

where $H_i(u)$ are quintic blending functions. The computation of $H_i(u)$ is described in [1]. The mapping from the geometric to parametric description of the curve is carried out using the following equations

$$\vec{P}_i = \vec{p}_i$$

$$\vec{P}'_i = m_i\hat{t}_i$$

$$\vec{P}''_i = m_i^2\vec{\kappa}_i + \alpha_i m_i\hat{t}_i$$

where $\vec{p}_i$, $\hat{t}_i$, $\vec{\kappa}_i$ are the position, tangent direction, and curvature vector at one end of the curve; m_i **is** the first derivative magnitude and a, completes the relationship between curvature and the second derivative. During the minimization the scalar m_i must be constrained to be positive:this **is** because if m_i became negative then P'_i would reverse direction. We impose this constraint by using mi2 rather than m_i in equations and **.**

A curve is made up of a sequence of vertices connected by quintic elements. G^2 curves are pieced together from these elements by sharing geometric specifications at the common points. Data structures associated with the vertices hold the point, tangent, and curvature information, while data structures associated with the elements hold the parameters $\alpha_{i,0}$, $\alpha_{i,1}$ and $m_{i,0}$, $m_{i,1}$ (Fig. 2). Each element is defined by three data structures, a pair of structures associated with the vertices at the element's endpoints and one structure associated with the element itself. By distributing the curve/element specification in this way, adjacent elements share vertex structures and are guaranteed to meet with curvature continuity. Note that discontinuities can be introduced by giving adjacent elements independent geometric specifications.

As we have described thus far, a single quintic Hermite segment or *element* is placed between every two sets of constraints. Because of the limited descriptive power of polynomial elements, a single element can only approximate the ideal minimum variation curve. To improve the approximation, multiple elements can be inserted between constraints. In practice, a single element per constraint pair is normally sufficient.

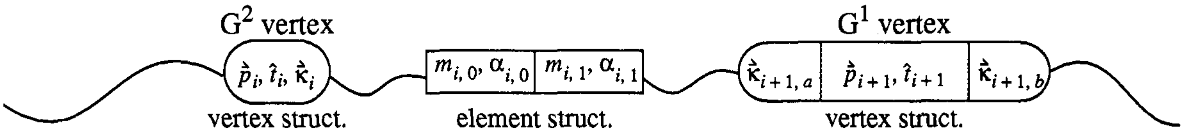

Figure 2. *Schematic View of Curve Representation. On the left is a G^2 joint where incident curves share the full geometric specification. On the right is a G^1 joint where incident curves have independent curvatures.*

3.2 Parametric Functionals

The MEC, MVC, SI-MEC, and SI-MVC functionals are defined in terms of integrals of a function over arc length. To evaluate these functionals and their gradients in the context of the parametric piece-wise polynomial curves described in section 3.1, the functionals must be converted to a parametric polynomial form and evaluated in a piecewise fashion. In this section we describe the transformation of the SI-MVC functional from in arc length based definition to its parametric form. The transformation of the other functionals is analogous. The value of the functional for the curve as a whole is computed as the sum of the values of the functional for each element. In the first conversion step the arc length based definition

$$\left(\int_0^l ds\right)^3 \left(\int_0^l \frac{d\vec{\kappa}}{ds}^2 ds\right) \qquad \textbf{SI-MVC}$$

is changed to the product of integrals of a function of the curve $\vec{C}(t)$ parameterized by t of the form

$$\int_\alpha^\beta f(\vec{C}(t))dt \ .$$

The bounds, α, β, of the integral are set to 0 and 1, since the Hermite representation is parameterized with t varying from 0 to 1. The differential with respect to s is converted to a differential in **t** Since

$$\frac{dt}{ds} = \frac{1}{\|\vec{C}'(t)\|} \text{ then } ds = \|\vec{C}'(t)\|dt$$

where $\quad \|\vec{C}'(t)\| = (\vec{C}'(t) \cdot \vec{C}'(t))^{1/2}$.

Next, the derivative of curvature with respect to arc length $\frac{d\vec{\kappa}}{ds}$ transforms to: $\frac{d\vec{\kappa}}{dt}\frac{dt}{ds}$. These two steps yield:

$$\int_0^l \frac{d\vec{\kappa}}{ds}^2 ds = \int_0^1 \left(\frac{d\kappa}{dt}\frac{dt}{ds}\right)^2 \|\vec{C}'(t)\|dt \qquad \int_0^l \frac{d\vec{\kappa}}{ds}^2 ds = \int_0^1 \frac{d\vec{\kappa}}{dt}^2 \frac{1}{\|\vec{C}'(t)\|}dt.$$

Lastly we find the expression for $\frac{d\vec{\kappa}}{dt}$ in terms of the parametric derivatives of the curve, $\vec{C}(t)$. The expression for curvature is

$$\vec{\kappa}(t) = \frac{\vec{C}'(t) \times \vec{C}''(t)}{\left\| \vec{C}'(t) \right\|^3}$$

Taking the derivative with respect to t yields

$$\frac{d\kappa}{dt} = \frac{v\,du - u\,dv}{v^2}$$

$$u = \vec{C}'(t) \times \vec{C}''(t), \qquad du = \vec{C}'(t) \times \vec{C}'''(t), \qquad \text{n.b. } \vec{C}''(t) \times \vec{C}''(t) = 0,$$

where

$$v = \left\| \vec{C}'(t) \right\|^3 = (\vec{C}'(t) \cdot \vec{C}'(t))^{3/2}$$

$$dv = 3\left\| \vec{C}'(t) \right\| (\vec{C}'(t) \cdot \vec{C}''(t))$$

3.3 Optimization Procedure

Cast as a multi-dimensional optimization problem, the curve is represented by a point in $\Re^n$ corresponding to its n degrees of freedom. The functional undergoing minimization is expressed as $f(x), x \in \Re^n$. Standard optimization techniques are used to minimize $f(z)$. Starting from a heuristically established point in $\Re^n$, Polak-Ribiere conjugate gradient descent [8] is used to traverse the space while reducing the objective function and ultimately arriving at a minimum.

4. Examples

In this section we compare the two new scale-invariant cost functions with their scale-dependent forerunners on a few relevant test cases that will bring out their salient differences most clearly. We demonstrate the different behavior of the four types of curves on the same external constraints and compare the resulting overall cost values for the four different curves according to all four different cost functionals. We also provide plots of the curves with curvature profiles attached; curvature profiles are an alternative to curvature plots. The profile width is proportional to curvature, normalized by arc length; curvature is constant when the profile's edge and the curve are parallel. Constraints are indicated by circles ● and triangles ▲; circles indicate positional constraints, and triangles indicate position-tangent constraints.

The quantitative measures provided by the four functionals, combined with the qualitative assessment of curvature distribution provide insight into the quality and robustness of the different curves. In the tables of functional values, grayed out entries refer to unstable curves. These values reflect the curve shape shown, resulting from stopping the optimization after a small number of iterations. The tabulated numbers are from optimizations terminated after the infinity norm of the gradient was reduced below 10^{-3}, the last couple of digits are not significant.

4.1 Three Points

As we have already mentioned MVC have exhibit greater continuity than MEC. As both curves pass through interpolatory constraints they exhibit discontinuities. Figure 3 illustrates curves passing through three position-tangent constraints. Note that both the MEC and SI-MEC have curvature discontinuities at the con-

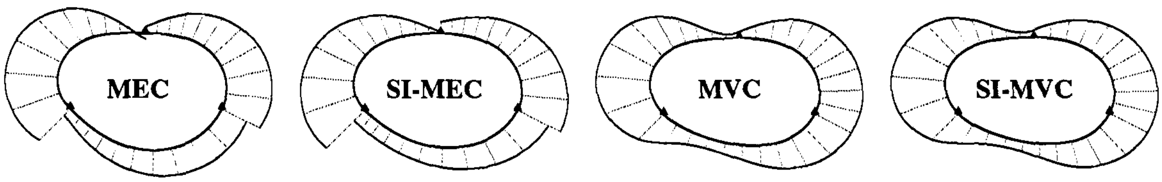

Curve Type	$\int ds \equiv L$	$\int \vec{\kappa}^2 ds$	$L \int \vec{\kappa}^2 ds$	$\int \left(\frac{d\vec{\kappa}}{ds}\right)^2 ds$	$L^3 \int \left(\frac{d\vec{\kappa}}{ds}\right)^2 ds$
MEC	12.56333	8.14509	49.03337	5.83229	1272.40928
SI-MEC	12.14680	8.24163	48.38989	4.20937	852.00564
MVC	11.89730	8.44883	49.88922	9.11292	1876.24165
SI-MVC	11.74763	8.51627	49.94979	9.19264	1854.78463

Figure 3. *Minimum energy curves specified with both position and tangent constraints have curvature discontinuities at the constraint points. Similarly, minimum variation curves have discontinuities of the derivative of curvature at points where curvature constraints are specified.*

straint points, while the MVC are G^3 continuous. A curvature constraint produces a G^3 discontinuity in an MVC. It is not possible to apply a curvature constraint to an MEC.

Three non-collinear points define a circle, and when interpolating three point constraints, a circle is assumed by all but the MEC. The MVC and SI-MVC naturally form circles since circles have constant curvature. The SI-MEC forms a circle since it naturally forms a circle when left completely unconstrained. Finally, if the triangle formed by the three positional constraints deviates too much from equilateral, the MEC will not even stay confined, but will escape through the largest side. Four points, if placed on the corners of a rectangle, are able to contain even the MEC; it will then assume a symmetrical shape that in the limit forms two copies of Horn's curve (Fig. 1), as the defining rectangle takes on more extreme aspect ratios. For these four constraints, the other three functionals continue to lead to circles.

4.2 Woodford's Data

This example is a simple interpolation problem first used by Woodford [9] in his development of an algorithm for computing a discrete MEC. We have added a natural spline to the suite of curves for comparison. The natural spline is a linearized approximation of an MEC and is used widely in CAD packages. We see the MVC and Sl-MVC have curvature distribution superior to the natural spline, MEC, and SI-MEC (Fig. 4). Also note that whereas these energy-based curves have zero curvature at their endpoints, the curves based on curvature variation have constant curvature.

4.3 Figure-8 Curves

We specify a simple figure-8 shape by 3 points along its symmetry axis (Fig. 5). Released after initialization, asymmetric but reasonable initial tangent directions are specified at these points to help define the desired curve. This leads to an asymmetrical initial condition from which both the MEC and the MVC can run away. By turning about the central point, these figures-8 shapes can grow to infinity, reducing curvature, while still passing through the three given constraint points. For the SI-MEC and SI-MVC where this force to lower cur-

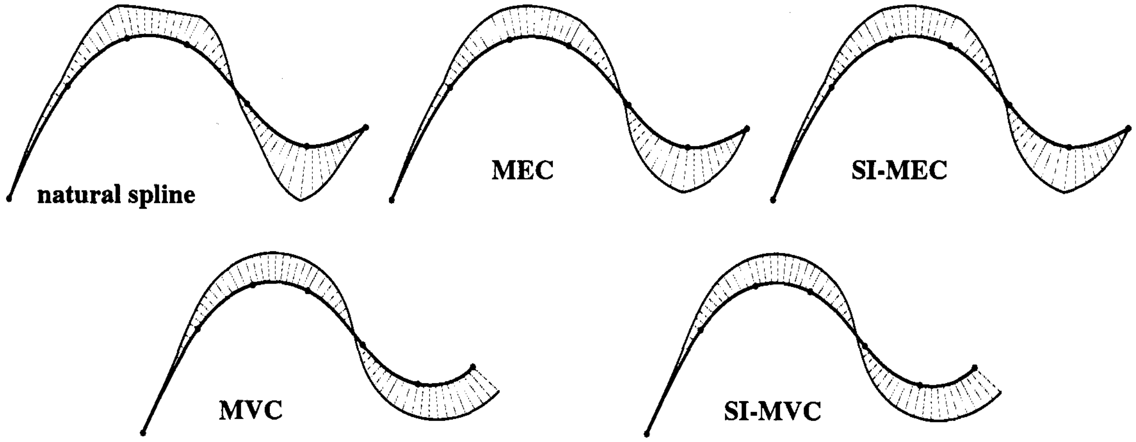

Curve Type	$\int ds \equiv L$	$\int \vec{\kappa}^2 ds$	$L \int \vec{\kappa}^2 ds$	$\int (\frac{d\vec{\kappa}}{ds})^2 ds$	$L^3 \int (\frac{d\vec{\kappa}}{ds})^2 ds$
MEC	12.15348	2.24767	18.51844	2.72641	1524.77207
SI-MEC	12.14947	2.24785	18.51680	2.76989	1548.31233
MVC	12.18984	2.55115	21.06973	1.15379	649.96890
SI-MVC	12.19133	2.54786	21.04185	1.15377	649.89717

Figure 4. *The curves fit to Woodford's data set illustrate the superior fairness and greater continuity of MVC functionals. Also note that, whereas the MEC start and finish with zero curvature, the MVC have constant curvature at their end points.*

vature is offset by an arc length scale factor, the end results are figure-8 shapes that closely resemble a lemniscate and which are only slightly larger than the initial starting shapes. Under symmetric starting conditions, specified by three position-tangent constraints, all four functionals lead to *lemnoid* shapes aligned with the original symmetry axes.

5. Unconstrained Curves

By comparing functionals without using a specific set of constraints we are certain to be free of any bias in the choice of constraints; this approach is even more relevant when comparing functionals for surface design. SI-functionals permit us to study the optimization of *closed* curves that are free of any external constraints beyond the initial definition of their topology. For every given curve shape, we can find the closest local optimum of the corresponding SI-MEC or SI-MVC, and thus determine what constitutes the fairest representative of that class of "topologically equivalent" shapes. In this context, shapes are "equivalent" if they can be smoothly deformed into one another without ever assuming singularities such as kinks and cusps; this is the case for all curves with the same *turning number.* The turning number is the integral over the change of the tangent direction angle taken for a whole curve traversal; it is zero for a smooth figure-8 and ± 1 for single-turn circles.

As we will see, it is possible to deform a closed curve sufficiently — without introducing any singularities — so that a gradient descent optimization procedure will travel down to a different local minimum. The SI-functionals then allow us to compare different locally minimal shapes in an objective manner, and to determine which one is the globally minimal shape relative to the particular cost functional. This opens up the possibility

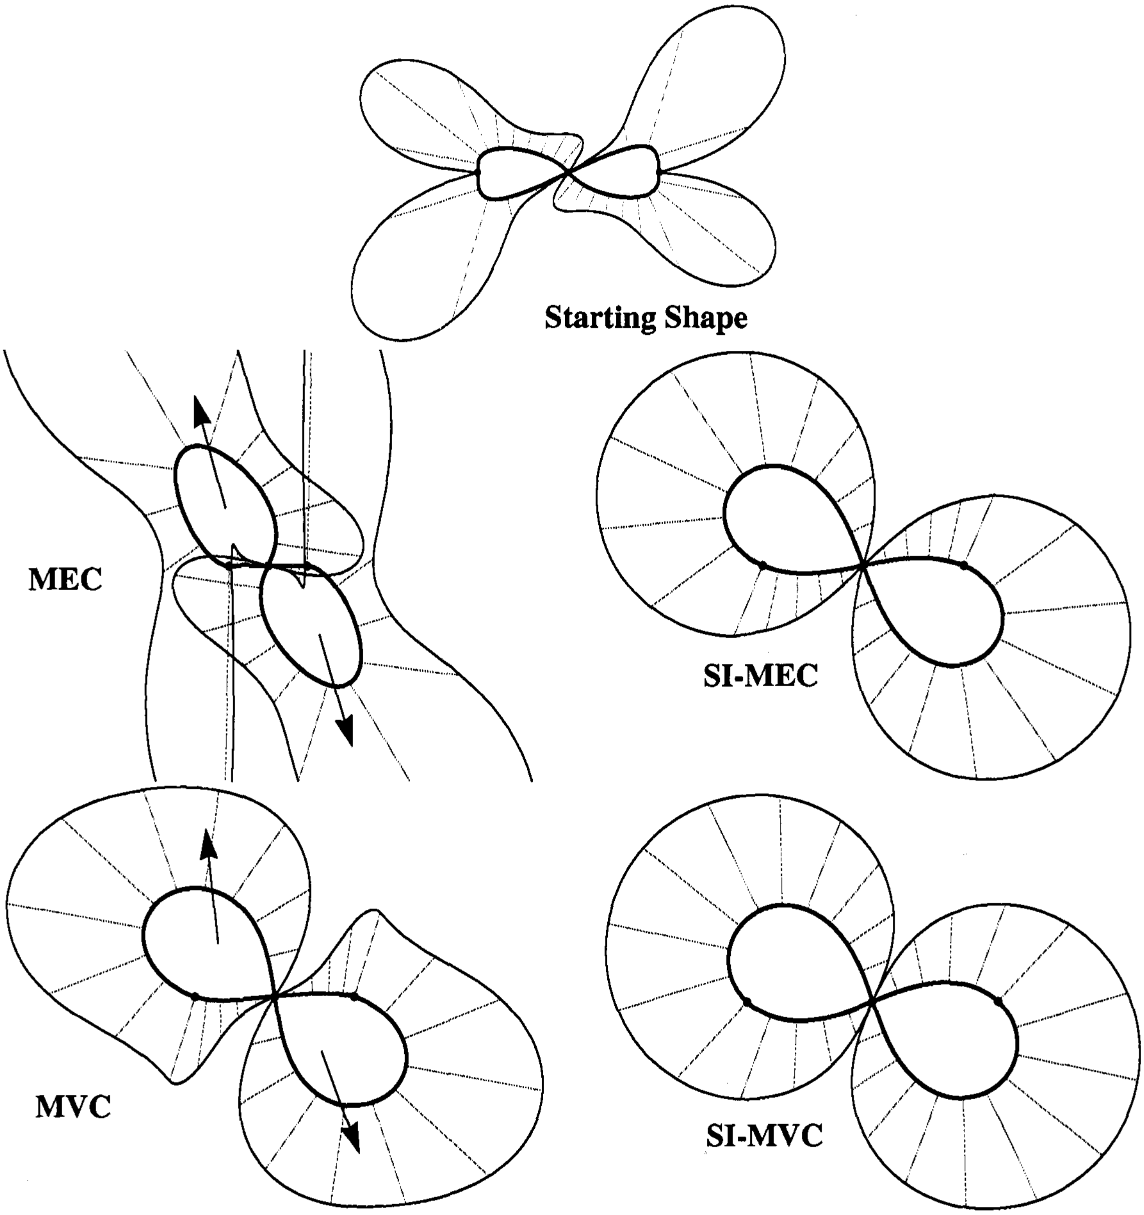

Curve Type	$\int ds \equiv L$	$\int \vec{\kappa}^2 ds$	$L \int \vec{\kappa}^2 ds$	$\int \left(\frac{d\vec{\kappa}}{ds}\right)^2 ds$	$L^3 \int \left(\frac{d\vec{\kappa}}{ds}\right)^2 ds$
Initial	42.53726	16.38924	164.79680	220.53971	224210.75847
MEC	565.82951	5.20216	188.47815	36.35917	1729200.14077
SI-MEC	84.28997	7.23245	112.44187	1.32306	4971.73079
MVC	143.26783	5.79170	114.45129	1.17128	9038.68746
SI-MVC	56.59292	8.76329	113.75806	2.05773	4501.25246

Figure 5. *Starting with a lop-sided initial shape, both the MEC and MVC escape to infinity. Their scale-invariant relatives remain stable.*

to search for the "absolutely best" shapes for each class of curves represented by different turning numbers. The answer to this quest turns out to be surprisingly simple and elegant.

For the scale-invariant curves, the constraints used to define the turning number of the curve are removed immediately after the initial curve has been constructed. Then, as the optimization proceeds without any external constraints, the "natural" unconstrained shape of the curve is revealed.

5.1 Elementary SI-MVC Shapes

The lemnoid shape and circles traced one or more times in either direction can represent curves with any possible turning number. These shapes form the global minima for our scale-invariant cost functions. Non-self-intersecting closed curves open up into circles. Curves with self-intersections but without inflection points end up in multiply traced circles as for instance the cloverleaf. (Fig. 6). However, there are some distinct, interesting local minima which frequently get assumed by complicated self-intersecting initial curves.

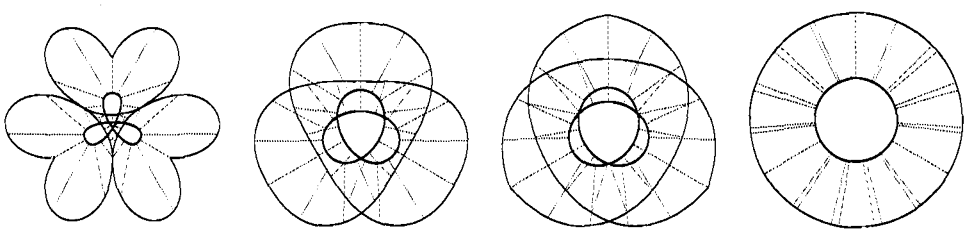

Figure 6. *Optimizing a cloverleaf curve into a double circle.*

5.2 Local and Global Minima

SI-MVC such as the one shown in Figure 7 can either open up into a circle (Fig. 7aa) or snap shut into a lemnoid with inner circle (Fig. 7bb); it depends on the exact initial conditions. The two final states are separated by a cost barrier in optimization cost space. The two initial shapes shown in Figure 7a & b lie on either side of this cost ridge and have monotonically decreasing cost paths to either of these two possible end states. Starting from Figure 7aa, the SI-MEC opens up into a circle, thus there is no SI-MEC barrier between these two shapes.

One can understand why the lemnoid with the inner circle forms a strong local minimum for the SI-MVC functional. Since the circle has constant curvature it only contributes it arc length to the scale invariant curvature variation functional. To unwind this shape into a simple circle would force some deformations that would increase the variation of curvature. One can also see that we can insert another circle into the second lobe of the lemnoid and get another stable minimum.

From the above one might suspect that there is a plethora of different shapes corresponding to local minima. This does not seem to be the case. We exclude from this discussion certain observed termination states with cusps. They are very sensitive to initial conditions and the curve representation and integration parameters, and are the result of numerical artifacts; these occurrences can be made to disappear by selecting a different number of integration points along each curve segment (e.g. 21 instead of 20). Aside from these spurious shapes, we have not yet found other stable local minima for the SI-MVC functional than the lemnoid with inner circles.

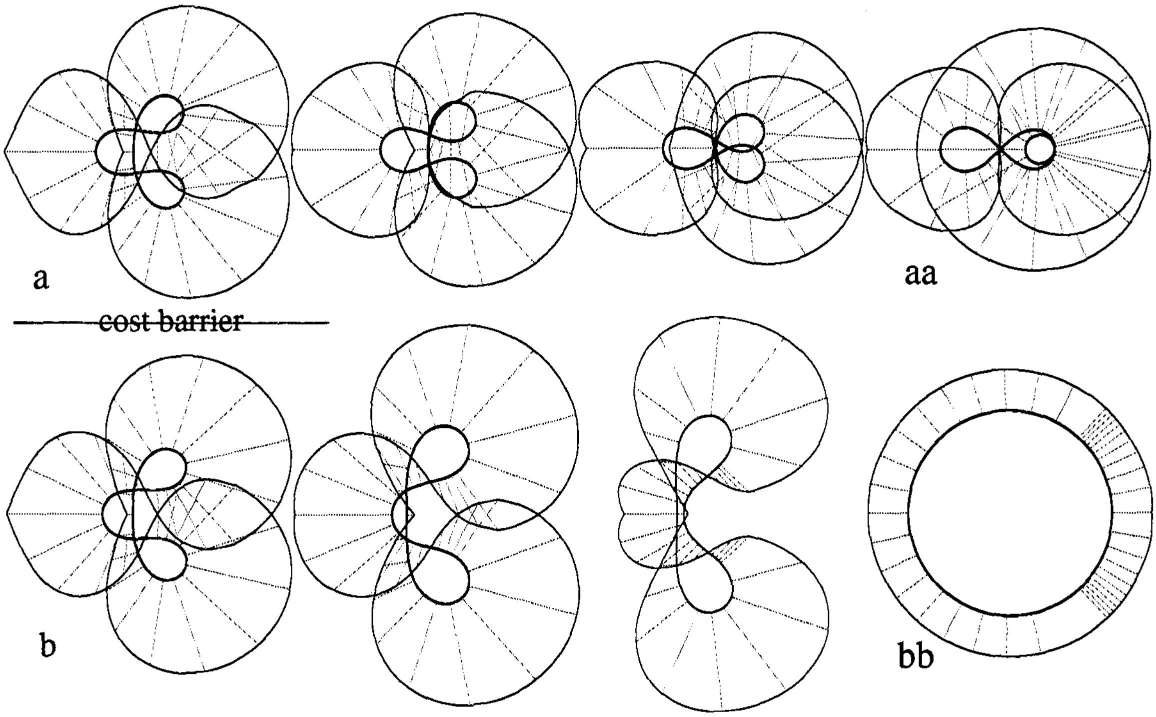

Figure 7. *Transformation between two SI-MVC curves of turning number 1.*

We have yet to find a local minimum for the SI-MEC. Since it incurs a direct penalty for any curved section, it has a greater tendency to unwind into the simplest possible shape, either a circle or a lemnoid, depending on the turning number of the initial curve.

6. Scale-Invariant Surfaces

Working from MVC we developed similar techniques for modeling minimum variation surfaces (MVS) [5]. The technique of converting a functional so that it is scale invariant is equally applicable to surface modeling. Minimum energy surfaces (MES) are naturally scale invariant [3]. The MVS functional requires a scale factor proportional to surface area:

$$(\int dA)\left(\int \left(\frac{d\kappa_2}{d\hat{e}_1}\right)^2 + \left(\frac{d\kappa_2}{d\hat{e}_2}\right)^2 dA\right).$$

We have begun the work to implement this new scale invariant functional, modeling the surfaces using quintic triangular and quadrilateral patches. We are also investigating other surface functionals based on curvature variation. Once satisfied with a scale invariant minimum variation surface, we will then catalogue the elementary shapes of minimum variation for the low-genus surfaces.

7. Conclusions

With this paper we introduce two new curves for computer-aided geometric design: the scale-invariant versions of the minimum-energy curve (SI-MEC) and of the minimum-variation curve (SI-MVC). These curves minimize cost functions that have been derived from the traditional functionals of the MEC and of the MVC by multiplying the respective functionals by suitable powers of the curve's arc length. While the first curve type is presented primarily for completeness and for comparison purposes, the latter seems to offer practical design solutions of superior fairness and stability.

The SI-MVC retains the attractive features of the MVC: its smoothness characterized by at least G^3-continuity, its lean look that avoids unnecessary bumps and wiggles, and its tendency to assume circular arcs wherever they are compatible with the given constraints. To that it adds the stability of a scale-invariant curve which permits it to assume a stable minimum even in situations where no absolute constraints confine the size of the solution curve.

We implemented the SI-MEC and SI-MVC using a representation, based on quintic Hermite polynomial segments, which is globally optimized by a finite-element non-linear optimization package. The user may specify any number of positional, tangential, and/or curvature constraints. The system uses heuristics to construct an initial approximation of a curve satisfying the given constraints. The system then globally minimizes the functional by varying all the non-fixed parameters of the underlying curve representation. Our approach is computationally expensive and complicated examples take a few minutes on the fastest work stations available. The large computational cost is justified by the quality of the results.

Our approach for achieving scale-invariant cost functions by multiplying with a suitable power of the length/area of the manifold to be designed also extends to surfaces. The computational cost of scale invariant surfaces is considerable, resulting in times to solution stretching into hours.

8. References

1. Gerald E. Farin. *Curves and Surfaces for Computer Aided Geometric Design, A Practical Guide.* Academic Press, San Diego, 1990.
2. B. K. P. Horn. "The Curve of Least Energy." *ACM Transactions on Mathematical Software,* 9 (December 1983): 441-460.
3. L. Hsu, R. Kusner and J. Sullivan. "Minimizing the Squared Mean Curvature Integral for Surface in Space Forms." 058-92, Mathematical Sciences Research Institute, Berkeley, California, June 1992.
4. Augustus Edward Hough Love. *A Treatise on the Mathematical Theory of Elasticity.* 4th ed. Dover, New York, 1944.
5. Henry P. Moreton and Carlo H. Sequin. "Surface Design with Minimum Energy Networks." In *Proceedings Symposium on Solid Modeling Foundations and CAD/CAM Applications,* ed. J. Rossignac and J. Turner, 291-301. ACM Press, New York, 1991.
6. Henry P. Moreton and Carlo H. Sequin. "Functional Optimization for Fair Surface Design." *Computer Graphics,* 26 (July 1992): 167-176.
7. Henry P. Moreton. "Minimum Curvature Variation Curves, Networks, and Surfaces for Fair Free-Form Shape Design." Ph.D. Thesis, University of California, Berkeley, 1992.
8. William H. Press, Brian P. Flannery, Saul A. Teukolsky, and WIlliam T. Vetterling. *Numerical Recipes in C.* Cambridge University Press, Cambridge, 1988.
9. C. H. Woodford. "Smooth Curve Interpolation." *BIT,* 9 (1969): 69-77.