

Chapter 1

Two-parameter splines

The broad goal of this work is to find the *best* spline, or at least provide a menu of best choices, each suited to a specific requirement. In attempting to characterize the space of all splines, including making a taxonomy of existing splines in the literature, a remarkable number share this property:

A two-parameter spline is defined as one for which each curve segment between two knots is uniquely determined by the two angles between tangent and chord at the endpoints of the segment.

Examples of two-parameter splines include the Minimum Energy Curve, the Euler spiral spline, Mehlum’s KURGLA 1, Hobby’s spline, IKARUS, and Séquin’s circle spline. Note that the count of parameters excludes the scaling, rotation, and translation required to get the curve segment to coincide with the two endpoints—those additional parameters

Any given two-parameter spline can be characterized by two properties: the family of curves parameterized by the endpoint tangent angles, and the algorithm used to determine the tangent angles. For example, the Euler spiral spline uses a curve family with linearly varying curvature, and uses a global solver to ensure G^2 continuity to determine the tangents.

Of the two-parameter splines listed above, the MEC and the Euler spiral are extensible. In the case of the MEC, this property follows directly from the variational statement - if adding a new on-curve point would produce a new curve with a lower cost functional, that curve would have satisfied the old constraints as well. In addition, both splines are based on cutting segments from a fixed *generating curve*, the rectangular elastica or the Euler spiral.

These properties are not simply coincidental. In fact, all two parameter, extensible splines correspond to a single generating curve, and any generating curve can be used as the basis of a two-parameter extensible spline. The two families, then, are essentially equivalent.

1.1 Extensible two-parameter splines have a generating curve

This section presents the argument that all extensible two-parameter splines have a generating curve.

The Minimum Energy Curve and the Euler spiral spline have a number of features in common, among them that all segments between two adjacent control points are fully determined by two parameters. This count of parameters does not include the scaling, rotation, and translation needed to make the two endpoints coincide—just the shape of the curve once the positions of the endpoints are fixed.

In fact, the parameter space for each such segment is described by the tangent angles at endpoints. Since our definition of “interpolating spline” requires G^1 continuity, determining a single tangent angle at each control point suffices. Thus, combining a two-parameter curve family with a technique for determining tangent angles at the control points yields an interpolating spline. This section explores the entire family of such splines, called “two-parameter splines” for short (non-interpolating splines won’t be considered, so the “interpolating” qualifier is often omitted for brevity).

| Curve type | Tangents | Continuity | Extensibility | Roundness | Robustness | Locality |
|----------------|---------------|---------------|---------------|-----------|------------|----------|
| MEC | G^2 | G^2 | +++ | -- | -- | + |
| Euler spiral | G^2 | G^2 | +++ | +++ | ++ | + |
| Hobby | $\approx G^2$ | $\approx G^2$ | - | - | +++ | + |
| Biacr (IKARUS) | Lagrange | G^1 | -- | +++ | ++ | + |
| Kurgla 1 | G^2 | G^2 | - | +++ | + | + |
| Circle | local | G^3 | --- | +++ | +++ | +++ |

Table 1.1 summarizes the properties of several splines in this two-parameter family. All of these are known and appear in the literature.

- MEC is the interpolating spline minimizing the total bending energy. Each segment between control points is a piece of a rectangular elastica, and the tangents are chosen so that G^2 continuity is preserved across each control point. (Section ??)
- Euler spiral is the interpolating spline for which each segment between control points is a section of an Euler spiral, i.e. curvature is linear with arclength. Like the MEC, the tangents are chosen for G^2 continuity. (Section ??)
- Hobby is a piecewise cubic spline developed for the Metafont system. It is best understood as a polynomial approximation to the Euler spiral spline, with better robustness but lower performance on other metrics such as smoothness and extensibility [1] (Section 1.8).
- Ikarus (Section ??).
- Kurgla 1 (Section 1.10).
- Circle spline (Section 1.11).

Note that there are also several splines in the literature designed to approximate the MEC (quite accurately in the case of small angles), but do not fit into the two-parameter framework. In general, cubic polynomial curves have four parameters. In a standard cubic spline, the curve segments can be any cubic polynomial, so the parameter space is four dimensional. Parameterizing by chordlength approximates the MEC more closely but does not reduce the dimensionality of the parameter space.

Similarly, the scale invariant MEC (dubbed SI-MEC in [9]) has segments that are piecewise elastica, and there is an additional parameter global to the spline (and associated with tension forces at the endpoints, in a physical model of the spline) that increases the dimension space of individual segments from the two of the pure MEC to the three admitted by the elastica.

While the previous chapter examined the elastica with constrained endpoints, a more general problem is the minimum energy curve that interpolates a sequence of control points. More formally, we seek, over the space of all curves that pass through the control points (x_i, y_i) in order, the one that minimizes the MEC energy functional, restated here:

$$E_{MEC}[\kappa(s)] = \int_0^l \kappa^2 ds \tag{1.1}$$

This problem is the mathematical idealization of a mechanical spline through control points fixed in position but free to rotate, and with an infinitely thin beam free to slide through them with zero friction.

The earliest discussion of this more general problem is probably Birkhoff and de Boor in 1964. Over the next few years, several authors published algorithms for computing these “nonlinear splines.” Probably the most illuminating and accessible discussion of the early nonlinear spline theory is the 1971 report of Forsythe and Lee[6]. A brief summary of their results follows:

- Each segment between knots is a MEC with $\lambda = 0.5$.
- Curvature is continuous across knots (G^2 continuity).

- Curvature is zero at endpoints (in the open case).
- The spline is extensible.

For each segment between knots, the elastica with λ fixed has two remaining parameters, corresponding to the tangent angles at the endpoints. Recall that Figure ?? shows instances over the entire two-dimensional space in which such solutions exist, as well as the boundary of that solution space. In particular, in the case of symmetrical endpoint tangents, the curve is *not* a circular arc, so the MEC spline does not have the roundness property. Forsythe and Lee found this unappealing, and propose adding an arclength penalty term to the functional:

$$E[\kappa(s)] = \int_0^l \kappa(s)^2 + c ds \tag{1.2}$$

As we now know, this change admits the solutions to the elastica equation other than $\lambda = 0.5$. Forsythe and Lee derive the value for c so that the solution is a circular arc (corresponding to $\lambda = 0$), and speculate “whether such an approach could be generalized is an open question.” Answering their question is a primary focus of this chapter.

1.2 Properties of the MEC spline




Figure 1.1: MEC roundness failure.

The most important property in which the MEC spline is lacking is roundness. For three control points placed in an equilateral triangle (shown in Figure 1.1, with the spline drawn in a thick line, and the circle going through the control points thin), the containing radius for the spline is approximately 1.035 times the radius of the circle going through the control points. With four points, this radius deviation decreases to about 1.009, and asymptotically converges as $O(n^4)$ with the number of points.

Another unpleasant property of the MEC spline is the failure to converge on a stable solution. This property is intimately related to the relatively small envelope enclosing the solution space—the maximal endpoint tangent for symmetric solutions to the $\lambda = 0.5$ is $\pi/2$, and not much better for any of the non-symmetric solutions.

1.3 SIMEC

Several attempts have tried to address these shortcomings in the MEC spline. One is the scale invariant minimum energy curve (or SI-MEC) proposed by Moreton and Séquin [9]. This spline is defined as the curve minimizing the SI-MEC functional:

$$E[\kappa(s)] = l \int_0^l \kappa(s)^2 ds \tag{1.3}$$

The functional is termed “scale invariant” because the absolute value is invariant as the entire curve is scaled up. By contrast, the MEC functional decreases with the reciprocal of the scaling constant. The MEC functional can be seen as “rewarding” increases in the length of the curve, which accounts for the roundness failure—the actual curve obtained is slightly longer than the circle, which accounts for the lower MEC functional. In the SI-MEC, the bias toward longer curves is removed, and, in fact, the spline is round when the control points are co-circular.

The SI-MEC minimizes the MEC functional for a given length (a curve with a lower MEC energy at the same length would also obviously have a lower SI-MEC cost functional). Thus, an appealing physical interpretation is that of a mechanical spline with a bit of tension pulling at the endpoints, making the overall curve slightly shorter. Such a curve is generally *not* two-parameter in the formulation above, and, in fact, all values of the elastica are possible (thus, it can be seen as a three parameter spline). The SI-MEC has an unpleasant global coupling (equivalent to finding the tension to apply to the mechanical spline) that makes a global solver considerably slower than the corresponding pure MEC [9].

The SI-MEC is extensible, which follows directly from its variational formulation as the curve minimizing an energy functional, such that it passes through all the control points. It shares G^2 continuity and similar locality properties with the MEC.

1.4 Piecewise SI-MEC

Another practical approach to improving the MEC is to use segments that each minimize the SI-MEC functional between any adjacent pair of control points (given the tangent constraints at the endpoints), and enforce G^2 continuity across the control points. A careful reading of Mehlum’s description [8] shows that this is the spline approximated by his KURGLA 1 algorithm.

The MEC, piecewise SI-MEC, and Euler spiral splines are all very closely related, and are essentially identical in the small-angle case. For larger bending angles at the control points, the difference is in the way that curvature varies:

- Curvature in the MEC increases linearly with the distance from the line tangent to the inflection point.
- Curvature in the piecewise SI-MEC increases linearly along the chord.
- Curvature in the Euler spiral spline increases linearly along the arclength.

*** really needs a figure ***

Since any dependence on the chord breaks extensibility (adding a new point changes the chords), the arclength is the most robust measure. For any given set of control points, the Euler spiral spline will have a slightly larger MEC functional, but almost all of that can be traced to the fact that the scaling of the MEC functional rewards a longer total curve length.

1.5 Cubic Lagrange (w/ length parameterization)

- + Reasonable approximation to MEC in small-angle case
 - + Not round (radius error == quartic??)
 - + Not extensible
 - + Locality same as MEC

1.6 IKARUS (biarc)

The spline used by Peter Karow’s IKARUS system [2] also falls into the category of two-parameter splines. It was intended for font design, and was particularly well suited for interpolating between “masters” of different weights.

The IKARUS spline uses two completely different mechanisms to achieve the two basic elements of the two parameter spline: the determination of tangents and the shapes of the curve segments between control points. To determine the tangents, first a cubic Lagrange polynomial (with normalization based on the chordlengths) is fit through the control points. Then the polynomial curve is discarded (but the tangents preserved), and, for each segment, a curve defined by biarcs is substituted. An arbitrary biarc has one additional degree of freedom, corresponding to the split point where the two arcs join, and IKARUS defines this somewhat arbitrarily to provide reasonable results.

The cubic Lagrange spline is defined as follows (see [5, p. 23] for more detail), given a sequence of input points (t_i, z_i) :

$$L_{i,3}(t) = \sum_{j=i-1}^{i+2} z_j \frac{\omega_{i-1,3}(t)}{(t-t_j)\omega'_{i-1,3}(t_j)}, \quad \omega_{i-1,3}(t) = \prod_{j=i-1}^{i+2} (t-t_j) \quad (1.4)$$

To obtain the Lagrange polynomial normalized by chordlengths, successive t_{i+1} so that $t_{i+1} - t_i = |z_{i+1} - z_i|$, in otherwords, the parameterization is the same as the chordlength (set $t_0 = 0$ for convenience, but the shape of the spline doesn’t depend on it).

Even though the Lagrange polynomial fit can easily incorporate inflection points, the biarc primitive cannot, so inflection points must be explicitly marked in the input.




Figure 1.2: Typical lettershape using IKARUS, with Spiro comparison.

The spline is not extensible, as adding additional points perturbs the chordlength parameterization. Further, roundness is fairly decent, as the biarc becomes a circular arc when the endpoint tangents are symmetrical. Continuity is only G^1 , but the biarc segments keep the spline from exhibiting extreme variations in curvature as is typical of purely polynomial-based splines. The spline is reported to work well when the control points are dense. Karow recommends one control point for every 30 degrees of arc [?]. A typical example is shown in Figure 1.2 (reproduced from [?]). For comparison, a Spiro spline for the same letter is also shown. As can be seen, approximately half as many control points are needed, and the smoothness and quality of the result is quite favorable.

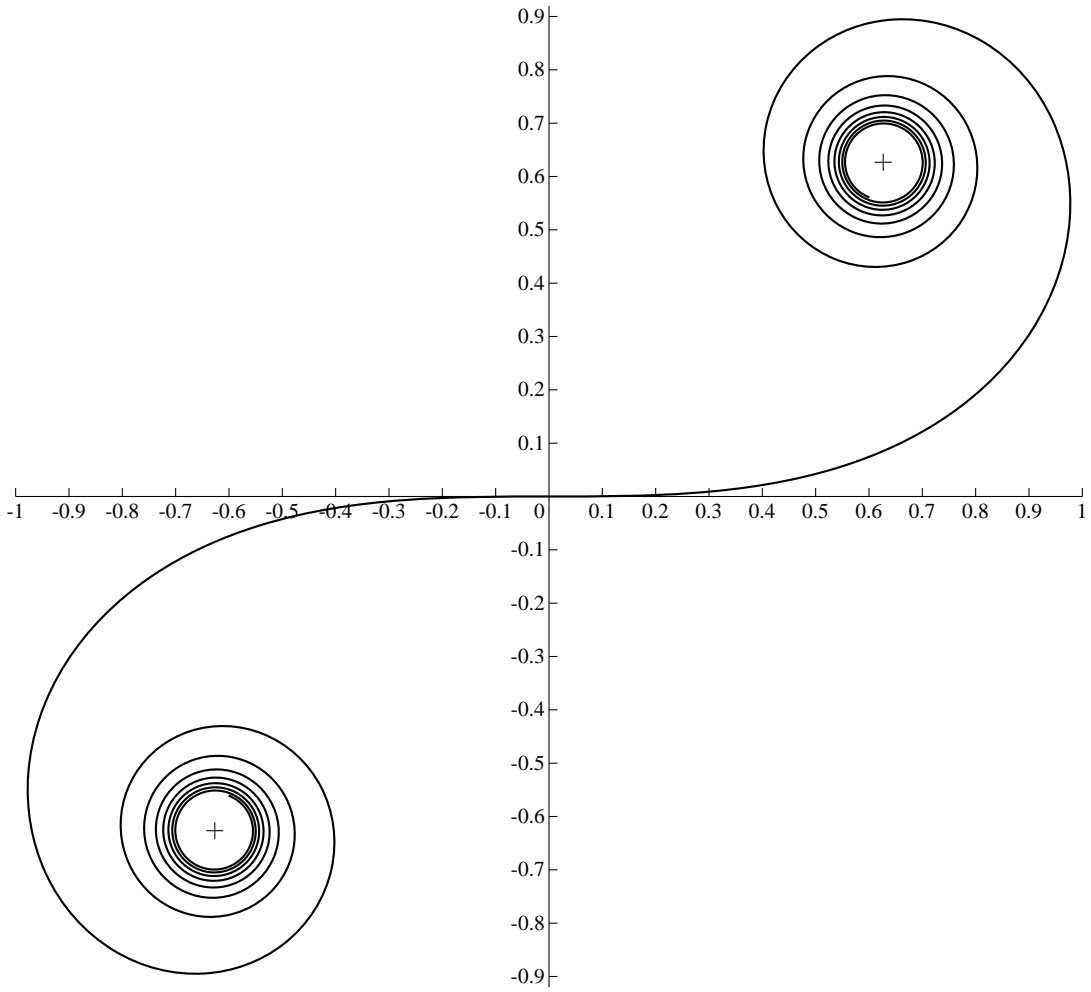


Figure 1.3: Euler's spiral

1.7 Euler's spiral

One particularly well-studied two-parameter spline is Euler's spiral, also known as the spiral of Cornu or the clothoid. More precisely, it is the spline formed by joining piecewise Euler's spiral segments with curvature (G2) continuity at the knots.

Euler's spiral, plotted in Figure 1.3 over the range $[-5, 5]$, is most easily characterized by its intrinsic definition: curvature is linear with arclength.

$$\kappa(s) = \frac{s}{2} \tag{1.5}$$

Integrating the intrinsic equation yields this equation, known as the Fresnel integrals (historically, one each for the real and complex part):

$$x(s) + iy(s) = \int_0^s e^{it^2} dt \tag{1.6}$$

This integral is also commonly expressed with a scale factor applied to the parameter. The resulting curve is similar; its size is $\sqrt{\pi} \frac{2}{\pi}$ times that of the first formulation, with the velocity similarly scaled so that it also has unit velocity.

$$C(s) + iS(s) = \int_0^s e^{i\pi t^2/2} dt \quad (1.7)$$

1.7.1 Euler’s spiral as a spline primitive

Euler’s spiral is readily adapted for use as a two-parameter primitive for a spline. Given tangent angles at the endpoints, find parameters t_0 and t_1 from which to “cut” a segment from the spiral, so that the resulting segment meets the endpoint constraints (for now, assume that such parameters can be found; an efficient algorithm will be given in Section ??).

The behavior of this primitive is shown in Figure 1.4. Compare to Figure ??, showing the corresponding behavior for the MEC spline. Note, among other things, that the Euler’s spiral spline primitive is defined and well-behaved over a much wider range of endpoint tangent angles.

1.7.2 Variational formulation of Euler’s spiral

The spline literature has generally been split into two camps: variational techniques, and the use of explicitly stated curves such as the Cornu spiral. Ironically, Euler’s spiral was born of the study of elastic springs just as was the elastica, but for the most part the literature does not examine Euler’s spiral from a variational point of view. This section ties these two threads together, showing that the Cornu spiral does indeed satisfy some fundamental variational properties.

First, as noted by Kimia et al[3], the Cornu spiral is among the solutions of the equation for the MVC functional, i.e. the curve that minimizes the MVC energy functional subject to various endpoint constraints:

$$E_{MVC}[\kappa(s)] = \int_0^l \left(\frac{d\kappa}{ds} \right)^2 ds \quad (1.8)$$

This result can be readily derived by treating the problem as an Euler variational problem expressed in terms of κ , i.e. $F(x, y, y') = y'^2$, with s for x and κ for y in the terms of Equation ??, as stated in Chapter ?. The corresponding Euler-Lagrange equation is particularly simple:

$$\kappa'' = 0 \quad (1.9)$$

The solution family is thus $\kappa = k_0 + k_1s$, which are simple similarity transformations of the basic Cornu spiral.

While the simplicity of this result is pleasing, it is unsatisfying for a variety of reasons. First, because the variational equation is written in terms of κ rather than θ , the boundary conditions are written in terms of curvature rather than tangent angle. Similarly, the additional terms corresponding to Equation ?? representing the *positional* endpoint constraints are missing. Recall that without these constraints, the only solution to the MEC equation is a straight line.

So this result only tells us that the Cornu spiral is the MVC-optimum curve satisfying *curvature* constraints on the endpoints, but with unconstrained endpoint positions and tangent angles. As we shall explore more deeply in Section ??, the solution to the MVC under other constraint configurations will in general not be a Cornu spiral, a point not clearly made in the presentation by Kimia et al.

A more fundamental problem is that the degree of continuity doesn’t match. A spline made of piecewise Cornu segments has G2 continuity, the same as a MEC spline, while the solution to the MVC functional with positional constraints for internal knots has G4 continuity.

So the Euler spiral is a reasonably good first-order approximation to the MEC functional, and is a special case in the solution space of the MVC functional, but might there be a functional for which it is actually the optimum? In a sense, this is the reverse of the usual variational problem, in which a function is sought that minimizes a given function. Here, we seek a functional which happens to be minimized by the given function, in this case the Euler spiral.

We conjecture that there is indeed a functional: the maximum rate of curvature change over the entire arclength.

$$E_{MAXVC}[\kappa(s)] = \max \left| \frac{d\kappa}{ds} \right| \quad (1.10)$$

This functional is related to the MVC; while the MVC is the 2-norm of curvature variation, E_{MAXVC} corresponds to the ∞ -norm.




Figure 1.4: Euler's spiral spline primitive over its parameter space.




Figure 1.5: Closed Euler example.

1.7.3 Euler’s spiral spline in use

How does the Euler spiral perform in the intended application domain of drawing glyph outlines? Experimentation yields encouraging results. One such example, a lowercase ‘c’ traced from 14 point metal Linotype Estienne, is shown in Figure 1.5. Estienne is a distinguished book face designed by the eminent English printer George W. Jones, and released by Linotype in 1930[7, p. 133]. It is not currently available in digital form, making it an attractive candidate for digitization.

The ‘c’ is a simple closed curve, reasonably smooth everywhere, but with regions of high curvature at both terminals. Below the glyph itself is a curvature plot, beginning at the knot marked by an arrow and continuing clockwise. The long flat section on the left is the inner contour, the sharp spike is the lower terminal, and the other long flat section (which is both longer and lower curvature) is the outer contour. Note that the curve has two inflection points, corresponding to zero crossings on the curvature plot and marked on the outline with short lines.

A more complex example is shown in Figure 1.6, in this case a lowercase ‘g’ from the same font. In this case, there are three closed curves; the outer shape and two “holes.” While the top hole is a simple smooth curve, the other curves contain both G2-continuous smooth knots and “corner points”, with only G0 continuity. Such curves can be factored into sections consisting of the corner points at the endpoints and the smooth knots in the interior, then joined together. Note that sections with no smooth knots are straight lines (there is one at the top of the “ear”), and sections with a single smooth knot are circular arcs (two of these make up the rest of the ear). More generally,




Figure 1.6: Complex Euler example.




Figure 1.7: Difficulty with straight-to-curved joins.

each spline segment bordered by a corner point is a circular arc, with zero variation of curvature.

Experience with drawing a significant number of glyphs is encouraging. The Euler spiral spline can represent a wide range of curves expressively and with an economy of points. Further, in an interactive editor, the correspondence between the input (the placement of knots and control points) seems considerably more intuitive than the standard Bézier curve editing techniques.

The trickiest curves to draw using pure Euler’s spiral splines are transitions between straight lines and curves, as typical in the shape of a “U.” The difficulty is illustrated in Figure 1.7. Essentially, the problem is that the segment adjoining the corner point tends to be a circular arc segment, rather than being constrained to be a straight line. It’s possible, by careful interactive placement, to decrease this curvature so that it’s almost a straight line, but then any changes to the curved segment will undo the effect of this careful placement. Similar problems occur if the points are determined by interpolation between different masters, or some other parametric variation. We shall return to the problem of straight-curve joins in Chapter ??.

1.8 Hobby’s spline

A significant interpolating spline, especially in the application domain of font design, is Hobby’s spline as implemented in Metafont[1]. It is perhaps easiest to understand this spline as a very computationally efficient approximation to the Euler spiral spline of the previous section.

The definition of the primitive segment, with endpoint tangent angles θ_0 and θ_1 , is as follows:

$$\begin{aligned}
a &= \sqrt{2} \\
b &= \frac{1}{16} \\
c &= (3 - \sqrt{5})/2 \\
\alpha &= a(\sin \theta_0 - b \sin \theta_1)(\sin \theta_1 - b \sin \theta_0)(\cos \theta_0 - \cos \theta_1) \\
\rho &= \frac{2+\alpha}{1+(1-c)\cos \theta_0+c\cos \theta_1} \\
\sigma &= \frac{2-\alpha}{1+(1-c)\cos \theta_1+c\cos \theta_0} \\
x_0, y_0 &= 0, 0 \\
x_3, y_3 &= 1, 0 \\
x_1, y_1 &= \frac{\rho}{3\tau} \cos \theta_0, \frac{\rho}{3\tau} \sin \theta_0 \\
x_2, y_2 &= 1 - \frac{\sigma}{3\tau} \cos \theta_1, \frac{\sigma}{3\tau} \sin \theta_1
\end{aligned} \tag{1.11}$$

The resulting (x_i, y_i) are the coordinates for a standard cubic Bézier curve. The adjustable parameter τ corresponds closely to the “tension” parameter of B-splines, and its default value is 1. In the Metafont sources for Computer Modern, the most common other value found is 0.9, and that is only for a few glyphs, including lowercase ‘c’[4, p. 311].

The behavior of the Hobby spline primitive over its parameter space can be seen in Figure 1.8 (see Figure 1.8 for comparison to the MEC). Note that it is defined for all values of θ_0 and θ_1 , but the figure only shows values up to 2.156 (123.5 degrees).

Similarly, the curvature plot is shown in Figure 1.9. For relatively small angles, is approximation to linear curvature is good, but sharp curvature peaks are apparent at larger angles.

While the Euler spiral spline is defined as the one which is piecewise Euler spiral and G2 continuous at the knots, the Hobby spline uses a slightly different criterion. Hobby introduces the concept of *mock curvature*, which is the linear approximation (taking only the first-order term of the exact equation) of endpoint curvature as a function of endpoint angles θ_0 and θ_1 . The global solution to the spline is defined as the assignment to the vector of tangent angles so that mock curvature is equal on either side of each knot.

A very important consequence is that solving the global spline is a *linear* problem. In fact, since each tangent angle affects two neighboring segments, and each segment in turn contributes to the mock curvature constraint at its two endpoints, the entire system of constraints can be represented as a tridiagonal linear equation, which is readily solvable in $O(n)$ using textbook techniques. Hobby also proves that the matrix is non-singular within a reasonable range of values for τ , including of course $\tau = 1$. Thus, Hobby’s is one of the few “advanced” splines for which there is a solution for *all* configurations of input points.

Evaluating the spline properties, we find that the Hobby spline does fairly well. For small angles, it closely approximates the Euler spiral spline. It is not round, but converges to roundness as $O(n^6)$ as opposed to the MEC, which is only $O(n^4)$. Similarly, the extensibility and continuity of curvature are only approximate, and only good at small angles. However, as noted above, it is an exceedingly robust spline. It is thus a good choice if computational resources are limited, robustness is key, and the input control points are closely enough spaced that angles are small.

1.9 Splines from fixed generating curves

The MEC has the property that all segments between control points can be cut from a single generating curve, the rectangular elastica (here, “cut” means selecting points on the generating curve, then performing the unique rotation, scaling and translation to make these endpoints match the corresponding control points). This property can be derived from the equilibrium-of-forces formulation of the elastica, and the fact that the constraints at the control points exert zero tension, but these facts are not intuitively obvious.

Similarly, in the Euler spiral spine, segments between control points are cut from the Euler spiral (“cut” may also include a mirror flip, because this curve lacks a mirror symmetry).

What do splines cut from a single generating curve have in common? Why do so few other splines have this property? Might there be some additional flexibility that can be gained by choosing segments from a curve family rather than a fixed generating curve?

The key to answering these questions is extensibility. In fact, the class of two-parameter extensible splines is exactly equal to the class of two-parameter splines cut from a generating curve.

+ Give argument that anything that’s 2-parameter and extensible must be cut from a fixed curve.




Figure 1.8: Hobby spline primitive over its parameter space.




Figure 1.9: Curvature plot of Hobby spline over its parameter space.

Other properties, such as roundness, can then be derived from properties of the generating curve. For example, for the spline to be round, the curve must tend to a circular shape in the limit. More precisely, $\lim_{s \rightarrow \text{inf}} \kappa'(s)/\kappa(s) = 0$. The Euler spiral meets this condition because κ tends to infinity, while κ' is fixed. Another possibility would be a curve tending to fixed curvature in the asymptote, in which case κ' would tend to 0. The MEC, by contrast, clearly does not have this property.

The envelope of endpoint chord angles (as plotted in Figure ?? for the MEC) also has an effect. A small envelope yields failure to converge. A generating curve without an inflection (such as a parabola) excludes all antisymmetric endpoint angle configurations from this envelope, even for tiny angles.

1.10 Piecewise SI-MEC splines, or Kurgla 1

The Autokon system of Mehlum [8] was inspired by the goal of simulating a real, mechanical spline, but its Kurgla 1 algorithm actually implemented something a bit different, with both advantages and disadvantages with respect to the MEC it was intended to mimic.

Mehlum based his design on the fact that the curvature of an elastica linear in distance from the line of action. His solver approximated the integration of this intrinsic equation as a sequence of circular arcs, the curvature of each determined by the distance from the line of action.

As we've seen, the general elastica has three degrees of freedom; two for the position and angle of the line of action with respect to the control point chord, and one for the amplitude of the curvature variation. Note that inflections occur at the line of action. In this formulation, choosing all these parameters to minimize the total bending energy is a tricky global optimization problem, and Kurgla 1 did not attempt it.

In particular, it nailed down the *angle* of the line of action, choosing it to be perpendicular to the chord. In this way, the parameter space was reduced from three to two, and the problem became much more tractable. In particular, the curve segments were determined by the tangents at the endpoints, and the remaining global problem of choosing tangents for G^2 continuity yields to simple Newton-style solution.

As shown in Section ??, an elastica segment with the line of action perpendicular to the chord is the one minimizing the SI-MEC energy for the angle constraints at those endpoints. Thus, the Kurgla 1 algorithm effectively computes a piecewise SI-MEC.

One consequence is that Kurgla 1 is round, unlike the pure MEC, since the SI-MEC describes a circular arc in the case of symmetrical endpoint angle constraints. Another consequence is considerably better robustness than the pure MEC, because the envelope of parameter space is approximately twice as large.

However, the price paid for these improvements is loss of extensibility; when a new on-curve point is added, it generally changes the chords, and thus the lines of action for the resulting elastica sections.

Thus, though it is based on the elastica, and has some desirable properties, Kurgla 1 is neither a true energy minimizing spline like the MEC, nor does its extensibility and robustness compare well to the Euler spiral spline. Indeed, Mehlum was aware of these limitations, and his subsequent Kurgla 2 algorithm is the earliest known implementation of the Euler spiral spline.¹

1.11 Circle splines

Interpolating splines have a tradeoff between extensibility and locality. The two-parameter splines described above favor extensibility. It is possible to make different choices, and favor locality instead. An excellent example is the “circle spline” of Sequín, Lee, and Yin[10].

In the circle spline, the tangent angle at each control point is chosen as the tangent of the circle going through the control point and its two neighbors. Each segment is then constructed (using a blending of these circles) from the tangents at the endpoints. Thus, moving a control point only affects four segments—two on either side.

The construction of the primitive curves is robust and smooth. In fact, the construction guarantees zero κ' at the endpoints, and that the curvature is identical to that of the circle used to determine tangents, so the resulting spline has G^3 continuity.

Unfortunately, even though the spline has good locality and G^3 continuity, it is not particularly smooth for all inputs, and adding additional on-curve control points can cause serious deviations from extensible, in many cases

¹Mehlum wrote to George Forsythe on July 9, 1971, mentioning “Fresnel integral splines” as recent work. The work was published in 1974 [8].




Figure 1.10: Circle spline extensibility failure.

adding additional inflections. Figure 1.10 clearly shows the lack of extensibility; the two paths differ only by the addition of the central point in the right-hand figure.

Bibliography

- [1] John D. Hobby. Smooth, easy to compute interpolating splines. Technical Report STAN-CS-85-1047, Stanford University, 1985.
- [2] Peter Karow. *Digital Formats for Typefaces*. URW Verlag, April 1987.
- [3] Benjamin B. Kimia, Ilana Frankel, and Ana-Maria Popescu. Euler spiral for shape completion. *Int. J. Comput. Vision*, 54(1-3):157–180, 2003.
- [4] Donald E. Knuth. *Computer Modern Typefaces*, volume E of *Computers and Typesetting*. Addison-Wesley, Reading, MA, USA, 1986.
- [5] Boris I. Kvasov. *Methods of shape-preserving spline approximation*. World Scientific, 2000.
- [6] Erastus H. Lee and George E. Forsythe. Variational study of nonlinear spline curves. *SIAM Rev.*, 15(1):120–133, 1973.
- [7] Mac McGrew. *American Metal Typefaces of the Twentieth Century*. Oak Knoll Books, 2nd, revised edition, 1993.
- [8] Even Mehlum. Nonlinear splines. *Computer Aided Geometric Design*, pages 173–205, 1974.
- [9] Henry P. Moreton and Carlo H. Séquin. Scale-invariant minimum-cost curves: Fair and robust design implements. *Computer Graphics Forum*, 12:473–484, 1993.
- [10] Carlo H. Séquin, Kiha Lee, and Jane Yen. Fair, g2- and c2-continuous circle splines for the interpolation of sparse data points. *Computer-Aided Design*, 37(2):201–211, 2005.