Splines and Curves of Least Energy

Chris Krumm, TropoStudio LLC
06/11/2024

Caveats: I design and build museum exhibits for a living. I design and build an occasional boat for fun. I like math and physics, although both make my head hurt. Analytic solutions are elegant, but why remember 'rules of integration' when numerical methods get you close enough? I felt this way even before PC's.

Claims: I'm pretty certain there are no general CAD or boat design programs splines that precisely represent homogenous beams with simple supports. I believe this may be one reason for the 'CAD doesn't match reality on the floor' complaints one hears about in boatbuilding.

Simplistic Beam Deflection: Mathematical Elements for Computer Graphics by David F. Rogers and J. Alan Adams was first published by McGraw Hill in 1976. The 2'nd Edition came out in 1989. It is still *the* best resource covering the math behind 2D and 3D CAD geometry. The 1'st version can be found online as a PDF. The 2'nd version can be found as a used book via Amazon, AbeBooks, eBay, etc.

From Chapter 5-3: Cubic Splines:

> The mathematical spline derives from its physical counterpart - the loftsman's spline. A physical spline is a long narrow strip of wood or plastic used by a loftsman to fair in curves between specified data points. The splines are shaped by lead weights called "ducks." By varying the number and position of the lead weights the spline can be made to pass through the specified data points such that the resulting curve appears smooth or "fair."
>
> If the physical spline is considered to be a thin elastic beam, then Eulers equation (cf Ref. 5-1) yields
>
> $$M(x) = \frac{EI}{R(x)}$$
>
> where M(x) is the bending moment, E is Young's modulus, I is the moment of inertia, and R(x) is the radius of curvature. For small deflections the radius of curvature R(x) may be replaced by 1/y", where the prime denotes differentiation with respect to x. Thus, we have
>
> $$y''(x) = \frac{M(x)}{EI}$$
>
> Assuming that the ducks act as simple supports, them M(x) is a linear function between the supports. Letting M(x) = A + Bx and integrating the above equation twice shows that the physical spline is described by cubic polynomials between supports.

The bending formulas are part of any intro to mechanics of materials. They apply to a beam of unform cross section and homogenous material, under *elastic deformation*, subject to *small deflections*. Note the emphases: 'elastic deformation' means the material will return to its original position when unloaded. Foam massaged into shape with a heat gun doesn't count. 'Small deflections' is subjective – closer to a loaded floor joist than a cedar strip bent around a boat mold.

Take a strip of material with a constant modulus of elasticity (E) and area moment of inertia (I), such as a wooden plank. Bend it around a series of mold stations and it is a beam with multiple point supports. Let each end run wild with no constraints past the bow and stern stations. The plank curvature will return to zero at both ends. The plank may be close to a curve you drew on the board or in CAD. It won't match exactly because it was not defined the same way.

Hull 'planking' material can be plastically deformed, deposited via additive manufacturing, or milled from a block, and the result can precisely match the CAD geometry. Even an elastically deformed strip can be forced to match a respective CAD curve with plenty of molds and tangency constraints - up until the strip deforms plasticly or breaks!

Actual Beam Deflection: To figure out the 'actual' shape a bent strip or rod takes requires looking into Curves of Least Energy. A beam of uniform cross section and material bent through simple supports will assume a curve that minimizes the total potential energy within the beam. Additional constraints or forces will increase the potential energy of the beam and change the shape.

Online papers regarding curves of least energy are abundant. Most extend the math behind cubic splines and basic beam-bending formulas without focusing on real-world experiments. The best derivation of the curve of least energy I've found is by BKP Horn, Professor of Computer Science and Engineering at MIT, in his 1983 paper "The Curve of Least Energy." A PDF is available at:
https://people.csail.mit.edu/bkph/papers/Least_Energy.pdf

Horn presents several curve options – ellipses, the Euler Spiral, and arc approximations – before settling on what I'll call the **'Horn Curve' - *the curve whose curvature varies with distance along the axis of symmetry***. Interestingly, Horn's career has been focused on machine vision, and he conjectures that the curve of least energy may be related to our sense of 'visual fairness.'

The Experiment: Before reading the BKP Horn paper, I thought **the Euler or Cornu Spiral - *the curve whose curvature varies linearly with arc length*** – represented the curve of least energy. That was wrong. Here is the experimental setup:

- A 0.125" dia. pultruded fiberglass rod is constrained horizontally at lower left and lower right by ball-bearing V-rollers. The pair of v-wheels at the top do not support the rod – they are there just to prevent it shooting up or down. The rod floats in between those 2 v-wheels - trust me on this.
- The rod is constrained only by opposing horizontal loads on the v-wheels. The rod is tangent to the y-axis at each end. The curvature of the rod at each end is

zero. The curvature of the rod at the apex is maximum. This is an example of a curve of least energy.

- The watermark refers to 19.875" from the horizontal baseline to the apex of the curve. It is 12" from the vertical axis to the left and right centerlines of the rod along the horizontal axis. The ratio of height/width for ½ of the symmetrical curve is 1.656:1. That ratio will come up again.
- The chart underlay is from an Excel spreadsheet I put together at first thinking the experiment curve would match a Euler Spiral. It didn't. The Euler Spiral peaks at the top of the photo, above the upper v-wheels. After discovering the Horn paper, I put together a worksheet to calculate the 'Horn Curve.' That is the set of points that lay over the fiberglass rod.

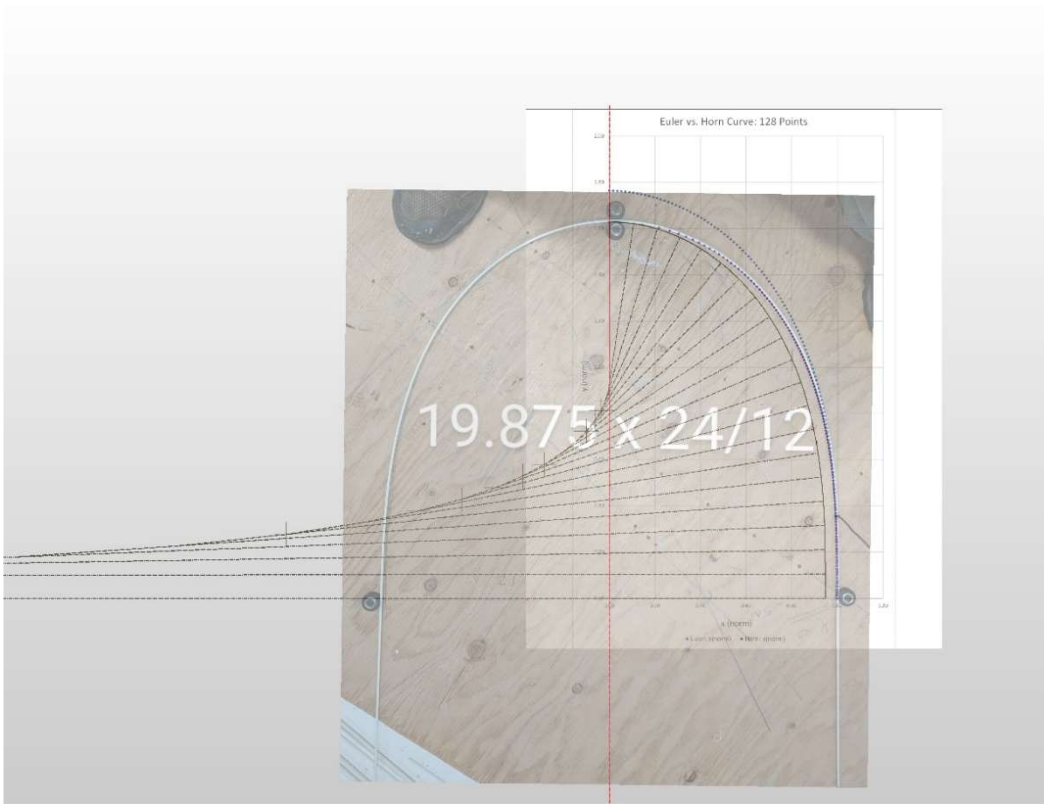

- The background CAD is a geometric layout of the Horn Curve using 20 sub-intervals. X ≠1 at Y=0 due to the low number of sub-intervals. The spreadsheet allows the user to select the number of sub-intervals and normalizes values to X=1. The more intervals used, the closer the approximation comes to Horn's analytic solution based on elliptic integrals.

Conclusion: The Euler or Cornu Spiral has more potential energy than the Horn Curve. The Euler Spiral only corresponds to the Horn curve at very low curvature. A simple beam of uniform cross section and material, bent through simple supports, will match the Horn Curve and have the lowest potential energy.

Below is a composite image of the 0.125” rod, the spreadsheet, and a ‘natural’ cubic spline curve through 3 points from my CAD program (KeyCreator). The cubic spline is dark green and underspecifies the other curves. The end conditions are ‘natural’ which means the curve is tangent to the y-axis and curvature is zero beyond the endpoints (the spline is not constrained beyond the endpoints).

It is a poor fit to the Horn Curve – the ‘true’ curve of least energy. Imagine defining a fantail stern in plan view with a batten through three points vs. using a ‘natural cubic spline’ in CAD through the same three points. They aren’t even close. Do the same with a 2^{nd}, 3^{rd} , or 4^{th} order B-spline in CAD. None will match the Horn Curve without adding more control points, modifying tangency conditions, or ‘tensioning’ the spline.

The image below is a composite of the 0.125” diameter fiberglass rod overlaid with a photo of a 0.188” diameter rod in the same jig. They align, which makes sense: $k = M(x)/EI$. Curvature, k, will remain the same for a homogenous beam if the material modulus of elasticity, E, stays the same while the bending moment, $M(x)$, and moment of area, I, increase proportionally.

Swapping spring scales or load cells for the v-wheels would show the proportional variation in load – apologies for not doing that.

The Excel Spreadsheet I put together generates Euler and Horn curves with a user-selectable number of sub-intervals ($8 \leq n \leq 2048$). X max is normalized to X=1 in all cases. Here is a comparison chart using 512 intervals:

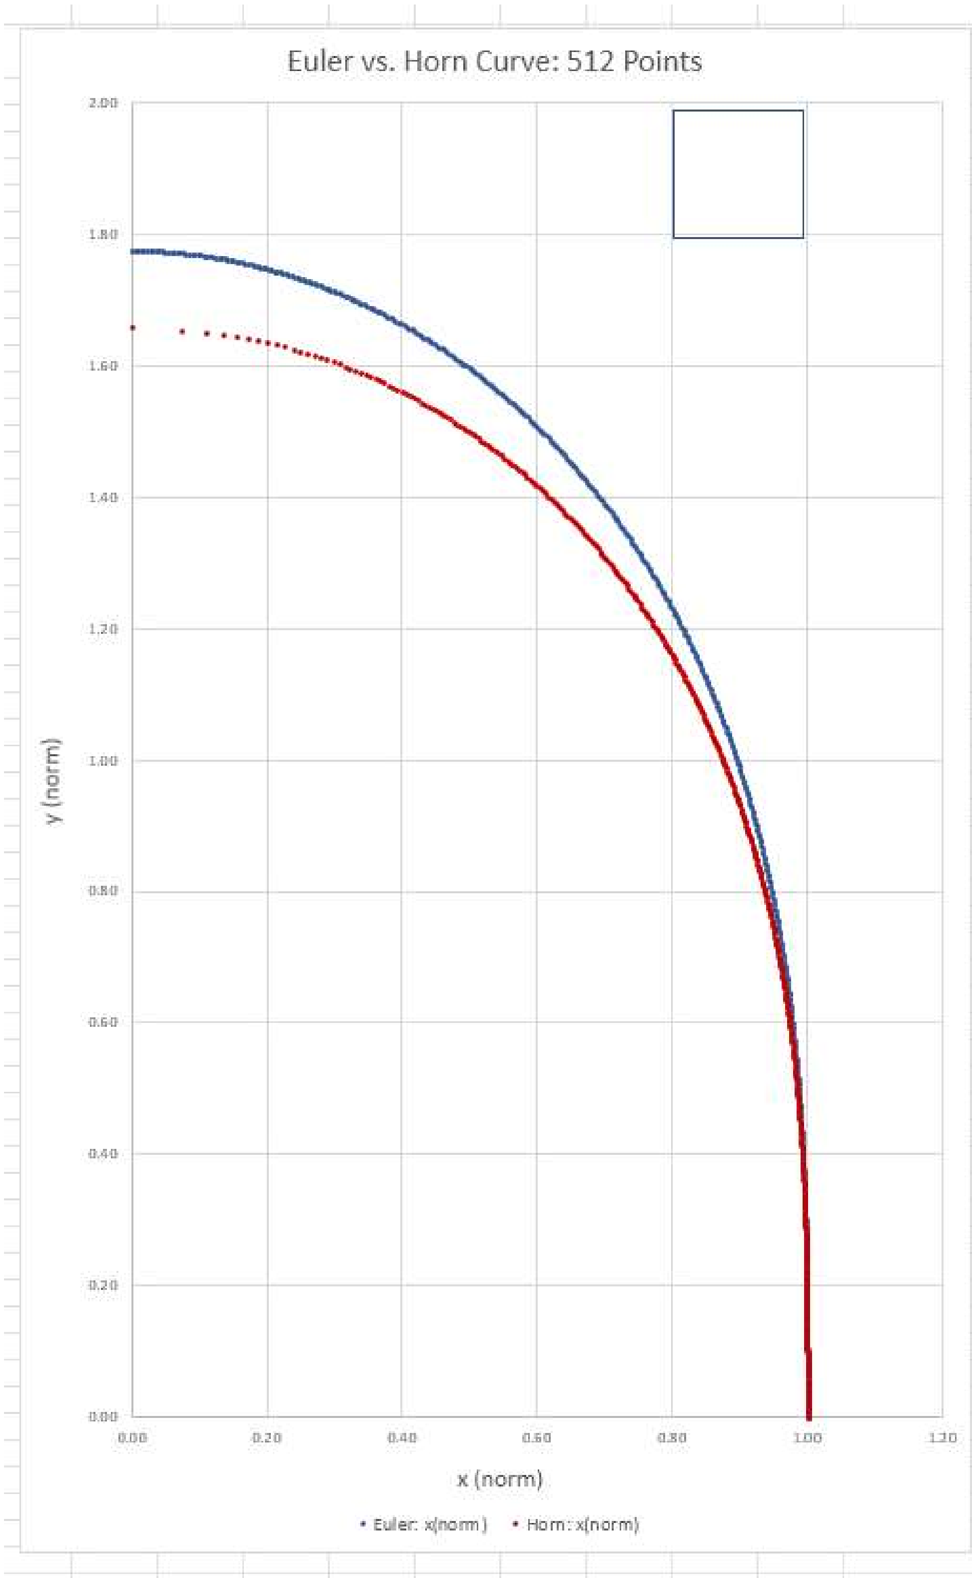

Here are few cells from the Euler Curve worksheet:

Euler Spiral: curvature varies linearly with arc length along the curve

				Curve Length Solver:	
Total Curve Length	S	2.27209299	Re-Calc Curve	FALSE	
Number of Sub-intervals	n	512		1	
Discretization Length	ds	0.00443768		32767	
Maximum Curvature	k_{max}	1.38000000		0	
Minimum Curvature	k_{min}	0.00000000			
Minimum Radius of Curvature	r_{min}	0.72463768			
Maximum Radius of Curvature	r_{max}	∞		SUM $d\theta_{RAD}$	SU

	SUM $E(S_n)$	SUM $E(S_n)_{norm}$
	1.44655294	1.44806335
y(norm)	$E(S_n)$	$E(S_n)_{norm}$
1.77556945	0.00845112	0.00845995

Note that Y max = 1.77556945 and Total Energy = 1.44808335.

Here are a few cells from the Horn Curve worksheet:

BKP Horn Curve: curvature varies linearly with distance along the axis of symmetry

Maximum Curve Height	Y	1.65747353	Re-Calc Curve	
Number of Sub-intervals	n	512		
Discretization Length	dy	0.00323725		
Maximum Curvature	k_{max}	1.19870000		
Minimum Curvature	k_{min}	0.00000000		
Minimum Radius of Curvature	r_{min}	0.83423709		
Maximum Radius of Curvature	r_{max}	∞		SUM dx(n)

	Sum E(Sn)	SUM $E(S_n)_{norm}$
	1.43286316	1.43330287
y(norm)	$E(S_n)$	$E(Sn)_{norm}$
1.65798218	0.00000000	0.00000000

Note that Y max = 1.65798128 and Total Energy = 1.43330287. Recall the 1.656:1 height/width ratio from the rod bending experiment. Close enough.

The Euler Curve is much easier to calculate than a Horn Curve over the same number of intervals. I looked at mapping the Euler curve on to the Horn Curve to maintain accuracy with a faster calculation time. Scaling down the Euler curve along the Y-axis followed by a cubic polynomial adjustment on the residual Y-value differences between the two curves looks pretty good. Note the yellow 'Euler Scaled +Cubic curve' overlaid on the red 'Horn curve.'

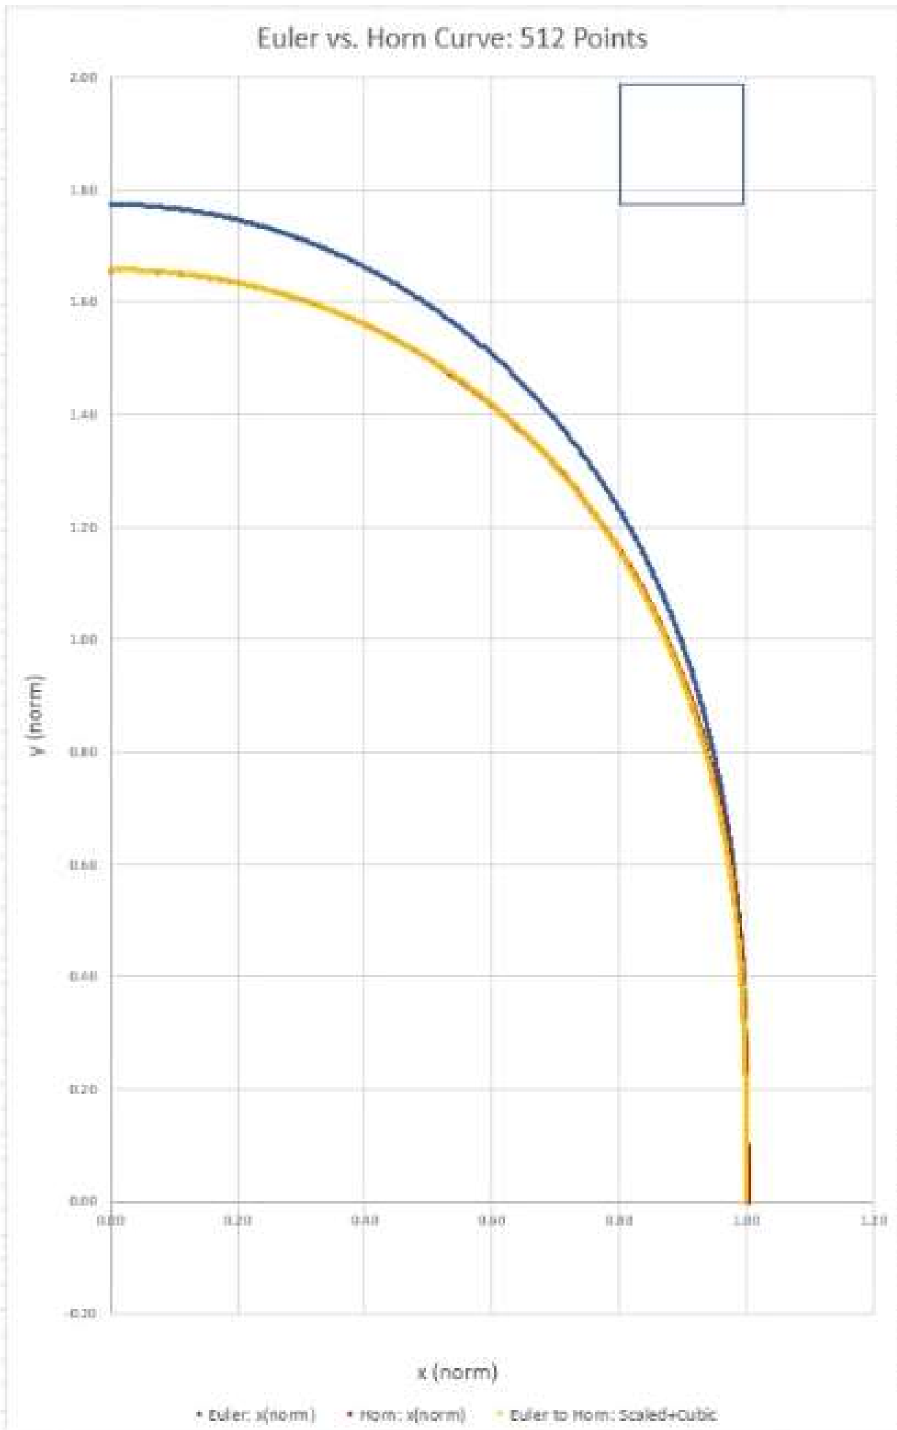

Here's a full worksheet, with the graph of y differences between the Euler and Horn data alongside a quick & dirty 4-point cubic fit to same. It's a poor fit from a regression standpoint, but it works.

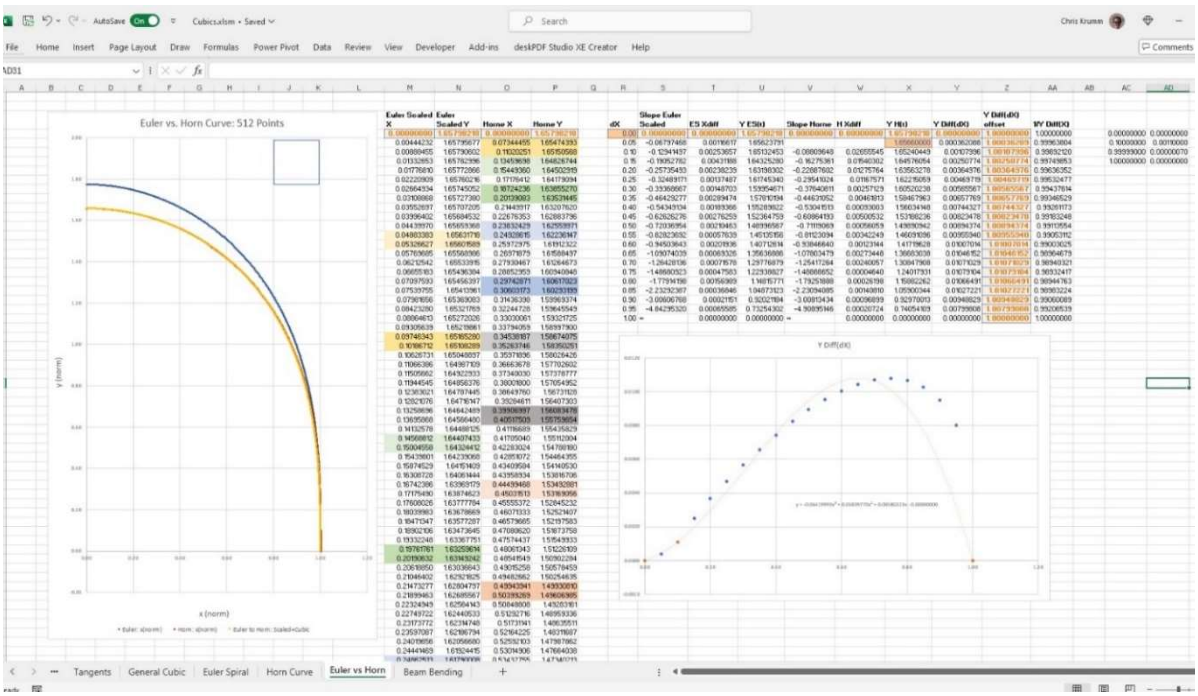

Additional Background: I recently discovered a monograph by R. Frisch-Fay, Lecturer in Civil Engineering at the University of New South Wales, titled Flexible Bars (Butterworth and Co., London; 1962). A PDF can be found at: https://bigoni.dicam.unitn.it/varie/flexible_bars_frisch-fay_1962.pdf. Frisch-Fay analyzes the deflections of thin rods in linear and nonlinear forms under many loadings. The solutions are based on elliptic integrals and predate Horn's work. I have no idea whether Horn was aware of Frisch-Fay's book. The math gives me a headache 😊.

Next Steps: There's a lot more to be done to make this useful for CAD or boat design.

- Extend the applicability to a 2D curve through 3 points. The curve would be comprised of 2 scaled sections of the Horn Curve with k=0 at each end and $k_1=k_2$ at the intermediate point.
- Extend the applicability to a 2D curve of least energy with 4 or more control points.
- Look at simple beams or rods curved in 3 dimensions. Torsion has been added to bending. Perhaps look at it as two 2D bends from orthogonal planes.
- Parametrically distribute longitudinal 'splines of least energy around 'nominally orthogonal' station curves to form a surface. The station curves may be any 2D family of curves. Conics are well suited, and the math isn't too difficult.