Minimizing the Squared Mean Curvature Integral for Surfaces in Space Forms

Lucas Hsu, Rob Kusner and John Sullivan

CONTENTS

We minimize a discrete version of the squared mean curvature integral for polyhedral surfaces in three-space, using Brakke's Surface Evolver [Brakke 1992]. Our experimental results support the conjecture that the smooth minimizers exist for each genus and are stereographic projections of certain minimal surfaces in the three-sphere.

1. INTRODUCTION

Elastic surfaces, such as biological membranes, resilient metal plates, and interfaces between polymers, are of fundamental interest in science. (See, for example, [Seifert 1991] and references therein.) A simple geometric model, proposed around 1810 by Sophie Germain, sets the elastic energy $E(S)$ of a surface S equal to the integral with respect to surface area of an even, symmetric function of the principal curvatures of S. The surface S may be embedded or immersed in three-space (typically $\mathbf{R}^3$, but possibly another three-manifold of constant curvature, such as $\mathbf{S}^3$ or $\mathbf{H}^3$), perhaps with volume or boundary constraints. We shall assume that the integrand is quadratic and therefore of the form $a + bH^2 + cG$, where H is the mean curvature and G the Gaussian curvature. Thus the energy is given by

$$E(S) = \int_S (a + bH^2 + cG)\, dA.$$

Physically, this formula is called Hooke's law; b and c here are "bending" energies, while a is a surface tension or "stretching" energy. (When the two sides of the elastic surface are distinguished, as in a polymer interface, our assumption of evenness may not be satisfied. This case can be handled by replacing H with $H - H_0$ in Hooke's law.)

This research was supported in part by NSF grants DMS-8505550 (Hsu) and DMS-8908064 (Kusner), by an NSF Postdoctoral Fellowship (Kusner), and by The Geometry Center (Sullivan). Computational facilities were provided, in part, by additional NSF and DOE grants, and by a generous loan from Silicon Graphics, Inc.

1058-6458/92 $0.50 per page

An equilibrium elastic surface S constitutes a critical point for $E(S)$, subject to the constraints of the problem. What this means is that, for any variation S_t of $S = S_0$, we have

$$\frac{d}{dt}\bigg|_{t=0} E(S_t) = 0.$$

Assuming S is smooth enough, this implies that S satisfies the Euler–Lagrange equation

$$\Delta H + 2(H^2 - (G - a/b))H = p,$$

where Δ is the Laplace–Beltrami operator of S and p is a constant, called *pressure*, which vanishes in the absence of a volume constraint. If we recall that $H = -\frac{1}{2}\langle N, \Delta X\rangle$, where X is the position of the surface S and N is its unit normal, we see that this is a fourth-order, nonlinear, elliptic partial differential equation. If the surface S has boundary, we should impose a pair of boundary conditions, such as the position of the boundary curve together with the surface normal vector along this curve. (In the case of elastic plates, this is commonly called a *clamped* boundary. At a *hinged* boundary, the curvature of S perpendicular to the boundary is specified instead. A third possibility is to specify H along the boundary; this is an average of the clamped and hinged cases.) We shall assume instead that S is a closed surface, with no boundary.

Observe that the coefficient c of G does not enter into the Euler–Lagrange equation. Poisson noticed this around 1815, decades before Gauss and Bonnet showed that the third term in $E(S)$ is actually a topological constant:

$$\int_S G\,dA = 2\pi\chi(S),$$

where $\chi(S)$ is the Euler characteristic of S. Thus it suffices to consider energies of the form

$$E(S) = \int_S (a + bH^2)\,dA.$$

If $b = 0$, the energy is just a multiple of surface area, for which the Euler–Lagrange equation reduces to the condition of constant mean curvature. We are not interested in this degenerate case, so we assume that $b > 0$ and, in fact, by scaling, that $b = 1$. We also assume that a is equal to the curvature of the ambient three-manifold (so $a = 1, 0, -1$ if we are working in $\mathbf{S}^3$, $\mathbf{R}^3$, $\mathbf{H}^3$). The energy is commonly denoted W in this case. If the ambient space is Euclidean—the case that has physical significance—this choice of a means there is no surface tension. The problem is then clearly scale-invariant: the energy is unchanged under Euclidean similarities. Thus, to model a physical pressure across the surface, we would now have to impose not just a volume constraint, but a constraint on the ratio between volume and surface area. We will assume there is no such constraint.

It is a remarkable fact that when a is equal to the ambient curvature, as we have assumed, the energy W is *conformally invariant*, that is, $W(S) = W(\mu(S))$ for any conformal transformation μ of the ambient space. In fact, if $\hat{S}$ is a surface in $\mathbf{S}^3$ and S is its stereographic projection to $\mathbf{R}^3$ (from a pole not on $\hat{S}$), we find that the two elastic energies

$$W(\hat{S}) = \int_{\hat{S}} (1 + H^2)\,dA,$$

computed by using the H and dA induced from $\mathbf{S}^3$, and

$$W(S) = \int_S H^2\,dA,$$

computed in Euclidean space, are equal. We call the inverse stereographic projection map from $\mathbf{R}^3$ to $\mathbf{S}^3$ *conformal compactification*, because the surface S may have ends at infinity that get compactified in $\hat{S}$. In this case, which occurs when the pole of projection lies on $\hat{S}$, the elastic energy of $\hat{S}$ exceeds that of S by an amount equal to $4\pi k$, where k is the number of ends of S.

Thus the symmetry group for W is the group of Möbius (conformal) transformations. From the viewpoint of $\mathbf{R}^3 \cup \{\infty\}$, this is the group generated by Euclidean similarities and the inversion $I(x) = x/|x|^2$ in the unit sphere.

Blaschke and Thomsen [Thomsen 1924] discovered in the 1920s that $W(S)$ can also be computed (up to a constant from the Gauss–Bonnet formula) as the area of the image of S under the *conformal Gauss map*. This map assigns to each point of S the oriented round two-sphere tangent to S and with the same mean curvature as S at that point. (The space of such two-spheres is naturally identified with the Lorentz four-sphere $\mathbf{S}^{3,1}$ in Minkowski five-space $\mathbf{R}^{4,1}$; coordinates in this five-space are Darboux's *pentaspherical coordinates* for the two-sphere in question [Darboux 1887]. Note that the

Möbius group acts on $\mathbf{S}^{3,1}$ via its fundamental linear representation as $O^+(4,1)$.)

Thus minimal surfaces in $\mathbf{S}^{3,1}$ that happen to be conformal Gauss images of surfaces in $\mathbf{S}^3$ yield critical surfaces for W. But the conformal Gauss image of a minimal surface S in $\mathbf{S}^3$ is simply the polar surface S^*, which is again minimal in $\mathbf{S}^3$ (and minimal in $\mathbf{S}^{3,1}$), so any such S is W-critical; since $H \equiv 0$, we see that $W(S)$ is the area of S. The problem of finding global solutions to this minimal surface problem was untouched for decades, perhaps because it is difficult to get a good existence theory for minimal surfaces in Lorentz manifolds, or because it is not clear how to handle the constraint of being the conformal Gauss image of a surface in $\mathbf{S}^3$.

Apparently unaware of this earlier work, Willmore [1965] observed that the round sphere attains the absolute minimum value of 4π among all compact surfaces, and is the only surface to do so. He also studied tori of revolution, which led him to conjecture that the Clifford torus in $\mathbf{S}^3$ with $W = 2\pi^2$ minimizes W among all tori. Implicitly, Willmore proposed the basic global problem of finding compact surfaces of a given topology minimizing W.

Lawson [1970] constructed compact minimal surfaces of every genus embedded in $\mathbf{S}^3$. These, by the discussion above, provided the first examples of higher-topology surfaces critical for W. Lawson also found immersed nonorientable minimal surfaces in $\mathbf{S}^3$ of every topological type except that of the projective plane $\mathbf{RP}^2$.

Bryant [1984] studied W-critical spheres, showing that they all arise from certain complete minimal surfaces in $\mathbf{R}^3$ via conformal compactification: they all have $W = 4\pi k$, where k is the number of ends of the corresponding complete minimal surface. Using this idea, Kusner [1987] found a family of W-critical projective planes, including an absolute minimizer, and Bryant [1988] classified all minimizers. These surfaces and the round sphere are the only explicitly known W-minimizing surfaces.

In the case of tori, Simon [1986] proved that a smooth embedded W-minimizer exists, while for every genus, Kusner [1989] showed that the W-minimizers, if they exist, are necessarily embedded. In particular, Kusner and Pinkall independently observed that there is a unique embedded Lawson minimal surface of genus g with $W < 8\pi$. Kusner conjectured that, for each genus, this surface gives the unique W-minimizer (up to conformal equivalence). From the above discussion, we see that this conjecture can naturally be split into two parts:

(i) Among surfaces of genus g, is W minimized by the minimal surface in $\mathbf{S}^3$ of genus g with smallest area? (The existence of such a smallest-area surface is guaranteed by the smooth compactness theorem of Choi and Schoen [1985].)

(ii) Is the Lawson surface alluded to earlier the minimal surface of smallest area? As g tends to infinity, there is theoretical evidence for this: in particular, the Lawson surfaces and the surfaces of smallest area both must converge as varifolds to the same limit as g increases, this limit being supported on the union of two orthogonal equatorial spheres, with $W = 8\pi$.

2. EXPERIMENTAL SETUP

Our interest in testing these conjectures led to the experiments we describe below. In its simplest form, the experimental task is to flow a given initial surface according to the downward gradient of W and to see where it settles. Generically, it should settle at a local minimum, and possibly the global minimum. To test W-stability, we can start near a critical point and observe whether we flow toward it or away.

Our experiments are now being carried out with Brakke's Surface Evolver [Brakke 1992]. A surface is modeled by a triangulated polyhedron in $\mathbf{R}^3$, and the energy is then a function of the n vertices. For any fixed triangulation, the polyhedron flows (in discrete time) in $\mathbf{R}^{3n}$ according to the negative gradient of this energy. We can also refine the triangulation as necessary.

Note that a polyhedron, viewed as a limit of smooth surfaces, has its mean curvature concentrated along the edges: this concentration means that W is infinite for the actual polyhedral surface. To use polyhedra to approximate smooth surfaces, we must define a new discrete energy w, which approximates the integral W over some smooth surface nearby.

The discrete energy w used in Brakke's algorithm is derived from a notion of mean curvature at each vertex v. Motion of this vertex will affect its *star*, the set of all incident triangles. Let a_v be one third of the area of this star, so that $\sum_v a_v$ is the total surface area. If we were moving the surface to decrease area, we would move the vertex v in the direction of the (logarithmic) gradient of the area a_v:

$$2h_v = \frac{-\nabla_v a_v}{a_v}.$$

The h_v defined here is the discrete mean curvature vector at the vertex v. As the mesh size of the triangulation gets uniformly small, this approximates the variational definition for the mean curvature vector near v of an approximating smooth surface.

(There has been some recent work at Princeton on this "polyhedral mean curvature" and its relation to smooth mean curvature in the context of finding minimal surfaces. Alice Underwood [1992] observed that if we start with a minimal surface and choose points on this surface in a nearly equilateral triangulation, h_v need not approach zero as the mesh size decreases. In fact, consider a vertex v on a minimal surface. If its neighboring vertices are not quite evenly spaced, there will surely be a nonzero h_v, as the neighbors may tend to be above or below v on average. If we take neighbors in a finer mesh around v, but still in the same pattern, the value h_v is nearly unchanged, because it scales properly as a curvature.

However, Underwood seems to find that close to any polyhedral surface with $h = 0$ there is a minimal surface. Similarly, our experiments suggest that when we minimize w, the resulting surface has w close to the value of W for a nearby smooth surface; perhaps the minimization process naturally leads to triangulations on which the problems of the last paragraph do not arise.)

Now we imagine taking h to be a constant h_v over the area a_v. Thus, to define our approximation to W, we set

$$w = \sum_v w_v, \qquad w_v = h_v^2 a_v.$$

To further test the accuracy of our discretization w, we experiment with a simple polyhedron with only one movable vertex. Consider a tetrahedron with a fixed equilateral triangle as its base and a free vertex v at the tip. We expect that the flow decreasing w_v will push v down into the plane of the base, resulting in a flat simplex. We will place v initially above the center of the base triangle, and by symmetry it will move vertically. In fact, if we start with a relatively flat simplex (with the top corner duller than a cube's corner), the flow proceeds as expected. But a cube corner is a saddle point for w_v, and if we start with a sharper corner, the flow is in the opposite direction, making the corner even sharper. The problem here is that, if v is well above the plane, $\nabla_v a_v$ is roughly independent of its height, so h_v and w_v decrease to zero as the corner gets sharper. It seems that narrow triangles (far from equilateral) can also give misleading values for h_v at their vertices.

Although our experiments give us great confidence in w as an estimate of W for "good" triangulations, we must be careful to get a fairly smooth polyhedron from which to start the evolution. If we start with a cube, triangulate the faces and naïvely start flowing by w, the cube does not run down toward a sphere, but instead grows a long "horn" at each corner. Fortunately, there are simple commands within the Evolver [Brakke 1992], such as vertex averaging (which will help remove a sharp corner), edge notching (which eliminates sharp edges) and equiangulation (which helps find nearly equilateral triangles), which permit us to prepare an initial polyhedron properly, making it smooth enough that we can trust the w-evolution. As we near a minimum value of w for any given triangulation, we can then use the Evolver's refine command to subdivide each triangle into four.

We will see later that a cube (properly prepared with the commands shown above, and refined as needed) flows toward a good approximation of a sphere (Figure 1), and a rough polyhedral torus flows toward a Clifford torus (Figure 4). In each case the final value of w is slightly less than the value of W for the smooth surface we are evidently approaching.

These preliminary experiments suggest that the discrete w-flow approximates the smooth W-flow well, but they also raise several basic computational and mathematical issues.

Mathematically, we might ask how good this approximation is. This is a basic numerical analysis problem. Our experiments with the sphere and torus suggest that a polyhedron that minimizes w for its combinatorial type will approximate a smooth W-minimizer, with $w < W$. It seems also that the error $1 - w/W$ is on the order of l^2/A, where A is the total surface area and l is the

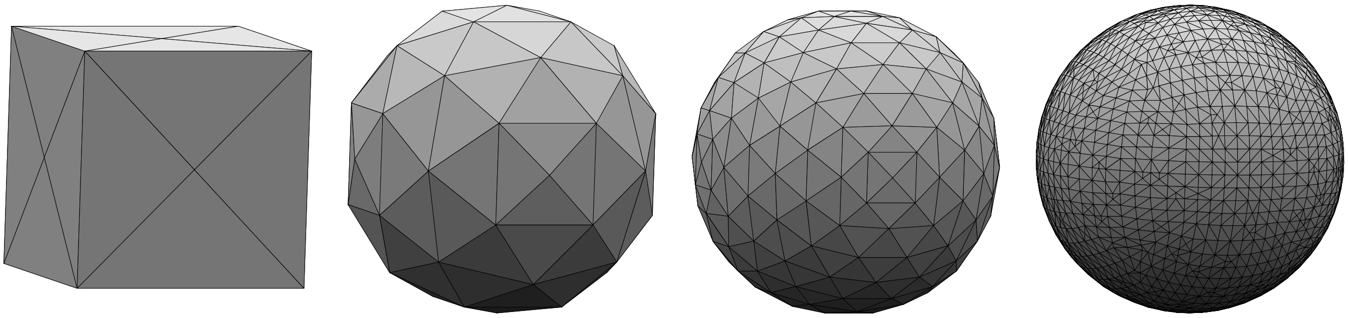

FIGURE 1. The evolution of a cube into a sphere (Test 1 in Section 3). The initial cube (left) must be carefully refined and rounded off to avoid having the corners grow. With 96 facets, we get a decent approximation to a sphere (middle left) corresponding to the first full line of Table 1. One further refinement, followed by motion to reduce the energy w, results in $w = 12.39$ (middle right). Further refinements (right, with 3624 facets) were done with "long-edge subdivision", which allows selective refinement where most needed, but here gives a less uniform pattern in the triangulation. This picture would already look like a smooth sphere if we omitted the edge lines. See also Table 1.

longest edge length. Of course, none of this has been proved; we don't even know what W-minimizing surface (or value) we are trying to approximate, except in the case of the sphere. Here an ideal result would be a rigorous proof that the existence of a polyhedral minimum for w guarantees a smooth minimum for W, and that these minimizing surfaces are suitably close.

For any computational work with surfaces, it would be nice to have a general-purpose program for creating initial polyhedra with a given topology, and perhaps with other desired properties, such as symmetries or moments of inertia. We have taken one step in this direction, creating the program Genus, which outputs a polyhedron of genus g to the Evolver (Figures 5, 9 and 12).

On an algorithmic level, we would like to automate the procedure of properly preparing and evolving polyhedra. Currently, Evolver users must keep careful watch, especially near the beginning of the evolution, and choose when to modify the triangulation with one of the commands mentioned above. We have had some success doing our test runs with an automatic script that applies these commands often enough to avoid trouble. This raises the basic issue of determining *a priori* what class of polyhedra are smooth enough to give sensible values of w, and proving that w can be minimized among these.

Although these are important issues, we cannot treat them properly here. Instead we report on our ongoing Evolver experiments and summarize our results to date.

3. MINIMIZING SPHERES

As already mentioned, it is known that the minimum of W among immersed spheres is 4π, achieved only by the round sphere (a totally umbilic embedding). Although the solution is known, it is useful to conduct experiments in this case, as a way to test the Evolver's algorithm and predict the dependence of the discretization error on the longest edge length. Three tests were made, starting from a cube, a perturbed parallelepiped, and a nonconvex surface.

Test 1 (see Figure 1). We begin with a unit cube, coarsely triangulated into 24 facets (left in the figure). To get a better triangulation, we use the Evolver's refine command `r`, which gives a triangulation with 96 facets. After 20 iterations under the w-flow, we obtain the polyhedral surface depicted in the middle left. A further refinement gives a triangulation with 384 facets, which after another 20 iterations give the spherelike surface in the middle right.

Another way to get a finer triangulation of the surface (which is especially good when the triangles are far from equilateral) is to use the Evolver's long-edge command `l` to subdivide all edges longer than a given upper bound. For this particular example, dividing edges longer than $\frac{1}{6}$ unit gives a triangulation with 1296 facets. After 80 iterations, we obtain a nearly spherical surface, as good as we can get with this many triangles. If we now refine again, further iteration gives the surface shown on the right in Figure 1. This procedure can be continued to obtain ever finer approximations to the sphere. The process is summarized in Table 1. Notice that the discretized energy w approaches the value $4\pi \approx 12.56$ as the triangulation of the surface gets better.

number of facets	fineness	approx. max. angle	$\log_{10}$ iterations	w
24				
96	29	29°	1	11.79
384	120	15°	1	12.39
1296	183	15°	2	12.51
5008	828	8°	2	12.54
12652	2207	8°	3	12.56

TABLE 1. The evolution of a cube into a sphere (Test 1). The meanings of the various quantities is described in detail below. Roughly, fineness is the best measure of the quality of a triangulation, as long as the maximum edge angle is small enough. For a good triangulation, we expect the value of w to be just under the elastic energy W of a nearby smooth surface. See also Figure 1.

Before proceeding, we explain the meanings of the various quantities that appear in our tables. The first column gives the number of triangles in the polyhedron. If we triangulate a surface of Euler characteristic χ with $2n$ triangles, there are exactly $3n$ edges and $n+\chi$ vertices. Roughly speaking, the larger the number of facets, the better the triangulation can approximate a smooth surface.

For a polyhedral surface, the Gauss map is constant on each face, and to approximate a smooth surface we should try to get these normal vectors spread evenly around the sphere. Thus, for a surface of genus g in $\mathbf{R}^3$, the number of facets (or vertices) needed in a polyhedral approximation should depend on the (unsigned) area of the Gauss image in $\mathbf{S}^2$, or total absolute curvature, which satisfies

$$\int |G|\, dA \geq 4\pi(1+g).$$

Thus the number of facets in a good triangulation should grow linearly with the genus.

We mentioned above that the discretization w is not well behaved on triangles that are far from equilateral, so we feel that a surface is best approximated by a mesh of nearly equilateral triangles of nearly uniform size. Thus, instead of counting the number of facets, we look for the longest edge length and compute how many triangles there would be if the total area were covered by equilateral triangles of this edge length. This and the considerations of the last paragraph lead us to define a scale-invariant *fineness* as

$$\frac{A}{(1+g)\, l^2},$$

where A is the total area and l is the maximum edge length. This gives a measure of triangulation quality that is independent of genus and that favors uniform equilateral triangles. This quantity is thus a better measure of how good the triangulation is than the number of facets alone was.

Table 1 suggests that the error $1 - w/W$ is approximately the reciprocal of the fineness; subsequent tests for genus zero and one also support this. Thus, to be confident that we have a good triangulation, we shall require that the fineness be at least about 1000.

Another measure of how well a polyhedron approximates a surface involves the angle between adjacent facet normals. As noted in the Introduction, large edge angles will result in a severe underestimation for the value w. So, for a good triangulation, we also demand that the maximum edge angle be less than 15 degrees. This maximal angle is reported in the third column of the tables, to within a factor of two. If the facet normals were spread out evenly, the maximum angle would go to zero as the reciprocal square root of the fineness.

The fourth column in our tables shows the order of magnitude of the number of iterations of gradient descent required to get the resulting near-equilibrium surface under the w-flow. An entry n in the table means that the number of iterations was within a factor of $\sqrt{10}$ of 10^n. In each iteration, motion is in the gradient direction; the distance moved in this one-dimensional space is chosen by the Evolver so as to minimize w. For the last two rows in each table, we used mostly evolution by the conjugate gradient method (see [Brakke 1992] for details). This method provides a faster rate of convergence to equilibrium once we are in an energy valley.

The last column gives the discretized energy w of the resulting surface. Note that in all the tables, only the entry in the last row of this column is really meaningful. When the triangulation is good, the relative error in w should be roughly one

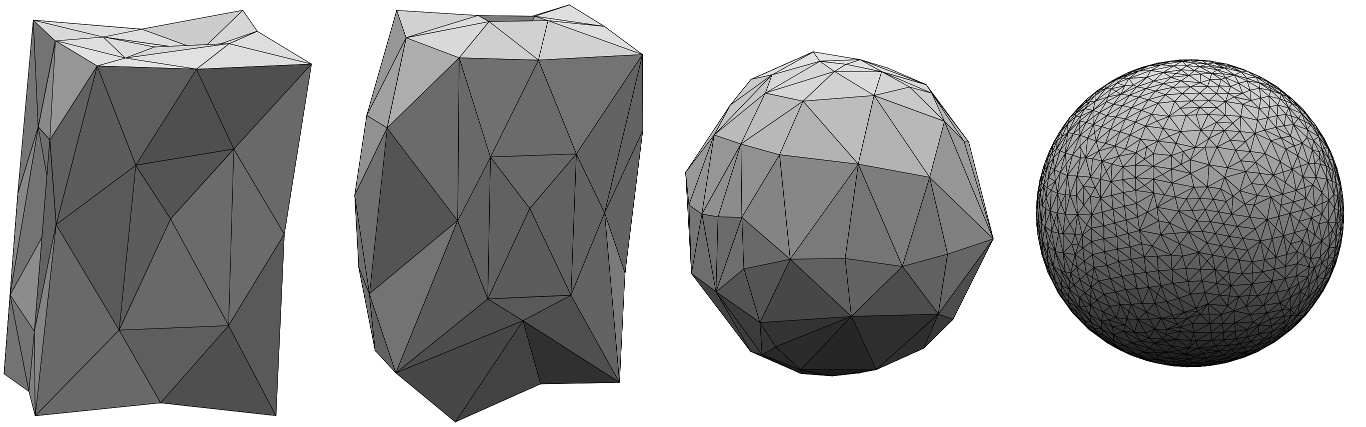

FIGURE 2. The evolution of a perturbed rectangular box into a sphere (Test 2). We start with a $1 \times 2 \times 3$ rectangular box and randomly perturb the vertices. We smooth off the resulting surface (left) with vertex averaging (middle left). This surface can be evolved (with some long-edge subdivision) to a surface (middle right) that starts to look spherical and corresponds to the first full line of Table 2. With further refinement, we get to a surface (right) with fineness 515, corresponding to the third full line in Table 2.

part in 1000, and there should be a nearby smooth W-critical surface; for coarser triangulations, the underestimation of W may be more severe.

Test 2 (see Figure 2). Here we begin with a rectangular box with sides of length 1, 2 and 3. Then, to break symmetry, we do a random perturbation of the surface, using the Evolver's jiggle commands `j` and `jj` to obtain the surface depicted on the left in the figure. We then proceed to evolve and refine the surface. The results are summarized in Table 2. As before, the surface evolves toward a sphere, with w approaching 4π as the triangulation becomes good.

number of facets	fineness	approx. max. angle	$\log_{10}$ iterations	w
96				
152	21	31°	1	11.97
1590	229	15°	2	12.51
3204	515	8°	2	12.54
5514	916	6°	2	12.55
14436	2543	6°	3	12.56

TABLE 2. The evolution of a perturbed box into a sphere (Test 2). See also Figure 2.

Test 3 (see Figure 3). Here we start with a nonconvex polyhedron of genus zero constructed by attaching to a cube a smaller, concentric cube, by means of a connecting rectangular cylinder. To dull the sharp facet angles of the initial polyhedron, we use the Evolver's refine command `r` followed by vertex averaging `VV` (left). Table 3 summarizes the evolution process for this example. From the figure and the table it is clear that the surface is evolving toward a round sphere with corresponding w approaching 4π.

number of facets	fineness	approx. max. angle	$\log_{10}$ iterations	w
88				
4408	400	16°	2	25.27
6918	620	12°	4	14.83
10174	1013	8°	4	12.55
15512	2277	8°	4	12.56

TABLE 3. The evolution of a nonconvex genus-zero surface into a sphere (Test 3). See also Figure 3.

Remarks. (a) The Evolver's jiggling commands `j` and `jj` allow us to test the stability of a putative w-minimizer. The result of Test 2 is consistent with the fact that the round sphere is W-stable. It would be interesting to verify experimentally if all the other W-critical spheres [Bryant 1984; Kusner 1987] are W-unstable.

(b) The w-evolution of Test 3 suggests that an embedded genus-zero surface with $W \leq 8\pi$ will remain embedded while evolving to a round sphere under the W-flow. However, it seems unlikely that the W-flow preserves embeddedness in general.

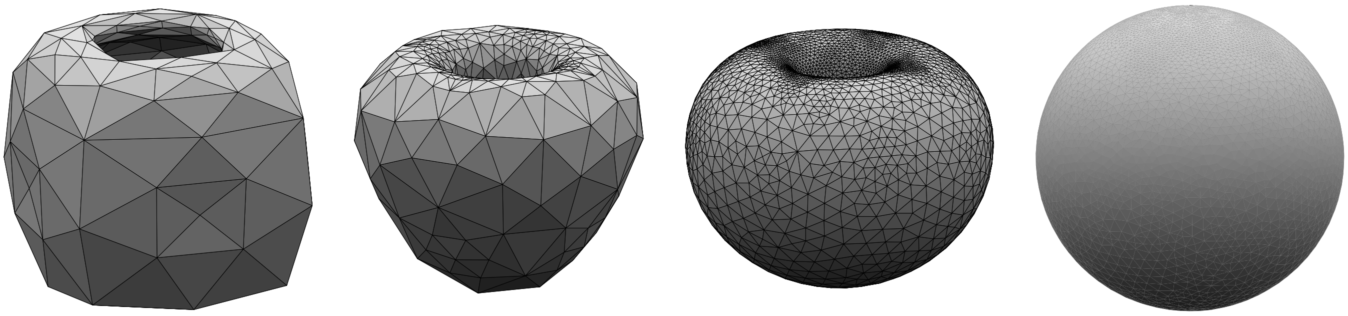

FIGURE 3. The evolution of a nonconvex genus-zero surface into a sphere (Test 3). We start with a hollowed-out cube, with an opening at the top. We round off the corners slightly (left), to get the evolution started. In the middle left panel, the neck is starting to pull through, after much refinement. In the middle right, there is still a big dimple in the sphere, which is why its energy w is over 20. Further motion, however, leads to a round sphere (right), with energy approaching 4π, as shown in Table 3.

4. MINIMIZING TORI

For genus-one surfaces, the outstanding problem is Willmore's conjecture, which states that $W \geq 2\pi^2$, with equality only for tori conformally equivalent to the minimal Clifford torus in $\mathbf{S}^3$. (This Clifford torus stereographically projects to the anchor ring generated by revolving a circle of radius r about an axis whose distance from the center of the circle is $\sqrt{2}r$.) Indeed, Willmore [1965] has shown that $W \geq 2\pi^2$ for all smooth tori of revolution. Simon [1986] proved that a smooth embedded torus exists realizing the infimum of W. Since the only embedded surface with $W = 4\pi$ is the round sphere, this implies that the minimum value of W among tori is strictly greater than 4π.

In this section we summarize a number of experiments that test the validity of Willmore's conjecture and Brakke's algorithm. We first study the evolution under the w-flow of approximate tori of revolution, such as the octagonal torus shown on the left in Figure 4. This surface evolves toward the Clifford torus shown on the right in the same figure, with w approaching $2\pi^2 \approx 19.73$ as the triangulation becomes good. Table 4 summarizes the evolution. Similar results, shown in Table 5, are obtained for the evolution of the "cubical" torus of Figure 5, built using the program Genus.

number of facets	fineness	approx. max. angle	$\log_{10}$ iterations	w
256				
1328	92	40°	2	19.26
5312	366	29°	2	19.65
7936	738	17°	3	19.70
16394	1337	13°	3	19.72

TABLE 4. The evolution of an octagonal torus into a Clifford torus. See also Figure 4.

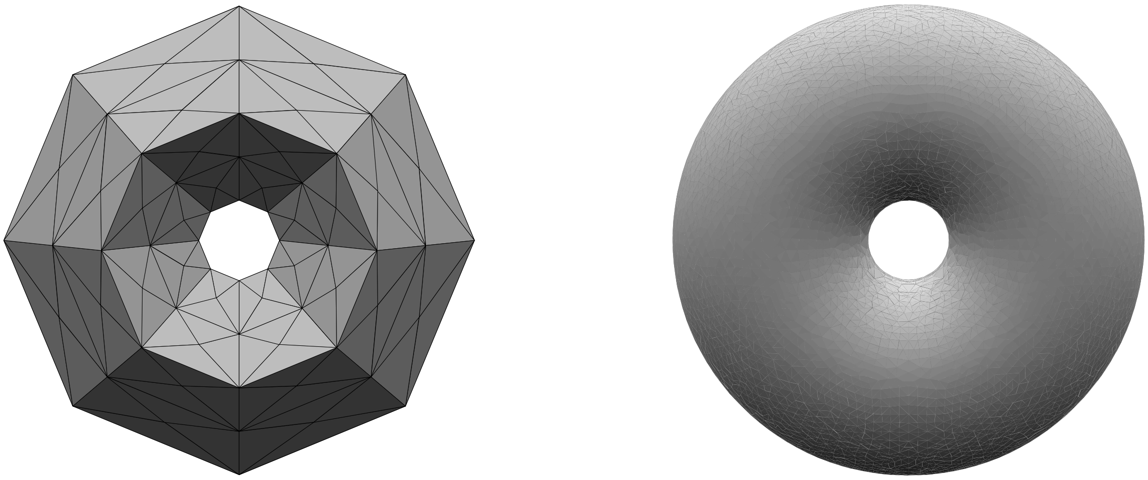

FIGURE 4. An initial octagonal torus (left) and the Clifford torus (right) into which it evolves. See also Table 4.

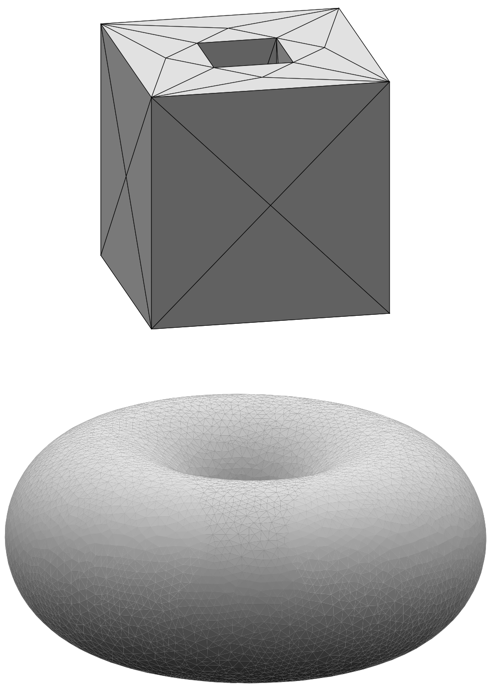

FIGURE 5. An initial cubical torus (top) evolves into a Clifford torus (bottom). See also Table 5.

number of facets	fineness	approx. max. angle	$\log_{10}$ iterations	w
64				
2754	142	20°	3	19.58
5230	110	13°	4	19.67
8112	600	13°	4	19.71
17328	1350	12°	4	19.72

TABLE 5. The evolution of an octagonal torus into a Clifford torus. See also Figure 5.

Remark. Yet another test of Brakke's algorithm is to establish the invariance of w under conformal transformations of $\mathbf{R}^3$. Since w is obviously invariant under Euclidean motions and scaling, it suffices to check invariance under sphere inversion. (We created a program Sphinv for this purpose.) In all our examples, this invariance property is satisfied with only small deviations. For instance, take the approximate Clifford torus evolved from the cubical torus (with $w \approx 19.72 < 2\pi^2$). If we invert its vertices in a unit sphere centered at one of the corners of the original cube, we get a Dupin surface (Figure 6), with $w \approx 19.81$. This surface evolves after about 100 iterations to a nearby one with $w \approx 19.72$ again.

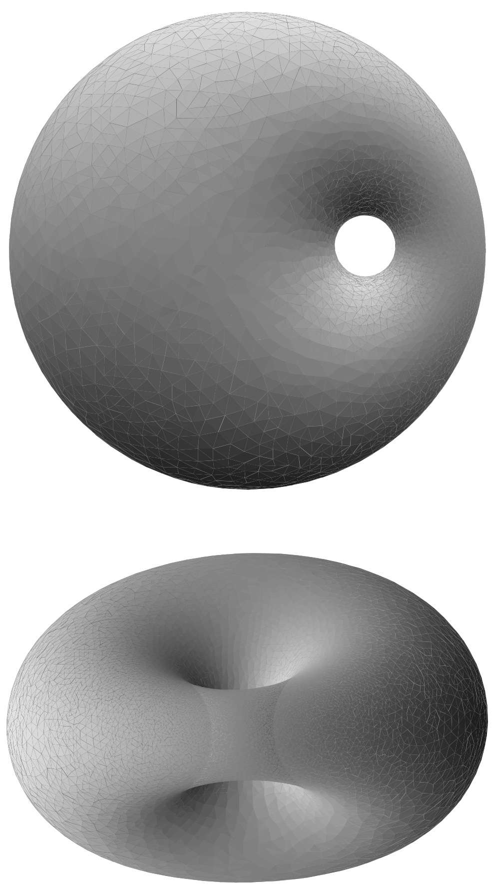

FIGURE 6. Two views of the Dupin surface obtained by inverting the Clifford torus of Figure 5 in a sphere centered at one of the corners of the the original cube (Figure 5, top).

Finally we consider deformations of the Clifford torus, such as the one shown on the left in Figure 7. The result of the evolution is shown on the right, and the statistics are given in Table 6. In every example we tried, the deformed surface appears to evolve toward a Dupin surface with w approaching $2\pi^2$. This corroborates the W-stability of the Clifford torus and its conformal images, which is known [Weiner 1978], and lends credibility to Willmore's conjecture.

5. HIGHER-GENUS MINIMIZING SURFACES

The generalization of Willmore's conjecture to the higher-genus setting [Kusner 1989] states that the stereographic projection to $\mathbf{R}^3$ of Lawson's minimal surface $\xi_{1,g}$ in $\mathbf{S}^3$ minimizes W among genus-g surfaces. In fact, $\xi_{1,1}$ is the Clifford torus—the conjectured minimizer in the case of tori. In contrast with the genus-one case, however, the existence of

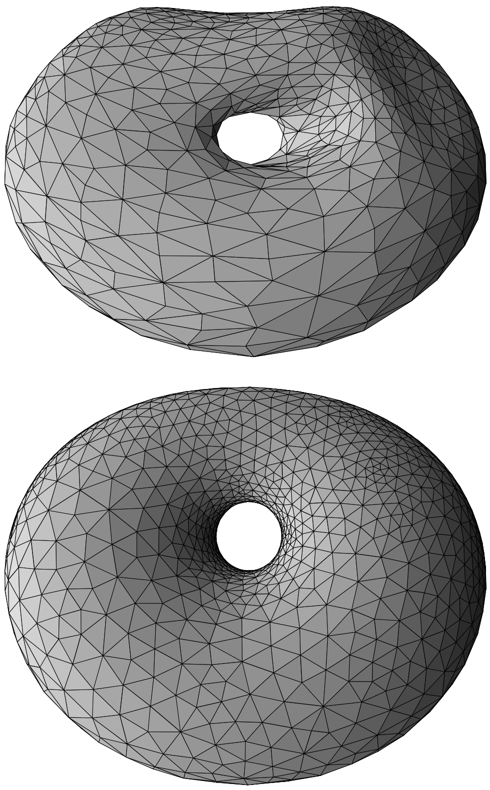

FIGURE 7. The deformed torus at the top evolves into the approximate Dupin surface at the bottom, which corresponds to the first full line of Table 6.

number of facets	fineness	approx. max. angle	$\log_{10}$ iterations	w
1490				
3520	194	16°	3	19.74
5986	425	15°	3	19.71
7416	739	15°	4	19.70
16820	1364	13°	4	19.73

TABLE 6. The evolution of a deformed torus into a Dupin surface, in a stability test for w for surfaces of genus one. See also Figure 7.

a smooth minimizer for $g \geq 2$ has yet to be established.

The Lawson surface $\xi_{1,g}$ has $(g+1)$-fold dihedral symmetry (and another mirror symmetry) in the sphere. There is a stereographic projection to $\mathbf{R}^3$ that preserves this symmetry, and we write $\xi_{1,g}$ also for the image of this projection, a symmetric surface in $\mathbf{R}^3$.

In our experiments we made a detailed study of genus-two and genus-three surfaces. We also made preliminary investigations into surfaces of genus four and five.

Our experiments suggest that a smooth minimizer exists for each genus, and is conformally Lawson's genus-g surface $\xi_{1,g}$.

Genus-Two Minimizers

We first study the evolution under the w-flow of a polyhedral approximation of Lawson's genus-two surface $\xi_{1,2}$, shown on the left in Figure 8. After a few thousand iterations, the surface appears to evolve toward $\xi_{1,2}$, with w approaching the value 21.89 (see Table 7 for a summary). The W-energy of each of the Lawson surfaces $\xi_{1,g}$ is known [Kusner 1989] to be less than $8\pi \approx 25.13$.

number of facets	fineness	approx. max. angle	$\log_{10}$ iterations	w
78				
2932	81	20°	2	21.77
4924	93	15°	4	21.78
8742	367	13°	4	21.86
12430	574	13°	4	21.88
19036	1013	13°	4	21.89

TABLE 7. The evolution toward the Lawson surface of genus two. See also Figure 8.

Next we begin with a genus-two surface shown on the left in Figure 9, obtained by fusing two cubical tori using the program Genus. The bottom row of the same figure shows the final product of the evolution. The statistics are given in Table 8.

number of facets	fineness	approx. max. angle	$\log_{10}$ iterations	w
120				
2794	114	35°	2	29.00
4168	156	20°	3	21.83
6226	281	15°	3	21.87
10192	539	15°	4	21.89
20384	956	13°	4	21.90

TABLE 8. The evolution from a cubical genus-two surface to a button surface, conformally equivalent to the Lawson surface. See also Figure 9.

Notice that the end product of the evolution in Figure 9 has w very close to that of $\xi_{1,2}$ (Figure 8, right). This suggests that these two surfaces are related. In fact, one can find a Möbius transformation mapping one surface into the other. This

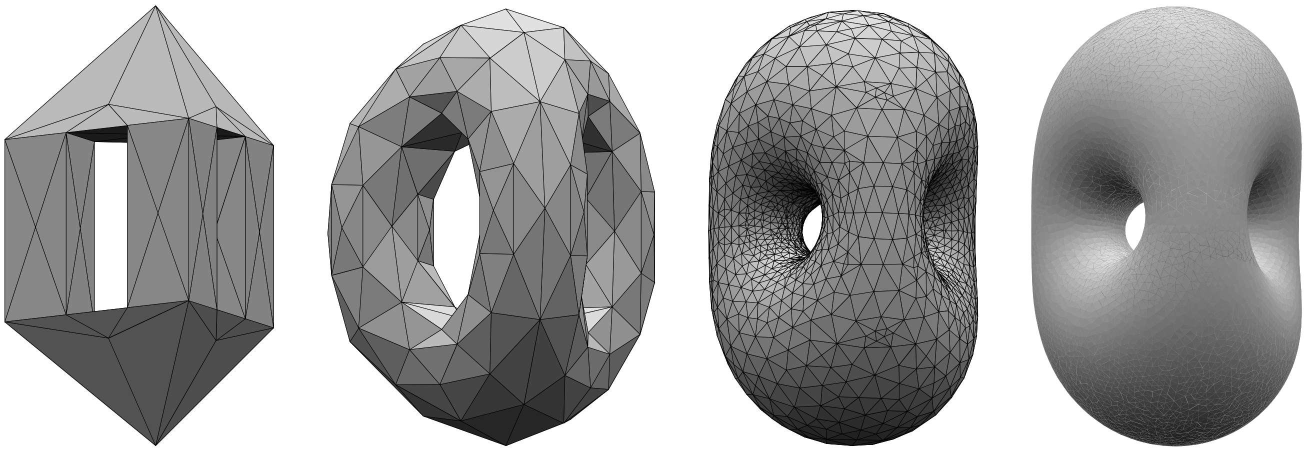

FIGURE 8. A polyhedral surface (left) with the symmetry of the genus-two Lawson surface $\xi_{1,2}$ can be rounded off and refined, giving the surface in the middle left panel. This will then evolve under the w-flow and further refinements (middle right and right), yielding a close approximation to $\xi_{1,2}$. See also Table 7.

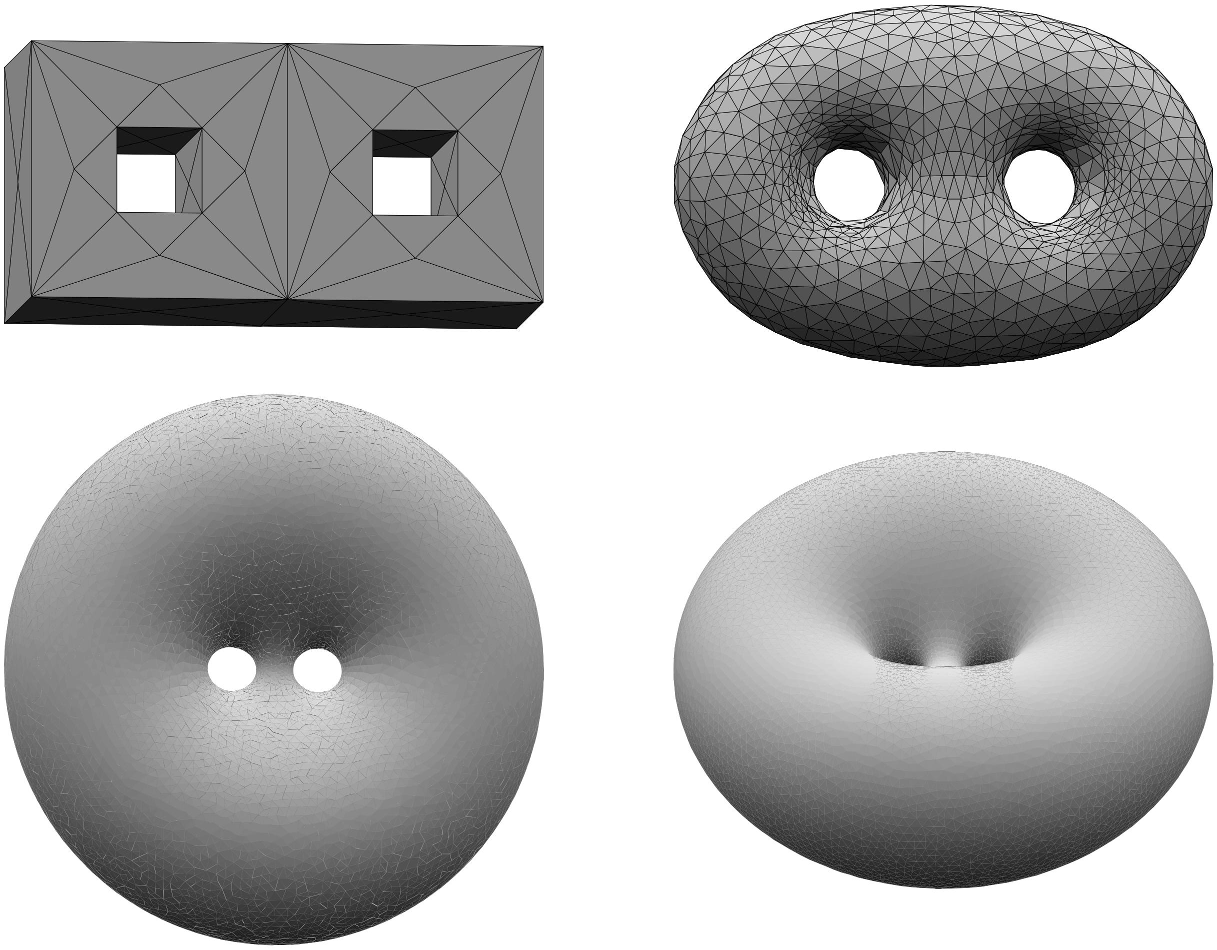

FIGURE 9. The program Genus builds a cubical genus-two surface (top left). A smoothed version (top right), corresponding to the first full line of Table 8, still has energy $w = 29$. Under the w-flow, the holes get smaller and move toward each other, giving the button surface shown in two views at the bottom. This surface is conformally equivalent to the Lawson surface shown on the right in Figure 8.