Computing Sphere Eversions

George Francis, John M. Sullivan, and Chris Hartman

Mathematics Department, University of Illinois, Urbana, IL 61801, USA

Abstract. We consider several tools for computing and visualizing sphere eversions. First, we discuss a family of rotationally symmetric eversions driven computationally by minimizing the Willmore bending energy. Next, we describe programs to compute and display the double locus of an immersed surface and to track this along a homotopy. Finally, we consider ways to implement computationally the various eversions originally drawn by hand; this requires interpolation of splined curves in time and space.

Introduction

In an earlier paper [14], we described a minimax sphere eversion, driven automatically by minimization of the Willmore bending energy for surfaces. Here, we consider three extensions of that work, which are of interest in the areas of optimal geometry, splined surfaces and regular homotopy theory. These fields are some of the many areas of geometry and computation which converge so nicely on the problem of finding new ways to evert the sphere. We consider this paper a continuation of [14], to which we refer the reader for background; further information on sphere eversions can be found in [20,10,24].

Symmetry has always been a great aid to understanding complicated geometrical structures, including sphere eversions. Twenty-five years ago, Bernard Morin proposed the family of tobacco-pouch eversions, with increasing rotational symmetry. In 1983, Rob Kusner observed that certain Willmore surfaces could serve as good halfway models for these eversions. Two years ago, we succeeded in implementing the first minimax eversion, driven by minimizing the Willmore energy in Ken Brakke's *Evolver*. This eversion, with two-fold symmetry, was topologically the simplest of the tobacco-pouch eversions. However, that computation did not make use of the symmetry. If instead we enforce the symmetry, and work with only one fundamental domain, we gain efficiency and the ability to compute the higher-order eversions, as we describe in Section 1.

An immersed surface can be profitably investigated by examining its self-intersections or *double locus*. In a regular homotopy, like a sphere eversion, we have a continuous family of double loci X_t at different times. In their seminal paper [1], Morin and Apéry demonstrated the significance of the *double locus surface* $X \subset \mathbb{R}^4$ obtained by stacking the X_t on top of each other. Events in the regular homotopy correspond to critical points for the height function t on X. For the eversion we presented in [14] (equivalent to one in [1]), counting these critical points shows that X has Euler characteristic -1 and thus is a

Dyck's surface (a sphere with three cross-caps) immersed in $\mathbb{R}^4$, which we want to visualize. This involves computing the double locus at each stage of the eversion, splicing these level curves together into a surface, and then viewing the resulting immersion in $\mathbb{R}^4$ by projection or slicing. We describe the software we used for these three steps in Section 2.

Once we have the software for the middle step (knitting a polyhedral surface together from a discrete sequence of plane curves representing its levels), we can extend it to also allow us to arrange a temporal succession of such layered surfaces into a regular homotopy. This is precisely the visualization exercise expected of the reader of certain classical sphere eversions like the one by Tony Phillips [24]. We want to design our program to work, in such a way that sculpting and rearranging the homotopy in time and space requires the manipulation of only a few control points, with real-time graphics feedback. Our solution uses Catmull-Rom splines in both space and time, as is described in Section 3.

Our first test case is a sphere eversion suggested by Bryce DeWitt [8]. He presented an array of level curves at different times; the conjecture that these could be interpolated by a regular homotopy of nonsingular surfaces has stood unchallenged for twenty years. We hope to resolve this question empirically with the computational tools described in this paper.

The authors wish to express their gratitude to all of the many students and colleagues who have contributed to the projects described herein, but special thanks go to Alex Bourd, Ken Brakke, Ulises Cervantes-Pimentel, Glenn Chappell, Rob Kusner, Stuart Levy, Dennis Roseman, Nick Schmitt, and Matthew Stiak.

1 Symmetric Eversions Driven by Willmore Energy

Earlier we described a minimax sphere eversion [14] computed numerically by minimizing a surface bending energy. In fact, that eversion seems to be only the first in an infinite family of similar eversions with higher-order symmetry. We now have methods to compute these further eversions, and have performed the computations for several examples.

Tobacco-pouch eversions. One way to describe a sphere eversion is to give an immersed sphere, called a *halfway model*, which can be turned inside-out by a rigid motion. In other words, the halfway model has a symmetry which interchanges its inside and outside. More precisely, for an appropriate choice of parameterization, antipodal points on the abstract domain sphere are mapped by the immersion to points related by this symmetry. Given such a halfway model, any regular homotopy which simplifies it down to the round sphere can be extended by symmetry to a sphere eversion.

The halfway models that have been used in this way can be divided into two classes. The first class includes double-covered immersions of a projective

plane like Boy's surface. Here the orientation-reversing symmetry is simply the identity map in space: antipodal points of the sphere cover the same point in $\mathbb{RP}^2$ and thus are at the same place in the halfway model. The early sphere eversions of Shapiro [12], Phillips [24] and Kuiper [17] used Boy's surface as a halfway model. In the second class, the halfway model has $2p$-fold rotational symmetry reversing orientation (and thus p-fold symmetry preserving orientation). The original Morin-Froissart halfway model was of this type, with $p = 2$.

In the early seventies, Morin proposed a family of sphere eversions for integers $p > 1$, later called the tobacco-pouch eversions.[1] In 1977, after Nelson Max kindly loaned one of us (Francis) his newly issued film [22], the students in a freshman honors topology seminar helped design an accurate, but only combinatorial, description of the tobacco-pouch eversions [9]. For p even, the halfway model used in these eversions is of the second class, with $2p$-fold rotational symmetry (reversing orientation). For p odd, it is of the first class, a projective plane with p-fold rotational symmetry. In both cases, the entire eversion can proceed maintaining p-fold rotational symmetry.

Morin soon found analytic expressions for the essential steps of these regular homotopies [23] (see [10, p. 116f]), which were further developed by Apéry [1]. Although these formulas are analytically elegant, they do not lead to nice pictures, so we are led to look for nicer (or even optimal) geometric forms for these same eversions, which can be computed automatically.

Willmore-critical spheres. An elastic bending energy for surfaces should be quadratic in the principal curvature, and if symmetric can be reduced by the Gauss-Bonnet theorem to the integral of mean curvature squared, $W = \int H^2\, dA$, known as the Willmore energy [26]. (See [15] for more about the history of this energy, and some early computer experiments minimizing it.)

In the 1980's, Bryant [6] showed that all critical points for this energy (which is Möbius-invariant) among immersed spheres arise as Möbius transformations of minimal surfaces in R^3 with k flat ends, and thus they can be described explicitly by the Weierstrass representation with data on the Riemann sphere. The Willmore energy of such a critical point is $W = 4\pi k$; aside from the round sphere (a global minimum at $W = 4\pi$) the lowest energy examples occur with $k = 4$.

Shortly thereafter, Kusner [18,19] found particular examples of such critical spheres with rotational symmetry, which he proposed (inspired by a draft of [9]) as particularly nice geometric realizations of the halfway models for the tobacco-pouch eversions. He described a minimal surface S_p as the image

[1] Because of their resemblance to the way the French *blague de tabac automatique* closes up.

of the (punctured) Riemann sphere under the map

$$S_p(w) = \mathrm{Re}\left(\frac{\left(i(w^{2p-1} - w),\ w^{2p-1} + w,\ i\frac{p-1}{p}(w^{2p}+1) \right)}{w^{2p} + \frac{2\sqrt{2p-1}}{p-1} w^p - 1} \right).$$

From this formula, it follows that the orientation-reversing $2p$-fold symmetry $w \mapsto e^{\pi i/p}/\bar{w}$ of the Riemann sphere becomes a rotational symmetry of the surface S_p by angle $\pi - \pi/p$ around the z-axis. Looking at the pth power of this symmetry highlights the distinction between p odd (when $S_p(-1/\bar{w}) = S_p(w)$ so we have a double-covered (punctured) projective plane), and p even (when $S_p(-w)$ is $S_p(w)$ rotated a half-turn around the z-axis).

To get a halfway model with the same rotational symmetry, we apply a Möbius transformation to S_p by inverting in a sphere centered at some point $(0, 0, s)$ along the z-axis. Because S_p passes through the origin (but no other point of the z-axis) we must choose $s \neq 0$ to get a compact image M_0^p. We see no reasons other than esthetic ones to pick any particular value for s; for low p, we have found $s \approx \frac{1}{3}$ gives an appealing halfway model. (Note that the sculpture at Oberwolfach described in [16] is the Boy's surface obtained from these formulas with $p = 3$ and $s = \frac{1}{2}$.)

Minimax symmetric eversions. Each of the immersed spheres M_0^p described above is a critical point for the bending energy W, with an orientation-reversing symmetry of order $2p$. In general, we expect the (Morse) index of a critical point to decrease as more symmetry is imposed. Here the index is not known theoretically, but the numerical experiments we have performed support a reasonable conjecture: if this $2p$-fold symmetry is enforced, M_0^p is stable (a local minimum for the energy), while if only the p-fold orientation-preserving subgroup is enforced, M_0^p is unstable with index one. (With no symmetry imposed, the index would be higher, depending on p.) That is, in the space of immersed spheres with p-fold rotational symmetry, the round sphere minimizes W, while M_0^p is the lowest saddle point (with $W = 8\pi p$).

Therefore, we propose to generate a minimax eversion with this p-fold symmetry, by flowing along a heteroclinic orbit for the gradient flow of W, starting at M_0^p and ending at the round sphere. (Of course this gives only the second half of the eversion; the first half is the same homotopy, reversed in time and orientation, and rotated by $\pi - \pi/p$.)

The gradient flow for W is a fourth-order parabolic flow, which is not well understood. Probably in some cases the flow could start with a smooth surface and pinch off a handle or otherwise produce a singularity. We have only numerical evidence for the fact that this does not happen in our case: computer experiments with Brakke's *Evolver* give an approximation to the flow we want, with no singularities.

It is interesting to note that sphere eversions have often produced with simplification of the halfway model as a guiding principle. For his original

eversion, Morin progressively simplified the double locus. Earlier tobacco-pouch eversions were guided by simplification of the apparent contours. It is remarkable that we get essentially the same eversion by reducing the bending energy W to simplify the surface.

Symmetric halfway models in the *Evolver*. Brakke's *Evolver*[2] is a general tool for geometric optimization, like minimization of surface area or bending energies [4]. There are several discretizations of the Willmore energy available in the *Evolver* [15]. In fact, we made use of these before to implement the minimax eversion with $p = 2$, which is equivalent to Morin's original eversion [14]. However, there we did not make use of the two-fold symmetry. This means that the full surface was computed at each stage, wasting effort, and that, as small numerical errors accumulate, the symmetry we start with is not perfectly preserved during the evolution.

However, the *Evolver* can also work with just one fundamental domain of a symmetric surface. For instance, with mirror symmetry, we simply constrain boundary vertices to lie in the bounding mirror planes (and insure that the energy computations deal with them appropriately). There is also a general-purpose symmetry-group mechanism, where edges from one fundamental domain to the next are marked with a "wrap" value in the group. However, with rotational symmetry, there are certain extra difficulties when the surface spans across the axis of rotation. Recently, with Brakke's help, we have implemented the special features needed to overcome these [5].

To implement one of the eversions, we first get a good triangulation of the halfway model, with symmetries. Our explicit formula for the surface is as a map from the Riemann sphere (which we identify with the unit sphere in space), and it is easier to picture the symmetries on this domain. Thus we start by implementing a sphere with $2p$-fold symmetry generated by a flip-rotation. This can be described with just two vertices (one at the pole and one on the equator), three edges connecting them and their images under the group, and two triangles (one in each hemisphere); the complete input file is shown in Fig. 1.

We can evolve this crude sphere to a nicely triangulated round sphere, for instance by minimizing W as we refine the triangulation. We then make use of an auxiliary program which applies the Weierstrass map to each vertex position; this was implemented in C++ so that the formula could be written in complex notation, almost exactly as it appeared above. (The subroutine to map a single point is shown in Fig. 2.)

We then only need to change one line in the file header to tell the *Evolver* that the generating symmetry is now rotation by $\pi - \pi/p$ instead of flip-

[2] Available at URL `http://www.geom.umn.edu/locate/evolver/`.

```
parameter rotation_order = 2*5
parameter cr = cos(2*pi/rotation_order)
parameter sr = sin(2*pi/rotation_order)
view_transform_generators 1
cr sr 0 0   -sr cr 0 0   0 0 -1 0   0 0 0 1
constraint 1 formula x
constraint 2 formula y
constraint 3 formula x^2+y^2+z^2-1
symmetry_group "flip_rotate"
quantity willmore energy modulus rotation_order/4/pi \
                 global_method star_eff_area_sq_mean_curvature

vertices
1 0 0 1 axial_point constraints 1 2
2 1 0 0

edges
1 1 2
2 2 2 wrap 1
3 1 2 wrap 1

faces
1 1 2 -3
2 3 2 -1

read
transform_expr sprintf "%.0fa",rotation_order-1;
set facet density 0;
raw_cells;
mob_offset := .35;
do_dump := {
        set vertex constraint 3; recalc; unset vertex constraint 3;
        dump sprintf "s%.0f.fe", rotation_order/2;
        system sprintf "weier %.0f %f 0 s%.0f.fe h%.0f.fe",
            rotation_order/2, mob_offset, rotation_order/2, rotation_order/2;
}
```

Fig. 1. The input file `sphere.fe` for a sphere with flip-rotation symmetry in the *Evolver*; only the first line (current set for $p = 5$) needs to be changed for different p, since the constraints, symmetries, and viewing transformations are all expressed in terms of this parameter.

```
void calc(int p, double s, double theta, double x[3], double y[3])
{
    double q,r;
    complex z,denom,i,zp,eit;
    i = complex(0,1);
    if (x[2]==1)
        {y[0]=y[1]=0; y[2]=1/t; return;}
    z = complex(x[0],x[1])/(1-x[2]);
    q = 2*sqrt(2*p-1)/(p-1);
    zp = pow(z,p);
    denom = -i*(zp*zp+q*zp-1);
    eit = exp(i*theta/180*M_PI);
    y[0] = real(eit*(zp*zp/z-z)/denom);
    y[1] = real(eit*-i*(zp*zp/z+z)/denom);
    y[2] = real(eit*(p-1)*(zp*zp+1)/denom/p) + s;
    r = y[0]*y[0]+y[1]*y[1]+y[2]*y[2];
    y[0] /= r; y[1] /= r; y[2] /= r;
}
```

Fig. 2. This C++ subroutine maps a point x on a round sphere to the corresponding point on the halfway model of order p. (If $\theta \neq 0$ we get an inversion of an associate minimal surface.) It is part of a standalone program that transforms *Evolver* input files.

rotation by π/p.[3] This gives an approximation of the halfway model with all its symmetry enforced. Using the standard interactive features of the *Evolver*, we can now even out the triangulation while keeping near a minimum for the discrete W. Once we are happy with the model, we can run a special *Evolver* script to duplicate this fundamental domain and thus enforce only half the symmetry (the p-fold rotation that will remain throughout the eversion).

Evolution scripts. We have described how to create an initial data file for the halfway model M_0^p, respecting the p-fold symmetry. The *Evolver* can compute the matrix of second derivatives for the energy W, and as we expected, we find it has a single negative eigenvalue. This eigendirection is the perturbation we should make to push off from the saddle point and start evolving downhill towards the round sphere. In fact, the `saddle` command in the *Evolver* makes exactly this kind of motion in the lowest eigenvalue direction.

The evolution now proceeds much as in the $p = 2$ case [14]. Theoretically, once we have pushed off the initial critical point, we could continue downhill simply with the conjugate gradient method. However, we discovered for $p = 2$ that this was computationally very slow both at the beginning and again near the end when $W \approx 8\pi$, and the surface has the shape of a gastrula. We found that using the `saddle` command repeatedly (between every 50 iterations of conjugate gradient) sped up the computation at these stages. For the eversions with higher p-fold symmetry, we have decided to use these regular `saddle` commands throughout the eversion, since it is not clear *a priori* when they would be needed.

Thus, our script tells the *Evolver* to repeat (as many times as needed) a saddle command (of a certain magnitude) followed by, say, 50 iterations of conjugate gradient and a retriangulation step. This last step, a subroutine `trian`, involves several *Evolver* primitives, refining long or sharply bent edges, weeding out small triangles, *etc.*, as described in [14].

We mentioned that we observed these evolutions to proceed without any singularities forming. However, we must be careful: if there are not enough triangles to model the necks and other highly curved regions of the surface, problems will develop. The `trian` step will create new triangles in regions where they are needed (as well as removing them from other regions which are no longer so delicate). But we need to carefully choose the constants in that routine. For instance, to recompute an eversion with a finer triangulation, we not only must start with more triangles in the halfway model, but also must ensure that `trian` is adjusted to try to keep the triangles smaller. Probably the most important constant to choose is the bound on how bent an edge can be: the maximum allowed dihedral angle between adjacent faces.

[3] Note that for p odd, we now have a double-covered projective plane. The generating symmetry is a rotation which seems to have order p, but we treat it as still having order $2p$ in the symmetry group, to get both sheets of the surface.

On an a suitable computer,[4] these *Evolver* computations take on the order of ten minutes. To study the resulting eversions with interactive graphics, we save a sequence of *topes*, to be fed into our viewer `illiVert`. Generally we save a tope each time through the loop described above. If we wanted to get finer resolution in time (for creating a high quality video) we should use a smaller saddle step and fewer iterations of conjugate gradient each time through the loop.

The resulting eversions. Roughly, these tobacco-pouch eversions all start by pushing the north pole of the sphere down (forming a gastrula) and then through the south pole, creating the first self-intersection, or double curve. Then, around the neck, we see p fingers push up, and intersect each other near the symmetry axis. The middle stages, approaching the halfway model, are harder to understand without powerful visual assistance.

For pictures of the minimax eversion with $p = 2$, we refer the reader to [14]. Several stages in the $p = 3$ eversion are shown in Color Plate 7.

It has often been thought that a sphere eversion must include a "twisting" phase (like the one seen clearly in Thurston's eversion [21] where the poles rotate a full turn in opposite directions). In our minimax eversions, there is no obvious twisting. It would be interesting to examine them more carefully, to see how they avoid this step that been used in nearly all earlier eversions.

In [14] we described extensively the double locus of the $p = 2$ minimax eversion. The halfway model is a Morin surface with one quadruple point; at nearby times we see the generic creation of such a point, with six double lines intersecting in four triple points forming the edges and vertices of a tetrahedron whose faces are four sheets of the surface. The sheets move inwards, and the tetrahedron shrinks down to the quadruple point; as the sheets continue to move, it then expands again, inverted.

For $p = 3$ we can give a similar description; near the halfway stage the double locus is like a four-fold cover of the double locus of Boy's surface. As the sheets come together, at the center we see a small cube shrinking, instantaneously creating a single sextuple point before expanding again, inverted. Fig. 3 shows the double locus at a stage near $t = 0$; note that there is one extra loop because the two sheets cross each other an additional time.

For general p, the halfway model has a $2p$-tuple point. At nearby times, the $2p$ sheets of surface are arranged like the faces of the polyhedron dual to a p-gonal antiprism; note that this polyhedron has parallel opposite faces if and only if p is odd. As these faces move inwards and through each other, the polyhedron shrinks to a point and then expands again. Again, it is helpful to focus on the double locus to gain understanding of the eversion.

[4] For example an SGI Onyx with ten R10000 processors running in parallel.

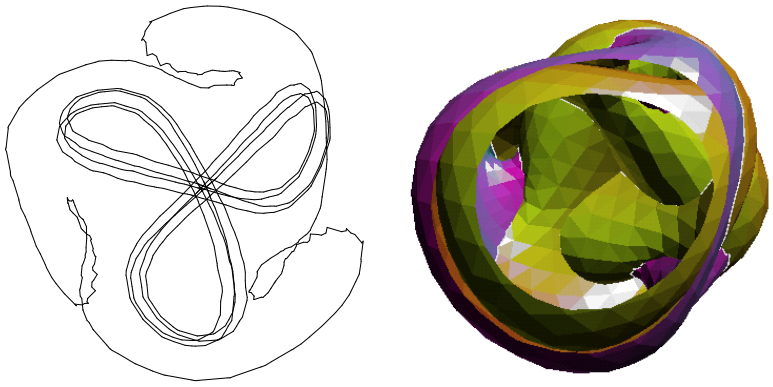

Fig. 3. Near $t = 0$ in the $p = 3$ minimax eversion, we see the double locus (left) and a cutaway view of the whole surface (right). At this stage, the double locus includes a four-fold cover of the three loops in the double locus of Boy's surface, and one extra curve. Compare Color Plate 7.

2 Visualizing the Double Locus Surface of an Eversion

The double locus surface of a sphere eversion is a two-dimensional surface immersed in $\mathbb{R}^4$ whose level sets show the self-intersections of different stages of the eversion. We have developed software to compute these level sets, to splice them together into a surface, and to view the resulting surface in projection.

The double locus surface. Suppose $h : \mathbb{S}^2 \times [-1, 1] \to \mathbb{R}^3$ is a regular homotopy which parameterizes an eversion of the sphere. The *double locus* of the eversion at time $t \in [-1, 1]$ is the subset of $\mathbb{R}^3$ where the immersed surface $M_t = h(\mathbb{S}^2 \times \{t\})$ self-intersects:

$$X_t = \{h(x, t) : h(x, t) = h(y, t) \text{ for some } y \neq x\}.$$

Generically, X_t consists of a collection of curves, which meet at isolated multiple points. However, for those eversions whose halfway model is a doubly covered projective plane, at $t = 0$ we have $X_0 = M_0$, meaning that every point is in the double locus.

We shall assume that for all but a finite set of critical times, the immersed surface M_t intersects itself *transversally*, which means, for all practical purposes, that the surface is in as general position as possible given that it is part of a homotopy. Each X_t consists of the double curves of M_t meeting at triple points, except at the critical values of t, where X_t changes connectivity. In a favorable case, the succession of curves X_t may be regarded as the level curves of a surface $X = \bigcup_t (X_t \times \{t\}) \subset \mathbb{R}^3 \times \mathbb{R} = \mathbb{R}^4$ which is itself immersed in 4-space, and which is called the *double locus surface* of the eversion.

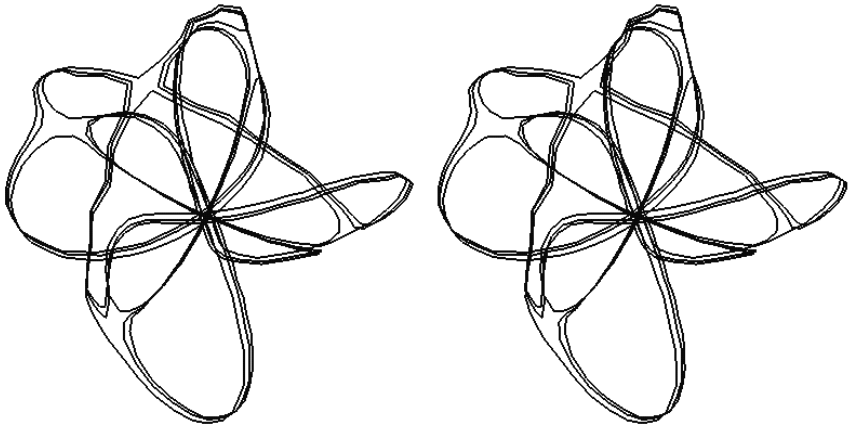

Fig. 4. The double loci of three adjacent topes near $t = 0$ in the $p = 2$ minimax eversion, drawn in cross-eyed stereo. The quadruple point of the halfway model is in the back center, and there are four saddles symmetrically around it, with a fifth in the front upper left. Compare Color Plate 8.

Morin and Apéry [1] introduced the double locus surface X because of the connection between events in the homotopy and critical points of the height function t on X in $\mathbb{R}^4$. For example, in the Morin-type eversion described earlier, the double locus of the halfway model M_0 is a bouquet of six loops emanating from the quadruple point, which cross pairwise at the five points where M_0 has double tangency [14, Fig. 4.5]. In a neighborhood of time $t = 0$ and in a neighborhood of each of these points, the situation can be modeled as follows. At the quadruple point, the four planes of a shrinking tetrahedron pass through each other at $t = 0$ and then regrow the tetrahedron; this is not a critical point of the height function t. At each of the other five places, the model is that of a saddle sinking below sea level (an *isthmus event*); these are saddle critical points for t on X. The double loci at three times near $t = 0$ are shown in Fig. 4.

Following through the second half of the eversion, there is a time when two pairs of triple points are destroyed. Here, X_t changes from being two immersed loops to being two embedded loops, but this is not a critical time. However, when these two loops disappear we get two critical times, which are the two local maxima of t on X. Symmetrically, in the first half of the eversion we find just the two local minima for t. On the whole surface X, therefore, t has these four extremal points plus the five saddle points at $t = 0$, so X has Euler characteristic -1. Thus it is the connected sum of a torus with a cross cap, or, by Dyck's theorem [10, p.101f], a sphere with 3 cross caps, which we call *Dyck's surface* [11].

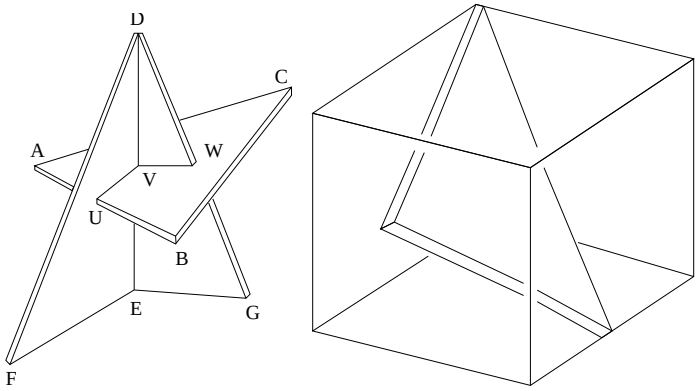

Fig. 5. At the left, the horizontal triangle ABC is pierced by an edge, DE, belonging to two triangles; these meet ABC (along the double curve UVW) in the two different possible ways. At the right, any triangle has at least one vertex at a corner, at least one on an edge, and at most one on a face, of its smallest (coordinate-aligned) bounding box.

Computing a double locus. We want to visualize the double locus surface of an eversion, like the particular immersion of Dyck's surface in $\mathbb{R}^4$ just described. The first step is to compute the double locus at any given time.

A numerical simulation of an eversion is more properly described by a discrete sequence (in time) of simplicial approximations to the immersed spheres $M_t = h(\mathbb{S}^2, t)$. We have coined the word *tope* for such a discrete stage in a homotopy to distinguish it more readily from a so-called *frame* in a computer animation. These topes are not parameterized spheres; in other words, we do not calculate numerical approximations of the function h itself, whose principal role is to simplify the topological description. For economy, we will henceforth use the same notation for smooth objects and their discrete polyhedral approximations.

For each tope M_t, the *Evolver* emits a list of vertices and facets (triangular faces), computed according to the script. We need to compute which facets intersect others. Generically, there are two ways that two triangles in space can meet along a segment: either the two boundaries link, with one edge of each cutting the interior of the other, or a corner of one triangle pierces the other (see Fig. 5). All degenerate cases are defaulted to one of these by using non-strict inequalities. Clearly, the primitive test which we must implement checks whether an edge DE intersects a triangle ABC, and if so, returns the intersection point V. In terms of linear algebra, the putative intersection point V has expansions as convex combinations of ABC and of DE:

$$A + u(B - A) + v(C - A) = V = D + t(E - D),$$

where t, u, v and $u + v$ are all between 0 and 1. We can express this as a matrix equation, to be solved within the constraints:

$$[t\ u\ v\ ] \begin{bmatrix} D-E \\ B-A \\ C-A \end{bmatrix} = [D-A]$$

This 3×3 matrix is (nearly) singular, $\delta = (D-E) \cdot (B-A) \times (C-A) \approx 0$, either if the segment FE is (nearly) parallel to the triangle ABC, or if the latter is (nearly) degenerate. For our visualization purposes it was sufficient to ignore these cases. We found it most useful to solve this system by Cramer's Rule. First, we compare the determinant $\tau = (D-A) \cdot (B-A) \times (C-A)$ with δ. If τ and δ have opposite signs, or if $|\tau| > |\delta|$, then D and E are on the same side of the plane ABC, so we abort the search. Otherwise, the segment does intersect the plane at $t = \tau/\delta$, so we compute the determinants in the numerators for u and v, and check the inequalities that guarantee that the intersection point V is inside the triangle.

For larger numbers of facets, it would be best to implement a divide-and-conquer strategy with oct-trees to decide which facets might intersect. For simplicity (and ease of parallelization) we simply used a double loop over all pairs of distinct facets F, F', checking if F' intersects any edge E of F. The only optimization we used was to immediately discard the pair E, F' if their (coordinate-aligned) bounding boxes (see Fig. 5) failed to intersect.

If we do discover an intersection point V of some E and F', we then proceed to find the other endpoint U of the segment UV along which F and F' intersect. U is the intersection either of some other edge of F with F', or of some edge of F' with F.

Of course, if F and F' intersect along UV, then some neighboring triangles intersect in adjoining segments. But we did not make use of this information in our search, and the output of this first stage was simply an unordered collection of segments. The next step was to arrange these into continuous curves, by identifying common vertices. For the $p = 2$ minimax eversion, the example we computed, at each stage the double locus X_t can be arranged into exactly two loops.

Splicing curves together into a surface. To mesh these polygons X_t in time, we first automatically resampled the curves so that each of them had the same number of vertices, approximately equally spaced around the curve. At this point, we needed to do two things by hand: to choose aligned base points (from tope to tope) on each loop, and to reorder the vertices correctly near the critical halfway time (where the height function on the double locus surface has saddle points).

Now that the we had a pair of 100-vertex cycles for each of 126 values of t, we could in principle look at a projection of the double locus surface from $\mathbb{R}^4$ to $\mathbb{R}^3$. The splining viewer, `illiLevel`, which was designed for viewing

layered surfaces with far fewer control points, and which we will discuss in Section 3, revealed an intolerably "over-sampled" surface which was very slow to render, and also corrugated to an ugly degree. Since we planned to examine the surface more thoroughly with the 4-D viewer `illiSlice`, we felt that little would be gained by resampling to let the splines do more of the work.

Instead, the meshed dataset was next translated into a list of triangular facets suitable for resubmission to the *Evolver*, which has subroutines for culling small (hence superfluous) facets, and for swapping diagonals to give more equiangular triangles. Dennis Roseman helped us reduce in this way a representation of X in terms of 86,912 facets into an eminently manageable model with just 3665 triangles.

Viewing the surface in four-dimensional space. Now we have data for a surface in $\mathbb{R}^4$, given as a collection of triangular facets. We can use Glenn Chappell's `illiSlice` (originally designed to investigate Roseman's knotted surfaces in 4-space [13]) to view it, by means of sections and projections. On an SGI Indigo Extreme II with High Impact graphics, a version written in C (with the `gl` library under Irix 5.3) can render 2000 triangles at the practical limit of 10 frames per second.[5]

Given the surface X in $\mathbb{R}^4 = \{(x, y, z, t)\}$ we can first of all rotate it about the xt, yt, and zt axes to give a new surface X^ρ in $\mathbb{R}^4$. (The other principal rotations are in the 3-D space of the viewer.) Secondly, given two values $t_0 < t_1$, `illiSlice` retriangulates X^ρ into three separate surfaces, the belt $X_0^\rho = X^\rho[t_0, t_1]$, and the distal parts $X_\pm^\rho$ beyond the belt. These are (separately) orthogonally projected to xyz-space, where they can be rotated and zoomed in the usual `illiView` style. The belt is colored in a range from blue to yellow to indicate which part is nearer to t_0 and which is nearer to t_1. The distal parts are less gaudily colored, and provided mainly as reference, especially for when we are in "slice mode" with $t_1 \approx t_0$. When the belt is wide enough to include all of the surface, we are in "projection mode" and perceive the fourth coordinate on the color scale.

Recall that, unrotated, X is already sliced by the X_t. As t_0 and t_1 are kept close, but moved together in time from the earliest to the last place on X, we see the homotopy of the double locus of the eversion. Rotated $90°$ in $\mathbb{R}^4$, the slicing and projecting becomes more interesting, and reveals a disturbingly complex immersion of Dyck's surface. The reason for this is not hard to discern. Recall that the *Evolver* script that reduces the Willmore bending energy (from 16π for the halfway model $X_0 = M_0^2$ to 4π for the round sphere) knows nothing about the double locus. In its evolution in time, the

[5] This is the same order of magnitude we expect from the CAVE wholly immersive virtual environment, powered by an Onyx with eight R10000 processors driving two Infinite Reality Engines which each render four images (two stereo walls) per frame.

eversion pulsates, perhaps moving the double locus about unnecessarily. Such excess motion is stacked into spatial complexity by passing from the $X_t \subset \mathbb{R}^3$ to $X \subset \mathbb{R}^4$. It is a wonder that the double locus turned out to be even as simple as it is.

3 Level Curve Methods for Everting Spheres

Many methods used for visualizing sphere eversions can properly be traced all the way back to the late nineteenth century. Then many geometers were preoccupied with the problem of demonstrating the existence (or its impossibility) of surfaces (with prescribed combinatorial or algebraic properties) in real or complex spaces. To be sure, the fact that a sphere could be turned inside out was the generally surprising consequence of a much later theorem by Smale [25] on classifying the space of immersions of surfaces up to regular homotopy. Yet the first concrete visualization of such a homotopy by Arnold Shapiro [12] was achieved with topological methods that had reached their maturity before the turn of the century. Even more to the point, Tony Phillips used precisely the same graphical methods for his celebrated sphere eversion (in the Scientific American cover story [24]) as Werner Boy [3] had used to describe his famous immersion of the real projective plane in $\mathbb{R}^3$ at the turn of the century.

Boy's surface via level sets. The problem Boy solved was to present a nonsingular (we now say 'immersed') surface in 3-space which, like the Klein bottle, was one-sided (nonorientable). However, unlike the Klein bottle, Boy's surface was a projective plane, with the smallest possible connectivity number for a closed nonorientable surface, namely Euler characteristic $\chi = 1$. Singular projective planes such as Steiner's Roman surface were well known, but the Klein bottle had been the simplest immersed nonorientable surface for decades. Ever since Möbius [16], we have known that χ for a nonsingular surface equals the number of maxima and minima minus the number of saddles, for any height function in general position (which we now call a Morse function). Thus the best Boy could hope for was to display a surface with precisely one of each kind of critical point.

The easiest way to remember Boy's levels, shown at the left in Fig. 6, is to begin with the saddle level, B^0, consisting of a figure-8 crossing a circle at two points, with the double point of the figure-8 the outside the circle. The circle is oriented so that it has tangent winding number (Gauss's *amplitudo*) $\tau = +1$. Either orientation of the figure-8 has $\tau = 0$. Following Boy, we take one of the two points where the figure-8 crosses the circle to be the saddle point, while the other is just another double point. We split the saddle point of B^0 in the two possible ways to give immersed circles, $B^{\pm}1$, above and below the saddle. With the orientation inherited from the circle in B^0, each has $\tau = 1$. Let B^1 have two positively oriented loops outside a negative

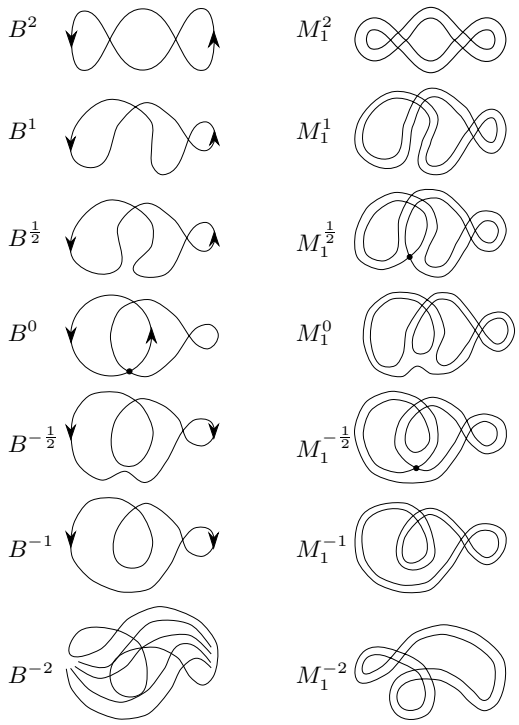

Fig. 6. Level sets of Boy's surface (left) and of a doubling M_1 (right); see text.

circuit, while B^{-1} has a positive loop inside, and a negative loop outside, a positively oriented circuit, as in Fig. 6. It is easy to see that the belt in between, $B[-1,1]$, is an immersed Möbius band with a hole. Hence, all that is left to do to complete Boy's surface is to attach one immersed disc to the top of B^1, and another to the bottom of B^{-1}.

Once B^1 has been straightened out a little (by an isotopy) into B^2 it is easy to see how to move it (by a regular homotopy) into a circle B^3. The belt $B[1,3]$ is thus an immersed cylinder with one double curve. It is harder to move B^{-1} into a circle. Indeed, there is no obvious way to do it. One way is to take hold of an arc opposite the inside loop, and move it across the loop, thereby forming a new positive loop while erasing the negative loop. The result is a curve B^{-2} which is isotopic to B^2, and we proceed as before to deform it to the circle B^{-3}. Whichever way the cylinder $B[-3,-1]$ is immersed, it must have one triple point by a theorem of Banchoff [2]. The double locus of $B[-3,3]$ is a bouquet of three loops emanating from this triple point.

Creating a sphere eversion from Boy's surface. We note that this Boy surface is a combinatorially optimal immersion $\beta : \mathbb{RP}^2 \to \mathbb{R}^3$ of the real projective plane, from the viewpoint of building the surface as a set of level curves of a Morse function. If we compose the double covering $\mathbb{S}^2 \to \mathbb{RP}^2$ (identifying antipodes) with β we obtain an immersion M_0 of the sphere, which covers every point of B twice. Imagine injecting air between the two sheets of M_0 to produce an immersed sphere M_1 in more general position. We shall have an explicit eversion of the sphere if we do two things. The easy part, which everyone started with until Morin's eversion, is to turn M_1 inside out, by moving it (along the normal lines of B) back through M_0 and out the other side to M_{-1}. The hard part is to construct a regular homotopy which moves the complicated immersed sphere M_1 into a round sphere $M_n = \mathbb{S}^2$, in a reasonable number of stages, so that each transition from M_i to M_{i+1} is immediately clear.

Note that the doubling M_1 of Boy's surface, using the same height function, now has two maxima, two minima, and two saddles. If we ask that no further critical levels develop and disappear along the homotopy M_t, this constraint makes the eversion harder to devise, but also easier to follow. This means that the eversion is to be *level-preserving* in the sense that it consists a two-parameter family of plane curves, $M_t^u, 1 \le t \le n, -m \le u \le m$, where M_1^u are horizontal slices through the inflated double of Boy's surface that we labeled M_1 above.

These slices of M_1 are shown at the right in Fig. 6. M_1^0 is the single immersed circle with $\tau = 1$ running along either side of the saddle level B^0 of Boy's surface. Immediately above and below are two saddles $M_1^{\pm\frac{1}{2}}$. The saddles break apart into two curves at $M_1^{\pm 1}$, which run parallel to $B^{\pm 1}$. Color Plate 9 shows these levels spliced together into $M_1[-1, +1]$.

Implementing DeWitt's suggested eversion. At the same Battelle Conference in 1967 where Morin began to develop Marcel Froissart's suggestion for a sphere eversion, the physicist Bryce DeWitt [8] suggested a different one. Sphere eversions were very much the topic of discussion at this conference, and Tony Phillips' drawings [24] were much debated.

In the approach to an eversion outlined above, we have to deform M_1 into some embedded sphere. There are precisely two axially symmetric embeddings of the sphere with two maxima, two minima and two saddles. One is shaped like a dumbbell, the other like a gastrula, both with axis horizontal. Phillips used the former and filled in all the intermediate steps. DeWitt proposed using the latter because is it actually easier to find regular homotopies $M_t^{\pm\frac{1}{2}}$ connecting the saddle levels on M_1 to the corresponding saddles on the gastrula than on the dumbbell. Any such homotopy extends to regular neighborhoods $M_t[\frac{1}{3}, \frac{2}{3}]$ and $M_t[-\frac{2}{3}, -\frac{1}{3}]$ of the saddles. The real difficulty lies in visualizing what to do in between, $M_t[-\frac{1}{3}, \frac{1}{3}]$, to meld these two surface homotopies into each other.

DeWitt designed a two-dimensional array of plane curves. At one side are the horizontal slices of Boy's surface, and on the other side is an immersed sphere which moves to the gastrula embedding in an obvious way. He connected certain critical levels by regular homotopies of the plane curves. The conjecture that these key elements could be interpolated by a continuous succession of surfaces which did not develop singularities during the homotopy has stood unchallenged for 20 years.

Our plan is to check this using suitable computer graphics tools. Furthermore, we plan to use the same general tools to realize Phillips' eversion and many other visualizations based on drawing and deforming complicated plane curves.

Splining curves in time and space. The abstract problem looks like this. Given a 2-parameter family of plane curves, M_{tu}, connect them into a succession of surfaces, M_t, and animate these with smoothly changing shapes. The approach we have taken for `illiLevel` is to use the same number of control points p_{ijk} for each level M_{tu}, and an efficient sampling of key levels $u = k$ and key topes $t = j$. We then obtain model surface homotopies that are easy to edit using piecewise cubic interpolating splines. Although these Catmull-Rom [7] splines do not enjoy great popularity among computer graphicists, we find that they serve our purpose well.

We begin with four points, p_0, p_1, p_2, p_3, in $\mathbb{R}^n$ (or in any vector space, even infinite dimensional), and consider the cubic polynomial $q(t) = m_0 + m_1 t + m_2 t^2 + m_3 t^3$, with coefficients chosen so that

$$q(0) = p_1, \quad q(1) = p_2, \qquad q'(0) = s_1(p_2 - p_0), \quad q'(1) = s_2(p_3 - p_1),$$

where the s_i are pseudo-speeds, described below.

Thus, to compute the coefficients of q we solve the linear system:

$$\begin{bmatrix} 1&0&0&0 \\ 0&1&0&0 \\ 1&1&1&1 \\ 0&1&2&3 \end{bmatrix} \begin{bmatrix} m_0 \\ m_1 \\ m_2 \\ m_3 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -s_1 & 0 & s_1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & -s_2 & 0 & s_1 \end{bmatrix} \begin{bmatrix} p_0 \\ p_1 \\ p_2 \\ p_3 \end{bmatrix},$$

or equivalently,

$$\begin{bmatrix} m_0 \\ m_1 \\ m_2 \\ m_3 \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -s_1 & 0 & s_1 & 0 \\ 2s_1 & (s_2-3) & (3-s_1) & -s_2 \\ -s_1 & (2-s_2) & (s_1-2) & s_2 \end{bmatrix} \begin{bmatrix} p_0 \\ p_1 \\ p_2 \\ p_3 \end{bmatrix}.$$

From this expression we see how to spline an arbitrarily long stream p_i of control points. We are free to choose a pseudo-speed s_i at each control point (except the first and the last). The interlacing of the tangents at the control points guarantees that the velocity is continuous at the corners. In practice,

we choose all pseudo-speeds to be equal, $s_i = s_*$, but vary s_* interactively. There are two distinguished values for the pseudo-speed, $s_* = 0$ and $s_* = \frac{1}{2}$. With $s_* = 0$, we find that $q(t) = p_1 + (3t^2 - 2t^3)(p_2 - p_1)$. Therefore, for this pseudo-speed the spline runs along the polygon connecting the control points, but with more interpolating points near the corners than in the middle.

In order to distribute the interpolating points more evenly (for evenly chosen values of $t \in [0, 1]$), we choose $s_* = \frac{1}{2}$. In this case, when the $p_i = p_0 + iu$ are uniformly distributed along a straight line, the cubic interpolation degenerates to a truly linear one, $q(t) = p_1 + tu$. In practice, values of s_* closer to 1 makes for better shaped curves and surfaces, though interpolating splines are never likely to form beautiful surfaces.

More interesting is the *tensorial* nature of these splines. For simplicity, consider first a four-by-four array of control points, p_{ij}. If we first spline in one direction, we obtain four cubic polynomials, $q_j(t) = t^a Q_{ai} p_{ij}$, where Q is the matrix determined above and we use the Einstein convention. If now we spline the quadruples produced by the interpolated values of t, using a different time parameter, s, we obtain cubic polynomials $r(s, t) = s^c K_{cj} q_j(t) = s^c t^a Q_{ai} K_{cj} p_{ij}$. We would get exactly the same result if we splined first in the other dimension. Note that we could choose pseudo-speeds independently for each dimension at each control point.

A graphical user interface (written by Matthew Stiak) lets us easily draw and move the level curves M_{tu} (as in Color Plate 9). It promises to make `illiLevel` into a versatile visualization tool that handles time and space in a uniform and interchangeable manner.

References

1. F. Apéry, *An algebraic halfway model for the eversion of the sphere*, Tohoku Math. J. **44** (1992), 103–150, with an appendix by B. Morin.
2. T. F. Banchoff, *Triple points and surgery of immersed surfaces*, Proc. Amer. Math. Soc. **46** (1974), 407–413.
3. W. Boy, *Über die Curvatura integra und die Topologie geschlossener Flächen*, Math. Ann. **57** (1903), 151–184.
4. K. A. Brakke, *The Surface Evolver*, Experimental Math. **1**:2 (1992), 141–165.
5. K. A. Brakke and J. M. Sullivan, *Using symmetry features of the surface evolver to study foams*, Visualization and Mathematics (H.-C. Hege and K. Polthier, eds.), Springer Verlag, Heidelberg, 1997, pp. 95–117.
6. R. Bryant, *A duality theorem for Willmore surfaces*, J. Differential Geometry **20** (1984), 23–53.
7. E. Catmull and R. Rom, *A class of interpolating splines*, Computer Aided Geometric Design (R. Barnhill and R. Riesenfeld, eds.), Academic Press, 1974, pp. 317–326.
8. B. DeWitt, *Eversion of the 2-sphere*, Battelle Rencontres, 1967, Benjamin, New York, 1967, pp. 546–557.
9. G. Francis, *Drawing surfaces and their deformations: The Tobacco pouch eversions of the sphere*, Math. Modelling **1** (1980), 273–281.
10. ———, *A topological picturebook*, Springer-Verlag, New York, 1987.
11. ———, *On knot-spanning surfaces: An illustrated essay on topological art*, The Visual Mind (M. Emmer, ed.), Leonardo Books: MIT Press, 1993.
12. G. Francis and B. Morin, *Arnold Shapiro's eversion of the sphere*, Math. Intelligencer **2** (1979), 200–203.
13. G. Francis, J. M. Sullivan, K. Brakke, R. Kusner, D. Roseman, A. Bourd, C. Hartman, G. Chappell, and J. Rubenstein, *LATERNA matheMAGICA*, Virtual Environments and Distributed Computing at SC'95: GII Testbed and HPC Challenge Applications on the I-WAY (H. Korab and M. D. Brown, eds.), ACM/IEEE Supercomputing'95, 1995.
14. G. Francis, J. M. Sullivan, R. B. Kusner, K. A. Brakke, C. Hartman, and G. Chappell, *The minimax sphere eversion*, Visualization and Mathematics (H.-C. Hege and K. Polthier, eds.), Springer Verlag, Heidelberg, 1997, pp. 3–20.
15. L. Hsu, R. Kusner, and J. M. Sullivan, *Minimizing the squared mean curvature integral for surfaces in space forms*, Experimental Mathematics **1**:3 (1992), 191–207.
16. H. Karcher and U. Pinkall, *Die Boysche Fläche in Oberwolfach*, Mitteilungen der DMV **1997**:1, 45–47.
17. N. Kuiper, *Convex immersions of closed surfaces in E^3*, Comm. Helv. **35** (1961), 85–92.

18. R. Kusner, *Conformal geometry and complete minimal surfaces*, Bull. Amer. Math. Soc. **17** (1987), 291–295.

19. _____, *Comparison surfaces for the Willmore problem*, Pacific J. Math. **138** (1989), 317–345.

20. S. Levy, *Making waves: A guide to the ideas behind Outside In*, AK Peters, Wellesley, MA, 1995.

21. S. Levy, D. Maxwell, and T. Munzner, *Outside In*, AK Peters, Wellesley, MA, 1994, narrated videotape (21 min) produced by the Geometry Center, University of Minnesota.

22. N. L. Max, *Turning a sphere inside out*, International Film Bureau, Chicago, 1977, narrated videotape (21 min).

23. B. Morin, *Équations du retournement de la sphére*, Comptes Rendus Acad. Sci. Paris **287** (1978), 879–882.

24. A. Phillips, *Turning a sphere inside out*, Sci. Amer. **214** (1966), 112–120.

25. S. Smale, *A classification of immersions of the two-sphere*, Trans. Amer. Math. Soc. **90** (1959), 281–290.

26. T. J. Willmore, *Note on embedded surfaces*, An. Stiint. Univ "Al. I. Cuza" Iasi Sect. I, a Mat. **11** (1965), 493–496.

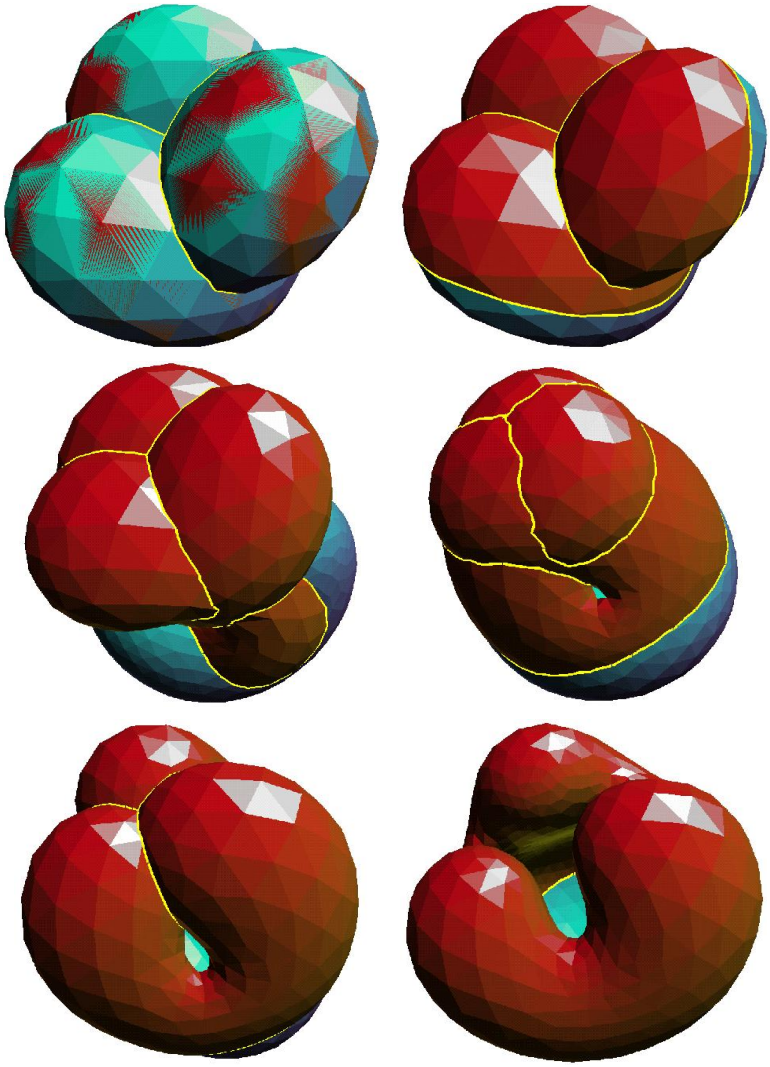

Fig. 7. Six stages in the $p = 3$ minimax eversion, starting in the top left with a double cover of a Boy's surface critical for the Willmore energy W. (Because the two sheets have opposite orientations, we see their red and blue sides; at this stage they are in the same place, so give mottled colors when rendered.) Next, the two sheets have pulled apart; here the white double locus between red surfaces is a four-fold cover of the Boy's surface double locus, while the part between red and blue is the extra curve seen in Fig. 3. In the middle row, the surface starts to be simplified by minimization of W. At right, a pair of triple points has just disappeared in front; because of the eversion's three-fold symmetry, this happens in three places. At the bottom, the three fingers have straightened and almost lost contact. In the last picture, the fingers are retracting; the double locus where the north pole pushes through the south pole will soon disappear.

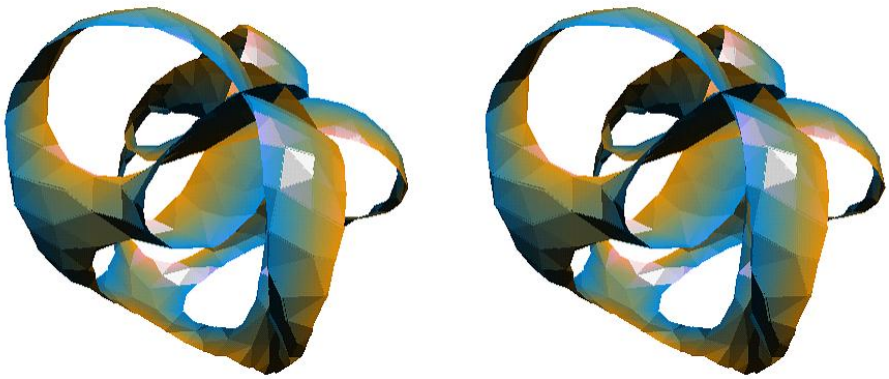

Fig. 8. A slice of the double locus surface for the $p = 2$ minimax eversion, near $t = 0$, with the t direction projected out, drawn in cross-eyed stereo. Six sheets of this surface cross where the halfway model has a quadruple point; the five isthmus events are saddles here. Compare Fig. 4, which shows three individual levels from which this surface was pieced together.

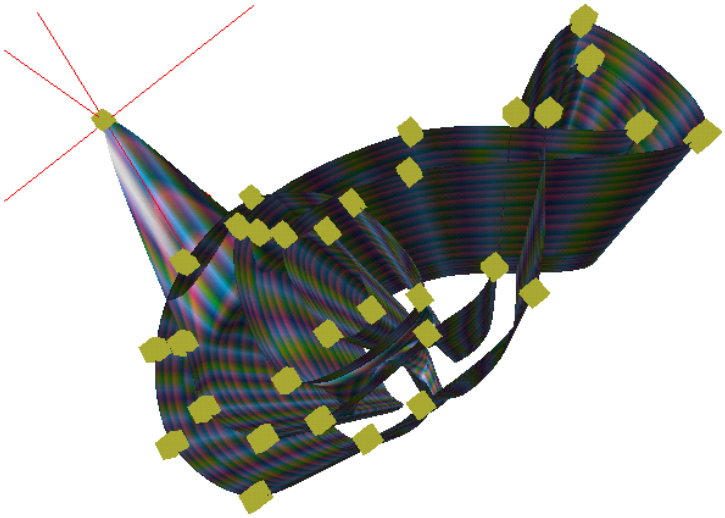

Fig. 9. The double cover of a Boy surface with sheets pulled apart, M_1, can be given by splines. Here we see the slab $M_1[-1, +1]$. The graphical user interface for `illilevel` lets us interactively move control points, as at the upper left.