

1992

The Surface Evolver

K. A. Brakke
Susquehanna University

Follow this and additional works at: http://scholarlycommons.susqu.edu/math_fac_pubs

Recommended Citation

K. Brakke, The Surface Evolver, *Experimental Mathematics* vol. 1 no. 2 (1992), 141-165.

This Article is brought to you for free and open access by Scholarly Commons. It has been accepted for inclusion in Mathematical Sciences Faculty Publications by an authorized administrator of Scholarly Commons. For more information, please contact sieczkiewicz@susqu.edu.

The Surface Evolver

Kenneth A. Brakke
Mathematics Department
Susquehanna University
Selinsgrove, PA 17870
email: brakke@susqu.edu

May 13, 1992

This version is close but not identical to the published version in *Experimental Mathematics*, **1** (1992) 141–165. Some figures have been regenerated, and the bibliography updated. Pagination is different, and the text may not be exactly identical.

Abstract

The Surface Evolver is a computer program that minimizes the energy of a surface subject to constraints. The surface is represented as a simplicial complex. The energy can include surface tension, gravity, and other forms. Constraints can be geometrical constraints on vertex positions or constraints on integrated quantities such as body volumes. The minimization is done by evolving the surface down the energy gradient. This paper describes the mathematical model used and the operations available to interactively modify the surface.

AMS classification (1990):

Primary:

49-04 Calculus of variations explicit machine computation and programs.

Secondary:

35-04 PDE explicit machine computation and programs.

49Q05 Minimal surfaces.

49Q20 Variational problems in a geometric measure-theoretic setting.

Keywords: minimal surfaces, calculus of variations.

Introduction

One of the fundamental problems in the calculus of variations is to find a surface minimizing some energy subject to constraints. A soap film on a wire frame minimizes its area subject to its boundary staying on the frame. A cluster of soap bubbles minimizes the total soap film area subject to enclosing fixed volumes in each bubble. A capillary surface minimizes the gravitational energy of a fluid in a vessel plus the surface energy of its free surface and its contact energy with the vessel walls. Other examples of surfaces are grain boundaries in metals, crystal facets, fluid interfaces, and cell membranes [Almgren 1982; Almgren and Taylor 1976]. Note that these surfaces need not be simply connected, need not be orientable (as in a Mobius band soap film), and need not be manifolds (as in bubble clusters).

The Surface Evolver is an interactive program for the study of surfaces shaped by surface tension and

other energies. The user specifies an initial surface, the constraints that the surface should satisfy throughout the evolution, and an energy function that depends on the surface. The Evolver then modifies the surface, subject to the given constraints, so as to minimize the energy. The user can intervene during the evolution, changing the surface's properties or applying certain operations to keep the evolution well-behaved.

The action of the Evolver is meant to model the process of evolution by mean curvature, which was studied in [Brakke, 1977] for surface tension energy in the context of varifolds and geometric measure theory. The energy in the Evolver can be a combination of surface tension, gravitational energy, squared mean curvature or user-defined surface integrals. The Evolver can handle complicated topology (as seen in real soap bubble clusters), volume constraints, boundary constraints, boundary contact angles, prescribed

mean curvature, crystalline integrands, gravity, and constraints expressed as surface integrals.

The main focus of the Surface Evolver is on two-dimensional surfaces in three-dimensional space, the so-called soap-film model. The principal data structures are set up with this in mind, but they are designed in such a way that the dimension of the ambient space and the "surfaces" of interest can be arbitrary. Thus, one-dimensional strings and higher-dimensional surfaces can also be handled. Moreover, the ambient space can be endowed with an arbitrary Riemannian metric, and even be a quotient space under a group action.

The Evolver has a graphical interface that allows the user to follow the evolution of the surface on the screen. The graphics can also be output to files in several formats, including PostScript.

The Surface Evolver program is freely available (see "Software Availability" at the end of this article) and is in use by a number of researchers. Some of the applications of the Evolver so far include modelling the shape of fuel in rocket tanks in low gravity [Tegart 1991], calculating areas for the Opaque Cube Problem [Brakke 1991b], computing capillary surfaces in cubes [Mittelmann] and in exotic containers [Callahan et al. 1991], simulating grain growth, studying grain boundaries pinned by inclusions, searching for partitions of space more efficient than Kelvin's tetrakaidecahedra, modelling the shape of molten solder on microcircuits [Racz et al.], studying polymer chain packing, and classifying minimal-surface singularities. Section 9 of this paper gives a proof, based on the use of the Evolver, that a conjectured area-minimizing cone in \mathbf{R}^4 is not area-minimizing.

The strength of the Surface Evolver program is in the breadth of problems it handles, rather than optimal treatment of some specific problem. It is under continuing development, and not every feature is described in this paper. Users are invited to suggest new features.

This paper describes the capabilities of the Surface Evolver so that readers can evaluate its usefulness in their research and so that users have a published description to cite. It is not a substitute for the Surface Evolver Manual [Brakke 1991a], although it does mention certain operational details so that users can find the corresponding features in the program or manual.

Acknowledgments.

The Surface Evolver is a result of my participation in The Geometry Center at the University of Minnesota (formally The National Science and Technology Research Center for Computation and Visualization of Geometric Structures, formerly the Geom-

etry Supercomputer Project), I would like to thank Fred Almgren, Jean Taylor, and Al Marden for making my participation in The Geometry Center possible and for their enthusiasm. I would also like to thank Charlie Gunn, Toby Orloff, Stuart Levy, and the rest of the staff of The Geometry Center for technical assistance, especially in graphics. Thanks also to John Sullivan, who has used the Evolver since its inception and has made many suggestions and found many bugs.

1 Surface models

1.1 Representing surfaces

Surfaces have been represented mathematically as graphs of functions, level sets of functions, images of maps, measures, simplicial complexes, polyhedral complexes, spline patches, etc. Each way has its strengths and weaknesses. The Surface Evolver uses a simplicial complex representation. This permits the representation of surfaces like soap films and bubbles, which may have arbitrary topologies and may not be orientable. It permits the exact specification of a surface in terms of a finite number of parameters, the vertex coordinates. See the figures accompanying this paper for some simple examples.

The main focus of the Surface Evolver is on two-dimensional surfaces in three-dimensional space, which is called the *soapfilm model*. The principal data structures are set up with this in mind, but they are so designed that the dimension of the ambient space can be arbitrary and one-dimensional "strings" can also be represented. Higher dimensional surfaces can also be handled to a limited degree.

The units of measurement are dimensionless. If the user wishes to model a specific physical problem, then all values should be in one consistent set of units such as cgs or MKS.

Section 3 describes what is available in the Evolver in terms of alternatives or elaborations to the basic soap-film model: "surface" of arbitrary dimension, ambient spaces with an arbitrary Riemannian metric, and so on.

1.2 Energies

Broadly speaking, the energies that the Evolver minimizes are any quantities that may be expressed as integrals over the surface. Foremost among them is surface tension. Soap films and interfaces between different fluids have an energy proportional to their area, which can also be regarded as a surface tension, or force per unit length. That is, across any line in the surface, there is a tension whose value is the

same as the surface energy density. In the Evolver, the user can specify a value of the surface tension for each facet.

Another common energy is gravitational potential energy, which can be written as a surface integral by means of the divergence theorem. Capillary surfaces may be modeled by including both surface tension and gravitational energy in the total energy. For more details, see section 4.

1.3 Constraints

Several types of constraints are available in the Evolver. Vertices may be fixed in place, reflecting the fact that the edge of a soap film should be attached to a wire. Vertices may be constrained to lie on smooth manifolds, reflecting the case when the edge of a soap film should lie on a wall, but is free to move on the wall. Edges and facets may likewise be constrained, which simply means that any vertices generated on them will inherit those constraints. Bodies may be constrained to have fixed volumes, as in the case of the volume of the column of liquid underneath a capillary surface. Section 5 explains the exact mathematical form of the various types of constraints.

1.4 Basic Operations

The fundamental operation of the Evolver is the iteration step, which reduces energy while obeying any constraints. A gradient descent method is used. The force at each vertex is the gradient of the total energy as a function of the position of that vertex. Every vertex is moved simultaneously by a global multiple of its force. This multiple, called the *scale factor*, can either be fixed by the user or be the factor that optimizes the decrease in energy. Details of all the options available for the iteration step are given in Section 6.

Several operations are available for manipulating the triangulation. *Refinement* is the subdivision of each facet into four similar facets, for better approximation of curved surfaces. *Equiangulation* readjusts the triangulation of a surface to make the facets as nearly equilateral as possible. *Vertex averaging* moves each vertex to the average position of its neighboring vertices. These and other operations, including some that change the topology of the surface, are more fully described in section 7.

2 Three examples of the Evolver in action

2.1 The Catenoid

The catenoid is the minimal surface whose boundary is two parallel rings not too far apart. It is an extremely simple surface, yet it illustrates some of the subtleties of evolving triangulated surfaces. Stages in its evolution are shown in Figure 1. The surface in the initial data file consists of six rectangles forming a cylinder between the two rings. In general, a data file contains only the minimum amount of information needed to correctly define the topology of the surface. When initially read in, the rectangles are automatically triangulated into facets (top left). The vertices and edges on the rings are fixed. The rings themselves are not shown. With so few facets, the initial surface cannot shrink, so the user refines the surface (top middle). Here, only one refinement is done to keep the facets large enough to be seen easily. Normally, there would be alternating stages of refinement and iteration. Note that the vertices created by subdividing the edges on the rings are themselves fixed on the rings. Equiangulation gives the much nicer triangulation shown top right, by switching the diagonals of some quadrilaterals to make the facets more equiangular. Fifty iterations with optimizing scale factor result in an area of 6.458483 (bottom left). At this point, each iteration is reducing the area by only .0000001, the triangles are all nearly equilateral, everything looks nice, and the innocent user might conclude the surface is very near its minimum. But this is really a saddle point of energy. Further iteration shows that the area change per iteration bottoms out about iteration 70, and by iteration 300 the area is down to 6.4336 (bottom right), near the true local minimum. We know it is a minimum because iteration produces no change, and the Evolver can calculate the Hessian matrix to be positive definite. One can see that the triangulation really wants to be twisted around so that there are edges following the lines of curvature.

If the two rings are too far apart, the neck of the catenoid will shrink down to a point, as shown in Figure 1. Upon iteration, the neck forms a ring of very short edges (top middle). These edges can be removed by identifying their endpoints, in a step (taken by the user) known as tiny edge weeding (§7.5). This produces a single vertex at the neck (top right). The Evolver can recognize the topology around the neck vertex as improper for a soap film and split the vertex into two (bottom left), when instructed to do so by the user. This called vertex popping (§7.7) The two parts of the surface then quickly collapse to disks (bottom right).




Figure 1: The evolution of a stable catenoid. Top left: The initial surface. The boundary wire circles are not shown. The rectangular faces of the datafile have been automatically triangulated. Top middle: After one refinement. Note how vertices on top and bottom edges follow the boundary wire circles. Top right: After equiangularization. Note the edges that have switched direction. Bottom left: After iterating fifty times. This is a saddle point in area. Bottom right: Ultimate endpoint of iteration, with edges following the lines of curvature, which are horizontal and vertical.




Figure 2: The evolution of an unstable catenoid. The neck pinches down to a point, and the surface splits into two pieces.




Figure 3: Capillary surface making contact with a wall. Left: the initial surface. The facets on the vertical wall (darker shading) have surface tension 0.5, while the ones on the horizontal free surface have surface tension 1.0. Thus the equilibrium contact angle is 60° . The free surface is bounded by a fixed wire in front and plane constraints on either side. Center left: After some evolution, the free edge has crept up the wall. Note that several wall facets are getting very narrow, which would soon cause problems in the evolution. Center right: After vertex averaging, the free edge has more room to migrate without overrunning vertices. Right: the final equilibrium surface.

2.2 Capillary Surface Meeting a Wall

Figure 3 shows the evolution of an example where we use constraints and varying surface tension to model a surface meeting a wall at a given contact angle – here 60 degrees. This is the situation for a capillary surface, but this example does not include volume constraints or gravity.

The initial configuration is shown on the left. The fluid surface is the light gray horizontal surface, which has one edge movably attached to the dark gray vertical wall. The junction of the surface and wall consists of edges belonging to both surface and wall facets. We give the wall facets half the surface tension of the surface facets, so the equilibrium contact angle is 60 degrees ($0.5 = \cos 60^\circ$). All vertices on the wall are constrained to stay in a vertical plane. The three vertices at the top of the wall are fixed in place, as are the vertices of the surface on the edge opposite the wall. The two lateral sides of the surface and wall are constrained to lie in vertical planes (not shown) perpendicular to the wall. These planes contain no facets and contribute no energy, so the equilibrium contact angle is 90 degrees.

After several iterations (center left), the contact line has moved up the wall, seeking the equilibrium contact angle. The shrinking of the wall facets more than offsets the stretching of the surface facets. The interior vertices of the wall do not move, since there are no net forces on them. This can give rise to problems if the contact line overruns interior wall vertices, as is about to happen. At this point, the user must intervene to adjust the triangulation, using vertex averaging, after which (center right) the contact line can continue to move up the wall. The final equilibrium state consists of a plane surface (far right).

The wall facets serve two purposes: to generate the correct contact angle and to help visualize the

surface. But they also are the source of problems as the surface moves up the wall, requiring repeated vertex averaging. They occupy more computer memory and calculation time than necessary. It is possible to omit them. In §4.5 I show how to use edge integrals to compute the equivalent energy of the facets. The visualization function could be served by having fixed facets on the wall that do not participate in the calculations and do not refine.

2.3 Grain Growth

When liquid metal solidifies, crystallization generally starts at many nuclei, with random orientations. The crystal lattices are mismatched where the grains meet, and the atoms along the grain boundaries are in a higher energy state than interior atoms. To a good approximation, the energy is independent of the orientation of the boundary. In the process of annealing, the metal is warmed enough for boundary atoms to switch from one lattice to the other through thermal motions, and the boundaries migrate at a rate proportional to its curvature, assuming that impurities or other obstacles do not interfere.

Figure 4 shows this process for a two-dimensional metal in a unit flat torus that initially crystallizes from 100 random nuclei, resulting in an initial grain configuration of 100 Voronoi cells (top left). The ultimate product of evolution consists of four unequal hexagonal grains. A video of the evolution is available in [Brakke 1992b].

Modeling the dynamics of the evolution requires using a fixed scale factor (the time step) much smaller than the optimizing scale factor. Here, a time step of 5×10^{-6} was used. The Evolver has several features used to automate the evolution, including automatic topology changes in the string model. With a feature called *autopopping*, any edge whose length




Figure 4: The evolution by mean curvature of 100 Voronoi cells in a two-dimensional flat torus. The time step is 5×10^{-6} . From left to right: The initial configuration; configuration after 200 steps, still with 100 cells; after 1600 steps, with 66 cells left; and after 5000 steps, with 35 cells left.

is projected to become less than the critical length is deleted and any improper vertices resulting are popped. The critical length is automatically set to the critical length for stability described in §6.4, which works out to 0.003 here. Only edges whose lengths are decreasing are tested in order not to eliminate the very short edges generated by vertex popping. Also, with *autochopping*, edges that become longer than a chosen cutoff length of 0.015 are automatically subdivided.

3 Surface Models

This section describes several variations on the basic soap-film model discussed in §1.1.

3.1 One-Dimensional "Surfaces": The String Model

The "surface tension" can be declared to reside in edges instead of facets. The surface then becomes a network of elastic strings, hence the term *string model* is applied to this mode of operation. The strings may reside in any dimensional space, but if the domain is two-dimensional, then the strings may bound regions. In this case, a region is defined with a facet structure and a body structure. The facet may have any number of sides. The body has just one facet on its boundary. The effect is to stretch the string network into a cylindrical surface of height 1, whence the mechanisms for surfaces can be applied with minimal changes. The grain growth example in §2.3 uses the string model.

3.2 Quadratic model

In an attempt to approximate curved surfaces better than by flat facets, there is a mode in which each facet is a quadratic spline patch. A midpoint is added to each edge, giving a total of six control

points. Each coordinate is then quadratically interpolated from these six points to form the surface. An edge becomes a curve that depends only on the control points on the edge, thus guaranteeing that neighboring facets meet without a gap. The disadvantages of the quadratic mode are that it is slower, surface area is calculated approximately by numerical integration, facets are displayed as if flat, and some Evolver features are not implemented.

3.3 Higher-Dimensional Surfaces

Higher-dimensional surfaces cannot be represented by the basic vertex-edge-facet scheme. There is a mode of operation that permits the facets making up the surface to be represented directly as simplices (vertex lists). This permits an arbitrary dimensional surface in an arbitrary dimensional space. However, many features are not implemented yet for this mode.

3.4 Quotient spaces

The ambient space can be a quotient space of \mathbf{R}^n under some symmetry group. The vertex coordinates are taken to be in a fundamental region, and each edge is marked with a group element to tell how its head vertex should be transformed (wrapped) relative to its tail. The user invents an integer representation for the group elements to be used in marking edges. The flat torus quotient space is built-in with some special notation in the datafile for specifying the fundamental parallelepiped and the wraps of the edges. Using some other quotient space requires the user to write C functions that handle group transformations and compositions.

The display of a surface in a quotient space poses some interesting problems, since the display space is Euclidean. The Evolver offers three options, two of which are illustrated in Figure 5, which shows a surface in a flat three-dimensional torus whose fundamental unit is a unit cube. The three options are:

Raw facets. Each facet is displayed as it is located in the fundamental domain. Location is based on the first vertex in the facet, with other vertices unwrapped as needed.

Connected bodies. For each body, all facets on the boundary of that body are gathered in a list. One point is chosen as a base point, and unwrapping of vertices spreads out from neighbor to neighbor facet until the entire body surface is done. This nicely displays each logical body as a single visual body. An example is shown in Figure 5 (left).

Clipping. For a torus domain, each facet is clipped to the fundamental parallelepiped. Pieces lying outside are wrapped around so they lie inside. This option clearly shows the fundamental domain and how it wraps around. An example is shown in Figure 5 (right).

The surface shown in Figure 5 bounds two tetrakaidecahedra. Lord Kelvin [Thompson 1887] conjectured that the optimal way to partition space into equal volume cells with least area is a packing of very slightly curved tetrakaidecahedra. Several people (including the author) have used the Evolver in attempts to beat Kelvin's partition, but nobody has succeeded yet.

3.5 Background metric

The ambient space can be endowed with a Riemannian metric. Only one coordinate patch is allowed, but quotient spaces are possible. Basically, the Evolver operates as usual in Euclidean coordinates, except that the metric is used for the calculation of edge lengths and facet areas. Edges are not taken to be geodesics, nor are facets geodesic surfaces; rather they keep their Euclidean shape. Surfaces are displayed as if their coordinates were in Euclidean space.

An example using a metric is discussed in section 9 below. Other possible ways to use a metric are: modeling a cylindrically symmetric surface by means of the string model, and implementing a spatially varying scalar surface energy, as in a surface whose own weight is not negligible.

The metric need not be positive definite. One can do minimal surfaces in a Minkowski metric, but it takes a little care.

3.6 Internal Representation

Each geometric element (vertex, edge, facet, body) is implemented as one data structure. An element is stored as an oriented entity, but may be referred to with this orientation or the inverse one. There is

a data type *element_id* that contains a pointer to an element structure and a relative orientation (normal or inverted). This type is used for all references to elements. The connectivity of the surface is specified by having links in each element structure to the next higher- or lower-dimensional elements it intersects. One design principle followed is that each element should contain links to a fixed number of other elements. Thus the body structure does not record bounding facets itself; rather, the facet structure has two slots to record which body (if any) is on each of its sides.

The surface connectivity is completed by introducing *facet-edge* structures, which are a simplification of the scheme described in [Dobkins and Laszlo 1987]. A facet-edge structure contains links to a facet and an edge on the facet's perimeter (with proper boundary orientation), plus links to the previous and next facet-edges around the facet and to the previous and next facet-edges around the edge. This permits a quick run-through of all facets containing a given edge, and also the representation of facets with an arbitrary number of sides, which is necessary in certain situations.

4 Energies

The Surface Evolver tries to minimize the total energy of a surface. This energy may have several components.

4.1 Surface tension

Soap films and interfaces between different fluids have an energy proportional to their area. In the Evolver, each facet has a surface tension of 1 unless otherwise specified. Different facets may have different surface tensions. It is possible to endow both facets and chosen edges with tension in order to model surfaces where singular curves have energy, as in [Morgan 1994].

Contact angles between free surfaces and walls, as in capillary problems, can be specified by introducing facets that are confined to the wall and have a different surface tension, as in the wall example in §2.2. Negative tensions are allowed, so all contact angles are possible. However, this method has the drawback that a moving free boundary on the wall can overrun wall facets, as noted in Figure 3. Another method of prescribing contact angles, described in §4.5, uses edge energy integrals.

There is no general mechanism yet to include the integral over the surface of a general scalar integrand which may depend on position and tangent plane orientation. However, the use of a Riemannian metric




Figure 5: Two of Kelvin’s tetrakaidecahedra in a flat torus. The fundamental region is a unit cube. A. The surface plotted as the boundaries of connected bodies. B. The surface clipped to the fundamental region, a unit cube.

on the ambient space can often achieve the same effect.

4.2 Crystalline integrands

The Evolver can model energies of crystalline surfaces. A crystalline surface energy density depends on the direction of the surface normal vector. Such a quantity is known as a *crystalline integrand* [Taylor 1983; 1988]. The energy density is given by the largest dot product of the surface normal with a set of vectors known as the *Wulff vectors*. Surface area can be regarded as a crystalline integrand whose Wulff vectors coinciding with the unit sphere. In the Evolver, a finite set of Wulff vectors may be specified, and the corresponding crystalline energy computed.

4.3 Surface integrals

A facet may contribute an energy resulting from integrating a vectorfield over the facet as a surface integral. Multiple vectorfields may be defined in the datafile as functions of the coordinates, and any facet may use any number of them. Besides surface energies, these can be used to do volume energies, such as gravitational energy, by using the Divergence Theorem to convert a volume integral into a surface integral. Integrals are done numerically using Gaussian quadrature.

4.4 Gravity

A body B having a density ρ contributes its gravitational energy to the total. The acceleration of gravity G is under user control. Letting ρ be the body density, the energy is defined as

$$E = G\rho \int \int \int_B z \, dV,$$

but is calculated by the Divergence Theorem as

$$E = G\rho \int \int_{\partial B} \frac{z^2}{2} \vec{k} \cdot d\vec{S}.$$

The integral is taken over each facet that bounds a body. If a facet bounds two bodies of different density, then the appropriate difference in density is used. Vertical facets or facets lying in the $z = 0$ plane make no contribution, and may be omitted if they are otherwise unneeded. Facets lying in constraints may be omitted if their contributions to the gravitational energy are contained in edge energy integrals.

Gravity is a special case of a surface integrand, but it is implemented internally for several reasons: it can be evaluated exactly without numerical integration; it is common enough to be worth saving the user the trouble of setting it up; and, in general, evaluation of user-defined integrand expressions is slower than that of compiled-in energies.

The built-in gravity does not apply to Riemannian metrics or quotient spaces; users must define their own gravitational energy integrands in such cases.

Gravity applies to bodies, not surfaces. Surfaces are weightless. If one did want heavy surfaces, one could use the metric mechanism to simulate a scalar surface integrand.

4.5 Edge integrals

An edge may contribute an energy resulting from integrating a vectorfield over the edge as a line integral. The objective of this is to let the free edges of a surface have energy. This is useful to control the contact angle of a surface on a wall. As mentioned earlier in section 3.1, the contact angle can be specified by giving an energy density to the wall on one side of the surface edge. Alternatively, an edge integral can be defined to give an energy equivalent to the wall energy by means of Stokes’ Theorem. This eliminates the need to coat the wall with facets. Likewise, edge integrals can be used to replace other facet energies, such as gravitational energy, for facets on constraints.

An edge integral is evaluated once for each edge – regardless of how many facets it is on – using the orientation given in the datafile.

4.6 Prescribed mean curvature and pressure

For an equilibrium surface with a constant surface tension, the mean curvature can be interpreted as proportional to the pressure difference across the surface. Therefore, prescribing the mean curvature is equivalent to prescribing the pressure difference. The Evolver permits the user to prescribe pressures in bodies. Pressure can be defined as the rate of change of energy with respect to volume, so the pressure feature is implemented by having each body with a prescribed pressure P contribute energy

$$E = -PV,$$

where V is the actual volume of the body. The energy is actually calculated by a surface integral

$$E = -P \int \int_{\partial B} z \vec{k} \cdot d\vec{S}.$$

The desired surface need not really be the entire boundary of a body; it may have fixed edges. The energy contributed by the omitted boundary is constant and so does not affect the shape of the surface.

This method will only work if the desired surface is stable at the prescribed pressure. If curvature decreases as volume increases, then the surface will either blow up or implode. For example, if a round soap bubble of surface tension T and initial radius R_0 is prescribed to have a pressure of P and $P > 2T/R_0$, then the pressure force will cause the bubble to expand, which reduces the curvature $1/R$, so the pressure can never be balanced by the curvature, and the bubble expands indefinitely.

4.7 Squared mean curvature

There are circumstances under which one wants the energy to include the integral over the surface of the squared mean curvature. For example, surfaces which minimize this integral are by definition *Willmore surfaces*. This presents a slight problem for a piecewise linear surface, as the mean curvature (in the form of first variation measure) is singular and concentrated on the edges. Its square integral is therefore always infinite. However, it is possible to come up with a usable approximation. An average mean curvature around each vertex can be calculated, and the integral of the square of this average counted as energy.

The definition of mean curvature used here is a variational one, and corresponds to the average of the sectional curvatures rather than their sum. The integral of squared mean curvature in the soapfilm model is calculated as follows: Each vertex v has a star of facets around it of total area A_v . The force on the vertex is

$$\vec{F}_v = -\frac{\partial A_v}{\partial v}.$$

Since each facet has three vertices, the area associated with v is $A_v/3$. Hence the average mean curvature at v is

$$\vec{h}_v = \frac{1}{2} \frac{\vec{F}_v}{A_v/3},$$

and this vertex's contribution to the total integral is

$$E_v = h_v^2 A_v/3 = \frac{3 F_v^2}{4 A_v}.$$

E_v can be written as an exact function of the vertex coordinates, so the gradient of E_v can be fed into the total force calculation.

The alternative to locating curvature at vertices is to locate it on the edges, where it really is, and average it over the neighboring facets. But this has the problem that a least area triangulated surface would have nonzero squared curvature, whereas in the vertex formulation it would have zero squared curvature.

Squared mean curvature is also implemented for the string model, but not for quadratic models. In the string model, let L_v be the sum of the lengths of the edges adjacent to v , so the force on a vertex is

$$\vec{F}_v = -\frac{\partial L_v}{\partial v}.$$

Each edge has two endpoints, so the length associated with v is $L_v/2$, so the curvature is

$$\vec{h}_v = \frac{\vec{F}_v}{L_v/2},$$

and this vertex's contribution to the total integral is

$$E_v = h_v^2 L_v/2 = \frac{2 F_v^2}{L_v}.$$

4.8 Gaps




Figure 6: Surface in a ring, showing the gap problem. A. The initial surface. The vertices are free to move along the boundary wire. B. After one iteration, the gaps have grown and the surface area shrunk.

Consider a soap film spanning a circular wire. The Evolver must approximate this surface with a collection of facets, as shown in Figure 6. The straight

edges of these facets cannot conform to the curved wire, and hence the computed area of the surface leaves out the gaps between the outer edges and the wire. If the vertices are free to move along the wire, then the Evolver will naturally try to minimize area by moving the outer vertices around so the gaps increase, as in Figure 6 right, ultimately resulting in a surface collapsed to a line. This is not good. Sometimes the vertices can be fixed on the wire, but other times this is not possible, for example when the surface is spanning the inside of a cylinder. Therefore there is provision for a “gap energy” to discourage this. A constraint of the type defined in §5.1 may be declared *convex* in the datafile. For an edge on such a constraint, an energy is calculated

$$E = \frac{k}{6} \left\| \vec{S} \times \vec{Q} \right\|$$

where \vec{S} is the edge vector and \vec{Q} is the projection of the edge tangent to the constraint at the tail vertex of the edge. The constant k is a global constant called the *gap constant*. A gap constant of 1 gives the best approximation to the actual area of the gap. A larger value of k minimizes gaps and gets vertices nicely spread out along the wire.

Another way to handle gaps is to define an edge integral (see §4.5) that is zero on the constraint and positive on the convex side of the constraint. Edge integrals are evaluated on interior points of the edge, so the bigger the gap, the bigger the integral. This encourages the edges to be equal length.

Of course, any addition of energy changes the problem slightly. But this energy decreases quadratically with the fineness of the triangulation, and so should not change the solution significantly. By changing the constant associated with this energy, one can see whether the problem is being significantly altered.

In actual practice, gap energy would seldom be used with a wire boundary, since it would be much simpler just to declare the vertices on the wire fixed. The real use for gap energy comes with surface edges on walls, where the vertices cannot be fixed.

5 Constraints

Energy minimization takes place subject to constraints of two types: constraints on the motion of vertices and constraints on the value of surface integrals. Vertices can be individually constrained by declaring them fixed, by confining them to level sets of functions (which is Evolver’s narrow meaning of the term *constraint*), or by defining their position in terms of parameters (Evolver *boundaries*). Surface integrals in general are called *quantities*, and the

particular case of body volumes is implemented internally.

5.1 Level-Set Constraints

A vertex may be confined to the zero level set of one or more functions. Such a function is called a *constraint* in Evolver terminology. It should be clear from context when the term *constraint* is used in this narrow sense or in the broader mathematical sense in this paper. The default is the narrow sense. Constraint functions are defined by the user in the initial datafile and numbered for reference. Vertices may be declared to be on one or more constraints simultaneously, but it is the user’s responsibility to ensure that the constraint function gradients at a vertex are linearly independent. There are also one-sided constraints, which means a vertex is restricted to the region where the constraint function has nonnegative (or nonpositive) values.

When a vertex is moved for whatever reason, Newton’s method is used to project it back to its constraints. There is a global constraint tolerance parameter the user can set to control the accuracy of constraint satisfaction. The coordinates of a vertex in the initial datafile do not have to exactly satisfy its constraints; it will be automatically projected. A vertex on constraints may also be declared fixed, which means that it will not move after its initial projection. If a constraint is modified during runtime by changing an adjustable parameter, all vertices are projected again so as to satisfy the new constraint.

Edges and facets may be declared to be on constraints, which means that all vertices generated by subdividing them will be on the same constraints.

A single constraint is the best way to attach a free edge of a surface to a wall. Two constraints confine a vertex to a curve; but if a one-dimensional wire is desired instead of a two-dimensional wall, then it may be better to use the parameterized boundary feature described below.

5.2 Boundaries

Evolver *boundaries* are one or two dimensional parameterized manifolds; they are an alternate way to constrain the position of vertices. A vertex on a boundary cannot also have constraints. Vertices, edges, and facets may be deemed to lie in a boundary. For a vertex, this means that the fundamental parameters of the vertex are the parameters of the boundary, and its coordinates are calculated from these. When a vertex on a boundary moves, the motion is projected back to parameter space and applied to the parameters. Edges and facets on a boundary bequeath the boundary to descendant vertices.

A delicate question is how to handle wrap-arounds on a boundary such as a circle or cylinder. Subdividing a boundary edge requires a midpoint, but taking the average parameters of the endpoints can give nonsense. Therefore the average coordinates are calculated, and that point projected on the boundary parameters as continued from one endpoint. The rings in the catenoid example in §2.1 are represented as one-parameter circles and show a case where the endpoint extrapolation is necessary in the refining operation.

A general guideline is to use constraints for two-dimensional walls and boundaries for one-dimensional wires. If one uses a boundary wire, the vertices and edges on the boundary can probably be declared to be fixed. Then the boundary becomes just a guide for refining the boundary edges.

5.3 Quantities

A *quantity* in the Evolver is a value that can be expressed as the sum of integrals of vectorfields over surfaces and edges. Quantities can be evaluated for informational purposes, or they can be used as constraints (in the mathematical sense). Body volume is an examples of a built-in quantity. Surface area is not a quantity in this sense, since it cannot be written as a vector integral.

User-defined quantities are supported – for example, one might use as quantities the center of mass, the moment of inertia, the magnetic flux, and so no. The key to many of these is the Divergence Theorem, which permits volume integrals to be written as surface integrals. One can also use a quantity in place of a body volume that is awkward to do with the built-in volume mechanisms, for example, when a background metric is used.

Each quantity is defined in the datafile by giving the surface integrand and which facets it is to be integrated over, and an edge integral and the edges it is to be integrated over. A quantity acts as a mathematical constraint when it is declared fixed and given a specified value. Its actual value may be displayed and its target value changed interactively. Quantities have the same mathematical form as surface and edge energy integrals discussed in section 3, but are used for constraints or information rather than as part of the objective function.

5.4 Volumes

The term *volume* is used for the highest dimensional measure of a region in n -space: area in \mathbf{R}^2 , volume in \mathbf{R}^3 , etc. A body may have a volume specified in the datafile, which then becomes a volume constraint.

The volume of a body B can be written as

$$V = \int \int \int_B 1 dV,$$

which by the Divergence Theorem can be written a surface integral:

$$V = \int \int_{\partial B} z\vec{k} \cdot d\vec{S}.$$

This integral is evaluated over all the boundary facets of a body.

The part of the boundary of a body lying on a constraint need not be given as facets. In that case, Stokes' Theorem can be used to convert the part of the surface integral on the constraint to a line integral over the edges where the body surface meets the constraint. The line integral integrands are given as part of the constraint definition in the datafile. These edge volume integrands can also be used to overcome the volume calculation problems caused by gaps between curved constraints and flat facets.

Volumes are a special case of quantities, and are implemented internally for Euclidean space and flat torus domains only. In general quotient spaces or in Riemannian metrics, it is up to the user to define volume constraints using the general quantity constraint mechanism.

5.5 Volumes in a torus domain

The volume of a body can be automatically calculated in a torus domain, but the wrapping of the edges across the faces of the fundamental region makes the calculation a bit tricky. Ideally, we would like to adjust the vertices by multiples of the fundamental region basis vectors to get a body whose volume we could find with regular Euclidean methods. Unfortunately, all we know are the edge wraps, i.e. the differences in the adjustments to endpoints of edges. But this turns out to be enough, if we are a little careful with the initial volumes in the datafile.

Let the facets of a particular body be indexed by m , and let v_{mi} , $i = 0, 1, 2$, be the vertices of facet m . Let \vec{A}_{mi} be the (unknown) vertex adjustment for vertex v_{mi} , and \vec{T}_{mi} be the (known) wrap vector (difference in endpoint adjustments) for edge i of facet m . Then

$$\begin{aligned} V &= \frac{1}{6} \sum_{facets} \sum_m (\vec{v}_{m0} + \vec{A}_{m0}) \\ &\quad \cdot (\vec{v}_{m1} + \vec{A}_{m1}) \times (\vec{v}_{m2} + \vec{A}_{m2}) \\ &= \frac{1}{6} (S_1 + S_2 + S_3 + S_4), \end{aligned} \quad (5.1)$$

where

$$\begin{aligned}
S_1 &= \sum_{\text{facets } m} \vec{v}_{m0} \cdot \vec{v}_{m1} \times \vec{v}_{m2}, \\
S_2 &= \sum'_{\text{facets } m} \vec{v}_{m0} \cdot \vec{v}_{m1} \times \vec{A}_{m2}, \\
S_3 &= \sum'_{\text{facets } m} \vec{v}_{m0} \cdot \vec{A}_{m1} \times \vec{A}_{m2}, \\
S_4 &= \sum_{\text{facets } m} \vec{A}_{m0} \cdot \vec{A}_{m1} \times \vec{A}_{m2}.
\end{aligned}$$

In S_2 and S_3 , the \sum' notation means that each facet is included three times in cyclic permutation, once with each vertex as base point.

The first of these sums is straightforward. The second sum can be regrouped with one term for each edge j , pairing the two facets j and j' for each edge together:

$$\begin{aligned}
S_2 &= \sum_{\text{edges } j} \vec{v}_{j0} \cdot \vec{v}_{j1} \times \vec{A}_{j2} + \vec{v}_{j1} \cdot \vec{v}_{j0} \times \vec{A}_{j'2} \\
&= \sum_{\text{edges } j} \vec{v}_{j0} \cdot \vec{v}_{j1} \times (\vec{A}_{j2} - \vec{A}_{j'2}) \\
&= \sum_{\text{edges } j} \vec{v}_{j0} \cdot \vec{v}_{j1} \times \frac{1}{2}(\vec{T}_{j'2} + \vec{T}_{j1} - \vec{T}_{j2} - \vec{T}_{j'1}).
\end{aligned}$$

We can regroup this into a sum which can be done facet by facet, again including each facet three times:

$$S_2 = \frac{1}{2} \sum'_{\text{facets } m} \vec{v}_{m0} \cdot \vec{v}_{m1} \times (\vec{T}_{m1} - \vec{T}_{m2}).$$

In S_3 , we group terms with a common vertex together, with the inner facet sum over facets with vertex k as base vertex:

$$\begin{aligned}
S_3 &= \sum_{\text{vertices } k} \left(\vec{v}_k \cdot \sum_{\text{facets } i} \vec{A}_{i1} \times \vec{A}_{i2} \right) \\
&= \sum_{\text{vertices } k} \left(\vec{v}_k \cdot \sum_{\text{facets } i} (\vec{A}_{i1} - \vec{A}_k) \times (\vec{A}_{i2} - \vec{A}_k) \right) \\
&= \sum_{\text{vertices } k} \left(\vec{v}_k \cdot \sum_{\text{facets } i} \vec{T}_{i0} \times -\vec{T}_{i2} \right) \\
&= \sum'_{\text{facets } m} \vec{v}_{m0} \cdot \vec{T}_{m2} \times \vec{T}_{m0},
\end{aligned}$$

which again can be done facet by facet. The step introducing A_k is valid since the A_{i1} are just a relabeling of the A_{i2} and so $\sum_i A_{i1} = \sum_i A_{i2}$.

The sum S_4 is a constant, and so only needs to be figured once. Also, it is a multiple of the fundamental

region volume V_c , so its contribution to the body volume is a multiple of $\frac{1}{6}V_c$ by equation (5.1). Therefore, if we assume the volume prescribed in the datafile is within $\frac{1}{12}V_c$ of the actual volume, we can calculate the other sums and figure out what the fourth sum should be.

The body volume gradient at vertex v can be easily found from the above sums, since the base vertex is dotted with other terms. The gradient is a sum over all facets with base vertex v :

$$\begin{aligned}
\frac{\partial V}{\partial v} &= \frac{1}{6} \left(\sum_{\text{facets } m \text{ on } v} \vec{v}_{m1} \times \vec{v}_{m2} \right. \\
&\quad \left. + \frac{1}{2} \vec{v}_{m1} \times (\vec{T}_{m1} - \vec{T}_{m2}) + \vec{T}_{m2} \times \vec{T}_{m0} \right).
\end{aligned}$$

6 Iteration

The heart of the Evolver is the iteration step that reduces energy while obeying any constraints. The surface is changed by moving the vertices. No changes in topology or triangulation are made. The idea is to calculate the force at each vertex and move the vertex in that direction, thus using a gradient descent method of minimization.

6.1 Force Calculation

The first step in an iteration is the calculation of the forces on the vertices. The total energy of the surface is viewed as a function of the coordinates of the vertices. The negative gradient of the energy as a function of the position of a single vertex gives the force on that vertex. Collectively, all the forces on all the vertices make up the negative of the total gradient of energy. No new approximations are introduced by the force calculation; the energy may be an approximate energy due to numerical integrations, but the gradient is the exact gradient of the approximate energy.

Vertices on constraints have their forces projected to the tangent spaces of the constraints. Vertices on boundaries have their forces mapped back to forces on the boundary parameters. Fixed vertices have their forces set to zero.

6.2 Volume and Quantity Constraints

The second step in an iteration is to enforce constraints on body volumes and other integrated quantities. This has two parts. The first part consists of correcting for any errors in the current values of the quantities. The second part consists of projecting the vertex forces to be orthogonal to the quantity gradients. Both parts use the gradients of the quantities

as functions of the vertices. Let \vec{W}_{vk} be the gradient of quantity k as a function of the position of vertex v . These gradients are projected on constraint level sets or boundaries or set to zero in the same manner as forces.

The value correction consists of applying a single step of Newton's method. Let the current excess value of quantity k be δ_k (which may be negative, of course). We assume the correcting motion \vec{R}_v at vertex v is of the form

$$\vec{R}_v = \sum_k c_k \vec{W}_{vk}$$

such that

$$\sum_v \vec{R}_v \cdot \vec{W}_{vk} = -\delta_k \quad \text{for each } k.$$

This leads to the following linear system for the c_k :

$$\sum_k c_k \sum_v \vec{W}_{vk} \cdot \vec{W}_{vk'} = -\delta_{k'}$$

for each quantity k' . This system is solved for the c_k , and hence for the motions \vec{R}_v . This motion is not done immediately, but is done as part of the overall motion described below. The quantity values are not perfectly corrected by this step, but over several iterations they should converge to the target values.

The second part is the projection of the forces. Let \vec{F}_v be the total force at vertex v . We want a projected force $\vec{F}_{v,proj}$ of the form

$$\vec{F}_{v,proj} = \vec{F}_v - \sum_k a_k \vec{W}_{kv}$$

such that

$$\sum_v \vec{F}_{v,proj} \cdot \vec{W}_{vk} = 0 \quad \text{for each quantity } k,$$

which leads to the linear system for a_k

$$\sum_k a_k \sum_v \vec{W}_{vk} \cdot \vec{W}_{vk'} = \sum_v \vec{F}_v \cdot \vec{W}_{vk'}$$

for each quantity k' . The coefficients a_k are the Lagrange multipliers for the quantity constraints, and can be interpreted as pressures for body volume constraints. Whenever the user asks for body volumes to be displayed, these pressure values are also shown.

6.3 Motion

Each vertex is moved by the quantity correction motion plus a scale factor times the force at the vertex. The scale factor is a global constant, the same for all vertices. The user may set the scale factor explicitly, or let the Evolver seek the optimal value. In

the latter mode, Evolver will successively double or halve the scale factor until a minimum in energy is bracketed. Then quadratic interpolation is used to estimate the optimum scale factor, and that value is used in the final motion.

Each time a motion is done, all vertices on constraints are projected back to their constraints by repeatedly applying Newton's method until the constraint function value is smaller than the constraint tolerance factor, which the user may set. If vertices subject to one-sided constraints are on the wrong side of the constraint, they are projected to the constraint. If such a vertex wants to move to the proper side of the constraint, it is freed from the constraint.

From experience, it seems that for two-dimensional surfaces driven by surface tension that the optimum scale factor is around 0.2 independent of the fineness of the triangulation as long as the surface is evolving without problems. The universality the scale factor is expected in this case since the scale factor is unitless in length for two-dimensional area only. When the scale factor dives toward zero, it is usually a sign of problems like an edge length becoming zero, or a facet area becoming zero. Unfortunately, the 0.2 value applies best to minimizing area; other models, such as the string model and squared mean curvature, have "normal" scale factors that vary with triangulation size and other factors, so it is hard to tell when the surface is evolving properly and when it is getting into trouble. In these circumstances, other methods have to be used, such as watching the surface visually and checking edge length and facet area histograms.

6.4 Motion by mean curvature

The mean curvature vectorfield \vec{h} of a surface S is defined to be the gradient of the area of S , in the sense that if S is deformed with an instantaneous velocity \vec{u} , the rate of change of surface area is

$$\frac{dA}{dt} = \int \int_{surface} \vec{u} \cdot \vec{h} dA.$$

By definition, $\vec{u} = -\vec{h}$ under motion by mean curvature, so that

$$\frac{dA}{dt} = - \int \int_{surface} h^2 dA. \quad (6.1)$$

By default, the gradients of quantities like energy and volume are calculated simply as the gradient of the quantity as a function of vertex position. This gives the force on a vertex, for example. But to simulate motion by mean curvature, it is necessary to have force per unit area instead. In the triangulation formulation, let A_v be the area of the star of facets around vertex v and let

$$\vec{F}_v = - \frac{\partial A_v}{\partial v}$$

be the force on vertex v . Then

$$\frac{dA}{dt} = \sum_v \vec{u}(v) \cdot -\vec{F}_v.$$

Take the area associated with a vertex to be one-third of the total areas of the facets surrounding the vertex, $dA = \frac{1}{3}A_v$. Since each facet has three vertices, this allocates all area. Hence as approximation to \vec{h} , we take

$$\vec{h} = -\frac{3\vec{F}_v}{A_v}.$$

If the Evolver is operated in “area normalization” mode, then the vertex motions are calculated using this formula. In this mode, the user should set the scale factor, which is the timestep for the evolution, to a constant small enough for the iteration to be a good approximation to the continuous evolution. Using an optimizing scale factor makes the time step too large.

The string model is more amenable to close scrutiny of how the Evolver does motion by mean curvature. In §2.3 an example was given of grain boundaries evolving in two dimensions. In approximating motion by mean curvature, there are two discretizations that must be made, in space and in time. The space discretization here consists of representing the grain boundaries as a set of line segments, and the time discretization consists of the iteration steps. There is an intermediate stage of discretization, that of the continuous time evolution of the discrete boundaries. The equation of motion for this problem has been chosen to approximate the continuous problem in a certain respect. Two-dimensional grain evolution has the property that the rate of change of area of a grain depends only on the number of its sides. If \vec{h} is the mean curvature of a grain G , so the boundary velocity is $-\vec{h}$, then the rate of change of area of a grain is

$$\frac{dA}{dt} = \int_{\partial G} -h ds = - \int_{\partial G} \frac{d\theta}{ds} ds = - \int_{\partial G} d\theta.$$

The total turning angle around the boundary of a grain is 2π , but each triple vertex contributes a turning angle of $\pi/3$, so for a grain of N sides, so

$$\frac{dA}{dt} = (N - 6) \frac{\pi}{3}.$$

The space-discrete problem preserves this property. The motion of each vertex is such that the area change of the grain due to the motion of that vertex is proportional to the turning angle at that vertex (or to the excess turning angle at triple vertices).

A property that is not exactly preserved is that the rate work is done sweeping out area is proportional to the rate of length loss,

$$\frac{dW}{dt} = \int_{\partial G} \vec{v} \cdot \vec{F} ds = \int_{\partial G} h^2 ds = -\frac{dL}{dt}.$$

The last equality is the string version of equation (6.1). To preserve this property would require solving a system of linear equations linking together the motions of all the vertices.

The space-discretized evolution is guaranteed dissipative, but the time-discretized evolution can suffer from instabilities. Consider a boundary made up of segments of length L zigzagging about a straight midline with small amplitude y . The velocity of a vertex will be $4y/L^2$, and if the time step is Δt , then the amplitude will grow if $4\Delta ty/L^2 > 2y$. Hence the maximum timestep allowable is $\Delta t = L^2/2$, or, conversely, the minimum edge length is $L = \sqrt{2\Delta t}$. Exact damping of the zigzag occurs when $\Delta t = L^2/4$.

6.5 Conjugate gradient

For minimizing a quadratic function, there is a technique called the *conjugate gradient method* [Press et al. 1988 §10.6] that can minimize much faster than gradient descent. With exact arithmetic, it minimizes an n -dimensional quadratic function in at most n iterations. This method does not follow the gradient downhill, but makes an adjustment using the past history of the minimization.

The Evolver uses the Fletcher-Reeves variant of the conjugate gradient method. At iteration step i , let S_i be the surface, E_i its energy, $\vec{F}_i(v)$ the force at vertex v as described above, and $\vec{h}_i(v)$ the “history vector” of v . Then

$$\vec{h}_i(v) = \vec{F}_i(v) + \gamma \vec{h}_{i-1}(v)$$

where

$$\gamma = \frac{\sum_v \vec{F}_i(v) \cdot \vec{F}_i(v)}{\sum_v \vec{F}_{i-1}(v) \cdot \vec{F}_{i-1}(v)}.$$

For the actual motion, a one-dimensional minimization is performed in the direction of \vec{h}_i , using the bracketing method described in §6.3. It is important that all volumes and constraints be enforced during the one-dimensional minimization, or else the method can go crazy.

the energy function of a surface is not exactly quadratic, but the method can still be applied, and sometimes it yields very good results. But sometimes it’s worse than regular iteration. The saddle point of energy in the catenoid example of §2.1 seems to confuse the conjugate gradient method. With conjugate gradient in effect, the saddle point is passed at iteration 17 and the area decreases again until iteration 30, when it reaches 6.4486. But at this point further iteration produces no change, and the conjugate gradient mode has to be turned off and on to erase the history vector. Once restarted, another 20 iterations will get the area down to 6.4334. This shows that the conjugate gradient mode can work much better than ordinary mode, but it can also have problems.

6.6 Hessian minimization

Minimization by gradient descent or even conjugate gradient can take many iterations. A more direct way to try to minimize is to calculate the Hessian matrix of second derivatives of energy and solve for the motion that gives zero gradient. This method is currently implemented only for the case where surface tension is the only energy and there are no constraints of any kind except fixed vertices. It assumes that the surface is close enough to a local minimum for the Hessian to be positive definite, which is not always true. When the method works, it can find the minimum energy to 15 decimal places in three or four iterations. But if the Hessian is not positive definite, the method blows up. It is easy to find examples with a saddle point of energy; see the catenoid example below. There are checks in place to ensure that the calculated motion does indeed reduce energy.

6.7 Diffusion

In real soap-bubble clusters, air can diffuse across the soap films, driven by pressure differences. Smaller bubbles tend to have higher curvature and hence higher pressure, so tend to shrink. One can watch a foam evolve over the course of minutes, changing its topology as bubbles disappear. The Evolver can simulate diffusion. If the diffusion mode is on, then at the start of an iteration target volume is transferred across each facet, in the amount equal to the area of the facet times the difference in body pressures times the global diffusion constant. The iteration step then corrects the actual volumes to the target volumes and does its normal energy minimization step. Topology changes are not done automatically yet; those are up to the user to do using the operations described in Section 7.

7 Surface Operation

This section describes the main commands available to the user, aside from the iteration step described in the previous section.

7.1 Refining

To *refine* a triangulation is to subdivide each facet to create a finer triangulation. The Evolver does this by creating new vertices at the midpoints of edges, which it then uses to subdivide each facet into four new facets, each similar to the original.

The first stage of refining is to subdivide all edges by inserting a midpoint. Hence all facets temporarily have six sides. For an edge on constraints, the

midpoint gets the same set of constraints, and is projected to them. For an edge on a boundary, the parameters of the midpoint are calculated by projecting the vector from the edge tail to the midpoint back into the parameter space and adding that to the tail parameters. This avoids averaging parameters of endpoints, which gives bad results when done with boundaries that wrap around themselves, such as circles. In the second stage, each facet is subdivided into four facets by connecting the new midpoints.

Certain attributes of new elements are inherited from the old elements from which they were created. The new facets inherit the surface tension of their parent facets. Fixity, constraints, and boundaries are always inherited by offspring of all dimensions. In a quotient space, some, but not all, new edges inherit symmetry group wrapping, so that the surface is correctly embedded.

Refinement can change surface area, energy, and volumes if there are curved constraints or boundaries.

In seeking the minimum energy, it is best to evolve with a coarse triangulation as far as possible. Each iteration can propagate a position adjustment only one edge at a time, so the finer the triangulation, the longer adjustments take to travel across the surface.

7.2 Equiangulation




Figure 7: The two adjacent facets on the left violate the equiangulation criterion, since we have $\theta_1 + \theta_2 > \pi$. Equiangulation flips the quadrilateral diagonal, making the triangles more nearly equilateral.

Triangulations work best when the facets are as close to equilateral (that is, equiangular) as possible for a given set of vertices. Given a set of vertices, how does one make a triangulation for those vertices that has triangles as nearly as possible equilateral? In the plane, the answer is the Delaunay triangulation, in which the circumcircle of each triangle contains no other vertex [Sibson 1978]. It is almost always unique. It can be constructed by local operations beginning with any triangulation. Consider any edge as

the diagonal of the quadrilateral formed by its adjacent triangles. If the angles of the two vertices off of the diagonal add to more than π , then the circum-circle criterion is violated, and the diagonal should be switched to form a replacement pair of triangles (Figure 7). When no more switches can be done, we have a Delauney triangulation.

Now suppose now we have a triangulation of a curved surface in space. For any edge with two adjacent facets, we switch the edge to the other diagonal of the skew quadrilateral if the sum of the angles at the off vertices is more than π . As in Figure 7, let a be the length of the common edge, b and c the lengths of the other sides of one triangle, and d and e the lengths of the other sides of the other triangle. Let θ_1 and θ_2 be the off angles. Then by the law of cosines,

$$a^2 = b^2 + c^2 - 2bc \cos \theta_1, \quad a^2 = d^2 + e^2 - 2de \cos \theta_2.$$

The condition $\theta_1 + \theta_2 > \pi$ is equivalent to $\cos \theta_1 + \cos \theta_2 < 0$. So we switch if

$$\frac{b^2 + c^2 - a^2}{bc} + \frac{d^2 + e^2 - a^2}{de} < 0.$$

The equianguation procedure over the whole surface may have to be repeated several times to get complete equianguation, but almost never more than three or four times. The process is guaranteed to terminate since a switch reduces the radii of the circumcircles and a finite number of vertices has a finite number of triangulations.

Equianguation can have an almost magical effect in improving a triangulation, and I highly recommend its regular use. It may temporarily increase area and change volumes, but the magnitudes of these effects are within the approximation error of using flat facets for a curved surface. A use of equianguation is shown in Figure 1.

7.3 Vertex averaging

An evolving surface can get into trouble if some of the vertices of the triangulation get too scrunched together, as shown in Figure 3 B. To get vertices to spread out, there is an operation called *vertex averaging*. For each vertex, this operation computes a new position as the area-weighted average of the centroids of the facets adjoining the vertex. Fixed vertices or vertices on boundaries are not moved, nor are vertices on triple or higher edges. Also, to keep the new surface as close as possible to the old one, volumes on both sides of the surface are preserved. If vertex v is on facets f_i with centroids \vec{x}_i , then the new position is calculated as

$$\vec{v}_{avg} = \frac{\sum_i \text{area}(f_i) \cdot \vec{x}_i}{\sum_i \text{area}(f_i)}.$$

The volume on one side of all the facets around the vertex calculated as a cone from the vertex is

$$V = \sum_{\text{facets } f} \vec{v} \cdot \vec{N}_f,$$

where \vec{N}_f is the facet normal representing its area.

The total normal \vec{N} is

$$\vec{N} = \sum_{\text{facets } f} \vec{N}_f.$$

To preserve volume, we subtract a multiple λ of the total normal from the average position:

$$(\vec{v}_{avg} - \lambda \vec{N}) \cdot \vec{N} = \vec{v} \cdot \vec{N},$$

so

$$\lambda = \frac{\vec{v}_{avg} \cdot \vec{N} - \vec{v} \cdot \vec{N}}{\vec{N} \cdot \vec{N}}.$$

Then the new vertex position is

$$\vec{v}_{new} = (\vec{v}_{avg} - \lambda \vec{N}).$$

Constrained vertices are then projected to their constraints.

Vertex averaging may slightly increase area, but this is usually offset by its benefits. It is useful in getting the vertices spread out evenly. Evolution can be awkward when facets are of very different sizes, since the same scale factor applies to the whole surface.

7.4 Notching edges

A surface can be locally highly curved, resulting in facets forming pronounced ridges along edges. One way to selectively refine the surface is to just refine around those edges whose adjacent facets are too far from parallel, putting a notch in the edge to make it more saddle-shaped. There is a command that lets the user do this with a cutoff angle of his choosing. The refinement is actually done by subdividing each adjacent facet by putting a new vertex in the center. Equianguation then completes the process. Formerly, the Evolver did notching by just subdividing the offending edges, but that tended to create lots of long skinny triangles and not always help matters. The new method seems to work better.

7.5 Edge and facet operations

One way to improve a triangulation is to simply eliminate all edges that have become too short. This operation is known as *tiny edge weeding*. Every edge shorter than a user-set cutoff length that can legitimately be removed is deleted by identifying its endpoints.

Sometimes there are very skinny triangles that should be eliminated, but which don't have a short edge to be found by tiny edge weeding. Therefore there is an operation called *area weeding* that removes triangles whose area is smaller than some cutoff. This procedure finds the shortest edge of a triangle and eliminates it by the same process as the regular tiny edge removal.

There is a command that will bisect all edges longer than a user-chosen length. All facets adjoining the edges are also subdivided into pairs of facets. If the new edges are still longer than the cutoff length, they are not further subdivided. It is suggested that this step be followed by equianguation.

Histograms of edge lengths and facet areas can be displayed in conjunction with any of these commands.

7.6 Annealing, or jiggling

Sometimes it may be desirable to perturb the surface to get it off a metastable position. There are both random and user-definable perturbations possible. Because of its similarity to the thermal perturbations responsible for annealing in metals, the characteristic magnitude of the perturbation is called the *temperature*.

Under a random perturbation, or *jiggle*, each coordinate of each non-fixed vertex is moved by $\delta x = TLg$ where g is a random value from the standard gaussian distribution (calculated from the sum of five random values from the uniform distribution on $[0,1]$), T is the current temperature, and L is a characteristic length that starts as the diameter of the surface and is cut in half at each refinement.

A *long jiggle* is a sinusoidal displacement of each vertex v by $\vec{A} \sin(\vec{v} \cdot \vec{w} + \psi)$. The amplitude \vec{A} , the wavevector \vec{w} , and the phase ψ may be specified by the user or chosen at random.

7.7 Popping Edges or Vertices

The Evolver does not change the topology of a surface on its own. However, there are many times when a naturally evolving surface will need to change its topology. For example, a neck might pinch out in a catenoid whose boundary rings are too far apart, or two growing metal grains might meet. Fortunately, the types of singularities possible in soapfilm-like surfaces in three-dimensional space were classified in [Taylor 1976] for uniform surface tension. Three surfaces may meet along a curve, or four triple curves may meet at a point. The Evolver has procedures known as *edge popping* and *vertex popping* to detect improper singularities and reduce them to proper types. These routines are designed only for surfaces with uniform

surface tension. There are many more types of singularities possible if the different component surfaces meeting at a singularity have different surface tensions.

Edge popping looks for edges with more than three facets that are not fixed and are not on boundaries or constraints. When found, such an edge is split longitudinally with a new facet in between. The two old facets with the smallest dihedral angle between them are attached to the new edge. This is repeated until only three facets are on the original edge. Each split is propagated along the multiple junction line as far as possible. If it is impossible to propagate the split beyond either endpoint, the edge is subdivided to provide a vertex which can be split.

Vertex popping assumes that all edges have at most three facets adjoining, so edge popping should be done first. The facet and edge structure around each vertex is analyzed to find which have the wrong topology. The analysis is done by looking at the network formed by the intersection of the facets and edges containing the vertex with the surface of a small sphere around the vertex. The numbers of sides of the cells of the network are counted. A simple plane vertex has two cells of one side each. A triple edge vertex has three cells of two sides each. A tetrahedral point has four cells with three sides each. Any other configuration is popped. The popping is done by replacing the vertex with a hollow formed by truncating each cell-cone except the cell with the largest solid angle. In case the network is disconnected, the solid angles of all the cells will add up to over 4π . Then the vertex is duplicated and the different components are assigned to different vertices. This lets necks shrink to zero thickness and pull apart.

In the string model, vertices with more than three edges are popped by finding the pair of edges making the least angle, pulling them out a short distance with a new vertex, and joining the new and old vertex with a short edge. This is repeated until the original vertex has only three edges.

Improper vertices may exist in the original datafile, or they may be introduced by short edge elimination. For example, the pinching neck in a catenoid must have all the short edges around its waist eliminated to pinch the waist down to one vertex, which can then be popped. The other operations described in this section (refining, vertex averaging, equianguation, notching) do not change the global topology and so do not introduce improper vertices. Improper vertices are not automatically detected unless the autopop feature is on.

7.8 Zooming

Sometimes the detail of a surface may require a closer look. The graphics display can magnify a surface, but that doesn't change the triangulation to follow the detail. There can be cases, as when a boundary wire passes through a soapfilm, where the detail around a point is on a scale 100,000 times smaller than the whole surface [Brakke 1992a]. Hence the Evolver contains a mechanism to zoom in on a vertex, throwing away the rest of the surface to save memory and time and to keep all triangle sizes reasonably close together.

The user specifies which vertex to zoom in on and the radius around it to keep. All vertices beyond the cutoff distance from the given vertex are deleted. Then all edges and facets containing any deleted vertices are deleted. Any remaining edge that had a facet deleted from it is made fixed to anchor the cut edges of the surfaces.

8 User Interface Details

This section describes the Evolver's user interface, including the initial data file, the command mechanism, and the graphics interface.

8.1 The Initial Data File

The initial configuration of a surface is read from a text file referred to as the *datafile*. The datafile has five sections: general definitions, vertices, edges, faces, and bodies. The catenoid datafile `cat.fe` is presented here to give the flavor. This file is slightly atypical in that none of the vertices are given directly by their coordinates. The `.fe` filename extension is a relic of early versions of the Evolver in which facet-edges had to be explicitly listed in the datafile; I continue to use it out of habit to identify datafiles.

```
// cat.fe
// Evolver datafile for catenoid.

// ring radius and height
// adjustable at runtime
PARAMETER radius = 1
PARAMETER height = 0.55

// upper ring, parameterized by p1
boundary 1 parameters 1
x1:   radius * cos(p1)
x2:   radius * sin(p1)
x3:   height

boundary 2 parameters 1
x1:   radius * cos(p1)
x2:   radius * sin(p1)
x3:   -height
```

```
vertices // second column = value of p1
1   0*pi/3 boundary 1 fixed
2   1*pi/3 boundary 1 fixed
3   2*pi/3 boundary 1 fixed
4   3*pi/3 boundary 1 fixed
5   4*pi/3 boundary 1 fixed
6   5*pi/3 boundary 1 fixed
7   0*pi/3 boundary 2 fixed
8   1*pi/3 boundary 2 fixed
9   2*pi/3 boundary 2 fixed
10  3*pi/3 boundary 2 fixed
11  4*pi/3 boundary 2 fixed
12  5*pi/3 boundary 2 fixed
```

```
edges // given by endpoint vertices
1   1 2 boundary 1 fixed
2   2 3 boundary 1 fixed
3   3 4 boundary 1 fixed
4   4 5 boundary 1 fixed
5   5 6 boundary 1 fixed
6   6 1 boundary 1 fixed
7   7 8 boundary 2 fixed
8   8 9 boundary 2 fixed
9   9 10 boundary 2 fixed
10  10 11 boundary 2 fixed
11  11 12 boundary 2 fixed
12  12 7 boundary 2 fixed
13  1 7
14  2 8
15  3 9
16  4 10
17  5 11
18  6 12
```

```
faces // given by oriented edge list
1   1 14 -7 -13
2   2 15 -8 -14
3   3 16 -9 -15
4   4 17 -10 -16
5   5 18 -11 -17
6   6 13 -12 -18
```

The datafile syntax provides several features for flexibility and ease of use. Simple macros can be defined to do text substitution. Compound expressions can be used anywhere a real number or a formula is expected. Normal arithmetic and standard functions are available. For functions that are evaluated during runtime (constraints, quantities, etc.) expressions are stored as syntax trees that are interpreted when the expression needs evaluation. If interpretation is too slow, user-defined functions may be written in C and compiled into the Evolver. Named variables may be declared and used in runtime expressions. Such variables (called *adjustable parameters*) can be changed interactively during runtime, and are useful for moving constraints and boundaries around, modifying the metric, changing contact angles, and so forth.

The definitions section describes everything not pertaining to particular geometric elements, such as

- declarations and initial values for adjustable parameters;

- the dimension of the surface and the dimension of the ambient space containing the surface;
- the specification of a quotient space, or, for a flat torus domain, the vectors defining the fundamental parallelepiped;
- Riemannian metric tensor components;
- for a crystalline surface energy, the name of the Wulff vector file;
- constraint function formulas, together with energy and volume integrands for edges on constraints;
- boundary definitions via formulas of coordinates in term of parameters;
- quantity integrands, with target values for constrained quantities;
- initial values for the gravitational constant, the diffusion constant, the weighting factor for the squared mean curvature in the energy, the Gaussian integration order, and linear or quadratic mode.

None of these items are required. If missing, the feature is not used or has a natural default value.

The vertices section lists the vertices, one per line. Each vertex is numbered for later reference, and is defined by its coordinates (or boundary parameters), which constraints or boundaries it is on, and whether it is fixed.

The edges section lists the edges, one per line, also numbered for reference. Each edge is defined by its tail and head vertex numbers, the constraints or boundaries it is on, the quantities it contributes to, and whether it is fixed. If a quotient space is being used, the group element for wrapping the head vertex to the proper place with respect to the tail vertex is also given.

The face section lists polygons forming the initial surface. Each polygon is given by its edge numbers in order around its circumference. Edges traversed in opposite direction from that given in the edges section are given as negative numbers. The polygons need not be planar, and they need not be triangles (which is why the section is “faces” instead of “facets”). The Evolver will immediately triangulate non-triangular faces by putting a new vertex at the average position of the original vertices and putting in edges from the new vertex to each original vertex. Each face may be on constraints, boundaries, or be fixed. It may be given a specific surface tension; the default is 1.0. It may be deemed to contribute to certain quantity integrals, and to have certain surface integrands contribute to the total energy.

In the bodies section, each body is defined by listing its bounding faces by number, negative numbers if the orientation of the face in the face list has an inward normal. There may be any number of faces in any order. Faces do not have to completely enclose a body; they are used to compute volume and other integrals, and if certain faces are not needed for that, they may be omitted. A body may be declared to have a fixed volume of a certain value. The actual initial volume need not be that exact value; the volume will be adjusted during the iteration process. A body may also be given a density, which will cause the total energy to include the gravitational potential energy of the body with that density.

8.2 Command Interface

The user command interface is built on a simple terminal-type model for maximum portability. The main prompt is “`Enter command:` ”. There are two types of commands. The first consists of one letter occasionally followed by a number; the second is an embryonic query language. Currently, queries are supported that list, display, refine, or delete elements by various criteria. Commands may be read from a file with the command “`read filename`”. The output of any command can be piped to a system command. Commands that change the surface or change the model will cause energies and volumes to be recalculated. Commands can be logged to a file for later repetition with the `read` command.

8.3 Graphics

It is possible to run the Surface Evolver without any graphics. The program uses the standard teletype-style interface for maximum portability. But it’s always nice to see the surface, and often essential to understanding what is happening. Unfortunately, every computer system has its own way of displaying graphics, and there is no universal standard. The Evolver isolates the system-dependent graphics to drawing two-dimensional triangles. This cuts down the effort in porting the Evolver to a new system to a reasonable level. The Evolver’s main graphics routine calculates the triangles to display, and calls the device-dependent subroutine to do the display. The device could be a screen display or a graphics output file writer.

Two classes of graphics devices are provided for: those which can do their own viewing transformations and hidden surface removal, and those which can’t. The former are simply provided a list of triangles with vertex coordinates in three dimensions. For the latter, the Evolver keeps an internal viewing transformation matrix, sorts the transformed trian-

gles from back to front, and feeds them to the display routine (painter's algorithm). The sorting algorithm will not subdivide intersecting triangles. If two triangles overlap, it will just find one point in the overlap and compare depths there. This can lead to some strange-looking displays for strange surfaces, but it works well for the types of surfaces the Evolver is meant for.

My favorite graphics system is the GeomView program on Iris workstations. GeomView is an interactive viewer that lets the user rotate, translate, zoom, etc. the surface by dragging a mouse cursor over the window in a trackball metaphor. A high-end Iris workstation can light, shade, and smoothly rotate a surface consisting of several thousand triangles. The Evolver uses a shared memory interface with GeomView, so that the display is automatically updated whenever the surface changes. GeomView was written at The Geometry Center, and is freely available (see "Software Availability" at the end of this article).

Other types of screen displays do not have such fancy view control as GeomView. Instead, there is a teletype-style graphics command interface that lets the user control the viewing angle and size of the display. This viewing transformation is also used for the graphics output files.

There are several graphics output file formats, most notably PostScript. There are also formats that list transformed or untransformed triangles as text, suitable for input to other programs.

8.4 Other Commands

It is possible to reset the values of many parameters during runtime, including the gravitational constant, body volumes, constrained quantity targets, the diffusion constant and the user-defined variables used in formulas for constraints, boundaries, quantities and metrics.

The current surface can be dumped to a text file in the same format as the datafile. This is the only way to save a file; there is no binary save format. The text format has the advantages that it is portable, editable and not too much larger than a binary file would be.

After minimizing energy at several levels of refinement, it is possible to extrapolate the energy to an infinitely fine refinement. The extrapolation uses the final energies of three successive refinements and assumes a power law approach to the ultimate minimum.

9 Application: The Hopf Cone Conjecture

In this section I present an example of a conjecture that was settled (negatively) through the use of the Evolver.

In contrast to the situation in \mathbf{R}^3 , the classification of area minimizing hypersurface cones in \mathbf{R}^4 is unknown. Frank Morgan once conjectured that a certain cone in \mathbf{R}^4 is absolutely area minimizing [Morgan 1986, p. 1278]. This example shows that it is not by having the Evolver generate a comparison surface with less area for the same boundary. It illustrates the use of a Riemannian metric to permit a three-dimension surface to represent a four-dimension surface by projecting out a symmetry.

Morgan's cone is based on the Hopf fibration of the 3-sphere \mathbf{S}^3 . Let \mathbf{S}^3 be parametrized by

$$0 \leq \alpha \leq \pi/2, \quad 0 \leq \beta \leq 2\pi, \quad 0 \leq \gamma \leq 2\pi$$

so that Euclidean coordinates are

$$\begin{aligned} x_1 &= \cos \alpha \cos \beta, & x_2 &= \cos \alpha \sin \beta, \\ x_3 &= \sin \alpha \cos \gamma, & x_4 &= \sin \alpha \sin \gamma. \end{aligned}$$

Then the boundary of Morgan's Hopf cone consists of the three surfaces

$$\begin{aligned} \beta - \gamma &= 0 && (\text{mod } 2\pi), \\ \beta - \gamma &= 2\pi/3 && (\text{mod } 2\pi), \\ \beta - \gamma &= 4\pi/3 && (\text{mod } 2\pi). \end{aligned}$$

These three surfaces have zero mean curvature and meet at 120° angles along the two orthogonal circles $x_1^2 + x_2^2 = 1$ and $x_3^2 + x_4^2 = 1$.

Theorem 9.1. *Morgan's Hopf cone is not absolutely area minimizing.*

Proof. The idea of the proof is to take the quotient space of \mathbf{R}^4 modulo the Hopf fibers \mathbf{S}^1 , giving \mathbf{R}^3 with a metric such that the area of a surface in \mathbf{R}^3 is the same as the 3-area of the lift of the surface back into \mathbf{R}^4 . The metric turns out to have a natural interpretation as a cone space, leading to a simple counterexample which can be verified by the Evolver.

The metric on \mathbf{R}^4 in Hopf spherical coordinates $(r, \alpha, \beta, \gamma)$ is

$$ds^2 = dr^2 + r^2 d\alpha^2 + r^2 \cos^2 \alpha d\beta^2 + r^2 \sin^2 \alpha d\gamma^2.$$

The coordinates of the quotient space \mathbf{R}^3 will be (r, α, θ) , where $\theta = \beta - \gamma$. The orthogonally projected metric is

$$ds^2 = dr^2 + r^2 d\alpha^2 + r^2 \sin^2 \alpha \cos^2 \alpha d\theta^2.$$

The lift of a point in \mathbf{R}^3 is a circle of circumference $2\pi r$. Hence we can make the 3-area of the lift of a surface have the 2π times the 2-area in \mathbf{R}^3 by multiplying the linear metric ds by \sqrt{r} , giving an effective metric of

$$ds^2 = r dr^2 + r^3 d\alpha^2 + r^3 \sin^2 \alpha \cos^2 \alpha d\theta^2.$$

This can be made to look more like the ordinary spherical coordinate metric by using coordinates $\rho = r^{3/2}/2$ and $\phi = 2\alpha$. Then

$$ds^2 = \frac{16}{9} d\rho^2 + \rho^2 d\phi^2 + \rho^2 \sin^2 \phi d\theta^2.$$

This is the Euclidean spherical coordinate metric except for the factor $16/9$, which is bigger than 1, making \mathbf{R}^3 into a cone space.




Figure 8: This surface is a counterexample to Morgan's Hopf cone conjecture. The three outer edges are equally spaced half-great-circles on the unit sphere. Note how the surface avoids the center of the sphere.

Morgan's Hopf cone projects to three planes meeting at 120° , which is known to be absolutely area minimizing in the standard metric of \mathbf{R}^3 . Its area inside the unit sphere (in the Hopf metric) is 2π . However, the cone factor makes it more expensive for a surface to go radially inward rather than go sideways. In a two-dimensional cone, it is easily seen (by unrolling the cone) that geodesics avoid the origin. A similar phenomenon happens here. The comparison surface was made by deforming the three planes by pushing the point at the origin out toward one of the boundaries. This configuration was fed into the Surface Evolver, which evolved the surface and produced the surface of Figure 8 with an area of 6.14, which is easily less than that of Morgan's Hopf cone. The numerical errors in this area are due to the numerical integration used to calculate the area of the facets and the gap between the curved boundary and the facets. Both of these can easily be estimated to be less than the improvement of 8 over the Hopf cone. ■

The Hopf cone that projects to the tetrahedral cone is also not minimizing. The Evolver gives 7.44

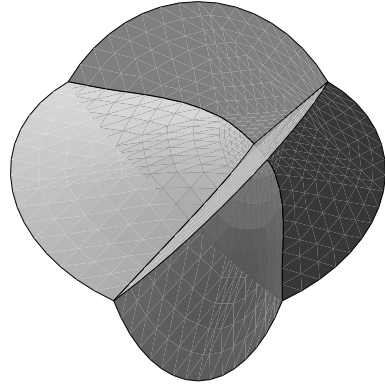


Figure 9: Comparison surface for the tetrahedral Hopf cone. The three outer edges lie on the unit sphere. Again, the surface avoids the center of the sphere.

for the area of the comparison surface Figure 9, while the area of the cone is $4 \cos^{-1}(-1/\sqrt{3}) = 7.6425$. The only other minimizing cone in standard \mathbf{R}^3 , the flat plane, also deforms to avoid the origin. Thus none of the area-minimizing cones in standard \mathbf{R}^3 lifts via the Hopf fibration to a minimizing cone in \mathbf{R}^4 .

10 Future directions

The Surface Evolver is under continual development. I welcome suggestions from users for new features. If they are reasonable, they will be added as time permits.

Some of the major items on my current list of things to do are:

- How close in various senses is an Evolver minimal surface to the true smooth minimal surface for a given problem?
- The Evolver gives an upper bound for the area of a minimal surface. The technique of *calibrations*, a generalization of the network theory min-cut max-flow duality, can give lower bounds. A near-minimal surface should be able to generate a near-maximal calibration. Hence a goal is to have the Evolver generate such calibrations.
- How close is an Evolver evolution by mean curvature to an ideal smooth evolution? Given an initial smooth surface, is it possible to construct an Evolver approximation that stays close to the ideal evolution? A more permissive notion of approximation would say that for each Evolver evolution there is an ideal smooth evolution that stays near it.

- The current method of approximation to motion by mean curvature needs further investigation. The gradient of energy is a covector, and motion is a vector. The conversion from covector to vector requires a metric or inner product. The inner product used by the Evolver is the Euclidean inner product at vertices, weighted by the vertex star area. Other inner products are possible and may have desirable properties.
- Instabilities of the type described in §6.4 often limit the size of the time step in an evolution. These instabilities need to be understood and methods have to be developed to speed evolution.
- Automatic triangulation management. Currently, users have to closely monitor the surface triangulation and intervene manually when it gets fouled up. I hope to be able to have any initial surface evolve for any length of time without any user intervention, as is now the case for string evolution as described in §2.3.
- An interactive graphical interface, which will let users select with a mouse the geometric elements (vertices, edges, etc.) they wish to work with, and it could be an interactive tool for the design of initial surfaces and datafiles.

11 Bibliography

- [Almgren 1982] F. Almgren, “Minimal surface forms,” *The Mathematical Intelligencer* **4**, no. 4 (1982), 164–172.
- [Almgren and Taylor 1976] F. Almgren and J. Taylor, “The geometry of soap films and soap bubbles,” *Scientific American*, July 1976, 82–93.
- [Brakke 1977] K. Brakke, *The motion of a surface by its mean curvature*, Princeton University Press, Princeton, NJ (1977).
- [Brakke 1991a] K. Brakke, *Surface Evolver Manual*, Research Report GCG 31, The Geometry Center. See “Software Availability” below for information on how to obtain it.
- [Brakke 1991b] K. Brakke, “The opaque cube problem,” in *Computing Optimal Geometries* (video and booklet edited by J. E. Taylor), American Mathematical Society, Providence RI, 1991.
- [Brakke 1992a] K. Brakke, “Minimal surfaces, corners, and wires,” *Journal of Geometric Analysis* **2** (1992), 11–36.
- [Brakke 1992b] K. A. Brakke, “Grain growth with the Surface Evolver,” in *Video Proceedings of the Workshop on Computational Crystal Growing* (edited by J. E. Taylor), American Mathematical Society, Providence, RI, 1992.
- [Callahan et al. 1992] M. Callahan, P. Concus and R. Finn, “Energy minimizing capillary surfaces for exotic containers,” in *Computing Optimal Geometries* (video and booklet edited by J. E. Taylor), American Mathematical Society, Providence RI, 1991.
- [Dobkin and Laszlo 1987] D. Dobkin and M. Laszlo, “Primitives for the manipulation of three-dimensional subdivisions,” technical report CS-TR-089-87, Department of Computer Science, Princeton University, Princeton, NJ, 1987.
- [Mittelman] H. Mittelman, “Symmetric capillary surfaces in a cube” *Math. Comp. Simulation* **35** (1993) 139–152.
- [Morgan 1986] F. Morgan, “Harnack-type mass bounds and Bernstein theorems for area-minimizing flat chains modulo ν ,” *Comm. Part. Diff. Eq.* **11**, (1986) 1257–1283.
- [Morgan 1994] F. Morgan, “Surfaces minimizing area plus length of singular curves”, *Proc. AMS* **120** (1994) 1153–1161.
- [Press et al. 1988] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling, *Numerical Recipes in C*, Cambridge University Press, New York, 1988.
- [Racz et al.] L. M. Racz, J. Szekely and K. A. Brakke, “A General Statement of the Problem and Description of a Proposed Method of Calculation for Some Meniscus Problems in Materials Processing”, *ISIJ International* **33** No. 2, (February 1993) 328–335.
- [Sibson 1978] R. Sibson, “Locally equiangular triangulations”, *Comput. J.* **21** (1978) 243–245.
- [Taylor 1976] J. E. Taylor, “The structure of singularities in soap-bubble-like and soap-film-like minimal surfaces,” *Ann. Math.* **103** (1976), 489–539.
- [Taylor 1983] J. E. Taylor, “Constructing crystalline minimal surfaces”, pp. 271–288, in *Seminar on Minimal Submanifolds* (edited by E. Bombieri), Annals of Math. Studies **105** Princeton University Press, Princeton, NJ, 1983.
- [Taylor 1988] J. E. Taylor, “Constructions and conjectures in crystalline nondifferentiable geometry,” *Proceedings of the Conference in Differential Geometry*, Rio de Janeiro, 1988, Pittman Publishing Ltd.

[Tegart 1991] J. Tegart, “Three-dimensional fluid interfaces in cylindrical containers,” AIAA paper AIAA-91-2174, 27th Joint Propulsion Conference, Sacramento, CA, June 1991.

[Thompson 1887] W. Thompson (Lord Kelvin), “On the division of space with minimum possible error,” *Acta Math.* **11**(1887), 121–134.

12 Availability

The Surface Evolver program is in the public domain, and there is no charge for it. It is written to be portable between systems. The only requirement is a C compiler. So far, it has been ported to Sun, Iris, NeXT, Xenix, and MS-DOS systems. The major effort in porting to a new system is writing a screen graphics interface. However, this is fairly simple, since the system-dependent routines need only display triangles. The program can also be run without any screen graphics, which makes it possible to run remotely.

A package containing source code, manual, and sample data files is available by anonymous ftp from geom.umn.edu (128.101.25.31) as `pub/evolver.tar.Z`. For those with NeXT computers, there is also a separate ftp archive `pub/evolver.next.tar.Z` which contains a NeXT executable Evolver and Interface Builder files. The Evolver is also available on floppy disk from the author. The manual in \TeX dvi format is included in the ftp archive. A hardcopy version of the manual can be requested separately from the author, or directly from The Geometry Center, 1300 South Second Street, Minneapolis, MN 55454. (Later note: the previous paragraph is obsolete. The Evolver home web page is now <http://www.susqu.edu/evolver>. Geomview is now available from <http://www.geomview.org>.)