

CAD Tools for Aesthetic Engineering

Carlo H. Séquin

CS Division, University of California, Berkeley, CA

ABSTRACT

The role of computers and of computer-aided design tools for the creation of geometrical shapes that will be judged primarily by aesthetic considerations is reviewed. Examples are the procedural generation of abstract geometrical sculpture or the shape optimization of constrained curves and surfaces with some global “cost” functional. Different possibilities for such “beauty functionals” are discussed. Moreover, rapid prototyping tools based on layered manufacturing now add a new dimension to the visualization of emerging designs. Finally, true interactivity of the CAD tools allows a more effective exploration of larger parts of the design space and can thereby result in an actual amplification of the creative process.

Keywords: Shape optimization, geometrical sculpture, sculpture generator, rapid prototyping.

1 INTRODUCTION

In this tutorial, we are concerned with computer-aided design tasks in which the final evaluation is mostly based on aesthetic criteria. While most engineers accept the fact that one needs to use computers to design jet engines, computer chips, or large institutional buildings, it is less clear whether computers are also useful in the design of artifacts that are judged mostly by their looks. In a traditional CAD setting, the computer primarily serves as a precise drafting and visualization tool, permitting the designer to view the emerging geometry from different angles and in different projections. A digital representation also makes it possible to carry out some analytical tasks such as determining volume or surface area of a part.

We will show that today the role of the computer goes much further. It actively supports the creation of geometric shapes by procedural means and can even optimize a surface by maximizing some “beauty functional.” It further can help to extend visualization aids for complex parts through the generation of rapid prototypes on layered manufacturing machines. Finally, it may even amplify the creative process itself by allowing the designer to quickly explore a much larger domain of design alternatives.

The objects used as examples in this tutorial are mostly abstract geometrical sculptural forms or mathematical visualization models (Fig.1). However, the principles and techniques discussed are readily applicable also to consumer products, or automotive parts and shapes. Creating maximally satisfactory forms for mathematical models or for geometric sculptures poses quite different requirements and constraints for any CAD tool than

developing an optimized airplane wing or designing the most powerful computer chip. Real-time interactivity becomes a crucial factor, when a designer's eye is the key evaluation instrument in the design loop.

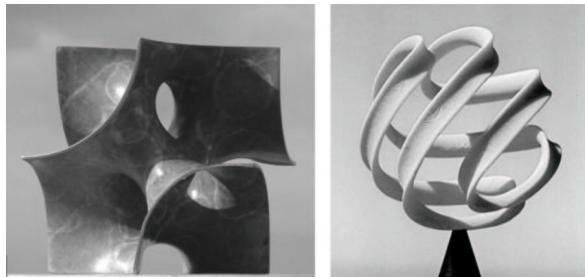


Fig. 1. Geometrical sculptures: (a) *Volution_5*, (b) *Altamont*.

This tutorial overview starts by looking at some generic tasks in curve and surface design, in particular, ongoing efforts for defining a “beauty functional” for procedurally optimizing shapes that are only partially constrained by the designer. It then discusses some research aimed at finding efficient implementations and approximations of such optimization functionals, so that they can be used at interactive design speeds. Next, we look a parameterized design paradigm that allows an artist to rapidly explore and compare many alternative versions of a geometrical shape. Finally, we make the point that a CAD tool that is well matched to the task at hand is much more than just a “drafting assistant” and can indeed become an amplifier for one's creative spark.

2 OPTIMIZATION OF SMOOTH SURFACES

Smooth surfaces play an important role in engineering and are a main application for many industrial CAD tools. Some surfaces are defined almost entirely by their functions; examples are ship hulls and airplane wings. Other surfaces combine a mixture of functional and aesthetic concerns, e.g., car bodies, coffee cups, flower vases... Finally, for some cases, aesthetics dominates the designer's concern, for instance in abstract geometric sculpture.

2.1 Beauty Functionals

For either situation, it can be argued that an ideal surface design system should allow a designer to specify all the boundary conditions and constraints and then provide the "best" surface under these circumstances. "Best" in the context of this tutorial would mean an optimization with respect to some intrinsic surface quality related to its aesthetic appeal. To be usable in a CAD tool, that quality has to be expressible in a functional or procedural form. Commonly, the characteristics associated with "beautiful" or "fair" surfaces imply smoothness - at least tangent-plane (G^1 -) continuity, but often also curvature (G^2 -) continuity. If the surface is covered with some textural pattern, then we have to demand more than just geometric continuity and also require smoothness of the parametrization, i.e., C^1 - or C^2 -continuity, respectively. Additional characteristics often cited in the definition of aesthetic shapes are symmetry and simplicity [1]. The first implies that symmetrical constraints should result in symmetrical solutions; and the second implies avoidance of unnecessary undulations or ripples.

All these properties are exhibited by *minimal surfaces*, i.e., by the shapes assumed by thin soap membranes spanning some given boundary (as long as the air pressure on both sides is the same). Experimentally, such shapes can be generated by dipping a warped wire loop into a soap solution. The lateral molecular membrane-forces will try to minimize overall surface area and thereby implicitly create a minimal saddle surface in which the mean curvature at every point of the surface assumes the value zero.

$$\text{MinimalSurface} \Rightarrow \kappa_1 + \kappa_2 = 0; \text{ everywhere} \quad (1)$$

An extension of this principle to include closed surfaces can be obtained by minimizing the total bending energy of the surface. In an abstraction and idealization that goes back to Bernoulli, the local bending energy of a thin filament or a thin sheet of stiff material is proportional to the square of the local curvature. The total bending energy of a shape then can be obtained as an arc-length or area integral of curvature squared over the whole shape.

$$\text{MinimumEnergyCurve} \Rightarrow \int \kappa^2 dL = \min \quad (2)$$

$$\text{MinimumEnergySurface} \Rightarrow \int \kappa_1^2 + \kappa_2^2 dA = \min \quad (3)$$

For closed surfaces, it turns out that minimizing bending energy is equivalent to minimizing mean curvature,

since the area integral of Gaussian curvature, $G=\kappa_1\kappa_2$, is a topological constant that depends only on the genus of the surface:

$$\int \kappa_1 \kappa_2 dA = 4\pi(1 - \text{genus}) \quad (4)$$

Four times the square of mean curvature, $H=(\kappa_1+\kappa_2)/2$, can also be written as:

$$(\kappa_1 + \kappa_2)^2 = \kappa_1^2 + \kappa_2^2 + 2\kappa_1\kappa_2 \quad (5)$$

Using mean curvature H as the energy functional is also known as Willmore energy [10], and the possible minimal-energy shapes for surfaces of different genus are well known [10]. For surfaces of genus 0, the minimal shape is, of course, a sphere, and it has a total bending energy of 4π regardless of its size, since the bending energy functional happens to be scale-invariant. For genus 1, bending energy is minimized in the *Clifford torus* in which the ratio of the two defining radii is equal to $\sqrt{2}$. For a higher genus, the energy minimizing shape is the *Lawson surface*, and the total Willmore energy for all these surfaces lies below a value of 8π .

For surfaces with a genus less than about 6 (Fig.2a), these minimal-energy shapes are quite pleasing to look at. With increasing genus, these surfaces approximate ever more closely two spheres intersecting along a circle of alternating tiny pillars and holes, reminiscent of the central portion in Scherk's second minimal surface [14] wrapped into a toroidal ring. Figure 2b shows some slices of this surface for the genus-11 case, revealing the shape of the obscured central parts. Most people do not think that this is an aesthetically optimal shape for the higher genus surfaces.

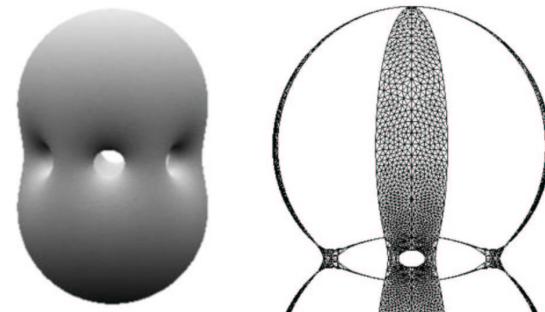


Fig. 2. Energy-minimizing Lawson surfaces: (a) the genus-5 case
(b) slices of the genus-11 surface.

It has been argued [12] that bending energy may not be the best "beauty functional." For the surfaces of higher genus, most people prefer a better balance between the toroidal handles and the holes between them. Also, if the perfect genus-0 shape is indeed a sphere, shouldn't the "penalty" (energy) function assume the value 0 for that shape? Thus we might obtain a better functional to evaluate the fairness of a curve or surface, if we try to minimize the integral over the 'change of curvature' squared, instead.

Moreton has created a first implementation of such a functional by integrating the squares of the derivatives of the principal curvatures in the directions of their respective principal directions [12].

$$\text{Minim.VariationCurve} \Rightarrow \int \frac{d\kappa^2}{de} dL = \min \quad (6)$$

$$\text{Minim.VariationSurface} \Rightarrow \int \frac{d\kappa_1^2}{de_1} + \frac{d\kappa_2^2}{de_2} dA = \min \quad (7)$$

In surfaces where the principal lines of curvature are exact circles, this *minimum-variation* (MV) functional evaluates to zero. Thus all *cyclides* (spheres, cylinders, cones, tori, and even horned tori) are “perfect” surfaces of zero MVS cost. To obtain some discrimination between tori that are too “skinny” and those that are too “fat”, we could also introduce the mixed derivative terms into the functional, i.e., $d\kappa_1/de_2$ and $d\kappa_2/de_1$. The consequences of introducing such variants into the minimum-variation functional have not been studied yet.



Fig. 3. Minimum variation surfaces: (a) the genus-2 case (b) a genus-5 surface with cubic symmetry enforced.

The first system to create *minimum-variation surfaces* (MVS) used bi-quintic quadrilateral Bézier patches stitched together so as to form the desired shapes [12]. All the degrees of freedom contained in the coordinates of the control points that are not specified by design constraints were then varied with the goal to minimize the overall cost function. The components of the energy gradient of all the available degrees of freedom were determined with finite differences, and a conjugate gradient descent method was used to move the system towards a local optimum. The area integral over the change of curvature was evaluated by Gauss-Legendre or by Lobatto quadrature, typically using about 400 sample points per Bézier patch. Penalty functions using Lagrange multipliers were employed in an inner optimization loop to enforce G^1 - and G^2 -continuity across the seams between adjacent patches. The system was very slow, using many hours for converging on even simple symmetrical shapes (Fig.3); but it produced beautiful results [12]. The challenge now exists to implement the evaluation of these cost functionals so that surfaces can be optimized at interactive rates.

2.2 Interactive Surface Optimization

Now, a decade later, what are the prospects for evaluating such functionals at the desired, almost instantaneous and truly interactive rate?

First, of course, computer power has increased by one to two orders of magnitude over the last decade, thus bringing us closer to our goal of full interactivity, even without any further innovations.

Second, and most importantly, subdivision surfaces have become mature and popular. They allow us to obtain surfaces with a reasonable degree of built-in continuity by their inherent construction, thus avoiding the very costly inner optimization loops that were used originally to guarantee smoothness at the seams. For instance, Catmull-Clark subdivision surfaces can offer G^1 -continuity everywhere and exhibit C^2 -continuity almost everywhere, except at extraordinary points where quadrilateral patches join with a valence different from 4.

Third, the inherently hierarchical organization of subdivision surfaces gives us the possibility to optimize the gross shape of the surface at a relatively coarse level, where only a small number of control points have to be adjusted. Then as we gradually refine the surface by increasing the level of subdivision, the number of degrees of freedom grows at a quadratic rate; but since the surface is already relatively close to the desired shape, the optimization procedure need not run for many iterations to achieve convergence.

Fourth, at the research frontier, experiments are now under way to find means to avoid the expensive numerical integration steps in the inner loop of the optimization. The aim is to find a discretized approximation of the salient surface characteristics, to obtain directly an estimate of the behavior of the cost functional that is good enough to guide the gradient descent optimization in the right direction.

2.3 The Basic Framework

As our basic framework, we use subdivision surfaces to represent the shapes to be optimized. Using finite differences based on incremental movements of the control vertices, a gradient vector for the chosen cost/energy functional is obtained and then used to evolve the surface iteratively towards a local cost minimum. After obtaining the minimum energy surface for a given mesh resolution, the mesh is subdivided to produce new vertices and therefore new parameters for optimization. In this general approach, we can vary the methods for calculating the actual optimization moves, trading off accuracy for speed.

As a baseline for comparing the various methods, we use exact evaluation of the subdivision surface [20], sampling the limit surface to obtain its geometric properties. Using differential geometry and numerical integration by Gauss-Legendre quadrature, we can compute with high accuracy a cost functional such as the bending energy. Using this energy computation in the above framework, we have obtained robust results that agree with the theoretically known energy minima for some highly symmetrical smooth

surfaces, such as spheres, tori, or the known energy minimizers of higher genus [10]. Since numerical integration and gradient calculations are computationally expensive, this method may take a few hours for surfaces like those depicted in Figures 2 and 3. However, it serves as an excellent benchmark for evaluating the following more approximate methods.

2.4 Approximating the Cost Functional

A first simplification calculates an approximate cost functional directly from the discrete mesh of control points of the subdivision surface, as is done, for instance, in [11][9]. We are exploring vertex-based as well as edge-based functionals that express the surface energy as a summation over the local energy at all the vertices or edges. These local energies are calculated with a discretized approximation, using polynomial expressions of vertex coordinates and/or dihedral angles along the edges. These simpler functionals are adequate to guide the gradient descent process in the same direction as a more exact functional evaluation would, but do so at significantly reduced cost and thus with higher speed.

An example of such an approach is used in Brakke's *Surface Evolver* [2]. Vertices are moved so as to locally minimize surface area. The local area considered is simply the sum of the areas of all the triangles surrounding a vertex, and the vertex is moved along the logarithmic gradient of that area (Fig.4).

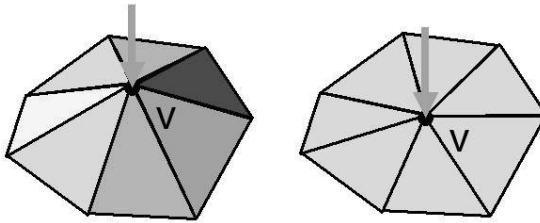


Fig. 4. Minimization of the area surrounding a vertex in Brakke's *Surface Evolver*.

In order to emulate functionals that rely on bending energy, we also have successfully used a formulation based on the dihedral angles along the edges of the subdivision polyhedron. For all edges we sum up the squares of the dihedral angles, weighted by the length of the edge, and normalized by the heights of the two attached triangles:

$$\text{TotalEnergy} = \sum_E \frac{\beta^2 \cdot \|e\|}{\|h_1\| + \|h_2\|} \quad (8)$$

For various test cases, ranging from spheres to more complex surfaces of genus 3, we have compared the shapes obtained in mere minutes with this discretized functional (Fig.5a) to previously calculated benchmark shapes, and we found the results to be in very good geometric agreement.

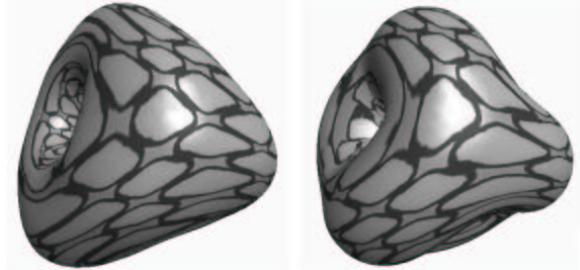


Fig. 5. Genus-3 surfaces: (a) MES obtained by minimizing a discretized bending energy; (b) MVS obtained by approximating minimum curvature variation with a direct vertex-move calculation.

2.5 Direct Vertex Move Calculations

A second simplification step tries to avoid also the gradient calculation based on finite differences. Instead we calculate directly the moves for the control vertices that promise to optimize the surface in the desired direction. As an example, we have developed a vertex-move procedure that aims to minimize the *variation* of curvature as attempted by [12]. For this purpose, we calculate for each edge in the control mesh a *change in turning angle* in the direction of the edge, and then aim to swivel the edge about a point on it so as to reduce this turning variation. Each vertex obtains a suggested move component from every edge attached to it, and it is then moved proportional to the mean of these components. Figure 5b shows a surface obtained by this direct method; the shape is very close to the shape found in 1992 after many hours of computation [12], but now it can be generated in just a few seconds!



Fig. 6. The desired future way of modeling car hoods with an interactive, constraint-based CAD system.

2.6 Interactive CAD Applications

With this speedup resulting from the use of discrete functionals and/or direct vertex-move calculations, we can envision a CAD system in the not-too-distant future, where the designer specifies boundary conditions and constraints for a surface panel (Fig.6), and then picks a suitable cost functional for a quick optimization of the surface. The designer may compare and contrast the results of using two or three different aesthetic functionals and choose the one that is most appropriate for the given application domain. The designer further can adjust some of the original

constraints or add new ones to force the surface to meet functional as well as aesthetic expectations. The role of the chosen functional is to take care of the details of the surface shape, e.g., to avoid geometric discontinuities or unneeded wrinkles and slope changes.

3 FAIR CURVES ON FAIR SURFACES

A second key CAD problem is the embedding of “beautiful” or “fair” curves onto the kind of optimized surface discussed above. For instance, one may need to draw a fair connecting line between two points on a smooth surface. The most direct such connection is a geodesic line, which exhibits no gratuitous lateral curvature. While it is easy to trace a directional geodesic ray on a smooth surface or on a finely tessellated polyhedral approximation thereof, it is a well-known hard problem to connect two points with the shortest geodesic path on a surface that exhibits many areas of positive and negative mean curvature.

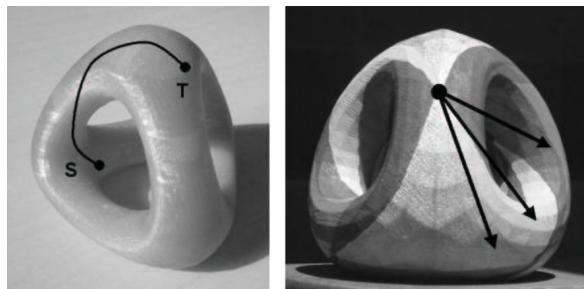


Fig. 7. (a) Geodesic line between two points.
(b) LVC curves with adjustable end-tangents.

Sometimes the geodesic line segment is too restrictive for design purposes; it offers no degrees of freedom or adjustable parameters to the designer. This limitation is particularly detrimental when multiple lines must radiate from the same point. In this situation, a designer would like to have some control over the initial tangent directions of these lines, perhaps to distribute them at equal angles around the point from which they emerge. For this purpose, a good alternative is a line for which its geodesic curvature is either constant or varies linearly as a function of arc length (Fig.8). Such LVC-curves offer the designer two parameters: the values of geodesic curvature at either end of the line segment. These can then be used to set the tangent directions at the two end-points (similar to the controls available in a Bézier curve in the plane). We have developed a scheme to efficiently calculate a good approximation to such LVC-curves on subdivision surfaces.

We will illustrate the use of this technique with an example from mathematical topology concerning a crossing-free embedding of a complicated non-planar graph on a surface of a suitably high genus. E.g., K_{12} , the complete (fully connected) graph of 12 nodes, requires a genus-6 surface for an embedding with no crossings, and the 66 edges of this graph will then divide the surface into

44 3-sided regions. To make pleasing-looking, easy-to-understand models of this partitioned surface, we want to make all edges as “fair” as possible, i.e., keep them nice and smooth with no unnecessary undulations. At the same time we would like to have the edges more or less evenly distributed around the nodes where they join. LVC-curves offer just the right amount of control for our purpose.

3.1 Our Approach

The designer starts by constructing a coarse polyhedral model of the needed genus-6 surface (Fig.10a). Choosing the oriented tetrahedral symmetry group for this surface and exploiting this symmetry to the fullest, the user only has to construct 1/12 of the surface, which can easily be done with 9 quadrilaterals or 18 triangles. The complete surface is then constructed by composing twelve copies of this fundamental domain with suitable rotations. On this surface, the user now places the nodes of the graph and draws piecewise linear connections between them (Fig.10a). If the graph also is given the same tetrahedral symmetry, then this work needs to be done only on the fundamental domain, i.e. on 1/12th of the surface.

Our algorithm starts from this polyhedral model. The triangle- or quad-mesh is the basis of a Loop or Catmull-Clark subdivision surface, and the piecewise linear paths between nodes will be converted into LVC-curve segments. The two refinement processes occur in parallel. For each generation of the subdivision process, the piecewise linear paths are modified so as to approximate a curve with linearly varying curvature (LVC).

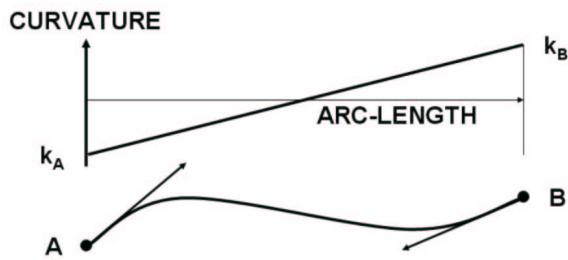


Fig. 8. (a) Path with linearly varying curvature (LVC) as a function of arc length; (b) this allows to control the end tangents separately.

Towards this goal, the vertices where the paths cross over the edges of the control mesh (Fig.9) are moved with a gradient descent method to approach the desired LVC-behavior. Specifically, each such vertex is moved along the edge on which it lies, so as to drive a discretized estimate of geodesic curvature at that point towards the mean of the geodesic curvature values at the two neighboring points on that path. A few dozen iterations of this optimization step are typically sufficient. After this curve optimization process has converged, the surface is subjected to another subdivision step. All linear path segments across all facets in the mesh are then split at the new subdivision edges, and all the path vertices are subjected again to the curve

optimization process. This general process loop is repeated until the desired degree of refinement has been reached. The technique works with many popular subdivision schemes.

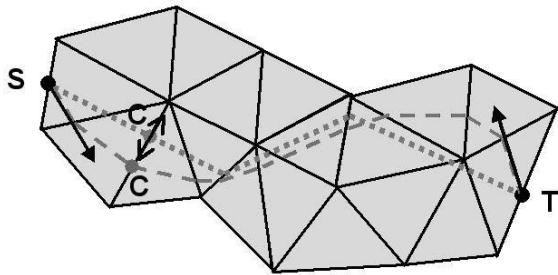


Fig. 9. Optimizing a discretized LVC curve linking S and T; the original path is the one with only three segments.

3.2 Results

The result of this process for the embedding of the K_{12} graph on a genus-6 surface of tetrahedral symmetry is shown in Figure 10b. The LVC curves have been enhanced to black bands to make them more visible, and the nodes of the graph are shown as small hemispheres. The 44 resulting 3-sided facets between the edges have been colored randomly. Thus we are able to provide a crisp visualization model for this difficult graph-embedding problem.

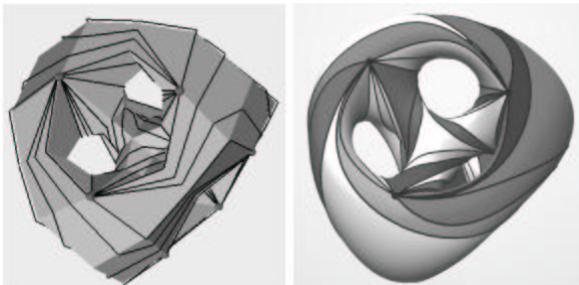


Fig. 10. (a) Initial piecewise linear paths on polyhedral model.
(b) Final optimized LVC curves on subdivision surface.

4 PARAMETERIZED SHAPE GENERATION

The design and implementation of geometrical sculpture is a relatively novel application domain for CAD, in which the techniques outlined above would be particularly useful. In 1995 I started to collaborate with Brent Collins, a wood sculptor who creates fascinating abstract geometrical shapes [3][4][8]. His work can be grouped into cycles that have a common recognizable constructive logic to them, and which exhibit a timeless beauty that captured my attention immediately when I first saw photographs of his work in *The Visual Mind* [7][8].

My interaction with Brent Collins was triggered by images of his *Hyperbolic Hexagon* (Fig.11a), which can be understood as a toroidal warp of a six-story segment of the core of Scherk's second minimal surface [14] (Fig.11b). In

our very first phone conversation, we discussed the question of what might happen if one were to take a *seven-story* segment of such a chain of cross-wise connected saddles and holes, and then bend it into a circular loop. We realized that the chain would have to be given an overall longitudinal twist of 90° so its ends could be joined smoothly. We further envisioned that interesting things might happen in this process: the surface may become single-sided, and its edges could join into a single continuous edge, forming a higher-order torus knot.

Since neither of us could visualize exactly what such a construction would look like, we both built little mock-up models from paper and tape (Séquin) or from pipe segments and wire meshing (Collins). In subsequent phone discussions, we expanded the scope of this paradigm. We asked ourselves, what would happen, if we gave the Scherk tower (Fig.11b) a stronger twist of, say, 270° , or of any additional 180° , which would allow the ends of the saddle chain to join smoothly. We also pondered what would a sculpture look like that uses third-order ("monkey") saddles, or even higher-order saddles, rather than the ordinary (biped) saddles of the original *Hyperbolic Hexagon*? What would be the proper amount of twist that such structures needed in order for the toroidal ring to close smoothly?

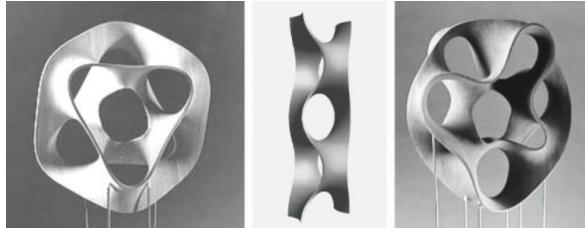


Fig. 11. (a) Collins' *Hyperbolic Hexagon*, (b) 4-story Scherk tower, (c) Collins' *Hyperbolic Heptagon*.

Constructing a realistic maquette of these relatively complex structures, precise enough for aesthetic evaluation, can be a rather labor-intensive process. During the first year of our collaboration, our ideas were coming forth at a rate much greater than what we could possibly realize in physical models. This led me to propose the use of the computer to generate visualizations of the various shapes considered, to judge their aesthetic qualities and to determine which ones might be worthwhile to implement as full-scale physical sculptures [15]. I started to develop a special-purpose computer program that could readily model these toroidal rings of Scherk's saddle chains, as well as all the generalizations that we had touched upon in our discussions. This led to *Sculpture Generator I*, which allowed me to create all these shapes interactively in real time by just choosing some parameter values on a set of sliders (Fig.12) [16].

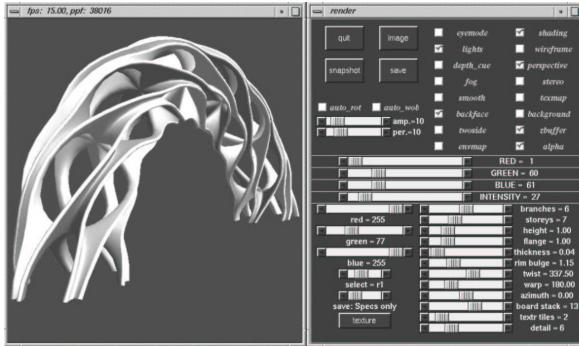
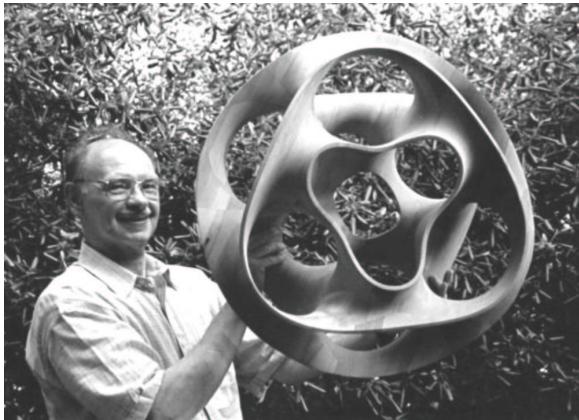


Fig. 12. Sculpture Generator I and its user interface

In the meantime, Collins had built the *Hyperbolic Heptagon* (Fig.11c), the twisted seven-story ring that we had first discussed on the phone. This two-foot wood sculpture showed us the potential of this paradigm of toroidal loops of saddle chains, and encouraged us to make additional sculptures of potentially much higher complexity. However, such sculptures would require more help from the computer than just the power of previewing the completed shape. Thus I enhanced my program with the capability to print out full-scale templates for the construction of these sculptures. The computer slices the designed geometry at specified intervals, typically 7/8 of an inch, and produces construction drawings for individual pre-cut boards from which the gross shape of the sculpture can then be assembled. Collins still has the freedom to fine-tune the detailed shape and to sand the surface to aesthetic perfection.

Fig. 13. Brent Collins holding *Hyperbolic Hexagon II*.

This eventually led to our first joint construction, the *Hyperbolic Hexagon II*, which features monkey saddles in place of the original biped saddles (Fig.13). It is possible that Collins could have created this shape on his own without the help of a computer. However, our next joint piece, the *Heptoroid*, a much more complex, twisted toroid, featuring fourth-order saddles (Fig.14a), would definitely

not have been feasible without the help of computer-aided template generation.

In a further extension of the Scherk-Collins paradigm, it was found, that Scherk's saddle-chain can be wound more than once around the toroidal ring. For a double loop, one needs to choose an odd number of stories, so that the saddles properly interlace on the first and second round. With an appropriate values for twist for the and flange-extensions, all self-intersection can be avoided (Fig.14b). With these generalizations of the original paradigm, intricate forms emerged whose relationship to the original *Hyperbolic Hexagon* are no longer self-evident.

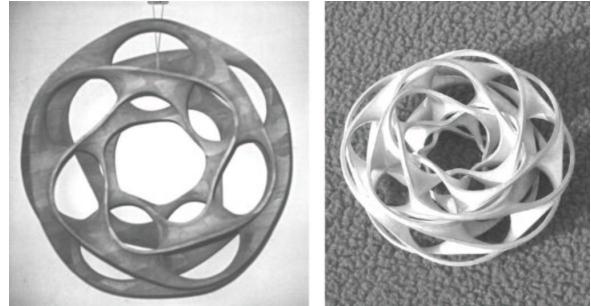


Fig. 14. (a) Heptoroid, from the collaboration with Brent Collins, and (b) doubly-wound quad Scherk-Collins toroid.

4.1 Capturing a Paradigm

In my interaction with Collins, an important new design task is added up front: I have to figure out what it is that I want my sculpture generator program to produce. This means that I first have to see a general underlying structure in a group of similar pieces in Collins' work and extract a common generating paradigm that can be captured in precise enough terms to be formulated as a computer program. This, by itself, is an intriguing and creative task. Moreover, if the paradigm is captured in a general enough form, it can then be extended to find additional beautiful shapes that have not yet been expressed in Collins' sculptures.

The question arises, whether a commercial CAD tool, such as *AutoCAD*, *SolidWorks*, or *ProEngineer*, would have been adequate to model Collins' sculptures. Indeed, with enough care, spline surface patches and sweeps could be assembled into a geometrical shape that would match one of Collins' creations. But this approach would be lacking the built-in implicit understanding of the constructive logic behind these pieces, which I wanted to generalize and enhance in order to produce many more sculptures of the same basic type. For that I needed stronger and more convenient procedural capabilities than those that commercial CAD tools had to offer. I chose C, C++, and OpenGL as the programming and graphics environments. The user interface originally relied on Mosaic and later on Tcl/Tk, in which my students had already developed many useful components, such as an interactive perspective viewing utility with stereo capabilities.

Capturing a sculpture as a program, forces me to understand its generating paradigm. In return, it offers precise geometry exploiting all inherent symmetries, as well as parametric adjustments of many aspects of the final shape. The latter turns out to be the crux of a powerful sculpture generator. If I build too few adjustable parameters into my program, then its expressibility is too limited to create many interesting sculptures. If there are too many parameters, then it becomes tedious to adjust them all to produce good-looking geometrical forms. Figuring out successful dependencies between the many different parameters in these sculptures and binding them to only a few adjustable sliders is the intriguing and creative challenge.

In practice it turned out that almost every sculpture family that I tackled, required a new program to be written. These programs became my virtual constructivist "sculpting tools." In the last few years, this virtual design environment has become more modular thanks to the SLIDE program library [19] created by Jordan Smith and enhanced with many useful modules for creating freeform surfaces by Jane Yen. Once a new program starts to generate an envisioned group of geometrical shapes, it often will take on a life of its own. In a playful interaction with various sliders that control the different shape parameters, and by occasional program extensions, new shapes are discovered that were not among the originally envisioned geometries. In this process the original paradigm may be extended or even redefined, and the computer thus becomes an active partner in the creative process of discovering and inventing novel aesthetic shapes [17].



Fig. 15. Hyper-sculpture: *Family of Twelve Scherk-Collins Trefoils* ($B=1$ to 4 from left to right; top: $W=2$; bottom $W=1$).

4.2 Illustrative Examples

In the *Family of Twelve Scherk-Collins Trefoils* (Fig.15), the space of parameter combinations is being explored for the range of saddles having from one to four "branches," and for single as well as double loops around the toroidal ring. The concept of a "saddle" has now been extended downwards to also include a single "branch" ($B=1$), which I

chose to be represented by a simple twisted band. For the case of the doubly wound loop ($W=2$), this band does self-intersect. For the single-branch case, the azimuth parameter has no relevant effect, and thus there are just single instances for the two cases, $W=1$ and $W=2$. For the cases with 2 and 3 branches, all possible constellations are exhibited, showing both (positive and negative) azimuth values (A_n, A_p) that give front-to-back symmetry for each case. For the fourth-order saddles ($B=4$) the structure becomes rather busy and starts to lose some of its aesthetic appeal; thus only a single azimuth value is shown for $W=1$ and $W=2$, respectively.

A graphical interface with individual sliders for each parameter allows the user of *Sculpture Generator I* to explore with ease the space of all Scherk-Collins toroids. For the twelve trefoils in this series (Fig.15), the width and thickness of the flanges was fine-tuned to optimize the aesthetic appeal of each particular trefoil by balancing the relative dimensions of the holes and branches and yielding a pleasing roundness – obviously a rather subjective process. The surface descriptions of the optimized shapes were then transmitted to a Fused Deposition Modeling machine [21] for prototyping of the twelve maquettes (see Section 4.3).

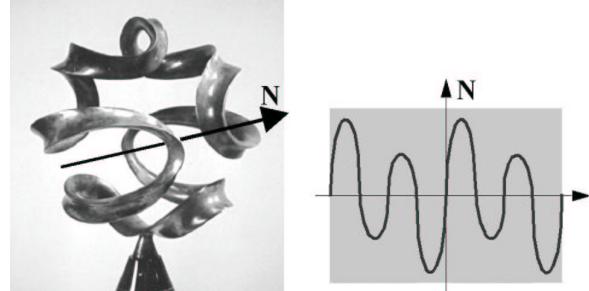


Fig. 16. (a) *Pax Mundi* by Brent Collins, and (b) its analysis as an amplitude-modulated sine-wave on the surface of a sphere.

As a second example I want to discuss a series of sculptures jointly called *Viae Globi*, since their shapes are reminiscent of the curvy pathways of an alpine road winding around large portions of a sphere. They were inspired by Brent Collins' *Pax Mundi* (Fig.16a; Fig.22 in [4]). When trying to find possible generative ideas behind this shape, I was reminded of other sculptures by Naum Gabo and by Robert Engman which also exhibit sweeping, meandering curves on the surface of a sphere, but of a simpler nature – more like the seams on a baseball. With this perspective in my mind, I could see *Pax Mundi* as comprising four periods of an amplitude-modulated sine-wave function wrapped around the equator of a globe (Fig.16b). Both, at the North pole (N) and at the South pole, a pair of opposite lobes closely approach each other, while the other pair of lobes is cut back in amplitude, so as not to cause any self-intersections. The generating framework in this case consisted of the specifications for the number of wave periods around the equator (say from 2 to 6) and for the amplitudes and widths (bulging) for each one of the lobes. This time, rather than writing a narrowly focused,

stand-alone generator program, I took a more modular approach based on the SLIDE program library [19].

The primary new piece of code that had to be written was a module to draw a nicely rounded meandering curve onto a sphere in accordance with the specified parameters. This curve would then serve as the path of a generalized sweep – a functionality that already existed in SLIDE. A second program module written in this context provided a parameterized description of a crescent-shaped cross section to match the profile used by Brent Collins. SLIDE already had the functionality to sweep any cross section along any sweep path, while controlling its azimuth and twist, as well as optionally varying the cross section by non-uniform affine scaling. The values for these latter parameters can be attached to any of the control points of the sweep path, and they are then interpolated by the same polynomial function that defines the spline for the sweep path.

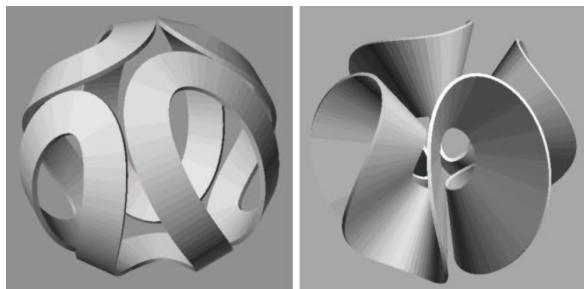


Fig. 17. Two different cross sections with different azimuth parameters swept along the same *Pax-Mundi*-like sweep path.

At a later time, the *Viae Globi* paradigm was further extended to allow the wave function to exhibit additional secondary wiggles of two or three times the base frequency. And finally, Jane Yen and later Kiha Lee even programmed special interactive curve editors, first based on a deCasteljau approach adapted to great circles on a sphere (Fig.18a), and later based on the very pleasing-looking interpolating circle splines [18] (Fig.18b). These tools allowed me to draw many intriguing free-form curves onto a sphere, which could then be imported into the SLIDE environment as sweep path definition. This approach resulted in more complex looking sculptures such as those depicted in Figures 1b and 20b.

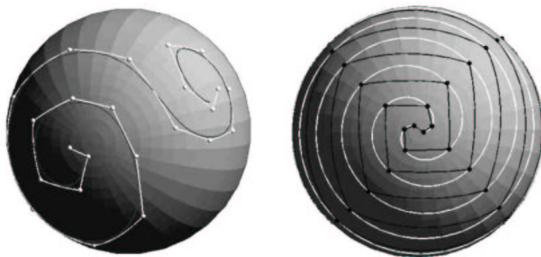


Fig. 18. Special purpose curve editors to make nicely rounded curves on a sphere: (a) based on an approximating deCasteljau method, and (b) based on interpolating circle splines.

Rapid Prototyping

The representation of geometrical shapes in a procedural form offers several advantages. Such designs can easily be optimized with the adjustment of a few parameters. More complex designs can be generated than could be crafted by traditional means. Interactive play with such parameterized programs extends the conceptual horizon of the designer and leads to new fertile insights. Finally, the procedurally generated output can readily be scaled to any size and can easily be targeted at various different interchange file formats. For instance, the simple, verbose, inefficient, but widely available .STL-format is accepted by most rapid prototyping machines and can thus be used to produce scale models by layered free-form fabrication.

In one such process, the Fused Deposition Modeling (FDM) process by Stratasys [21], the boundary representation of the sculpture is geometrically sliced into thin layers, 0.01 inches thick. These layers are “painted” individually, one on top of another, by a computer-controlled nozzle, which dispenses the ABS thermoplastic modeling material in a semi-liquid state at 270° centigrade, until the precise three-dimensional shape has been re-created (Fig.19).

In spite of the availability of ever more sophisticated rendering and visualization programs, physical 3D models play an important role in many design efforts. They are crucial to evaluate the tactile aspects of components such as the handle of a tool or the knobs on an appliance. They are needed to verify the proper functioning of a mechanism or the proper mating of parts in a modular assembly. But even for purely aesthetic artifacts, such as geometric sculptures, prototype maquettes, which can be readily inspected from all sides under varying lighting conditions, often reveal opportunities for design improvements.

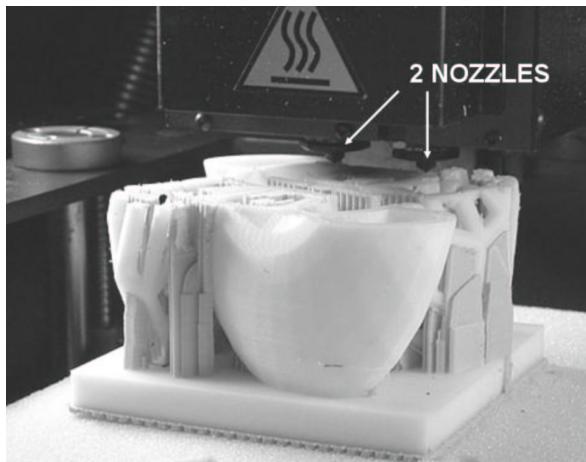


Fig. 19. Rapid prototyping of maquettes by layered manufacturing; look into the FDM machine by Stratasys.

It turns out that with some care the ABS plastic maquettes emerging from the FDM machine can be used

directly as the expendable masters for an investment casting process. This is particularly useful for very intricate and fragile geometries, where it would be difficult and impractical to make a mold from the original maquette in order to cast secondary wax masters for the investment casting process. Thus the ABS master is repeatedly dipped into a "Plaster of Paris" (silica) slurry, until a hard shell of some reasonable thickness has been formed. This shell is provided with a few drainage and venting holes, and is then heated to about 1100°C, where the plastic first melts and then evaporates rather cleanly. After some cleaning with compressed air, the shell is pre-heated again and can then be filled with molten bronze or with some other convenient casting metal. Figure 20 shows examples of sculptures that were cast in this manner by Steve Reinmuth [13]; he also provided them with the special patina that turns these shapes into true works of art.



Fig. 20. Bronze sculptures: (a) *Cohesion* – a Scherk-Collins toroid; (b) *Maloja* – a swept path on a sphere from the *Viae Globi* cycle.

4.3 Large-Scale Sculpture

When a sculpture is scaled from a desk model to a large size suitable for a public space, or when the material for its realization changes, the design will often have to be adjusted in subtle ways and cannot just be scaled uniformly. Cross-sectional profiles may have to be thinned or enlarged, flanges may have to be adjusted in thickness, and edges may have to be rounded differently. In this situation it is again a big advantage to have a suitably parameterized description of the geometrical form.

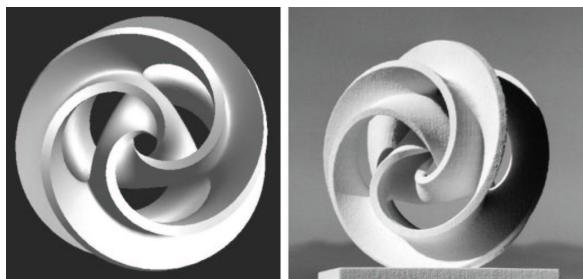


Fig. 21. Monkey Trefoil: (a) from Sculpture Generator I, and (b) fine-tuned into a maquette for a 12-foot snow sculpture.

This point was driven home quite clearly in the fall of 2002, when Collins and Séquin were invited on short notice to provide a design for the 13th Annual International Snow-sculpturing Championships in Breckenridge, Colorado. First the *Sculpture Generator I* was employed to create a couple of conceptual ideas (Fig.21a) for review by Stan Wagon [22], the experienced leader of our team. Based on his feedback, we could very quickly choose a set of parameters that would balance visual impact, complexity, and the potential for actually being realizable in snow. In a second refinement phase we then fine-tuned the parameters to optimally match the sculpture to the overall dimensions of the snow blocks (10' × 10' × 12 feet tall) made available for the competition. This final CAD description was then used to fabricate a scaled-down maquette on a rapid prototyping machine using a layered manufacturing technique (Fig.21b). The CAD representation also came in handy to make several orthogonal projections (Fig.22a) and cross sectional cuts as blueprints for on-site use during construction.

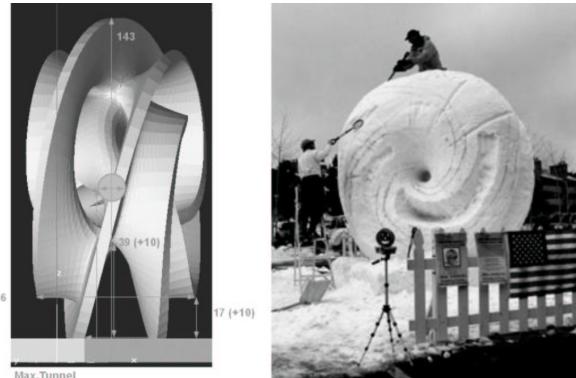


Fig. 22. Construction of snow sculpture: (a) blue-print projection from its side, (b) the definition of the flanges on the basic toroid.

The snow sculpting teams had four days to turn a 20-ton block of compact snow into some dramatic or whimsical display. To assure the regularity of our desired toroidal shape, we first reduced our block to a 6-foot thick vertical slab. On the faces of this slab we drew concentric circles defining outer perimeter, the major radius, and the small central hole of a perfect torus. We also made a half-circle plywood template with the minor radius of the torus. This template could be "swept" around the outer perimeter to guarantee uniform thickness and roundness of the torus. On this torus surface we could then mark the final visible edges of the spiral flanges (Fig.22b). From there, we proceeded with free-hand sculpting to create the desired shape. Our *Whirled White Web* turned out rather nicely (Fig.23) and was awarded the silver medal in that competition.

Fig. 23. Snowsculpture: *Whirled White Web*: (a) side, (b) front.

While it was fascinating for me as a designer to see one of my creations executed at the 12 foot scale, it leaves the question open, how one might go about realizing such shapes in a more durable material. (The *Whirled White Web* snow sculpture collapsed 45 minutes after judging had ended, due to unseasonably warm weather conditions).

One possible approach also resembles a layered manufacturing process. A large numerically controlled milling machine carves 1-2 feet thick sections of the sculpture out of high-density Styrofoam plates, and these sections are then manually assembled and glued together. Larger pieces of the sculpture could be used to form molds in which secondary wax masters can then be cast. Alternatively, the whole assembled Styrofoam sculpture could be provided with a more durable coating – similar to the surface of a surfboard. Thin layers of an Epoxy resin and of thin fiberglass matting are applied to the smoothed surface of the Styrofoam sculpture, until a hard weather-proof shell has been formed. Optionally this surface can also be painted with some outdoor enamel paints. This is the way some of the models in Disney World or such advertisement icons as the “Doggie Diner” puppy head have been constructed.

5 CAD TOOLS FOR AESTHETIC DESIGN

The utilities described in this paper exist in a research environment but have not yet made it into main-stream commercial CAD tools. What might the future of such tools look like?

5.1 Slow Adoption of New Ideas

Commercial CAD tools always lag one to two decades behind the first demonstration of a new idea in a research environment. Part of the reason is that customers as well as CAD tool vendors are very conservative and naturally resist changes. Learning to use a new CAD tool constitutes an investment of a fair amount of time, and most people are reluctant to learn a new tool, unless it is clear that it will offer significant advantages, or when the “old way” of doing things clearly becomes inadequate – as was the case for integrated circuit design during the 1980’s.

What are the main driving forces for new CAD tools for aesthetic engineering? Which market segments are large enough to bring about changes? The film industry is clearly one of those forces. Ever since *Toy Story*, Pixar and other computer graphics based animation firms have started to focus on more efficient and more robust representations to dynamic computer models that can be efficiently animated, shaded, and rendered. This has led to a resurgence of subdivision surfaces, and their extensions and modifications so that they can handle sharp features such as knife edges and creases.

Lack of standards is another problem that is slowing down the adoption of new ideas in the market place. The situation is reminiscent of the 1980’s when it became clear that designers wanted and needed NURBS in their tool kits. But support, particularly hardware-assisted support, was slow in coming. Companies like SGI, who would have been properly positioned to add such support to their graphics processor, were reluctant to pick one of the many possible NURBS variants that were used by different design teams and CAD companies. It was only when one particular format became sufficiently dominant that the use of NURBS and their support through proper tools took off.

Among the subdivision surfaces, Catmull-Clark surfaces and the triangle-based Loop surfaces now seem to gain ever more widespread acceptance. With the introduction of exact evaluation techniques [20], even in the presence of sharp features [6], these surfaces should also make their way from the film industry to the mechanical design community. Because subdivision surfaces can readily guarantee water-tight boundary descriptions of even very complex shapes of high genus, they are a good match for rapid prototyping by layered manufacturing.

5.2 Tools for Early Conceptual Design

The user interface of a CAD tool is often its weakest, least developed link. Often the user interface primarily reflects what is going on inside the computer, i.e., the data structures and procedures that were most easily programmed onto our classical computer processor architectures. But this may not be the easiest and most natural way for a designer to express herself or to approach a particular task.

In the early conceptual phases of design, the user interface plays a particularly important role. When one starts with a brand new concept, say, ‘*a bridge in the shape of a Moebius band*,’ it is often difficult to enter that first defining shape. 3D sketching tools, as far as they are available, are mostly inadequate. Many artists thus rather use clay, wire, scotch-tape, cardboard, or styrofoam, to make a first conceptual mock-up of a new geometric idea.

Experienced designers claim that creative thinking is greatly aided by some physical, bodily activity, such as pacing up and down the hallway, doodling with pencil and paper, bending wires, or folding and gluing paper pieces. Thus effective design ideation seems to involve more than just the eyes and perhaps a (3D?) stylus or other pointing

devices. How can future CAD tools provide a richer, more physical interface?

For quite some time, haptics has been seen as a promising technology that might revolutionize the user interface. However, this technology has yet to live up to its potential and it is still far from the needs of artistic designers. Commercially available systems typically include some haptic arm with a stylus that provides three or more degrees of freedom of force feedback. Often this device is used in a virtual reality (VR) setting, where a stereo display gives the appearance that the object under study is floating in free space in front of the display screen, where it can readily be reached with the haptic stylus. If the graphics display and the haptic model are properly collocated, then it indeed may feel like one is actually touching a 3D object suspended in space. Proper force-feedback, when the designer is modifying the shape by pushing or pulling parts of its surface, can further give the impression that one is actually bending or denting a thin piece of sheet metal. Some authors have even gone so far as to try to implement a *virtual clay* modeling system that allows the user to quickly make free-form shapes and distort them with free-form modeling operations [23][5].



Fig. 24. Virtual reality design station with a haptic arm, including a simulated appearance of the object under study.

But let's pause for a minute and evaluate such systems from the point of view of an artist who is used to working with real clay. What a difference! Human hands have more than a dozen separate "actuators" each and hundreds of "sensors" distributed densely all over the palms and the fingers. All of these elements work in concert when a piece of clay is squeezed and shaped into the desired form. Feedback about the current shape of the clay and particularly about its surface properties comes as much from the sensors in the hands as it comes from the eyes of the artist. Just consider, how impoverished clay modeling would be, if at any one time you could only touch a single point on the surface of your work piece! Perhaps DataGloves [] or some similar gadget will provide a better interface in the future – but this technology still has a long way to go before it meets the artist's needs.

However, computers and CAD tools do have other advantages. For most untrained people it would be difficult to make a nice, graceful clay vase; the results might typically be lopsided with un-even wall thickness. However, on a typical CAD package, such users could readily shape a simple, pleasing looking spline curve for the desired cross section and then apply a rotational sweep in order to get a good-looking result very quickly. Clearly the CAD environment has something to offer that we would not want to miss.

5.3 A Dream System

So, what is it that I would like to see in an "ideal" CAD system, useful for the initial design of, say, abstract geometric sculpture or free-form shapes for consumer products? Such a system should combine the best of both the virtual CAD environment and of the real physical world. As virtual elements, any construction parts possess infinite strength; they can be glued together easily, and just as easily be disassembled again. Of course, virtual constructions are not subject to gravity, and thus there is no need for any scaffolding – or for any glue for that matter.

However, some hands-on interaction seems rather desirable. Sweeping a hand through space to define a curve or a profile is often the most natural way to express one's intent. On other occasions, one would like to use a piece of physical material, such as an elastic steel blade, or a piece of heavy velvet cloth to define a shape that is then governed by the intrinsic properties of the chosen material. These physical artifacts would temporarily be collocated in the virtual context of the emerging design and would then be captured by some simple vision system, perhaps using structured light or some fast scanning process. The captured shape is then made available as a new geometric node in the design tree, and can be subjected to all the usual manipulations in the virtual design space. Once entered into the system, these shapes could then be assigned new – possibly fictitious – materials properties. Beams might bend like steel wires (MEC); surfaces may stretch like soap films (MES), or they could be subjected to some optimization process under the influence of one of several artificial "beauty" functionals, e.g., a minimum-variation functional (MVC, MVS).

Even limited haptics capability can serve a role. At the very least, it can provide some resistance when a user is drawing on an already existing shape to mark a region for modification (extrusion, indentation) or for the precise attachment of another part. In assembly tasks, haptics can provide important feedback when two parts come into physical contact.

Finally, we should put to good use the ever increasing compute power available in a designer's work station. When desirable, one should be able to effortlessly engage a simulator of the appropriate type. With its help the designer could then analyze and possibly optimize the design. For instance, an artist designing a large bronze sculpture may want to determine how strong it has to be made to withstand

gravity and wind forces, while at the same time keeping the amount of bronze needed as low as possible. Ideally, the corresponding simulations could be switched on seamlessly in the design environment if and when the artist is ready to deal with those issues.

6 CONCLUSIONS

Computer-aided design tools are gradually also becoming more suitable for aesthetic engineering and for artistic shape optimization. Subdivision surfaces are being introduced to commercial CAD tools, and systems for exact evaluation and for optimization of such surfaces with a variety of tailor-made energy functionals have been demonstrated in research labs. Computers are now powerful enough to subject these shapes to various simulations, ranging from simple stress analysis to the evaluation of their sound qualities. As the "mechanics" of design receives ever better support, the creative aspects of design and thus the early conceptual phases of design will need more attention. This is the domain where CAD tools need the most improvement. A more direct physical user interface would clearly be desirable, but haptics technology may take a while to live up to its expectations.

The key concern, in whatever is being offered for the conceptual design phase, is that it be fast enough to allow true real-time interactivity. Any tool that slows down the designer's creative thinking process will not be used. But when the designer can play with ease and explore effortlessly many new conceptual ideas, the CAD environment can truly become an amplifier for one's creative impulses.

7 ACKNOWLEDGMENTS

I would like to acknowledge the technical contributions by Pushkar Joshi and by Ling Xiao to the prototype CAD tools discussed in this paper. This work was supported by MICRO research grant 03-077, "Collaborative Design Environment," and by the CITRIS Institute, one of the California Institutes for Science and Innovation (CISI).

8 REFERENCES

- [1] Birkhoff, G. D., *Aesthetic Measure*. Harvard University Press, Cambridge, Mass. (1933).
- [2] Brakke, K., The Surface Evolver. *Experimental Mathematics*, vol. 1, no. 2, pp. 141-165, (1992)
- [3] Collins, B., Evolving an Aesthetic of Surface Economy in Sculpture. *Leonardo*, Vol. 30, No. 2, pp. 85-88, (1997).
- [4] Collins, B., Finding an Integral Equation of Design and Mathematics. *Proc. BRIDGES'98*, Winfield KS, pp. 21-28 (1998).
- [5] Dachille, F., Qin, H., and Kaufman, A. E., A novel haptics-based interface and sculpting system for physics-based geometric design. *Computer-Aided Design* 33, 5, pp 403-420, (2001).
- [6] DeRose, T., Kass, M., and Truong, T., Subdivision Surfaces in Character Animation. *Proc. of ACM SIGGRAPH'98*, pp 85-94, (1998).
- [7] Emmer, M., ed., *Visual Mathematics*, special issue, *Leonardo*, Vol 25, No. 3,4, pp. 313-320, (1992); reprinted in M. Emmer, ed., *The Visual Mind*, MIT press (1993).
- [8] Francis, G. K. and Collins, B., On Knot-Spanning Surfaces: An Illustrated Essay on Topological Art. *Leonardo*, Vol 25, No. 3,4, pp. 313-320, (1992).
- [9] Grinspun, E., Hirani, A., Desbrun, M. and Schröder, P., Discrete Shells. *Proc. of ACM Eurographics Symposium of Computer Animation* 2003.
- [10] Hsu, L., Kusner, R., and Sullivan, J., Minimizing the squared mean curvature integral for surfaces in space forms. *Experimental Math* 1, 3, pp. 191-207, (1992).
- [11] Meyer, M., Desbrun, M., Schröder, P., Barr, A. H., Discrete differential-geometry operators for triangulated 2-manifolds. In Hans-Christian Hege and Konrad Polthier, editors, *Visualization and Mathematics III*, pages 35-57. Springer-Verlag, Heidelberg, 2003.
- [12] Moreton, H., and Séquin, C. H., Functional optimization for fair surface design. *Proc. of ACM SIGGRAPH'92*, pp. 167-176. (1992).
- [13] Reinmuth, S., Bronze Casting of Geometrical Sculpture. *Proc. of MOSAIC*, Seattle, WA, pp 209-214, (2000).
- [14] Scherk, H. F., Bemerkungen über die kleinste Fläche innerhalb gegebener Grenzen. *J. Reine Angew. Math. (Crelle's Journal)* 13, pp 185-208 (1835).
- [15] Séquin, C. H., Virtual Prototyping of Scherk-Collins Saddle Rings. *Leonardo*, Vol. 30, No. 2, pp. 89-96. (1997).
- [16] Séquin, C. H., Interactive Generation of Scherk-Collins Sculpture. *Proc. Symp. on Interactive 3D Graphics*, Providence RI, pp. 163-166, (1997).
- [17] Séquin, C. H., Computer-Augmented Inspiration. *Proc. ISAMA'99*, San Sebastian, pp. 419-428, (1999).
- [18] Séquin, C. H., Lee, K., and Yen, J., Fair G²- and C²-Continuous Circle Splines for the Interpolation of Sparse Data Points. Accepted for publication: *Computer-Aided Design* (2004).
- [19] SLIDE design environment: <http://www.cs.berkeley.edu/~ug/slides/docs/slides/spec/> (2004).
- [20] Stam, J., Exact Evaluation of Catmull-Clark Subdivision Surfaces at Arbitrary Parameter Values. *Proc. of ACM SIGGRAPH'98*, pp. 395-404, (1998).
- [21] Stratasys Corp., Fused Deposition Modeling machines: <http://www.stratasys.com/> (2004).
- [22] Wagon, S., Refs. to snow-sculpting on home page: <http://www.stanwagon.com/> (2004).
- [23] K. T. McDonnell, H. Qin, and R. A. Wlodarczyk. Virtual Clay: A real-time sculpting system with haptic toolkits. In *Proc. 2001 ACM Symposium on Interactive 3D Graphics*, pp 179-190, (2001).