



Physics-based Surface Modeling using Quasi-Static Liquids

Kurt W. Swanson¹, Kenneth A. Brakke² and David E. Breen³

¹Drexel University, kurt.w.swanson@drexel.edu

²Susquehanna University, brakke@susqu.edu

³Drexel University, david@cs.drexel.edu

ABSTRACT

This paper presents a physics-based modeling approach for the creation of high quality surfaces for application in CAGD. Physics-based modeling is commonly used in animation and scientific modeling, and simulates realistic dynamic motion for computer graphics. We propose a physics-based modeling of liquid surface motion as a means to generate complex geometry. Regions of a model are defined to be in a solid phase and other parts are in a liquid phase. The liquid surface is mobile and moves in response to physics-based forces resulting in a smooth, minimum energy surface. The objects created by this method are referred to as Temporal Computational Objects (TCO's) and provide a capability that is fully integrated with standard CAD, FEA and CAM utilities. We demonstrate the usefulness of the approach with a wide range of examples that produce high quality physics-based surface models that have been analyzed and fabricated.

Keywords: shape design, transformation, morphing, physically based design.

DOI: 10.3722/cadaps.2009.759-768

1. INTRODUCTION

The work described here presents a class of surfaces that behave as quasi-static liquids and that have been systematically applied to form engineered articles. We employ an energy minimization process that is used to deform a liquid membrane-type surface. At the same time the surface is subjected to a wide variety of external forces and constraints. Our results include physical articles that have been modeled, analyzed and fabricated. Parts generated by this method are unique in form and appear organic, as they mimic physics-based processes. A clear benefit in structural efficiency, weight, number of parts is achieved over conventional modeling methods. The initial input datasets - sparse initial boundary conditions (SIBC) - are several orders of magnitude smaller than common CAD datasets. It appears that there is no strong correlation between the initial dataset size and the final object complexity. The complexity of TCO's is produced via computational evolution and not from input dataset size. Complex shapes are derived from SIBCs via a physics-based computational process. The objects created by this method are referred to as Temporal Computational Objects (TCO's) and provide a capability that is fully integrated with standard CAD, FEA and CAM utilities.

Liquid surfaces have been extensively studied in the past, but we have overcome the difficulties in adapting them to CAD and fabrication [20]. We have considered, and abandoned, the common shape modeling approach of elastic membrane perturbation methods. A liquid surface is used instead for geometry generation. Regions of the surface can be fixed (solid) or free (liquid) with the solid phase providing the boundaries and the liquid phase having mobility. This basic method is expanded with

the introduction of scalar functions, pressure and other energies, that provide user control of individual surfaces. We also demonstrate a method of object formation by joining dissimilar surfaces. Multiple examples of production artifacts are then shown including a femur and talus to show a connection between biological forms and TCO forms.

A TCO surface form is inclusive of the common surfaces: planes, cylinders and spheres. For example a TCO surface with coplanar edges may form a square - absent pressure; all the points of a liquid film on a coplanar boundary will be coplanar. This is a valuable result since planar surfaces are a very important surface form. An assembly of such surfaces can be arranged into many useful structures such as an 'I'-beam. It can be noted, that most of the common engineering surfaces, cylindrical columns for supports, or spheres and cylinders for pressure vessels, conform to the notion of constant mean curvature. Surfaces of constant mean curvature have been characterized by Laplace, Plateau [16], Clerk Maxwell [13], Gauss, Delaunay, Lagrange, Wittemore, Riemann, Weierstrass, and Schwartz. Many papers have been written on the subject, some over a century ago, with the clearest examples provided by Clerk Maxwell and Plateau. The primary constant mean curvature surface forms that we generate are the plane, cylinder, sphere, unduloid, catenoid, and nodoid. These can be used together in the same object in general, but are separated by fixed boundaries.

A TCO model is a representation of a liquid - either a liquid-to-gas meniscus, or a soap film. This means that the surface, while having many of the forms of an elastic membrane, will not support shear or bending. Also, while elastic materials have an elastic limit of deformation; a TCO modeled as a liquid membrane has no limit of deformation. Of course, the final object is fabricated as a solid having solid structural properties and is used within its elastic limits. While this may appear as a limitation, the opposite is true. TCO's engender a great simplification of the process of shape modeling and have outstanding properties for engineering surfaces and are almost limitless in form. These properties include reduced weight and part count (reduced assembly cost), reduction in stress concentrations, and multifunctional parts. We explore these benefits in more detail in section 5.

2. BACKGROUND

The TCO approach can be compared favorably to elastic membrane deformation methods that have been explored by Celniker and Gossard [4]. These approaches use an elastic deformation model that includes physical properties of bending and shear. Elastic deformation models (and the proposed TCO models) have advantages over solid modeling in that many shapes can be created from character lines in what they term a 'ShapeWright Paradigm'. Within the paradigm character lines are skinned and then deformed with the objective function of the deformation being an energy reduction of the surface. The typical potential energy applied to the surface is pressure. The original use of energy-based shape modeling is attributed to Schweikert [19], who introduced the method to Computer-Aided-Design with splines in tension. Nielson [14] elaborated on Schweikert's process of minimizing an energy function over the surface of an object. Subsequently, an energy-based process has been used to improve the ability to modify a surface by varying the parameters of the model [6, 7, 8, 9, 15, 18]. The common attribute of these methods is the use of a quasi-linear elastic model to deform the surface and/or the underlying curves. Welch and Witkin have expanded this basic work to facilitate user control of surfaces [21, 22, 23]. More recent developments include natural forces and time varying functions using DNURBS [5, 17], and variational methods for subdivision surfaces [10, 11]. Minimal mean-curvature-variation flow [24] and discrete Willmore flows [1] are some of the latest methods introduced. TCO's do not employ spline based surface generation methods. Instead TCO's use a quasi-static liquid evolution process to produce unique surface forms.

3. SURFACE COMPUTATION

A TCO is constructed as a tessellated simplicial complex, union of triangles, or simply a mesh. No higher order representation is proposed here or required. The principal motion of a TCO surface is by its mean curvature and the mean curvature motion results from surface tension energy. This is a very natural, physical property that affects the motion of fluid membranes. The TCO membrane interacts with its environment in response to pressure, tension and other energies. The TCO is initially defined with sparse geometry. It is then refined by the subdivision of its facets. The facets move in response to surface tension and pressure forces while obeying constraints such as volume and fixed boundaries.

The objects therefore have intrinsic forces and an extrinsic environment that are used to control the form during generation.

3.1 Surface Motion

A TCO surface moves by vertex displacement. Each movement reduces the surface's energy while obeying constraints. No changes in topology or triangulation are made during vertex motion iteration. (Other operations affect these attributes.) The method calculates the force at each vertex and moves the vertex in that direction, using a gradient descent method of minimization. A brief description of the method is provided here. More information can be obtained in [2].

The motion of the surface is modeled on the physics of soap films and liquid surfaces. Here, surface tension tends to minimize surface area. For liquid surfaces equilibrium is reached when pressure on the surface equals surface tension times mean curvature. Surface tension of liquids is measured in energy per unit area, so total energy of the surface is proportional to its total area. Therefore TCO surface shape evolution is defined as a local area minimization process.

Given that a TCO surface is represented by a mesh, 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. In the triangulation formulation, let A_v be the area of the star of facets around vertex v . The force that minimizes area locally is then defined as

$$\vec{F}_v = - \left(\frac{\partial A_v}{\partial x_v}, \frac{\partial A_v}{\partial y_v}, \frac{\partial A_v}{\partial z_v} \right), \quad (3.1)$$

where x_v , y_v , z_v are the x, y and z components of vertex v . Given the formula for the area of a triangle as a function of its vertices, a closed form of the partial derivatives may be derived and evaluated.

Each vertex is moved by a scale factor times the force at the vertex, projected to preserved constrained quantities such as volume. See reference [2] for more information on constraint enforcement. The scale factor is the time step for motion. The scale factor is a global constant, the same for all vertices. The scale factor can be set explicitly, or the system will seek an optimal value. In the latter mode, the system will successively double or halve the scale factor until it brackets a factor that gives the greatest energy decrease per computational iteration. Then quadratic interpolation is used to estimate the optimum scale factor, and that value is used in the final motion.

3.2 Mesh Operation

Mesh operations are required to adjust the spacing and uniformity of the mesh for computational convenience. As a mesh expands or contracts, the mesh features, such as aspect ratio, become skewed and surface motion becomes difficult. Mesh smoothing such as Delaunay triangulation, vertex averaging and global mesh refinement serve to adjust the mesh to improve motion computation. Mesh operations are not employed for smoothing as this is the result of surface motion as described in section 3.1. The complete description of the process can be found in the reference material [2], if more information on basic mesh operations is desired.

3.3 Element Division, Refinement and Mobility

TCO's are formed by a physics-based process and can be quite complex. They must embody a sufficient complexity to enable fluidic motion, e.g. describe a scalar field environment and define the computation steps, to evolve. A concise description of complexity itself is somewhat subjective. Lloyd [12] describes complexity as: algorithmic; computational; organizational; logical depth; etc. (as many as forty seven categories). Kolmogorov [3] describes the complexity of an object as the length of the shortest computer program that can reproduce the object. These concepts are intriguing, but for TCO

complexity we only wish to describe the processing times for a typical TCO and to determine the scaling requirements. How big can the object be? How many elements are needed to differentiate desired features? What are the computational constraints?

The TCO feature density and motion computation are determined by how the object is to be used. A TCO may be used to create objects at a variety of physical scales. An appropriate feature size for biologic tissue is 20 μm , for a NC machined part 3 mm, and a casting 1 cm. The application may require several different levels of refinements such as one for FEA and another for fabrication. The resolution of a TCO may be adjusted to meet the requirements of these different follow-on applications. One measure of TCO complexity; therefore, is simply the number of facets that comprise an object.

A standard definition of a cube, with eight vertices and six facets does not contain enough information needed for evolution – the range of available transforms is quite limited. A TCO must have a large number of elements with uniform properties to enable surface motion. By dividing facets, the number of base elements increases, and the complexity of the TCO increases.




Fig. 1: The refinement of the surface of a unit cube to an 8 X 8 X 8 cube provides mobility and allows the transformation into a sphere as shown on the right.

The cube and sphere illustrated in Fig. 1 have a triangular simplicial complex organization. The initial cube, on the left, has little mobility, and is simply defined by its height, width and length. After refinement, the cube surface is mobile and transforms into a sphere on the right. Therefore, the refinement of the initial form has allowed motion to occur in an energy-based environment of the TCO.

The unit cube provides a measure of the complexity for a wide range of TCO objects and computational requirements. The refinement of a TCO depends on the downstream resolution requirements, e.g. viewing, analysis, manufacture or data storage. The object represented in Fig. 2 is a TCO and was created from inception to fabrication as a mesh (no solid model). The general edge length requirements for downstream use are as follows:

- 10 mm for a FEA;
- 5 mm for rapid prototyping,
- 2 mm for surface lofting.

A complex TCO can be deformed from a cube, but we generally start with geometry established in a CAD tool to simplify object initialization. 1,736 facets were used to start the generation process for the object in Fig. 2. At the end of generation, the object had 240,000 facets. This is the same mesh complexity as a unit cube with refinement of 141 facets along each edge.

3.4 Computational Difficulty

A subdivision of a unit cube from six facets to 240k facets is a sufficient number of elements for the TCO form shown in Fig. 2. The number of elements is one measure of complexity, another is the computation required to move all the facets to a minimum energy state. Mesh refinement, mesh operations and motion operations occur together during the evolution to the TCO. Gross motion is less computationally intensive if done with fewer elements, so most of the large deformation is

accomplished with a coarse mesh. Much computational effort is expended making small local adjustments to smooth the surface.




Fig. 2: This TCO required 240K facets and 115K FEA tetrahedral elements. The refinement requirements are varied: coarse for FEM analysis; medium for STL fabrication; fine for back-lofting.

The TCO of Fig. 2 was measured for area and element count before and after generation. The area reduction is a result of the energy minimization function of a liquid surface. Each computational operation was counted - Floating Point Operations (FLOP) and Millions of Instructions (MI) - were measured for the generation of each surface segment and then tabulated. The two charts of Fig. 3 show the pre- and post-generation element counts for a single surface. The right graph in Fig. 4 shows the relative sizes of CATIA, IGES and TCO datasets that represent the TCO shown in Fig. 2. The left graph shows the computational requirements for this 57 lb object defined by some 240,000 facets. The charts show a large increase in the number of elements. This corresponds to a requirement for mobility. At the same time the surface area is reduced. The surface area reduction is directly proportional to a reduction in surface tension energy. The size of the tetrahedral model is included although this was generated in a downstream FEA utility.

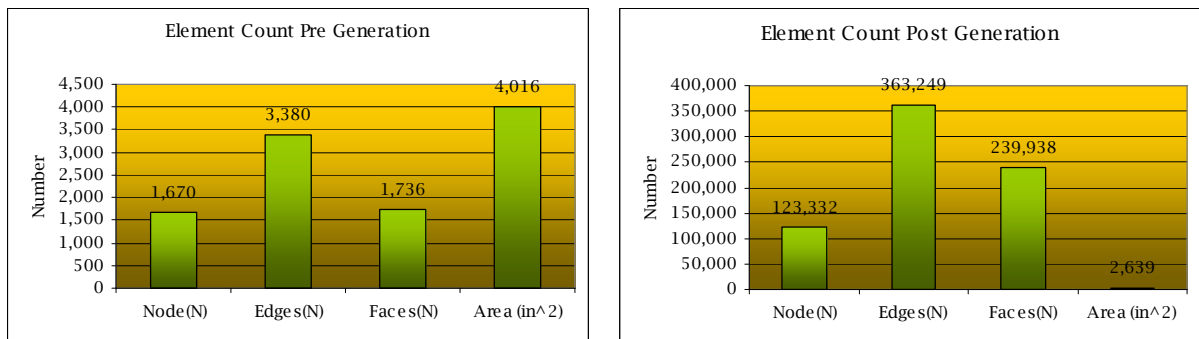


Fig. 3: These two charts show the before and after condition of refinement and evolution for the object of Fig. 2. The area change is in response to the surface tension energy minimization process of a TCO.

The object of Fig. 2 is comprised of multiple surfaces of varying complexity. Some are planar and some have compound curvature components, the latter requiring more computation time. Using the scale listed in Fig. 4, the computational time for a complex surface would be ten seconds using an Intel Xeon processor. The observed times are generally 20 - 30 seconds for generation. The entire surface is a collection of only six discrete data files and the combined execution time for the complete surface is approximately sixty seconds.




Fig. 4: These two charts show the requirements for the object of study Fig 2. The 57 lb object required a total of 2.8×10^{11} FLOP + MIP counts for surface generation. The chart on the left shows the processor requirements to refine and move a surface to an equilibrium state. The chart on the right shows the relative dataset sizes for CATIA, IGES and TCO datasets.

The original dataset requirements of a TCO model are extremely small, several orders of magnitude smaller, in comparison to a CATIA surface model or an IGES surface model, as shown in Fig. 4. This TCO dataset is only 42.5 kilobytes in size. The differences are magnified by the fact that the CATIA and IGES models are essentially static representations. The TCO model has a mobile surface and represents innumerable possible states. The TCO model has a multi-scale capability as the degree of refinement can be increased or decreased depending on usage. The method outlined here can determine the computational requirements for very large complex structures, which in general can be computed very rapidly.

4. STRUCTURAL APPLICATION AND SUITABILITY

A characteristic property of a liquid meniscus is uniform surface tension. For a sphere subject to pressure the strain is everywhere uniform. Another way of looking at these surfaces is that the form is independent of the magnitude of the applied loads. So while a sphere may enlarge slightly due to increased pressure - it is still a sphere. If the material were solid the surface would remain a sphere within the elastic limits of the material. As the pressure increases the internal stress and strain increases uniformly in relation to the Laplace-Young equation,

$$P = T(1/R_1 + 1/R_2). \tag{8.1}$$

This equation is central to describing the relationship between surface curvature, pressure P and tension T, where R_1 and R_2 are the principal radii of curvature of the surface. (e.g. for a sphere $R_1 = R_2$ and $P = T/2R$.) The resulting spherical surface form is considered to be optimal because the strain energy is uniform throughout the vessel. The same can be said for a cylinder. These forms are convenient for use - they are mathematically well known and easy to create and assemble into objects. For many applications these are suitable and we need go no further. If we were to limit ourselves to only pressure vessels we might be inclined to believe that there are only two suitable forms - cylinders and spheres -- but this is indeed not the case. There are many forms that are available and they can be easily characterized and will serve a similar role.

The object shown in Fig. 5 was generated from MRI data of a talus (foot) bone. The outer bounding curves are formed at a discontinuity between the three ends of the bone (not shown) and the free surface or side of the bone shown in blue. The inner surface was generated by scaling the outer curves and provides a foliate surface to mimic the internal boundary of the cortical bone. The TCO surface is generated as a liquid film to span the three boundary curves. By definition the liquid surface has uniform tension throughout, as does the inner foliate surface. If an additional curve is added between the curve pairs, shown on the right in Fig. 5, then a third surface will form as a foliate. That is, the

surface is nested completely within the outer and inner surfaces. An innumerable series of surfaces can be generated in the same manner. All will fall within the bounding surfaces and have uniform tension. If the material were changed to a solid, a compressive load can be applied uniformly at the curve boundaries and the response will be a uniform compressive load in the solid.




Fig. 5: Surfaces are generated as foliates nesting on fixed boundaries. The two left panels are generated from a MRI scan of a talus (foot bone) having three boundaries.

A qualitative evaluation of the TCO objects created to date demonstrates that dramatic improvements in cost and weight can be and have been obtained. Some of these have been order of magnitude improvements. Shown in Fig. 6, are a solid model on the left and a TCO on the right. The TCO approach has an order of magnitude reduction in parts and weight over the predecessor design. Even though we observed this for a few applications, it is unreasonable to expect similar improvements for all applications. For highly optimized engineering designs, the improvements may provide only single digit percentages of improvement. Since TCO objects tend to look radically different than conventional part design, appearance and familiarity may be a primary reason to discount a comparative advantage in function. TCO objects that have been fabricated perform extremely well in comparison to solid modeling structures that they replace. In some cases the TCO presented the only viable response to design trade studies for strength, stiffness and cost.




Fig. 6: Shows a typical solid model approach on the left and a TCO based design on the right.

5. RESULTS

The use of TCO-generated surfaces has resulted in many studies and some fabricated articles. The first article was a series of struts as depicted in Fig. 7. A common requirement for the three struts was a square base and a cylinder/spherical end. The orientations of the ends are not necessarily aligned with the base. For stiffness and strength a torque tube provides the proper approach; however, integrating such a tube to the end conditions is very challenging in a conventional approach. A TCO

surface spanned the highly irregular boundaries very easily. In this instance the TCO surface was 're-used' to design multiple similar parts - rather than to recreate a new design for each unique occurrence. The lower end termination was moved to the desired location and the surface was recomputed.




Fig. 7: This series depicts reuse of the surface for three variations. All the results are superimposed in the last panel for comparison.

The proximal and distal ends of a femur are shown in Fig. 8. The information required to form this TCO was taken from an MRI dataset. Discontinuities are evident at the head and neck of the femur. A boundary was added here and at all other discontinuities in the form. Once the boundaries were defined, TCO forms were created to produce surfaces similar to the actual bone: the femoral neck is a catenoid surface; the femoral head was a sphere; the diaphysis is quasi-cylindrical as in Fig. 7 and so forth. It was only necessary to find the discontinuities on the bone in order to accurately recreate the surfaces of the femur. The TCO surfaces correspond to the MRI dataset very closely.

Shown in the upper left corner of Fig. 8 is an aircraft fuselage study. The diameter of the fuselage is twenty feet and the length is twenty feet as represented in a CAD system. An isogrid pattern of stiffening is created by a triangular lattice of nodes shown in green; a cylindrical outer shell is shown in blue; interior window panels are shown in purple. The surface smoothness is limited only by computer memory storage capacity. This study demonstrates the capability of creating large intricate structural forms that do not include planar elements but are comprised mainly of curved forms and to accomplish them quickly by computation. If the outer envelope were not cylindrical, the computation and final result would not differ.

The multi-color part having two arms presents the outcome of the analysis of the fabricated part shown in the center bottom of Fig. 8. This part had a stiffness requirement with an eccentric load. The load on one arm was greater than the other. Since the deflection had to be the same for both arms the approach taken was to move the center support in the direction of the arm having the highest load. The center support was moved and the surfaces were easily recomputed and re-analyzed, until the deflection for both arms was limited to within a narrow acceptable range.

The series of green panels at the middle of Fig. 8 is a highly complex casting with multiple functions. Two 'H' beams are blended into the base shell by using a unique liquid capability of three intersecting films. In this way the vertical 'H' beam is structurally integrated into the shell. This part uses pressure and volume constraints, singularity functions [20] and modification of surface tension in some regions to achieve the final result.

6. CONCLUSION

TCO designed objects are functionally excellent, with uniform strain distribution and outstanding stiffness properties. Order of magnitude reduction in part counts were observed in some applications. The computation of the articles is rapid. The process is scalable to large complex objects and small intricate objects. Re-use of design for a family of parts has been demonstrated. The process has been

fully integrated with CAD, FEA and CAM capabilities. TCO's enable the creation of complex and diverse forms uncommon to solid modeling. A TCO physics-based surface modeling based on a quasi-static liquid is an improvement over elastic membrane methods.




Fig. 8: These are a small sample of a larger assortment of study applications. The results of TCO models demonstrate a diversity of forms from large complex structures to small biologic forms.

5. REFERENCES

- [1] Bobenko, A.; Schröder, P: Discrete Willmore Flow, Proc. Eurographics Symposium on Geometry Processing, 2005, 101-110.
- [2] Brakke, K.: The Surface Evolver, Experimental Mathematics, 1(2), 1992, 141-165.
- [3] Burgin, M.: Generalized Kolmogorov complexity and duality in theory of computations, Notices of the Russian Academy of Sciences, 25(3), 1982, 19-23.
- [4] Celniker, G.; Gossard, D.: Deformable curve and surface finite-elements for free form shape design, (Proc. ACM SIGGRAPH) Computer Graphics, 25(4), 1991, 257-266.

- [5] Du, H.; Qin, H.: Free-Form Geometric Modeling by Integrating Parametric and Implicit PDEs, *IEEE Transactions on Visualization and Computer Graphics*, 13(3), 2007, 549 - 561.
- [6] Farin, G.; Sapidis, N.: Shape Representation of Sculpted Objects: the Fairness of Curves and Surfaces, *Proceeding of Sea Grant Conference*, MIT, October, 1988.
- [7] Hagen, H.; Schulze, G.: Automatic Smoothing with Geometric Surface Patches, *Computer Aided Geometric Design*, 4, 1987, 231-236.
- [8] Kallay, M.; Ravani, B.: Optimal twist vectors as a tool for interpolation a Network of curves with a minimal energy surface, *Computer-Aided Geometric Design*, 7, 1990, 465-473.
- [9] Kjellander, J. A.: Smoothing of bi-cubic parametric surfaces, *Computer-Aided Design*, 15, 1983, 288-293.
- [10] Kobbelt, L.; Schröder, P.: A multi-resolution framework for variational subdivision, *ACM Transactions on Graphics*, 17(4), 1998, 209-237.
- [11] Kobbelt, L.: Discrete fairing and variational subdivision for freeform surface design, *The Visual Computer*, 16(3-4), 2000, 142-158.
- [12] Lloyd, S.: *Programming the Universe*, Knopf, Mar 14, 2006.
- [13] Maxwell, C.: On Theory of Rolling Curves, *Trans. R.S.E.* xvi, 1894, 519-540.
- [14] Nielson, G. M.: Some piecewise polynomial alternatives to splines in tension, in Bamhill, RE and Riesenfeld, R. F. (eds.), *Computer Aided Geometric Design*, Academic Press, 1974.
- [15] Nowack, H.; Reese, D.: Design and fairing of ship surfaces, in Barnhill R. E. and Boehm, W. (eds.), *Surfaces in CAGD*, N-Holland Amsterdam, 1983, 121-134.
- [16] Plateau, J.: *Experimental and Theoretical Statics of Liquids subject to Molecular Forces Only*, 1873, [English Trans, <http://www.susqu.edu/facstaff/b/brakke/PlateauBook/PlateauBook.html>].
- [17] Qin, H.; Terzopoulos, D.: D-NURBS: physics-based framework for geometric design, *IEEE Transactions on Visualization and Computer Graphics*, 2(1), 1996, 85 - 96.
- [18] Sapidis, N.; Farin, G.: Automatic fairing algorithm for B-spline curves, *Computer-Aided Design*, 22(2), 1990, 121-129.
- [19] Schweikert, D. G.: An interpolation curve using a spline in tension, *Journal of Math and Phys.*, 45, 1966, 312-317.
- [20] Swanson, K.; Brakke, K.; Breen D.: *Temporal Computational Objects: A Process for Dynamic Surface Generation*, *Advances in Visual Computing*, Springer, 2008, 959-969.
- [21] Welch, W.; Witkin, A.: Free-form shape design using triangulated surfaces, *Proc. SIGGRAPH*, 1994, 247-256.
- [22] Welch, W.; Witkin, A.: Variational Surface Modeling, (*Proc. SIGGRAPH*) *Computer Graphics*, 26(2), 1992, 157-166.
- [23] Witkin, A. P.; Heckbert, P. S.: Using particles to sample and control implicit surfaces, *Proc. SIGGRAPH*, 1994, 269-277.
- [24] Xu, G.; Zhang, Q.: G2 surface modeling using minimal mean-curvature variation flow, *Computer-Aided Design*, 39, 2007, 342-351.