Chapter 1

Introduction

1.1 Motivation

There is an endless variety of shapes in the world; from the simple Euclidean shapes such as rectangles, spheres, and cylinders, through the smooth swept curves of a boat hull, to the fractal nature of clouds, mountains, and trees. If the computer is to assist us in our endeavors to represent the world, it is important that we can represent and manipulate the wide range of shape.

The requirements for a suitable shape representation and for tools to manipulate such representations vary according to application. Industrial designers require exact analytical control of shape. In automotive, aerospace, and marine vehicle design, sculptors require control of the "curvature" and "fairness" of a surface. Character animators in the entertainment industry require "expressive" control of their creations. The special effects industry requires emulations of natural effects to create rich and complex visual elements. Robotics require spatial shape representations that can be derived from the environment and then used to navigate through the environment. Medical scientists require surfaces that segment data collected from non-invasive sensing techniques, such as CAT and MRI scans. The physical sciences require surfaces that provide an "optimal fit" to empirical data. Traditional computer based tools for representing and manipulating shapes both limit and enhance the ability to create specific classes of shapes.

In computer graphics, a common trait among shape representations is that the global structure is explicitly defined. In the case of the classical primitives (spheres, cubes, superquadrics) the structure is defined by the underlying parameterized equations. In the case of splines and polygons, structure is defined by manually connecting surface patches. And in the case of skeletal implicit surfaces, the structure is defined by the underlying skeleton structures. While these shaping representations are excellent at creating a wide variety of shapes, large changes in the surface structure and changes that do not preserve genus require significant user interaction.

In this dissertation, we investigate the problem of developing a shape representation in which the structure of a shape is inferred, rather than explicitly specified. By removing the restrictions on structure, we open the door to new ways of both creating and manipulating shapes. Our investigation takes us away from a purely geometric approach to modeling to one that combines geometry and physics.

Our goal is to develop a new shape representation model, based not on an underlying parametric representation, but rather on systems of self-organizing elements from which shape and structure emerge. We illustrate our approach by applying it to three common tasks: free-form shape modeling, computer assisted animation, and surface reconstruction. To this end we have three auxiliary goals:

- To create and manipulate complex shapes, not by the traditional method of manually defining the arrangement of surface patches and continuity conditions, but rather through an intuitive and flexible sculpting metaphor where "physical" tools are used to manipulate synthetic materials.
- To show that these shapes can be animated over time, not by key framing, but by mimicking the physical behavior of various materials.
- To reconstruct and extract surfaces of arbitrary topology from three dimensional data.

1.2 Application Areas

Our three auxiliary goals apply to the following areas.

1.2.1 Shape Modeling

There are two basic approaches to the modeling of shapes. The first is to create a shape to meet a known set of specifications. Another approach, more exploratory in nature, is to evolve a shape from simple to complex, until it is aesthetically pleasing. In this case the focus is not on measurable analytic qualities but on subjective aesthetic qualities. Spline based interpolation, extrusion, and constructive solid geometry, enforce analytical constraints on a shape, aiding in the objective design approach. The success of analytically based tools are in part because the constraints of each technique assist in defining a particular class of surfaces. However, if used to create shapes outside of the intended domain, the same constraints may hinder the creation process by imposing unnecessary limitations. For the exploratory approach, the design tools must be flexible in the sense that the constraints on the *design process* should be minimized. Since numerous computer based tools exist to assist in the creation of shapes with strict geometrical properties, good for precise analytical control of a surface, let us instead consider what properties are useful to designers who follow a more exploratory approach.

The occupation of the sculptor, whether working in wood, clay, plaster, wax, glass, or stone, has existed for centuries. The medium that these artists have chosen to realize their creations have been chosen for their long term durability, and for the material characteristics which allow the artist to shape and reshape the material to its final form. Perhaps some of the needs of the exploratory design process can be observed from considering the materials and tools sculptors have traditionally chosen.

Sculptors often design their shapes using a malleable material, such as wax, which can be easily sculpted with a small collection of tools. The shape can be built up or carved away, thereby providing flexibility to the creation process. After the design is complete, the sculpture is realized in the more expensive and less flexible materials of cast metal and stone. Sculptors use glass to create smoothly curving surfaces. By heating the glass, the artist changes the physical properties so that it can be shaped without breaking. A common technique is to stretch the glass, either by directly pulling the glass or using the force of breath to "blow" the glass into a smooth flowing shape. In industrial applications, smoothly curved lines and surfaces are often designed by bending strips of metal, exploiting the inherent nature of the material. The ceramic sculptor has perhaps the widest range of sculpting methods, which is due to the nature of clay. Beginning with a lump of clay they can stretch the clay, deform it with pressure, break or cut pieces of clay off of the main body, join pieces of clay together adding handles and loops, and build up features by incrementally adding clay to an existing form. If we can incorporate some of these qualities into our shape design algorithms we will provide the user with powerful metaphors to use in sculpting shapes on the computer.

This premise is supported by the fact that several current computer techniques mimic techniques used by sculptors of physical materials. In computer graphics, spline curves are used to enforce the continuity conditions of metal splines traditionally used in the ship building industry. Surfaces of revolution are used to create curved surfaces symmetric about an axis, like shapes thrown by a potter at the wheel or created by a carpenter at a lathe. The ability to add and subtract geometric objects from one another are implemented by constructive solid geometry (CSG) and implicit surface formulations. Still, many of the techniques used by the sculptors of real world materials are unavailable for the computer based sculptor; and when they do exist, they usually do not occur within the same shape representation.

For a metaphor to be effective we do not need a formal mapping from the metaphor source to the target (Madsen, 1994). For example, the piecewise spline model is based on the behavior of metal splines, yet the computer spline does not need to be as cumbersome to manipulate and store as long metal strips. To effectively use the metaphor of sculpting with clay we do not necessarily want to incorporate the fact that clay must be kept moist while it is being formed to prevent the clay from cracking and crumbling. Likewise in the molten glass metaphor we do not have to simulate the fragile nature of cooled glass. In creating a metaphor, undesirable qualities should be ignored, while desirable qualities enhanced.

For shape design and rapid prototyping applications, we propose an interactive system, based on a sculpting metaphor, which alleviates the designer's need to think about the underlying representation or be limited by its choice (Sachs, Roberts and Stoops, 1991), similar to how a child can sculpt clay while ignorant of the clay's chemical structure. Borrowing from traditional sculpting materials, we propose creating synthetic shaping materials that can be manipulated in a variety of ways. The materials should provide the user the ability to cut, merge, and join shapes; to deform shapes; to carve away and add new material; and to shape and reshape the material with a basic set of tools.

1.2.2 Computer Assisted Animation

In animation, a story is brought to life as a stream of images and sound. Creating the images requires capturing all of the visual changes that occur in the story, including the shape, color, transparency, structure, texture, time-vary position (motion-dynamics), and lighting of objects (Foley et al., 1990). Not surprisingly, animation is expensive and time consuming. The quantity of visual information that must be created, manipulated, and stored is enormous. While short animations, such as television cartoons and commercials, are produced in only a few months, longer feature length animations require three to five years of effort. It can be argued whether or not the computer has sped up the process of creating animation, but it cannot be argued that it has had no impact. The computer has come to assist the artistic process, by both reducing the time spent per image and expanding the range of images routinely created.

Three dimensional computer animation is the art of modeling shape changing over time. To express intent and reaction, the animator needs ultimate control over the shape. Because of this, specifying the object shape at key frames has been the dominant method used. Interpolating between key frames parallels the job of the inbetweening animator in traditional animation. Inverse kinematics is coming to play a bigger role in animation, because instead of simply interpolating, it computes the movement of jointed figures. Recently, physics based simulations have been making their way into the animation and special effects businesses. In these simulations, control is traded for dynamic complexity and visual richness, making them attractive alternatives to more time consuming techniques.

For simulating secondary action, the effects animator needs not only a system which aids him in representing shapes, but also the flexibility to change the shapes and their behaviors. We propose a physically inspired approach that allows one to sculpt a variety of shapes, to animate the shapes over time through a physical simulation, and to varying the physical properties associated with the shape.

1.2.3 Surface Reconstruction

Large quantities of data are routinely generated by computer simulations and collected by non invasive sensing devices (Ney, Fishman and Magid, 1990; Stytz, Frieder and Frieder, 1991; Higgins, Chung and Ritman, 1990; Wolfe Jr. and Liu, 1988; Baker, 1988). The sheer quantity of the data collected presents a fundamental problem when it comes to interpretation. The human visual system is considered a high bandwidth channel for receiving information (Haber and McNabb, 1990; Hibbard and Santek, 1989), making visualization a powerful data presentation tool. Often specific sections of the data are of interest and displaying all of the data, as in many volume rendering techniques, distracts the user from the meaningful features. Such features can be extracted by fitting surfaces to them.

Other interesting uses of surface fitting are to recreate surfaces scanned by laser range finders or from video cameras (Szeliski, 1991; Szeliski, 1993). In the future, geometric modeling may be simplified by entering geometric descriptions that consist of "showing" the object to the computer. This would allow users to send "3D faxes" by scanning an object in one place and interactively viewing a full 3D model at a remote location (Carlbom et al., 1992).

Unknown object topology poses difficult challenges to surface reconstruction that have until recently been largely ignored in the vision literature. Unfortunately, vision systems that must derive quantitative models of complex real-world objects from multiple views cannot avoid the issue of unpredictable topological structure. This is also an important concern in the related fields of biomedical and geological imaging where there is a need to analyze three dimensional arrays of volumetric density or reflectivity data. These arrays are like blocks of marble with meaningful embedded structures. For further analysis, the often complex and topologically unpredictable surfaces of these structures must be extracted and represented as compact geometric models (McInerney and Terzopoulos, 1997). To solve these problems we propose to use a system of self organizing primitive shape elements which react to both the content of the data and neighboring primitive elements. In such a system the topology is not defined by a surface parametrization, but by the relative spatial position of the primitive shape elements.

1.3 Challenges

In creating our new shape representation technique, we face several challenges:

Changes in Structure

Traditional shape modeling techniques are primarily based on the user manually defining the structure of the object. Thus, gross deformations and changes in genus require significant user interaction. The challenge is to create a model where large changes in shape can be made with minimal user interaction. For free-form modeling, we want to be able to easily sculpt a shape without having to consider the underlying structure. For computer assisted animation, the shape structure should change in a natural manner as dictated by the simulation. For surface reconstruction, the technique should be able to infer unknown surface structure from 3D data sets.

Synthetic Materials

Our modeling paradigm is based on the metaphor of synthetic materials endowed with physical properties. For free-form modeling, the synthetic materials should have properties similar to materials used in designing sculptures, such as wax, clay, and glass. For computer assisted animation, the synthetic materials should be able to exhibit a wide range of dynamic and elastic properties and interact with their environment. For surface reconstruction, the synthetic materials should be influenced by the data sets and be able to simulate the optimal fitting properties such as those of a membrane or thin plate.

Customization

To obtain a general shape representation, we need to customize the behavior of the synthetic materials. Some desirable properties include: conservation of volume, conservation of surface area, mass, elasticity, stiffness, and fluidity. For different applications we should be able to select and enforce desirable properties.

Efficiency

The computational cost of the model should be tractable. As the complexity of the shape grows, the time and space costs should not grow exponentially.

1.4 Methodology

We address the challenges with the following approaches.

Criteria

The measure of quality we use for our shape representation is not based on the realism and accuracy of the synthetic material to real world materials, but whether the synthetic material exhibits properties and behavior that assist the application. For example, in a free-form modeling application based on the metaphor of sculpting a malleable material with tools, the synthetic material does not need to exhibit all the properties of say, clay or wax, only a select few which support the metaphor.

Structure

We address the issue of defining structure using discrete shape elements that can be arranged and rearranged into new shapes. The arrangement and spatial proximity between elements defines the surface structure and hence the topology and genus. Useful arrangements of the shape elements is achieved by using a self organizing system. The ability of the shaping elements to self organize, allows us to create shapes that are not based on parametric patches or a prescribed structure. The global structure of the shape emerges from the local physics based interactions of elements.

Particle Systems

We base our modeling representation on particle systems and Newtonian dynamics. Each particle is a discrete shape element; either a volume element or a surface element. Interactions between particles are based on energy functions and local geometric properties allowing for the creation of synthetic materials, that interact with their environment. In this dissertation we solve for a second order dynamical system for applications requiring realism such as animation. For other problems, such as solving an optimization criteria, a first order system is adequate and is a simplification of the more complex second order dynamical particle system we present.

Surface Descriptions

Surface descriptions are based on the arrangement of the shaping elements. For shapes described by volume elements, we use an implicit surface formulation. For shapes described by surface elements, we generate a polygonization based on the spatial proximity between elements.

Efficiency

We maintain efficient computation by insuring that the computational costs do not grow excessively for large systems. We convert a general $O(N^2)$ particle system computation into $O(N \log N)$ expected computation for surfaces, and $O(N)$ expected computation for volumes, by reformulating global computations into local computations. Key to this approach is being able to rewrite energy functions in local terms of each shape element, and the use of hierarchical tree structures to partition space into local regions.

Customization

Particle systems are customized to a given application by designing new inter-particle potentials, application specific forces, and specialized particle creation heuristics. For example, for modeling volumes we use a potential energy function which encourages the equal spacing and preservation of volume, and velocity based force functions to model inelastic deformations. For modeling surfaces we introduce curvature based potential energies to encourage particles to arrange into surfaces. By varying the weightings of these potentials we can encourage more or less smoothness of the surface. Particle creation heuristics are designed to allow surface modeling operations which "stretch" the surface, and surface reconstruction applications to interpolate between sparse 3D point samples. Forces are designed to provide interaction with objects external to the particle system, such as colliding with objects in an animation, interacting with shaping tools in free-form modeling, and attracting particles to "edges" in surface reconstruction applications.

1.5 Contributions

We close this chapter by summarizing our contributions.

1.5.1 Dynamically Coupled Particle System

This dissertation presents a dynamically coupled particle system for shape representation and manipulation. This research has been published in the computer graphics literature (Tonnesen, 1991; Szeliski and Tonnesen, 1992) and in the computer vision literature (Szeliski, Tonnesen and Terzopoulos, 1993b; Szeliski, Tonnesen and Terzopoulos, 1993a).

A novel feature of this approach to modeling shape is the use of a particle dynamics simulation in which particles interact through potential energy functions to create volumes and surfaces. The use of a spatially symmetric potential energy encourages particles to arrange into tightly packed orderings, good for modeling volumes. This basic model is then extended to develop oriented particles, a particle system which encourages particles to arrange into surfaces rather than volumes.

In the oriented particle model, each particle represents an oriented trihedral coordinate frame. Based on these frames, we design new interaction potentials which favor locally planar or locally spherical arrangements of particles. Thus, the oriented particles support smoothness constraints similar to those inherent in the deformation energies of popular, elastic surface models. Unlike spline patches or parameterized models, the model *does not* attempt to enforce *analytical* continuity conditions, such as C^1 tangent or C^2 curvature continuity, over the surface. Instead of a parametric approach to shape modeling, our research uses collections of primitive self-organizing elements from which geometric structure evolves as the system moves to local minimal energy configurations. Continuous surface descriptions are created by extending the concept of two dimensional Delaunay triangulation to generate surface triangulations of arbitrary topology, embedded in three dimensional space.

1.5.2 Free-form Modeling

Our approach to free-form modeling provides the following abilities:

- to cut, merge, and join shapes,
- to remove and add material,
- to shape material using a small collection of tools,
- to add handles and loops without having to respecify the basic surface structure,
- to add curvature discontinuities by locally ignoring smoothness constraints,
- and to "heat" and "cool" the model to vary the material characteristics.

The modeling paradigm we present is based on a sculpting metaphor in which the user begins with a volume or sheet of material and arrives at a final shape through incremental sculpting operations In Figure 1.1 we show the result of deforming a sphere into a torus using two shaping tools. Such properties allow a designer to create a variety of geometric shapes without the need to be concerned with the underlying structure or parametrization of the shape.

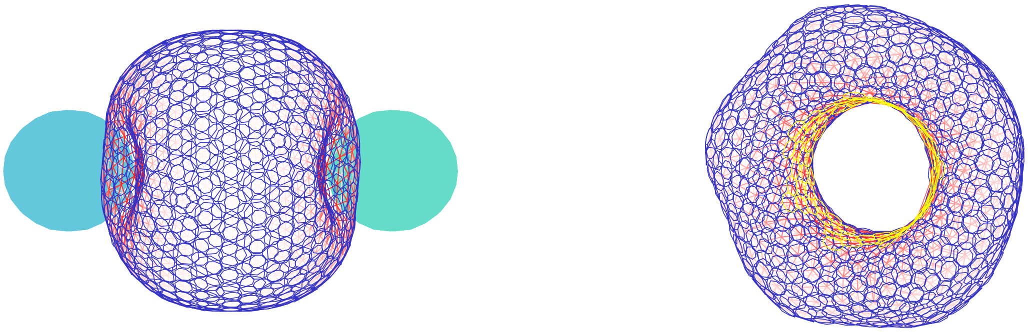

Figure 1.1: Free-form surface modeling

Deformation from sphere to torus using spherical shaping tools.

1.5.3 Computer Assisted Animation

When applied to three-dimensional animation, these same properties allow an animator to interactively sculpt the 3D models which can then be animated according to the associated physical characteristics. In addition, the physical characteristics may be modified during a scene to suit changes dictated by the story. For example, the technique can simulate from solid to fluid behavior, infinitely stretchable material, and materials that rip, tear, and that can be seamlessly merged together. Once created, the model may be influenced by a larger environment modeled for the animation. The manipulations can be global forces such as gravity, local forces derived from collisions with other objects, forces generated from hand gestures, scripted movement, or vector force fields defined in space. In Figure 1.2 we illustrate a selection of different physical properties, such as cloth draping (Figure 1.2a), plastic surface deformation (Figure 1.2b), and tearing (Figure 1.2c). The differences in physics were easily created by varying the weights of the potential energy functions.

1.5.4 Surface Reconstruction

When applied to surface fitting, the surface model and reconstruction methods developed retain the topological flexibility of the local patch methods, while constructing globally coherent surface models that can evolve consistently with time-varying data and forces. When reconstructing an object of arbitrary topology, the particles can be made to “flow” over the data, extracting and conforming to meaningful surfaces. The process is roughly analogous to pouring a viscous liquid over an object and wetting its surface. In Figure 1.3 we show the reconstruction of the surface from a CT scan of a plastic “phantom” vertebra model (decimated to 120 × 128 × 52 resolution). This smooth, triangulated model contains 6,650 particles and 13,829 triangles, and was created by seeding a single particle and extending the surface along high 3D edge

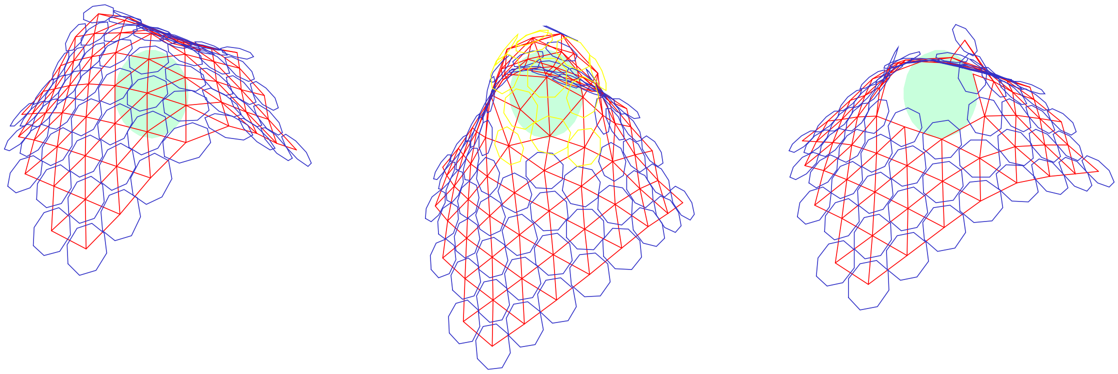

(a) (b) (c)

Figure 1.2: Computer assisted animation

Variable surface behavior: (a) cloth draping, (b) plastic deformation, and (c) tearing.

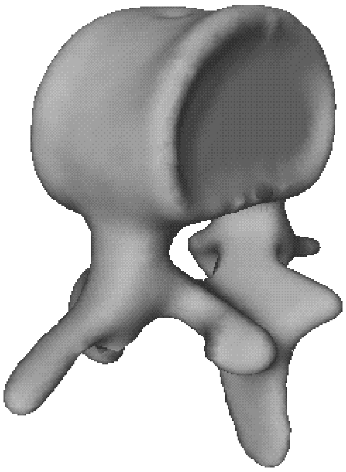

Figure 1.3: Surface reconstruction

Shaded 3-D Reconstruction of a vertebra from 120 × 128 × 52 CT volume data.

values until a closed surface was obtained.

The approach has two components. The first is the dynamic particle system which discovers topological and geometric surface structure implicit in the data. The second component is an efficient triangulation scheme which connects the particles into a continuous global surface model that is consistent with the particle structure. The evolving global model supports the automatic extension of existing surfaces with few restrictions on connectivity, the joining of surfaces to form larger continuous surfaces, and the splitting of surfaces along arbitrary discontinuities as they are detected.

1.5.5 Results

The main results of this thesis include:

- A new surface model representation that we call "oriented particle systems". This model combines particle systems, differential geometry, and physics into a single model. The model uses potential energy functions for minimizing the normal curvature, Gaussian curvature, and the twist of a surface, for determining the principal directions of surface curvature and lines of curvature, and for emulating the stretching and bending energies of physical objects.

- A free-form modeling paradigm in which surfaces are sculpted using tools in analogy to sculpting with physical tools. Basic modeling operations include merging, cutting, extending surfaces, and specifying curvature discontinuities.

- A heat transfer model for dynamically coupled particles derived from the macroscopic heat equation.

- A method of varying the stiffness of Lennard-Jones coupled particle models based on thermal energy while maintaining a conservation of total system energy. Combined with the heat transfer model, this extends the range of behavior exhibited to be from solid to fluid. It also allows for local variations in malleability of a model thus providing sculptors more control over modeling operations.

- A triangulation algorithm for creating continuous surface descriptions of arbitrary topology from 3D point sets.

- An analysis of the computational and memory requirements of dynamically coupled particle systems. This includes an analysis of neighboring finding techniques and an analysis of the stability and accuracy of the explicit Leapfrog integration scheme for the Lennard-Jones potential.

- Surface reconstruction algorithms for both interpolating 3D point sets and optimal surface fitting of 3D data. The algorithms were designed for both reconstructing open and closed surface from sparse 3D point sets, polygonizing iso-surface functions, and segmenting 3D volumetric data, e.g. CT scan data.

- Visuals simulations of volumes and surfaces interacting with external objects and forces, free-form modeling, and surface reconstruction.

1.6 Thesis Overview

In Chapter 2 we review the background material and related work. In Chapter 3 we develop a spatially coupled particle system model of deformable volumes. In Chapter 4 we extend the model of Chapter 3 to "oriented particles" which prefer to arrange into surfaces rather than volumes. In Chapter 5 we provide mappings from the discrete nature of a particle system to continuous surface descriptions. In Chapter 6 we further extend the behavior of the surface and volume models to be functions of thermal energy. In Chapter 7 we discuss implementation and efficiency issues. In Chapter 8 we apply our model to the applications of physics based animation, interactive free-form modeling of surfaces, and surface reconstruction of 3D data. In Chapter 9 we conclude with a summary and final remarks.

In Appendix A we review the differential geometry of surfaces. In Appendix B we review Newtonian dynamics. In Appendix C we derive equations for computing the gradient of energy functions based on the Euclidean norm, products of the norm with normal vectors, and weighting functions. In Appendix D we derive forces and torques for the co-planarity, co-normality, co-circularity, and Lennard-Jones potential functions. In Appendix E we provide a finite element analysis of the surface energies.