The ThingWorld Modeling System: Virtual Sculpting By Modal Forces *

A. Pentland, I. Essa, M. Friedmann, B. Horowitz, S. Sclaroff

Vision and Modeling Group, E15-387, The Media Lab, M.I.T.
20 Ames St., Cambridge MA 02138

Abstract

We describe a real-time solid modeling system that is based on the physical analogy of forming clay by applying forces. The system is implemented by simulating real materials as they react to user-supplied forces. Unlike other physically-based modeling approaches, the Thingworld system allows the user to restrict forming action to simple global deformations during the initial "roughing in" phase of modeling, and then later concern themselves with detailing. The Thingworld system also allows users to automatically model existing objects by using measurements taken from the object's surface. These measurements are used to generate artificial forces that mold the computer model much as a human would mold a clay model. Timed examples for constructing solid models are shown.

1 INTRODUCTION

A Computer Aided Design system must have much the same relationship to a design situation that people do; that is, it must be able to obtain, describe, and manipulate information about shape in a manner analgous to the way people do. Thus, a *shared representation* is crucial to effective, natural man-machine interaction; unfortunately, little is known about people's mental representation of shape. We do know, however, that people are very familiar with physical processes, and so computer analogies to physical processes have long been considered promising metaphors for modeling systems [1,2,4,5]. Perhaps the most attractive of these physical analogies is *sculpting in clay*, as clay is widely considered to be the fastest traditional modeling medium. Further, it is widely accepted that modeling is normally an iterative process, starting with *roughing in* and proceeding through *detailing*. As a consequence, it seems that a good computer modeler must allow us to adjust the level of detail during the modeling process.

These observations form the motivation for the Thingworld modeling system. The system presents the user with
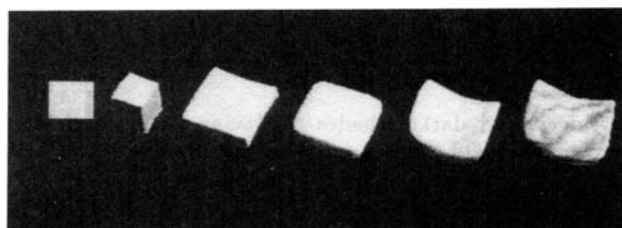
Figure 1: Sequence of operations used to form a single part.

virtual lumps of clay, and allows the user to apply virtual forces to this clay in order to form it. The resulting parts may then be fused together to form larger objects, or used as "virtual knives" to cut holes in existing parts. The system allows a variable level of detail, so that over time modeling may proceed from rough to fine forming operations.

Perhaps the most novel aspects of the Thingworld system are its use of a physical, force-based interface in a real-time modeling framework, its ability to vary the level of detail being modeled, and the ability to automatically generate models from measurements of real objects. These abilities stem primarily from the use of a modal representation of object dynamics and deformation [2] rather than the more common finite element representation, as the modal method provides both a hundred-fold speed-up in the calculation of the deforming object's dynamics, and a way of progressively adjusting the level of modeling detail.

2 MODELING USING MODES

By describing object dynamics using modes, the system can very efficiently simulate the deformation an object experiences in response to any force. Perhaps more importantly, the use of modes provides an easy way to control the level of detail. By incorporating only low-frequency modes, the object's deformation can be limited to simple, global changes: bending, tapering, squishing, pinching, and so forth. This allows the user to apply virtual forces and have the object deform *smoothly*, without sharp kinks or creases (as when a standard FEM model is used) or fingerprints (as with real clay).

After roughing in the design, additional detailing may be required. At this stage, the user can switch to a modeler specifically designed for local surface detailing. Thingworld uses a surface detailing interface that is similar to a paint system. The user "paints" on a flattened, 2-D elevation map of the surface, with brighter shades of grey producing peaks
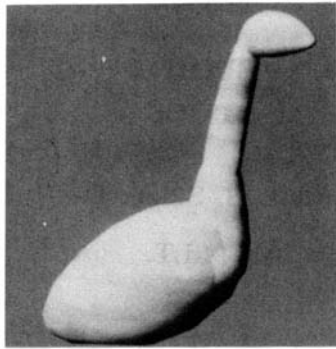
143

Figure 2: A solid model of a goose automatically produced from laser rangefinder data.

Figure 3: (a) chair: 5 minutes, (b) casting: 23 minutes, (c) dragon: about 30 minutes, (d) man and tree: about 1 hour.

and ridges, and darker shades producing valleys and pits. The coordinates of this map roughly correspond to the lines of principal curvature on the surface, so that it is easy to produce details that blend naturally with the overall shape of the object.

2.1 Automatic Modeling

In some situations we have measurements from a real object, and would like to create a computer model of that object. This is done by fitting one of the Thingworld deformable part models to these measured points. To start model fitting, the user interactively places a deformable model in approximate position over each "part" of an image of the object to be modeled. From these starting positions, the error between the model and measured point data can be numerically minimized, producing a part-wise 3-D model that closely fits the measured points. The fitting of model to data is accomplished by making use of the model's intrinsic dynamic behavior. That is, one simply gives each data point an artificial "gravity," and forces generated by these gravity fields cause the model to dynamically deform to fit the data; such fitting can be completed in only a few seconds on a SPARCstation 1[1]. The goose model shown in Figure 2 was produced in this manner, using data from a laser rangefinder. For more details see reference [3].

2.2 Performance Examples

In the current Thingworld system the user specifies forces by use of slider controls, which vary the amount of pinching, squishing, bending, etc., force. The computer screen then displays the deformed object in two orthographic views, with an update rate of 15 frames per second on a SPARCstation 1. Objects are rendered as line drawings, with hidden and intersecting lines and vertices removed. No special purpose hardware is employed. Figure 1 shows the progression of forming operations: First a cube is selected as the base "stock" form. This cube is then progressively oriented (rigid body modes), flattened and streched (linear modes), and finally bent and rounded (second order modes). The surface is then detailed to produce the final shape. Total elapsed time: under 30 seconds.

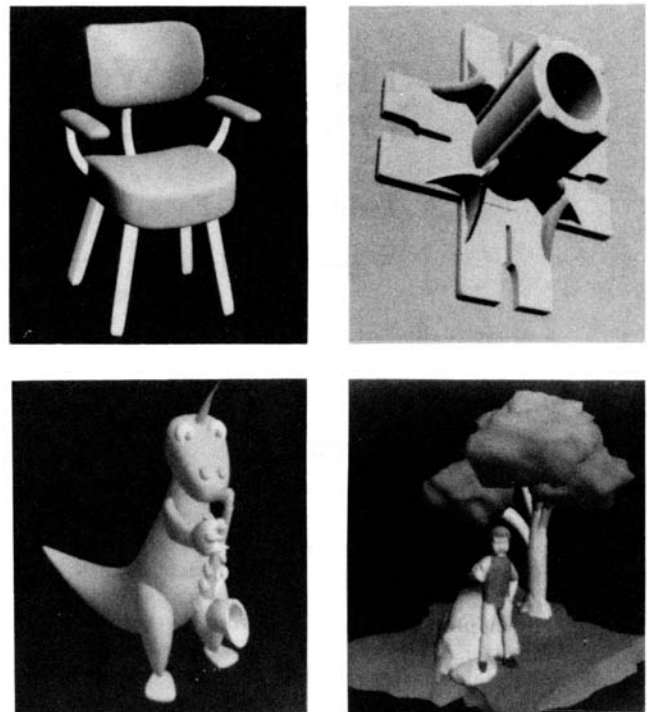Figure 3 shows four examples of solid models built using the Thingworld system[2]. The times given in the caption are total user times from start to finish, including filleting, surface detailing, and so forth. Experiments comparing Thingworld to people using clay — the fastest traditional modeling medium — have demonstrated that for low-accuracy, roughing in modeling Thingworld and clay require approximately the same amount of time, and for more detailed models Thingworld was substantially faster.

REFERENCES

[1] Borning, A., (1979), Thinglab — A Constraint-Oriented Simulation Laboratory. SSL-79-3, Xerox PARC, Palo Alto, CA.

[2] Pentland, A., and Williams, J., (1989) Good Vibrations: Modal Analysis for Graphics and Animation, *ACM Computer Graphics*, (Siggraph 89) Vol. 23, No. 4, pp. 215-223.

[3] Pentland, A., (1988) Automatic Extraction of Deformable Part Models, *M.I.T. Media Lab Vision Science Technical Report No. 104*, to appear in *Int'l Journal of Computer Vision*.

[4] Terzopoulos, D., Platt, J., Barr, A., and Fleischer, K., (1987), Elastically Deformable Models, *ACM Computer Graphics*, (Siggraph 87) Vol. 21, No. 4, pp. 205-214.

[5] Witkin, A., Fleischer, K., and Barr, A., (1987), Energy Constraints on Parameterized Models, *ACM Computer Graphics*, (Siggraph 87) Vol. 21, No. 4, pp. 225-232.

---

[1]SPARCstation is a trademark of Sun Microsystems
[2]or it's predecessor, SuperSketch

144