Critique: Deformable Objects

"Versatile and Efficient Techniques for Simulating Cloth and Other Deformable Objects", P. Volino, M. Courchesne, and N. M. Thalmann. *Proc. Siggraph* 1995, pp. 137-144. http://miralabwww.unige.ch/ARTICLES/SIGG95.html

Daniel Russel

February 27, 2000

1 Topic

The paper presents methods for simulating general flexible surfaces, with special emphasis on modeling textiles. The collision detection and resolution algorithms presented work on general triangular meshes and do not depend on orientation information, allowing them to work in situations where older methods could not.

2 Summary

The paper extends and improves the collision detection and resolution techniques presented in their earlier paper [1] to allow multiple flexible objects. They also present other details of their system which is optimized for simulations involving fabrics, especially clothing, but these are less interesting. While they only dealt with flexible surfaces, many of their collision detection techniques can be extended to volumes.

In their simulator, objects are represented using general triangle meshes with no regularity constraints. The internal dynamics of the objects are pretty straight forward, the main interesting parts of that being their tweaks to make the materials behave like textiles. These will be covered later. They use simple equations of motion and keep the time step very small in order to facilitate collision detection and handling. In order to mitigate some of the overhead of the small time step, the simulation supports dynamically

[1] "Efficient Self-Collision Detection of Smoothly Discretised Surface Animations using Geometrical Shape Regularity" P. Volino, N. Magnenat Thalmann, *Computer Graphics Forum (Eurographics Proc.)* 13(3). pp 155-166. 1994

varying the time step for the whole simulation, as well ask skipping updates of boring regions.

The small time step allows them to handle the collision interactions in a truly discontinuous manner and avoid ad-hoc collision forces. During each iteration, the objects positions are first updated without paying attention to collision constraints. Then, collisions are detected and the positions of the vertices are corrected to satisfy the geometric (objects cannot interpenetrate) and dynamic (conservation of momentum and inelastic collision) constraints. If there are multiple collisions that affect one another they iterate over the collisions until all constraints are resolved. This saves them from having to introduce exponential repulsion forces which are unphysical and can become extremely large.

The collisions are detected though a combination of standard bounding box techniques (between objects) and the curvature based technique they present in their other paper[1] (for self-collisions). Correctly resolving collisions can be tricky in complicated cases and their methods for doing so are probably the most interesting part. They keep track of the orientation of close vertex/triangle pairs and use this as their first method for determining which objects have penetrated one another. They also project the points back in time when the orientation is not previously known. Finally, since most collisions involve an area rather than single nodes, they locally walk over the region of interference and combine vertex information to orient the region as a whole. They don't provide any guarantees of the completeness of the methods, but they seem confident that they work in practice. They also mention techniques for updating the vertex proximity and orientation information as surfaces slide over one another, another commom case with textiles.

They make a number of somewhat ad-hoc modifications to the standard dynamics models to better mimic textiles. In order to facilitate tight folding, a problem since their techniques don't allow triangles to bend, they have a variable Young's modulus which allows easier compressior than stretching. In order to prevent discretization effects from making the bending strain blow up (since it is inversely proportional to curvature), they cap the bending strain and have a special handler for high strain cases. Also, when the model undergoes a great deal of stretching, they switch in to a plastic mode where the equilibrium dimensions change (and they allow the equilibria to gradually relax back to the original values when the high strain is removed). They artificially spread energy throughout the model when there are rapid vibrations characteristic of numerical instability. Which was an unusual idea.

3 Strengths and Weaknesses

The techniques presented produce very nice looking animations of difficult models (like the ribbon, or the clothes dryer). I don't have any basis for judging their computation time though.

A significant limitation is that collisions are all forced to be inelastic by the collision handling routines. While this is fine for textiles, it does not handle more rigid surfaces well. Simply making the collisions elastic could cause convergence problems, but this is not an area that they address. Another severe limitation is the requirement that the time step be small in order to avoid collisions which can't be accurately resolved. It allows their very clean collision resolution technique and allows them to avoid approximating collisions with continuous forces, but it severely limits what techniques can be used to speed up calculations.

The writing style of the paper left something to be desired and made unnecessarily hard to figure out at times. Of course some of this is doubtlessly due to they all being native French speakers. More importantly, they did not adequately (at least for me) define some variables used or adequately explain a few of their pictures. While these could be standard useages, and it was possible to figure everything out, a few more sentences would have gone a long way. I thought they spent too much time explaining their physics model, which was unexciting and their framework for assembling models, and the expense of their collision detection algorithms which were rather sparsely explained.

4 Questions

1. How does the running time compare to other techniques?

2. How well do the results extend to flexible volumes rather than surfaces? I think they would all work, but they may not be the most inefficient.

3. Is there some more systematic way to tweak the physics model to get the desired results?