Collision Approximation for Real-Time Cloth Simulation

Kristen Neal*
Mythic Entertainment
University of Virginia

1 Introduction

Dynamic cloth simulation has been a topic of interest within the graphics community for the past ten years. However, most of this research has focused on producing physically accurate results [Baraff et al. 2003]. Techniques such as these are useful for generating static images, but their computation time prohibits use in real-time simulation. While efforts on image based collision detection produced accelerated cloth simulation techniques, their implementation resulted in only a few frames per second [Vassilev et al. 2001] This type of simulation is still too slow to be incorporated within a large interactive applications such as computer games.

The predominate limiting factor for most cloth simulations is the computation time needed to model accurate collision detection and response. This is somewhat ironic, as people have a poor intuitive understanding of how cloth actually behaves. It is nearly impossible to predict how a piece of cloth will fold as it is draped over an object, or what patterns it will follow while flapping in the wind. With this understanding of human perception, approximations can be made to accelerate collision detection without sacrificing visual plausibility. The CollisionGrid system presented here uses the height field representation of terrain to approximate collision detection and response between cloth and world geometry.

2 CollisionGrid Architecture

The CollisionGrid system uses the height field representation for terrain to approximate collisions between a cloth mesh and world objects. An individual CollisionGrid defines the geometry of an object (or objects) in single axial direction. The CollisionGrid data specifies a grid sampling of this geometry that can be specified to any resolution of detail by modifying both the grid size and spacing. Through the use of multiple collision grids, any 3D surface can be easily reproduced.

While traditional collision detection algorithms compare position of each vertex in a cloth mesh to the vertices of all other objects in the scene, the CollisionGrid system simplifies collision detection into a simple projection and lookup. The time required to collide a cloth mesh against the grid depends only on the number of nodes within the cloth mesh, and is independent of size of the collision grid. Collision detection is done by projecting the position of each cloth node onto the CollisionGrid plane and bilinearly interpolating the nearest four data points to determine the height of the surface. The value is compared to the height of the cloth node along the collision axis to determine if a collision has occurred. Collisions are resolved by manually modifying the position and velocity of individual cloth nodes. Consequently, the CollisionGrid architecture is completely independent from the internal cloth dynamics.

The data values in a CollisionGrid (or set of grids) can be modified in predetermined sequences or on the fly from another dynamics engine to produce animation. Animating collision surfaces in this way simplifies the complex collisions, such collisions between a cloak and the back of a persons legs as they are running, into the simple grid lookup described before.

*e-mail:kneal@mythicentertainment.com

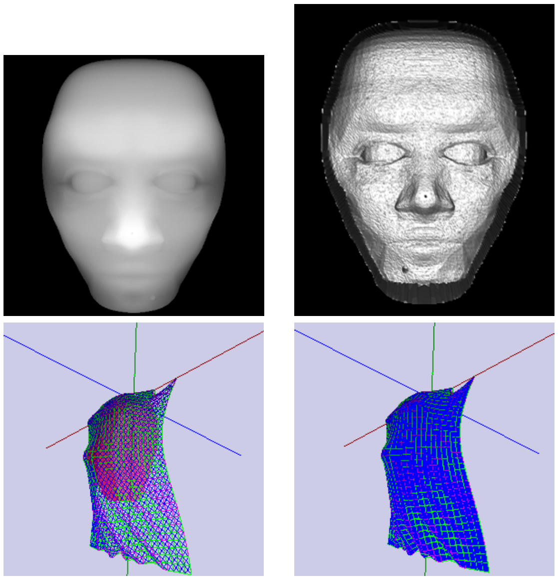

Figure 1: Top Left: A height map of a face model rendered in 3D Studio Max. Top Right: The resulting CollisionGrid produced by the height map data, rendered as a triangle mesh with a ray tracer to show the geometry approximation. Bottom Left and Right: An OpenGL application showing a piece of cloth draped over the same CollisionGrid. The cloth is drawn as a wire frame (Left) and triangle mesh (Right).

3 Results and Discussion

The simplicity, both in concept and in implementation, is the main advantage of the CollisionGrid system. CollisionGrids are easily generated from rendered height maps, and therefore provide an easy way to approximate the surfaces of any arbitrary geometry. The CollisionGrid system has also successfully demonstrated that collision detection for interactive cloth simulation can be greatly simplified without sacrificing plausible visual accuracy in the simulation. This technique is ideally suited for interactive applications such as computer games where high performance is required but exact physical accuracy is not necessary.

References

BARAFF, D., WITKIN, A., AND KASS, M. 2003. Untangling cloth. In *Proceedings of SIGGRAPH 2003*, 862–870.

VASSILEV, T., SPANLANG, B., AND CHYSANTHOU, Y. 2001. Fast cloth animation on walking avatars. In *Proceedings of EUROGRAPHICS 2001*, Blackwell Publishers.