Deformable Bodies: Mesh Generation and Simulation

Neil Molino
Stanford University

April 3, 2003

The generation of meshes for deformable bodies and the subsequent simulation of their motion are used in a number of areas including fracture, haptics, solid modeling, surgical simulations, and the modeling of biological tissue including the brain, the knee and even fatty tissue. Robust methods for generating these meshes are in high demand. Of particular interest to us are simulations of highly deformable bodies such as cloth (skin) and the muscle and fatty tissues commonly encountered in graphics, biomechanics, or virtual surgery. We propose a new mesh generation paradigm that gives particular consideration to conditioning the mesh for high deformation. We further discuss innovations for the simulation of deformable objects including a finite volume method for volumetric simulation, collision resolution, and internal dynamics modeling for cloth.

Our new mesh generation algorithm produces both high quality elements and a mesh that is well conditioned for subsequent large deformations. We primarily focus on tetrahedral meshes, but use this algorithm for surface meshing as well. We use a signed distance function defined on a grid in order to represent the object geometry. In the volumetric case, after tiling space with a uniform lattice based on crystallography, we identify a subset of these tetrahedra that adequately fill the space occupied by the object. Then we use the signed distance function or other user defined criteria to guide a red green mesh subdivision algorithm that results in a candidate mesh with the appropriate levels of detail. After this, both the signed distance function and topological considerations are used to prune the mesh as close to the desired shape as possible while keeping the mesh flexible for large deformations. Finally, we compress the mesh to tightly fit the object boundary using either masses and springs, the finite element method, or an optimization approach to relax the positions of both the interior and boundary nodes. The resulting mesh is well-suited for simulation since it is highly structured, has robust connectivity in the face of deformation, and is readily refined if deemed necessary during a subsequent simulation.

Mesh generation is not only a broad field, but is in some sense many fields each concerned with the creation of meshes that conform to quality measures specific to the application at hand. The requirements for fluid flow and heat transfer, where the mesh is not deformed, and for small deformation solids where the mesh is barely deformed, can be quite different from those for simulating soft biological tissue that may undergo large deformations. For example, while an optimal mesh for a fluid flow simulation should include anisotropically compressed elements in boundary layers, these highly stretched cells tend to be ill-conditioned when a mesh deforms significantly as is typical for soft bodies. Either the mesh is softer in the thin direction and the cell has a tendency to invert, or the mesh is stiffer in the thin direction and the simulation becomes very costly as the time step shrinks with higher stiffness and smaller element cross-section. Thus, although our method has been designed to provide a high degree of adaptivity both to resolve the geometry and to guarantee quality simulation results, we neither consider nor desire anisotropically stretched elements. Also, since highly deformable bodies tend to be devoid of sharp features such as edges and corners, we do not consider boundary feature preservation here. However, we believe that our algorithm can be extended to treat this case and plan to pursue it as future work.

Our main concern is to generate a mesh that will be robust when subsequently subject to large deformations. For example, although we obviously want an adaptive mesh with smaller elements in areas where more detail is desired, it is even more important to have a mesh that can be adapted during the simulation since these regions will change. Motivated by crystallography, we use a body-centered cubic (BCC) mesh that is highly structured and produces similar (in the precise geometric sense) tetrahedra under regular refinement. This allows us to adaptively refine both while generating the mesh and during the subsequent simulation.

Beginning with a uniform tiling of space, we then use a signed distance function representation of the geometry to guide the creation of the adaptive mesh, the deletion of elements that are not needed to represent the object of interest, and the compression of the mesh necessary to match the object boundaries. This compression stage can be carried out using either a mass spring system, a finite element method, or an optimization based approach. One advantage of using a physically based compression algorithm is that it gives some indication of how a mesh is likely to respond to the deformations it will experience during simulation. This is in contrast to many traditional methods that may produce an initial mesh with good quality measures, but also with hidden deficiencies that can be revealed during simulation leading to poor accuracy or element collapse.

In addition to testing the mesh with the physics based compression algorithm, we identify and guarantee several topological properties that ensure the mesh is sufficient for high deformation, for example a tetrahedron with all four of its edges on the boundary or an internal edge with both nodes on the boundary can limit the accuracy of the simulation, or worse, allow collapse.

We have generated tetrahedral meshes with this algorithm and used them in the simulation of highly deformable objects with a variety of constitutive models using several simulation techniques. The meshes retained their quality throughout all of these simulations. One new simulation technique that we tried is the finite volume method (FVM), which is more intuitive than the finite element method (FEM), since it relies on a geometrical rather than variational framework. We show that FVM allows one to interpret the stress inside a tetrahedron as a simple “multidimensional force” pushing on each face. Moreover, this interpretation leads to a heuristic method for calculating the force on each node, which is as simple to implement and comprehend as masses and springs. In the finite volume spirit, we also present a geometric rather than interpolating function definition of strain. We illustrate that FVM places no restrictions on the simulation of volumetric objects subject to large deformations, in particular using a quasi-incompressible, transversely isotropic, hyperelastic constitutive model for simulating contracting muscle tissue. B-spline solids are used to model fiber directions, and the muscle activation levels are derived from key frame animations.

Significant effort has been placed into accelerating FEM calculations including recent approaches that precompute and cache various quantities, modal analysis, and approximations to local rotations. In spite of significant effort into alternative (and related) methods for the robust simulation of deformable bodies, FVM has been largely ignored. Two aspects of FVM make it extremely attractive. First, it has a firm basis in geometry as opposed to the FEM variational setting. This not only increases the level of intuition and the simplicity of implementation, but also increases the likelihood that aggressive but robust optimizations might be found. Second, there is a large community of researchers using these methods to model solid materials subject to very high deformations. For example, it has been used with subcell pressures to eliminate the artificial hour-glass type motions that can arise in materials with strongly diagonally dominant stress tensors, such as incompressible biological materials.

We are particularly interested in the simulation of both active and passive muscle tissue. Biological tissues typically undergo rather large nonlinear deformations, and thus they create a stringent test for any simulation method. Moreover, they tend to have complex material models with quasi-incompressibility, transverse anisotropy, hyperelasticity, and both active and passive components. In this paper, we use a state of the art constitutive model, B-spline solids to represent fiber directions, and derive muscle activations from key frame animation.

The general theme of our mesh generation algorithm—tile ambient space as regularly as possible, select a subset of elements that are nicely connected and roughly conform to the object, then deform them to match the boundary—is applicable in any dimension and on general manifolds. We have used this idea to triangulate two-dimensional manifolds with boundary. We begin with a complete quality mesh of the object (actually created as the boundary of a tetrahedral mesh from our method, with additional edge-swapping and smoothing), and a level set indicating areas of the surface to trim. We keep a subset of the surface triangles inside the trimming level set and compress to the trimming boundary, making sure to stay on the surface. This can also be done for parametric surfaces with a triangulation in parameter space and a trimming level set defined either in three-dimensional space or in the parameter space. Quality two-dimensional surface meshes with boundary are especially important for cloth simulation. We, therefore, used cloth modeling as a testbed for our surface meshes. Two of the significant concerns in cloth simulation are collision resolution and internal dynamics modeling. A quality mesh free of slivers is important for well-conditioned numerical algorithms for both of these issues, particularly internal dynamics.

Collisions are a major bottleneck in cloth simulation. Since all points are on the surface, all points (and there will be many thousands for reasonably resolved simulations) may potentially collide with each other and the environment in any given time step. We argue that it is essential to maintain the constraint that the cloth mesh can never intersect itself. This can be efficiently achieved by a novel combination of fast and numerically well-conditioned penalty forces with a fail-safe geometric method. For dealing with collisions with large volumetric objects in the environment we demonstrate a technique for pushing the cloth to the surface of the objects while still preserving wrinkles and other features. Motivated by Newmark central schemes we demonstrate a time integration algorithm that decouples the collision dynamics from the interior cloth dynamics.

Cloth internal dynamics can be broken up into in-plane deformations (stretching and shearing) and bending. Bending is visually the most fundamental characteristic of cloth, yet it is not well understood. Working from basic principles we derive the only reasonable family of bending spring models that operate on triangle meshes, and show results with the simplest and most efficient member of this family. Our model easily permits non-flat rest poses, which allow the simulation of non-flat manifolds (e.g. skin) as well as pre-sculpted wrinkles and folds in cloth to maintain the desired artistic look of a garment during motion.

When simulating highly deformable objects, using a mesh designed for that purpose is obviously very important. Our new meshing algorithm produces such a mesh and has been tested in various simulations. As future work, we plan to explore the interplay between the mesh generation phase and the simulation phase for deformable bodes, for example allowing the mesh to slip during the simulation.