

3D Adaptive Multi-Resolution Triangulations Based on Physical Compression

Ricardo Marroquim

Paulo Roma Cavalcanti

Claudio Esperança

ricardo@lcg.ufrj.br

roma@lcg.ufrj.br

esperanc@lcg.ufrj.br

COPPE-Sistemas, Universidade Federal do Rio de Janeiro

Cidade Universitária - Ilha do Fundão

Caixa Postal 68530, CEP 21945-970

Rio de Janeiro, RJ, Brasil

***Abstract.** This paper presents a method for generating multi-resolution triangulations of non-manifold 3D objects composed of several regions with arbitrary geometry. The process of adapting the triangulation to the boundary of the object is based on physical compression, more specifically, a mass-spring system. The final triangulation usually has no degenerated triangles and provides an approximation of the boundary based on a chosen resolution.*

***Keywords:** Multi-resolution triangulations, physical systems*

1. Introduction

Numerical simulations have become an essential step in the development of engineering products, or in the prediction of the behaviour of physical phenomena, such as weather conditions, oil generation and migration, tide movement, earthquakes, etc.

In order to be able to carry out any simulation it is necessary to discretize the domain of interest, onto which a set of equations, describing physical laws, must be solved. This discretization process is generically referred to as the triangulation of the domain.

Although several triangulation algorithms have been proposed in the past three decades, these algorithms are meant mainly to deal with mechanical parts, produced by CAD systems. However, there are numerous important applications where the domain presents no symmetry at all, and boundaries are not well or clearly established.

Geological models, for instance, are often composed of several regions (forming the geological layers), and possess an irregular geometry, where surfaces can meet at very small angles, a situation known as a “pinch-out”. The geometry of the surfaces (horizons and faults) are obtained from seismic data, by a process which is not automatic and demands the interpretation of a geologist. As a consequence, the whole process of creating the model is not exact, and tends to be error prone.

Traditional triangulation algorithms can be classified according to the way they work. Delaunay based algorithms recover the boundaries after triangulating the convex hull of the pointset [4]. Ruppert and Shewchuck [13, 14, 15] insert a vertex at the circumcenter of a bad element to improve the mesh quality. In practice, if the original model already contains small angles, their algorithms have problems in converging, specially in 3D. Advancing front algorithms [9], on the other hand, start at the boundaries and proceed toward the center of the model, inserting new vertices to create good shaped elements. The trick is to merge the different fronts without generating any inconsistency. The success of the algorithm is highly dependent on the simplicity of the geometry of the boundary. Packing sphere algorithms [2, 8, 7] generate a distribution of vertices, forcing edges and faces to be present in the final mesh, which can be then generated, by a Delaunay like algorithm.

These traditional triangulation algorithms are hard to implement, have difficulty to deal with models having complex and irregular boundaries, and do not support multi-resolution. Even when a triangulation algorithm does a good job meshing this kind of model, it may generate small angles in the final mesh, which usually cause problems when running numerical simulations. Furthermore, if the original model already contains small angles these methods cannot eliminate them.

This paper presents an adaptive algorithm to triangulate non-manifold, multi-region models, which approximate their irregular boundaries instead of trying to match them exactly. Our main concern is to generate a good mesh, although the boundary is just an approximation, as close as possible to the original one.

Physically based triangulation algorithms are not new. Molino [11] has used a similar approach to obtain numerical meshes. Nonetheless, despite his good results, he considered only single region manifold objects. In this paper we extend Molino’s work to be able to deal with non-manifold multi-region objects, which are necessary to address a broader set of applications.

The algorithm has two main steps. In the first step, a grid is created over the object, adaptively, in order to subdivide the surrounding space. Some techniques to efficiently create the subdivision are also described. The representation of the input data is explained in Section 3., and Section 4. describes the distance function used to speed up the geometrical search in the algorithm.

The second step, described in Section 5., conforms the vertices of the grid to the boundary




Figure 1: A sliced sphere showing larger tetrahedra in the interior.

of the object through a compression scheme, while trying to maintain good shaped elements.

Our algorithm has several advantages over traditional triangulation techniques, such as: (1) being easy to implement, while offering multi-resolution at the same time; (2) being capable of dealing with models possessing complex and irregular boundaries; and (3) not being overly affected by numerical issues.

2. Spatial Subdivision

The first step of the algorithm is to subdivide the surrounding space of the model. An adaptive Binary Multi-triangulation (BMT) is used to hold the subdivision [10]. The grid obtained from the subdivision is also the base simplicial complex for the 3D triangulation.

The adaptive property of this mesh allows a refinement based on a given criterion. The one chosen here is the distance from a tetrahedron to a boundary surface. BMT adaptive meshes support multi-resolution naturally, and also have the property of producing a gradual transition of resolution, therefore leaving no gap between levels in the subdivision process. Figure 1 shows a sphere with finer resolution near the boundary, and coarser in the interior.

To subdivide a particular tetrahedron, the distance from its centroid to the boundary must be smaller than its edges average length. Also, each tetrahedron may only be subdivided a certain number of times. This maximum subdivision level is set by the user. An advantage of the BMT mesh is that it generates elements with good angles.

3. Input Model

Three-dimensional multi-region models are in general non-manifold. Therefore, to be properly represented, they require quite complex data-structures, such as the radial-edge data structure (RED) [16]. The RED data structure allows one to obtain any adjacency relationship in constant time.

Our models are created by using the *CGC* system [5], and the target are geological applications, which define a set of horizons and faults. Faults represent discontinuities (cracks) into the subsurface of the earth, and compart the model into a set of regions. A group of regions with similar attributes forms a geological layer. The geometry of geological models are intrinsically




Figure 2: 3D Gulf of Mexico Basin.

complex because the surfaces have been folded many times along the time, as can be seen in Figure 2. The model representing a sedimentary basin may extend from several hundred kilometers in x and y , to just a few hundred meters in z . Near the boundary a layer may be only 30 meters in depth.

4. Distance Function

Throughout the algorithm, the distance from a point in space to a boundary surface is computed several times. Therefore, a fast distance function is required. We make use of an octree to speed up the search.

The octree root node is defined by the model's bounding box coordinates. To make a compact structure, internal and leaf node coordinates are not stored, but computed on the fly while traversing the tree. Thus, the only attributes maintained in an internal node are pointers to its eight children, while a leaf node stores only the list of items intersecting its box.

The list of items in this case is the number of faces intersecting or into the node's box. If the list has more elements than a predetermined limit, the node is further subdivided. However, to prevent the refinement criterion from entering an infinite subdivision loop, the height of the tree is also limited. If there is a vertex with more incident faces than the imposed bound, the tree node containing it would be subdivided endlessly.

To find the distance from a given point to the boundary of the model, its octree is traversed starting at the root node. The estimated result distance is initialized to a suitable huge value. When a leaf node is visited, the distance to the closest face is returned. When an internal node (octant) is visited, the children octants are sorted in ascending order of distance from the query point. Each of these octants is then visited recursively, but only if the corresponding distance is smaller than the best estimate so far. If the traversal returns a distance estimate that is smaller than the overall best estimate, this is updated. In Figure 3, it is shown a pseudo-code for this algorithm.

5. Mesh Compression

To conform the refined mesh to the boundary of the model a physical mass-spring system is used. The idea is to select a subset of the mesh (called a d -mesh) and compress it to fit the boundary of the model. The set of vertices on the boundary of the d -mesh are marked to be projected towards the boundary of the model. When these vertices start to move, the forces created are propagated to the other vertices, rearranging all of the interior tetrahedra. Springs

```

global float  $best \leftarrow \infty$ 
procedure distance (point  $p$ , octreeNode  $T$ )
  if  $T$  is a leaf node then
    return the distance between  $p$  and the closest face in  $T$ 
  else
    sort the sons of  $T$  in ascending distance from  $p$ 
    for each son  $S$  (in sorted order) do
      let  $d$  be the distance between  $p$  and  $S$ 's box
      if  $d < best$  then
         $D \leftarrow \text{distance}(p, S)$ 
        if  $D < best$  then  $best \leftarrow D$ 
    return  $best$ 

```

Figure 3: Pseudo-code for the octree distance function.

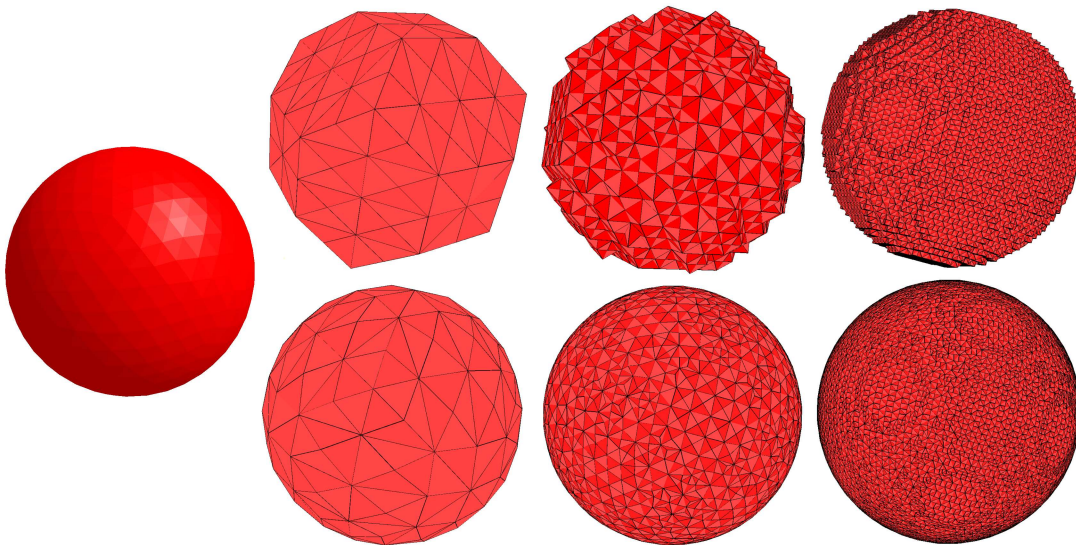


Figure 4: The d -mesh before and after compression for three different resolution levels (original sphere on the left).

are placed on each tetrahedron to hold their shape during the compression. Figure 4 shows the d -mesh and final triangulation of the sphere in three different resolutions.

5.1 d -Mesh

A set of vertices must be chosen to be pushed towards the boundary. One solution is to classify all tetrahedra as being *inside*, *outside* or *intersecting* the boundary of the model, and choose all exterior vertices of the intersecting tetrahedra. Analogously, all interior vertices could have been chosen either. However, merely choosing a set of vertices is not sufficient, since there can be inconsistencies during the compression. For example, some tetrahedra might have all four vertices onto the boundary of the model, and chances are they will be crushed against its surface. Also, some selected vertices may be far away from the boundary, creating a great disturbance when displacing them towards the boundary.

Ideally, the d -mesh should be close enough to the boundary of the model. Following [11],




Figure 5: A 2D representation of the d -mesh. White vertices represent the enveloped set. Vertices marked to be projected towards boundary of the model are drawn in black. The dark black line represents the boundary of d -mesh. The model's boundary is drawn in blue. The dashed lines are the discarded edges.

the d -mesh is composed of all tetrahedra incident to an enveloped set. The enveloped set contains vertices with all incident edges at least 25% inside the model. This guarantees that all tetrahedra of the d -mesh have a minimum of one vertex inside the model.

All vertices of the d -mesh that do not belong to the enveloped set are marked, and form a set of vertices to be pushed towards the boundary. This procedure insures that no tetrahedron has its four vertices marked. This set of vertices is a rough approximation of the boundary of the model. Tetrahedra that do not belong to the d -mesh are discarded.

The set of marked vertices contains some exterior and interior vertices, but they are all close to the boundary of the model, as can be seen in Figure 5. The more a vertex has to move to reach the boundary, the greater the force it creates on the mesh. Very high compression forces cause disturbances in the mesh that may lead to bad shaped elements.

There are two cases that have to be addressed to provide a smooth compression. First, an internal edge of the d -mesh, with both vertices marked, may be crushed during compression. An internal edge, as shown in Figure 6, is identified when all tetrahedra of its *star*, i.e., tetrahedra incident to the edge, belong to the d -mesh. Second, edges on the boundary of the d -mesh with more than two incident boundary faces are identified as being *non-manifold*. These internal and non-manifold edges are subdivided in a first and single pass, and the d -mesh is computed once again. When subdividing an internal edge, the new vertex is automatically inserted in the enveloped set. Because the edge is in the interior of the mesh, this new vertex is also guaranteed to be in its interior.

In subsequent passes the mesh is checked again for remaining inadequate edges. For each one of them, the vertex with the smaller distance to the boundary is inserted into the enveloped set. Because the distance is negative inside the model, if both vertices are in the interior, the deeper one is chosen. After each pass the d -mesh is recomputed, and the process is repeated until no more inadequate edges are found.

5.2 Multi-region Models

When working with multi-region models (possessing more than one region), different enveloped sets are created for each region. Figure 8 shows four different separated regions of the salt dome model. The enveloped set contains only vertices inside each region, therefore guaranteeing that different sets do not intersect. However, when computing the d -mesh, i.e., elements incident on the enveloped set, a tetrahedron might have vertices in different sets. Furthermore, this may lead to tetrahedra having all four vertices marked to be projected towards the boundary of the model, causing them to be crushed when compressed.




Figure 6: Red edges represent the interior edges of the d -mesh connecting two marked vertices. The boundary of the d -mesh is drawn with a thick black line, and the boundary of the model in blue. To facilitate the visualization the d -mesh is drawn in 2D, but the 3D case is analogous.

To solve this problem, it is necessary to have a delimiting surface between adjacent regions. This surface contains the marked vertices, and, after the compression stage, shall approximate the boundary between the regions.

When computing the enveloped set, the same rules are applied as described in section 5.1. A vertex must have all incident edges at least 25% inside its region to belong to the set. However, to find the delimiting surfaces, the d -mesh is only computed for some regions. A simple criterion based on the identification of the regions can be used. When selecting the marked vertices on the boundary between two regions, only the region with the smallest identifier is considered. Solving the boundary of the d -mesh for a region automatically solves it for the adjacent region, as can be seen in Figure 7.

5.3 Spring Configuration

A very important issue is how to place the springs on the physical system. Badly placed springs may not hold the shape of elements, or worse, may not prevent the tetrahedra from collapsing.

Two types of springs are used: edge and repulsion springs. An edge spring is placed in each d -mesh edge. Three repulsion springs are placed on each tetrahedron, connecting a vertex to its projection onto the plane of the opposite face. This configuration proved to prevent collapses efficiently [6].

Each edge has a reference to its corresponding spring, while each face has references to two repulsion springs. An edge spring is connected by two real vertices. However, repulsion springs connect a real to a virtual vertex, thus requiring some auxiliary vertices to represent the projection of a real vertex.

To compute spring forces two different equations are used: one for edge springs and the other for repulsion springs (projection edges). The edge spring equation is the Hooke's Law:

$$F = k_s * (|\Delta x| - x_o) + k_d \left(\frac{\Delta v \cdot \Delta x}{|x|} \right) * \frac{\Delta x}{|\Delta x|},$$

where k_s and k_d are the elasticity and damping constant, x_o is the rest length and Δx and Δv are the vertex position and velocity differences.

For the repulsion spring there is a non-linear equation, as explained in [6]. These non-linear springs apply an infinity force when their lengths go to zero.

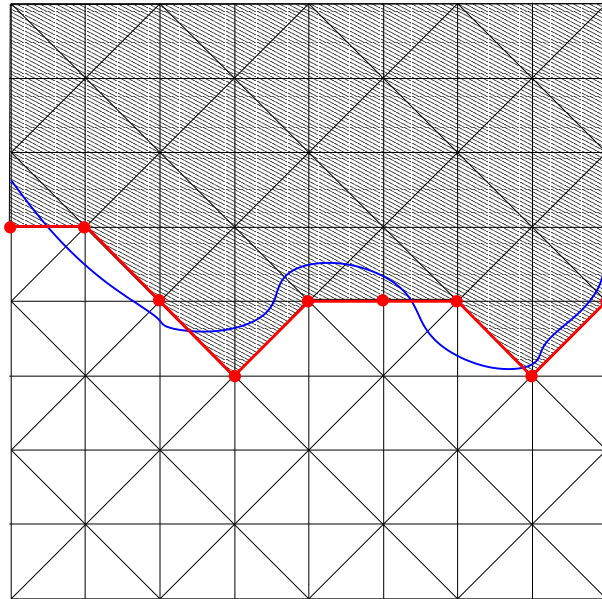


Figure 7: 2D view of the delimiting surface (drawn in red) between two regions, gray and white. The blue curve represents the model's boundary. The red vertices are those marked to be projected towards the blue curve.




Figure 8: On the left, a wireframe and a solid view of the Salt Dome. On the right, the regions of the model have been exploded.




Figure 9: Incident springs to vertex v_0 . Edge springs are drawn in blue, and repulsion springs in red. The other three edge and repulsion springs for this tetrahedron are not shown here.

$$F = k_s * \left(|\Delta x| - \frac{x_o}{|\Delta x|} \right) * \frac{\Delta x}{|\Delta x|}.$$

5.4 Particle System

To simulate the mass-spring oscillatory movement, the mesh is treated as a particle system [17]. Each vertex represents a particle in the system, and particles are associated to spring endpoints. Some temporary particles may be introduced to represent auxiliary vertices. These particles only exist when computing the force associated to repulsion springs, having only local influence in the system.

Each particle has the following attributes: position, velocity, force, and mass. The system iterates in time steps, updating at each cycle, the particles' position and velocity. At each time step there is an iteration through all edges to compute the forces at their particles. The force on each vertex is the sum of the forces applied by all incident springs.

Also, the particles are given different mass values. Those near the boundary are made heavier than those far from it. The reason is to give more robustness to the small tetrahedra near the boundary, while giving the bigger interior ones more freedom to rearrange. This avoids crushing the small mesh elements when the interior of the mesh offers too much resistance. The mass of a particle is computed as the average of all incident tetrahedra masses. The mass of a tetrahedron is computed as the ratio between its level and the maximum subdivision level.

The duration of one step can be redefined by the user. Usually, small steps are better because they create less oscillatory movement on the springs. However, smaller time steps result in more iterations, and consequently longer simulation running times.

Each system iteration has three main phases:

1. computation of the force accumulators for each vertex;
2. calculation of derivative values;
3. computation of new positions and velocities.

5.5 Force Accumulators

The force accumulated in each vertex is the sum of all contributions from all incident springs. Auxiliary vertices are not considered, but they contribute to the force accumulated

on the other endpoint of the repulsion spring.

The first step is to reset all force accumulators, with a simple iteration through all particles. Next, the edges and faces are examined in order to compute the forces applied on their corresponding springs.

5.6 Derivative Values

In this step, each particles' position and velocity derivatives are calculated, i.e., its velocity and acceleration, respectively. Here, an auxiliary array is created to hold all the values. The vector size is $6n$, where n is the number of particles in the system. The derivative vector is scaled for the time of an iteration step.

Another array is used to hold the particles' current position and velocity. The vector size is the same of the derivative vector. To compute the new attribute values the two vectors are added, and the particles are updated with the resulting values.

5.7 Compression Criteria

Spring rest lengths are set to their initial length, meaning that initially the particle system is completely stable. Unless some disturbance is applied on the system, it will not change. To start the compression, marked vertices of the d -mesh need to be pushed towards the boundary, see 5.1.

The initial velocity direction of a marked vertex is computed as the average of all incident boundary face normals. If the vertex is in the exterior of the mesh surface, the normals facing the interior of the incident tetrahedra are used. Otherwise, in case of an interior vertex, the normals facing outwards are used. This direction precludes a vertex from landing near edges or other vertices, spreading them apart fairly well.

Marked vertices also behave differently from the others in the system. Specifically, their force accumulator are projected onto the tangential plane of their direction. After each time step their velocity direction is recomputed according to the incident face normals. In summary, the marked vertices travel towards the mesh's surface while rearranging themselves in the tangential direction. If these vertices were allowed to interact freely with the system, the internal forces of the mesh would eventually return it to the initial state.

5.8 Limiting the Deformation Rate

To better control the mesh during the simulation, some bounds are imposed onto the movement of the particles. These bounds are motivated by the techniques used in cloth simulation, and limit the strain and strain rate [3, 1, 12]. First, no spring may compress a tetrahedron more than 60% of its initial height. Second, in a time step, no particle may move more than 10% of the current height of any of its incident tetrahedra.

When a vertex surpasses the strain bound, it is halted during the current time step. The other vertex new positions are iteratively recomputed until all tetrahedra abide the restriction. These artificial bounds create a higher degree of robustness, while compressing the mesh. They also help to obtain better quality meshes.

By this criterion some vertices being pushed towards the boundary may be continuously halted, and never reach it. This happens more often with low refined meshes, where the vertices have to travel large distances, and the size variation, of the tetrahedra near the boundary and those in the interior, is small. This procedure is valid in cases where a perfect match of the boundary is not necessary. The simulation finishes when all marked vertices have either reached the boundary of the model, or have been halted.




Figure 10: Rough approximation of the Stanford bunny with 6.6K tetrahedra.




Figure 11: Triangulation of a tiger with approximately 160K tetrahedra.

6. Results and Conclusions

Some preliminary results are presented in this section. A rough approximation of the Stanford bunny is shown in Figure 10. This triangulation possesses less than 7K tetrahedra, with minimum dihedral angle of approximately 10° . The total processing time was 25 minutes, on a pentium 3 computer, 1000Hz, 512Mb of RAM.

The triangulation of the tiger, shown in Figure 11, possesses 160K tetrahedra. The minimum dihedral angle is 11° , and the processing time was 2 hours. Figure 12 shows two slices of the models. In Figure 13, it is shown two views of Spock's head, and the respective d -mesh and final triangulation.

We have presented an algorithm for triangulating multi-region, non-manifold 3D models. Although the boundaries of the triangulation may differ from the boundary of the original model, there are several applications where this is acceptable.

Our algorithm has many advantages over traditional triangulation techniques, such as: (1) being easy to implement, and offering, at the same time, multi-resolution; (2) being capable of dealing with models possessing complex and irregular boundaries; and (3) being not affected in a high degree by numerical issues.




Figure 12: Tiger model sliced.

Acknowledgments

We would like to thank the CAPES and CNPq agencies and, specially, the financial support received from Cenpes/Petrobras, through the COPPETEC project Pesc 3678.

REFERENCES

- [1] David Baraff and Andrew Witkin. Large steps in cloth simulation. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, pages 43–54. ACM Press, 1998.
- [2] Marshall Bern and David Eppstein. Quadrilateral meshing by circle packing. *International Journal of Computational Geometry and Applications*, 10(4), August 2000.
- [3] Robert Bridson, Ronald Fedkiw, and John Anderson. Robust treatment of collisions, contact and friction for cloth animation. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 594–603. ACM Press, 2002.
- [4] Paulo Cavalcanti and Ullisses Mello. Three-dimensional constrained Delaunay triangulation: A minimalist approach. In *8th International Meshing Roundtable*, pages 119–129, Lake Tahoe, California, October 1999. Sandia National Laboratories.
- [5] Paulo Roma Cavalcanti, Paulo C.P. Carvalho, and Luiz F. Martha. Nonmanifold modeling: An approach based on spatial subdivisions. *Computer-Aided Design*, 29(3):209–220, March 1997.
- [6] Lee Cooper and Steve Maddock. Preventing collapse within mass-spring-damper models of deformable objects. In *The Fifth International Conference in Central Europe on Computer Graphics and Visualization*, pages 70–78, Plzen, Czech Republic, February 1997.
- [7] Xiang-Yang Li, Shang-Hua Teng, and Alper Üngör. Biting: Advancing front meets sphere packing. In *2nd Symposium on Trends in Unstructured Mesh Generation*, University of Colorado, Boulder, August 1999.
- [8] Xiang-Yang Li, Shang-Hua Teng, and Alper Üngör. Biting spheres in 3d. In *Proc. 8th Int. Meshing Roundtable*, pages 85–95, South Lake Tahoe, CA, October 1999.




Figure 13: Front and side view of Spock's head; d -mesh before compression; and final triangulation with 230K tetrahedra.

- [9] Dimitri J. Mavriplis. An advancing front Delaunay triangulation algorithm designed for robustness. 117(1):90–101, 1995.
- [10] Vinicius Mello, Luiz Velho, Paulo Roma Cavalcanti, and Claudio Silva. Bmt: A generic programming approach to multiresolution spatial decompositions. In H. Hege and K. Polthier, editors, *Visualization and Mathematics III*, pages 337–360. Springer-Verlag, Berlin, 2003.
- [11] Neil Molino, Robert Bridson, Joseph Teran, and Ronald Fedkiw. A crystalline, red green strategy for meshing highly deformable objects with tetrahedra. In *12th Int. Meshing Roundtable*, pages 103–114, Santa Fe, New Mexico, September 2003.
- [12] Xavier Provot. Deformation constraints in a mass-spring model to describe rigid cloth behavior. In Wayne A. Davis and Przemyslaw Prusinkiewicz, editors, *Graphics Interface '95*, pages 147–154. Canadian Human-Computer Communications Society, 1995.
- [13] Jim Ruppert. *Results on Triangulation and High Quality Mesh Generation*. PhD thesis, University of California at Berkeley, 1992.
- [14] Jim Ruppert. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *J. Algorithms*, 18(3):548–585, 1995.
- [15] Jonathan Richard Shewchuk. *Delaunay Refinement Mesh Generation*. PhD thesis, Department of Computer Science, Carnegie Mellon University, 1997.
- [16] Kevin Weiler. The radial-edge structure: A topological representation for non-manifold geometric boundary representations. In *Geometric Modeling for CAD Applications*, pages 3–36. North-Holland, 1988.
- [17] Andrew Witkin. An introduction to physically based modeling: Particle system dynamics. ACM Siggraph Course Notes, 1994.