Dynamic Meshes for Deformable Surfaces (posters_0040)

Fernando de Goes, IC-UNICAMP
Felipe P. G. Bergo, IC-UNICAMP
Alexandre X. Falcão, IC-UNICAMP
Siome Goldenstein, IC-UNICAMP
Luiz Velho, IMPA

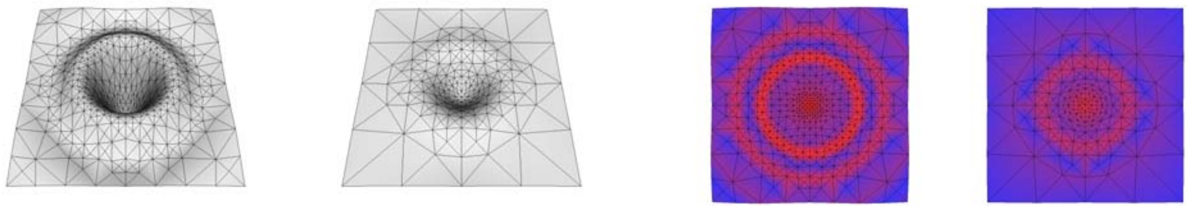

Figure 1: Deformations of a height function. Left side: perspective view. Right side: adaptation over the curvature.

Mesh adaptation is an area of intense research because of its importance. However, most of the work in this area has concentrated on adapted meshes for static models. The problem of creating dynamic meshes for deformable objects received less attention up to now. Although some results obtained for adaptation of static meshes can be used for dynamic meshes, truly effective computation can only be achieved by exploiting the specific nature of deformable objects.

We propose a framework for generating adapted dynamic meshes for deformable surfaces. We construct an initial mesh by refinement of a coarse base mesh. During the simulation, the current mesh is maintained adapted to the underlying deformable surface by local modifications to the mesh geometry and structure. The adaptation is based on the evolution of the surface. Nonetheless, it is decoupled of the simulation process, since we estimate the dynamics from stochastic samples that are evaluated at each time-step. Furthermore, our scheme is general and can be applied to parametric and implicit surfaces alike.

There are several criteria to determine the quality of a polygonal representation of a continuous surface. First, the piecewise linear approximation given by the mesh should be within a tolerance (according to some *error metric*). Second, the *mesh size* (i.e., the number of elements) should be small. Third, the *shape* of polygons (i.e., aspect ratio, orientation) should be bounded and adapted to surface features. Fourth, the *structure* of the mesh (i.e., degree of vertices) should be as regular as possible – which in the case of triangle meshes mean valences close to six.

Note that the criteria above are interdependent and may conflict with each other. In essence, the construction of a good mesh can be posed as a constrained optimization problem. One strategy to solve such problem is through a mesh adaptation process that balances the different criteria according to restrictions.

The adaptation process controls mesh geometry and topology to achieve its goals. Geometric control determines the displacement of vertex positions – normal or tangential to the surface. Topology control determines vertex connectivity by structural operations that may also affect the resolution of the mesh, such as vertex insertion and removal.

Our adaptation framework combines structural and geometric operations in order to maintain the quality of a dynamic mesh while the underlying surface is deforming. It keeps the surface approximation under some prescribed tolerance using a small number of elements that are well shaped and aligned to features. Furthermore, the resulting mesh structure has bounded and graded regularity.

An overview of the whole process is as follows: Given a deformable surface $S(t)$ and a base mesh M that has the same topology of S, we assume the ability to sample points on S at each time step t of the simulation, and that S does not change topology. The adaptation is an iterative algorithm: it has an initialization phase and a main loop. In the initialization, the base mesh is adaptively refined to create a mesh M_0 approximating S at $t = 0$. In the loop, for each time step $t = i$ of the simulation, the current mesh M_i is adapted by structural and geometric operations using a sampling of S_i.

In order to determine the correct level of approximation we need to evaluate an integral of the error over surface regions. Exact computation might be prohibitive, so, an alternative is *stochastic sampling*. Stochastic sampling selects points in the region of interest, and keeps the density of samples per area constant with uniform distribution. The quality of the integral's approximation (the sum of the samples' errors) is related to the sampling density.

As the model deforms we apply the following two-step procedure:

Structural Operations – Simplification / Refinement:
First the surface is sampled at the current time step. Then, the mesh resolution is modified based on the error criteria. We remove vertices in regions with error below the threshold and subdivide edges in regions with error above it. We employ the stellar operations edge *split* and *weld* to change the structure of the mesh. Their action is respectively to refine / coarsen a basic region of the mesh by subdividing two triangles that share an edge / simplifying four triangles incident to a vertex.

Geometric Operations – Vertex Position Correction:
During the structural adaptation, the vertex positions are set to reflect the current state of the deformable surface. However, after the geometry update there is no guarantee that all the triangles of the mesh will remain well shaped. We want a mesh with *good parametrization*. This happens when neighbor triangles have edges with similar length. We activate this by performing an intrinsic Laplacian mesh smoothing that is curvature and feature sensitive. The procedure is as follows: First, new vertices are projected onto the surface and existing vertices updated for the current time step. Then, all vertices are displaced on the tangent plane of S using the local curvature of the surface.

Our method works with implicit and parametric models. In Figure 1, we show the mesh adaptation of a deformable *sinc* height function. The curvature-sensitive adjustment plays a crucial role to keep the ridges of the *sinc* function sharp and well defined, while the error criteria controls the adaptation everywhere else.