Resampling Adaptive Cloth Simulations onto Fixed-Topology Meshes

George Brown
University of Minnesota

Armin Samii
Adobe Systems Inc.

James F. O'Brien
University of California, Berkeley

Rahul Narain
University of Minnesota

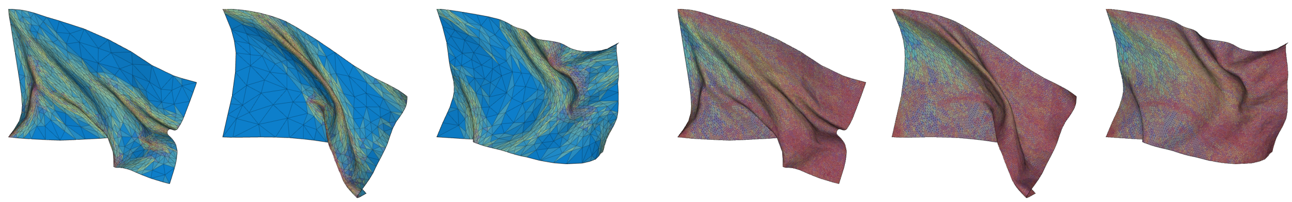

Figure 1: *Left: Three frames from an adaptive simulation of a flag. Right: The optimal fixed-topology mesh resolves all details in the animation, including those in frames not shown here. Colors indicate triangle size and shape (blue: large, red: small, yellow: anisotropic).*

Abstract

We describe a method for converting an adaptively remeshed simulation of cloth into an animated mesh with fixed topology. The topology of the mesh may be specified by the user or computed automatically. In the latter case, we present a method for computing the optimal output mesh, that is, a mesh with spatially varying resolution which is fine enough to resolve all the detail present in the animation. This technique allows adaptive simulations to be easily used in applications that expect fixed-topology animated meshes.

1 Introduction

One of the main challenges for realistic cloth simulation is the need for high-resolution discretizations in order to resolve fine detail such as localized wrinkles and folds. Adaptive remeshing [Narain et al. 2012] addresses this challenge by dynamically refining and coarsening the mesh to focus computational effort in regions where detail is present or is likely to arise. However, simulations using adaptive remeshing have been considered hard to integrate into many existing production pipelines because of the changing mesh topology at every time step. Recent techniques for high-quality interactive animation have employed precomputed simulation data, but they also assume a fixed topology and have not been able to take advantage of adaptive offline simulations.

2 Our Approach

We work with cloth models represented as triangle meshes, where each vertex carries both its current world-space position $\mathbf{x} \in \mathbb{R}^3$, and its coordinate $\mathbf{u} \in \mathbb{R}^2$ in an undeformed reference space. Equivalently, such a mesh may be thought of as a triangulation T of a flat domain $U \subset \mathbb{R}^2$ along with an assignment of the vertices of T to world-space positions in $\mathbb{R}^3$. The input to our algorithm is a sequence of meshes $M_1, M_2, \ldots$ produced by an adaptive simulator; because the simulator performs remeshing, these meshes correspond to different triangulations of the same domain U. To produce a single fixed-topology animated mesh $\overline{M}$, we have to (i) choose a fixed triangulation $\overline{T}$ of U, and (ii) resample the geometry of the input meshes $M_1, M_2, \ldots$ onto the vertices of $\overline{T}$.

For now, suppose $\overline{T}$ is known, for example as specified by the user. Resampling the geometry for any input mesh M_i may be achieved as follows. Each vertex of $\overline{T}$ is a point $\mathbf{u} \in \mathbb{R}^2$. We find the face that contains $\mathbf{u}$ in the input mesh's triangulation T_i, and perform interpolation to compute the corresponding world-space position, for example using exact evaluation of subdivision surfaces [Stam 1999]. This gives the world-space position of the vertex in $\overline{M}$.

It is also possible to determine the optimal mesh topology required to resolve all the features of the input animation. Adaptive remeshing uses a sizing tensor field $\mathbf{M}$ defined on vertices to characterize the resolution of the mesh. Given a number of sizing tensors $\mathbf{M}_1, \mathbf{M}_2, \ldots$, Narain et al. [2013] describe a procedure to find the sizing tensor which results in at least as much refinement as any of them; we denote this procedure by $\max(\mathbf{M}_1, \mathbf{M}_2, \ldots)$. We use this procedure to determine the optimal mesh as follows. Starting with $\overline{T} \leftarrow T_1$, the triangulation from the first input mesh, we iterate over each subsequent mesh $i = 2, 3, \ldots$. On each iteration, we

- update the sizing tensors at each vertex $\mathbf{u} \in \overline{T}$ via
$$\mathbf{M}(\mathbf{u}) \leftarrow \max(\mathbf{M}(\mathbf{u}), \mathbf{M}_i(\mathbf{u})),$$
where $\mathbf{M}_i(\mathbf{u})$ is obtained by interpolating the sizing field of the input triangulation T_i at $\mathbf{u}$, then
- remesh $\overline{T}$ using the updated sizing field $\mathbf{M}$.

Thus, at each step, $\overline{T}$ is refined at least as much as the input T_i, and is never coarsened. After all the input meshes have been processed, we use the final $\overline{T}$ for resampling as in the previous paragraph.

This approach allows adaptivity to be treated as a purely internal mechanism that enables efficient simulation, while the user only needs to see a fixed-topology mesh. We hope that these developments will accelerate the adoption of adaptive simulation techniques.

References

NARAIN, R., SAMII, A., AND O'BRIEN, J. F. 2012. Adaptive anisotropic remeshing for cloth simulation. *ACM Transactions on Graphics 31*, 6, 147:1–10.

NARAIN, R., PFAFF, T., AND O'BRIEN, J. F. 2013. Folding and crumpling adaptive sheets. *ACM Transactions on Graphics 32*, 4, 51:1–8.

STAM, J. 1999. Evaluation of loop subdivision surfaces. In *SIGGRAPH'99 Course Notes*, ACM.