

Supplementary Document for Hyper-Reduced Projective Dynamics

Comparison to Simplification Schemes

In the following we want to highlight the difference of our hyper reduction scheme to an approach in which the mesh is simplified, fully simulated and then the deformation is mapped to the original mesh using an *up-sampling matrix*.

Specifically, we simplify the surface of the original mesh using an algorithm similar to the method by Garland and Heckbert [1997], implemented such that all vertices of the simplified mesh lie on vertices of the original mesh. Using this correspondence, we create a matrix that maps displacements of the vertices of the simplified mesh to displacements of the vertices of the original mesh. There are many options to choose this matrix, and we found that no choice fixes the problems that we will address below. For the following experiment, we use our subspace construction, using the vertices of the simplified mesh as the samples on the original mesh from which the weights are computed (as described above, they are in direct correspondence). Note that we cannot add rotational degrees of freedom, since these are unknowns that cannot be determined from the vertex positions of the simplified mesh directly. This means, that the position of each vertex on the original mesh is a weighted combination of the positions of the vertices of the simplified mesh, where the weights are based on proximity to those samples.

In Figure 1 and in a supplementary video, we show (frames of) the simulations using our hyper reduction scheme and the simplification scheme described above. For our method we use the parameters listed in Table 1. To make the comparison as fair as possible, we used the same number of degrees of freedom for both our method and the method described above. That is, we simplify the surface mesh, such that the tetrahedralized version has 720 vertices (and 1490 tetrahedrons). This results in both methods having almost identical computation timings per iteration: Our method has some overhead compared to a full simulation on the same number of degrees of freedom due to solving the fitting problem (9), while the simulation on the simplified mesh has to evaluate significantly more constraint projections (1490 as opposed to 1000 for our method).

The experiment clearly demonstrates the issues with approaches that rely on mesh simplification. Firstly, the up-sampled results of the simulation on the simplified mesh suffer from strong linearization artifacts. Since no rotational degrees of freedom are solved for, and are thus unavailable, the linear relationship between the vertices of the simplified mesh and the original mesh becomes obvious. In contrast, we carefully choose degrees of freedom that can capture both local and global rotations and translations and solve the variational problem directly in these coordinates. This leads to an optimal solution to the equations of motion within the subspace, as opposed to an approximation that relies on a

simulation on a simplified mesh. Secondly, all details that are no longer visible or well represented in the simplified mesh, will not be part of the simulation and do not show any elastic behavior. Instead they are simply displayed as stiff extensions of the rough features still visible on the simplified mesh. This is a fundamental difference to our approach of approximating the dynamics of the original mesh via the proposed constraint projections fitting method. This can be seen when closely the smaller features of the squid mesh that still show individual and elastic dynamics in our hyper reduced simulation.

Note, that our subspace construction is certainly not the best method to generate a matrix for direct vertex displacements up-sampling and other approaches will yield better results. One problem of our approach is that we use the same support radius for the influence of each vertex of the simplified mesh. Therefore it is difficult to find a good trade-off between local influence and global coverage of the weights. The construction of a good up-sampling matrix is not trivial and not the goal of our approach. Moreover, the issues addressed above will hold for any type of linear up-sampling approach: detail that is not available on the simplified mesh can never take part in the reduced dynamics and a linear up-sampling approach can never be artifact free.

References

Michael Garland and Paul S. Heckbert. Surface simplification using quadric error metrics. *Proc. ACM SIGGRAPH*, pages 209–216, 1997.




Figure 1: Top: The original squid mesh with the edges and vertices of the simplified mesh embedded. Bottom: Frames of the simulation using our hyper reduced scheme (left) and of the simulation using the simplified mesh, up-sampled to the vertices of the original mesh (right).