Base Mesh Construction using Global Parametrization

Francisco Ganacim André Maximo Luiz Velho
VISGRAF Lab, IMPA, Rio de Janeiro, Brazil

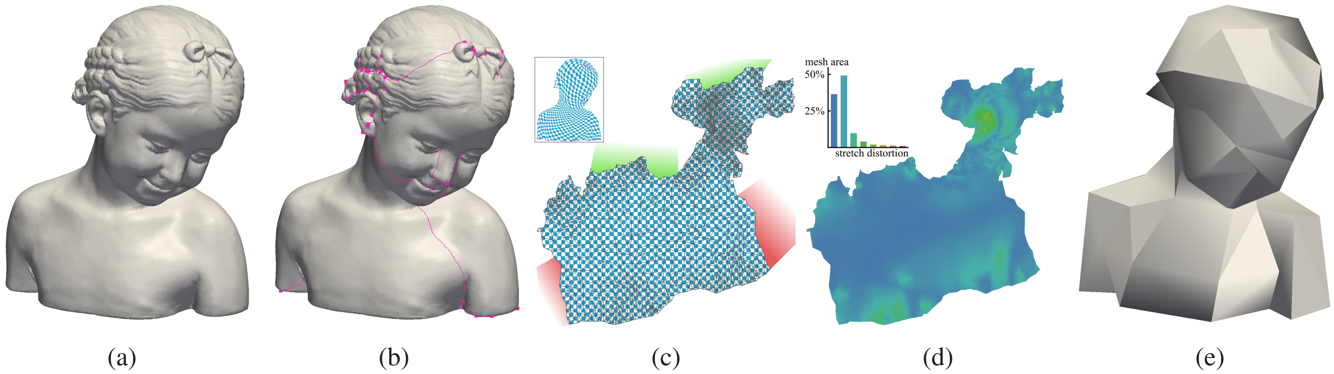

(a) (b) (c) (d) (e)

Figure 1: *The original mesh (a) is cut using seams (b) containing cone singularities (purple points); the mesh is parametrized (c) with connected boundaries (e.g. red and green regions); the base mesh (e) is generated using the metric distortion on the parametrization (d).*

Introduction and Motivation Base mesh construction from a dense-polygon mesh is often used to reduce the complexity of geometry processing problems. In the base or control mesh, each face corresponds to a region on the original surface and is used to encode its geometry. This encoding can involve a different representation of the surface, e.g. using displacement field and subdivision surfaces [Lee et al. 2000], or can be a more direct representation, e.g. through charts [Sander et al. 2003]. In the former example, the control mesh is constructed using edge-collapse simplification, and in the latter by an iterative seed-placement and chart-growth optimization process. Although both methods strive to optimize this construction, the first implies a sequence of local operations, lacking a global strategy, and the second iterates over greedy choices, which may not converge to a global solution.

In this work, we present an algorithm to construct base meshes that enables a global analysis of the problem. First, we compute a non-optimal global parametrization of the original mesh (see Figure 1), allowing distortions and using seams and cone singularities. Then, we map the boundaries of the parametrization pairwise, excluding the cone points and possible boundaries of the original surface. Finally, we place a few vertices on top of the global parametrization and compute a 2D Delaunay triangulation, yielding our base mesh. The analysis of where to place the base mesh vertices is done once, without iterations, and with global knowledge of the surface.

Our Approach Global parametrization is often done as the final goal – e.g. in texture mapping – not as the initial step to build an auxiliary data structure. In order to construct our base mesh, we initially compute a global parametrization of the surface to then triangulate the parametrized planar domain to obtain a base mesh and associated atlas structure. The method is as follows: first, we cut the surface open into a simply connected domain with boundary by carefully placing seams into the mesh (cf. Figure 1(b)). Next, we use a modified angle-based flattening algorithm with boundary restrictions, where mesh vertices lying on boundary seams are treated as internal vertices except for a few vertices marked as cone singularities. Note that all vertices that are not located at cone singularities must respect the planarity constraint. As a consequence, the total curvature of the surface is distributed at the cone points.

While performing the seam-cut step, we create new vertices on the seam pointing to its original vertex. This simple pointer strategy is used to identify vertices on the parametrization that are the same on the original mesh. The edges on seams are also identified by the pointers on its vertices. An important remark is that although vertices can be replicated several times, e.g. in seams bifurcations, the edges are mapped pairwise on the parametrization border (cf. Figure 1(c)). This mapping is only possible because of a special subset of unconnected vertices on the seams called cone singularities, that is, vertices allowed to have angle distortion. Although these vertices do not respect the planarity (or angle-consistency) constraint, the edges around them respect the length-consistency constraint and, thus, the scale is preserved across boundaries.

The last step of our algorithm is to create and place a small number of base vertices over the global parametrization and triangulate them to construct a new coarser mesh. We compute the length distortion ratio for all edges and place the vertices on the least stretched areas (cf. Figure 1(d)). The angular distortion is used to avoid placing base vertices on top of (or near to) cone singularities. The triangulation is a Delaunay triangulation on the 2D parameter domain and uses the correspondence across boundaries stored on edges. The final base mesh is obtained by getting the base vertices from the parametrization back to 3D maintaining the triangulation.

Future Directions Our approach creates only base meshes, but it has the potential to create meshes at any resolution. An interesting future direction is to investigate how these meshes interact and define a set of production rules to describe multi-resolution meshes.

Acknowledgements We acknowledge the grant provided by the Brazilian agency CNPq and the AIM@SHAPE repository.

References

LEE, A., MORETON, H., AND HOPPE, H. 2000. Displaced Subdivision Surfaces. In *Proceedings of ACM SIGGRAPH '00*, 85–94.

SANDER, P. V., WOOD, Z. J., GORTLER, S. J., SNYDER, J., AND HOPPE, H. 2003. Multi-Chart Geometry Images. In *Proceedings of Eurographics/ACM SIGGRAPH SGP '03*, 146–155.

SIGGRAPH 2012, Los Angeles, California, August 5 – 9, 2012.
ISBN 978-1-4503-1435-0/12/0008