

Unconstrained Spherical Parameterization

Ilja Friedel

Peter Schröder
Caltech

Mathieu Desbrun

Abstract

We introduce a novel approach for the construction of spherical parameterizations based on energy minimization. The energies are derived in a *general manner* from classic formulations well known in the planar parameterization setting (*e.g.*, conformal, Tutte, area, stretch energies, *etc.*), based on the following principles: the energy should (1) be a measure of spherical triangles; (2) treat energies independently of the triangle location on the sphere; and (3) converge to the continuous energy *from above* under refinement. Based on these considerations we give a very simple non-linear modification of standard formulas that fulfills all these requirements. The method avoids the usual collapse of flat energies when they are transferred to the spherical setting without additional constraints (*e.g.*, fixing three or more vertices). Our *unconstrained* energy minimization problem is amenable to the use of standard solvers. Consequently the implementation effort is minimal while still achieving excellent robustness and performance through the use of widely available numerical minimization software.

1 Introduction

There is by now a rich literature on the construction of energy-based parameterizations for surface meshes (for a recent survey, we refer the reader to [Floater and Hormann 2005]). While much of this work has focused on the planar case, *i.e.*, the mapping of a topological disk region of a given mesh to the plane, spherical parameterizations have been singled out as a special case occurring frequently enough in practice to warrant their own methods [Gotsman et al. 2003; Gu et al. 2004; Haker et al. 2000; Sheffer et al. 2004; Praun and Hoppe 2003]. These approaches often extend a specific method known in the planar setting to the sphere: Praun and Hoppe [Praun and Hoppe 2003], for example, use the method of Sander *et al.* [Sander et al. 2001]. Unfortunately, the intrinsic non-linearity of these extensions coupled with the lack of boundary vertices to anchor the parameterization seriously hinder practical implementation: tailored solvers are often used in conjunction with several vertex constraints in order to obtain non-degenerate solutions. Alas, these fixes do not lead to optimal parameterizations as constraints often introduce severe distortion.

Consequently, our approach begins with the observation that it would be desirable to construct a general procedure to take existing methods (we are deliberately agnostic as to the particular “weights” being used) from the planar parameterization case and adapt them to the sphere. In this paper, we present a straightforward technique for easing the computation of spherical parameterizations by *a simple modification* of traditional planar parameterization methods: our spherical energies differ from the usual planar quadratic energies only in the multiplication by a simple factor based on the inverse distance d_{min} of each triangle (A, B, C) from the sphere center. Examples include the following:

$$\begin{aligned} E_{Tutte} &= d_{min}^{-2} \cdot ((x_A - x_B)^2 + (x_B - x_C)^2 + (x_A - x_C)^2) \\ E_{Dirichlet} &= d_{min}^{-2} \cdot (\cot(\alpha) \cdot (x_B - x_C)^2 + \cot(\beta) \cdot (x_A - x_C)^2 + \cot(\gamma) \cdot (x_A - x_B)^2) \\ E_{Area} &= d_{min}^{-2} \cdot \frac{Area(A, B, C)^2}{InputArea} \end{aligned}$$

More importantly, our approach allows a robust minimization of the resulting non-linear energy *without* any artificial constraints or custom solvers. This allows us to use standard, high-performance minimization software. We

demonstrate the practicality of our technique with a number of examples employing different energies on various meshes of significant size.

The main motivation for scaling the planar energy by d_{min}^{-2} is to obtain an *upper bound* of the spherical integrals. Intuitively this can be done by measuring the energy of each triangle *after* transforming it into the tangent space of the sphere. As we show next, *this extra term removes the usual need for repeated spherical rejections and for unnatural point constraints.*

2 Analysis

For most meshes, the factor d_{min}^{-2} is very close to 1. So why is such a minor correction necessary? Let us examine what happens if we minimize the classic flat energies

$$E_{spring} = w_{AB} \cdot (x_A - x_B)^2 + w_{BC} \cdot (x_B - x_C)^2 + w_{AC} \cdot (x_A - x_C)^2$$

with spring coefficients w 's as discussed in [Floater and Hormann 2005] and variables x mapped on the sphere. A sequence of non-linear minimization steps typically looks like in Figure 1. During this process, the energy is

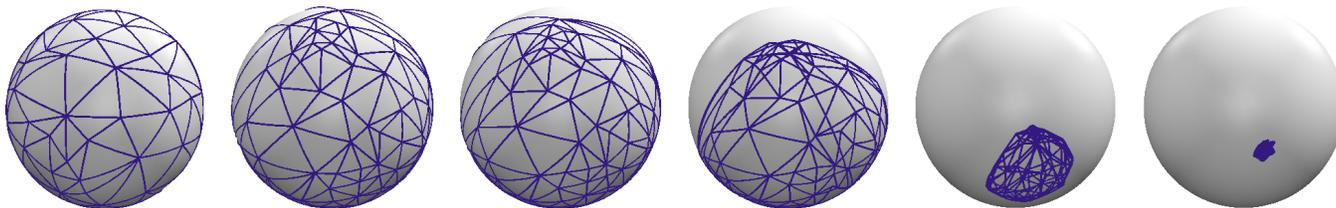
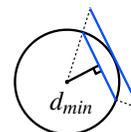


Figure 1: As the iterations proceed in the solver a triangle starts growing. Finally it slips over the “equator”, eventually shrinking the entire mesh to a point.

reduced in each iteration step, finally reaching its minimum at zero. This failure can be consistently observed. Our conclusion is that the spherical spring energy has no minimum at the expected configuration: instead the minimizer moves down a continuous slope leading to a collapsed configuration. Note that this behavior for the special case of the Dirichlet energy *cannot* be explained by the invariance of the parameterization to Möbius transformations [Gu et al. 2004]: while this energy has such a property in the continuous setting, this is no longer true in the discrete setting where the sphere is decomposed into simplicial cells.

One common fix to avoid a complete collapse consists of constraining three or more vertices. In practice the number of required vertices wildly varies with the input mesh. However, each additional constraint introduces extra distortion. For these reasons it is desirable to construct a method that does not require any additional constraints.

We decided to analyze the situation for various energies, like $E_{Dirichlet}$ and E_{Area} as cited in [Floater and Hormann 2005]. The planar energies *always underestimate* the integrals over spherical triangles. This is easy to see for energies based on areas: the area of a planar triangle cutting through the sphere is always smaller than the area of the corresponding spherical triangle. This is problematic, because the error increases disproportionately with triangle size. Used in a minimization process this is a recipe for disaster: the minimizer can find a way to decrease the energy by increasing the size of the triangle with the largest error, creating slippage. One way to avoid this situation is to account for the additional stretch which is caused by transforming planar into spherical triangles. This permits the design of spherical energies that are accurate for small triangles but otherwise *always overestimate* the continuous energy. We show that this can be achieved by using the central projection (or gnomonic map) which projects each flat triangle outwards until it is essentially tangential, as done by the division with d_{min} . A further consequence of this construction is the creation of infinite energy barriers for hemispherical triangles.



2.1 Spherical Dirichlet Energy

Pinkall and Polthier [Pinkall and Polthier 1993] wrote the Dirichlet energy for discrete conformal mappings between triangles as

$$E_{Dirichlet}(f) = \int_R \mathbf{trace}(Df^T Df). \quad (1)$$

The energy is an operator on maps $f : R \mapsto S$, that map the input mesh triangles R to domain triangles S . We denote with Df the matrix of partial derivatives (Jacobian) of f . In the work of Pinkall and Polthier the case of *planar* mesh triangles and *planar* domain triangles was considered. Consequently the map h was considered piecewise linear, with constant Jacobian over each triangle. Within this setting, rewriting the integral of the constant trace with respect to angles α, β and γ directly yields the celebrated cotangent weights.

However, for *spherical domain* triangles S^{sphere} , the energy has to be measured on the composed map $h = G^{-1} \circ f$, where $f : R \mapsto S^{flat}$ is the affine map from Pinkall and Polthier [Pinkall and Polthier 1993] and $G^{-1}(x) = \frac{x}{\|x\|}$ is the inverse of the gnomonic map. The choice of this map is not arbitrary, as it transforms spherical into planar triangles. (This is easily seen as it maps geodesics on the sphere into geodesics on planes [Snyder and Voxland 1989].)

An Upper Bound We can now deduce from the previous expression that:

$$\begin{aligned} E_{Dirichlet}(h) &= \int_R \mathbf{trace}(D(G^{-1} \circ f)^T D(G^{-1} \circ f)) \\ &\leq \int \mathbf{trace}\left(\frac{Df^T}{d_{min}} \cdot \frac{Df}{d_{min}}\right) = \frac{E_{Dirichlet}(f)}{d_{min}^2}. \end{aligned} \quad (2)$$

To obtain this inequality we used the monotonicity of the Dirichlet operator, and denoted as $d_{min} = \min\|f\|$ the minimum distance of triangle S^{flat} to the center of the unit sphere (see inset in previous paragraph). For acute triangles this minimum is achieved at the circumcenter of the triangle; for obtuse triangles the minimum distance is instead achieved at the midpoint of the longest edge. Finally rewriting Equation 2 as

$$E_{Dirichlet}(h) \leq d_{min}^{-2} \cdot (\cot(\alpha) \cdot (x_B - x_C)^2 + \cot(\beta) \cdot (x_A - x_C)^2 + \cot(\gamma) \cdot (x_A - x_B)^2) / 2$$

shows the familiar cotangent weights in our upper bound of the spherical Dirichlet energy.

Approximation Error Notice that we now have both lower *and* upper bounds on the spherical Dirichlet energy

$$E_{Dirichlet}(f) \leq E_{Dirichlet}(h) \leq d_{min}^{-2} \cdot E_{Dirichlet}(f).$$

The distance d_{min} is linked to the radius r_{mc} of the min-containment circle of the flat triangle by $d_{min}^2 = 1 - r_{mc}^2$. Using Taylor series expansion for small radii we get $d_{min}^{-2} = 1 + O(r_{mc}^2)$ with a *non-negative* error term. The implication is that the approximation error of the new operator is equivalent to methods using a flat triangle approximation; but $d_{min}^{-2} \cdot E_{Dirichlet}(f)$ will approach the real spherical energy *from above*. This simultaneously keeps min-containment circles sizes under control: as a triangle size approaches the entire hemisphere, $d_{min}^{-2} \rightarrow \infty$. This prevents collapses to the trivial solution even in the absence of any constraints.

2.2 Distortion Control

Many planar energies, like *MIPS* [Hormann and Greiner 2000] and *Stretch* [Sander et al. 2001], try to trade off angle and area distortion. It has been observed, for instance in [Floater and Hormann 2005], that the MIPS energy can be

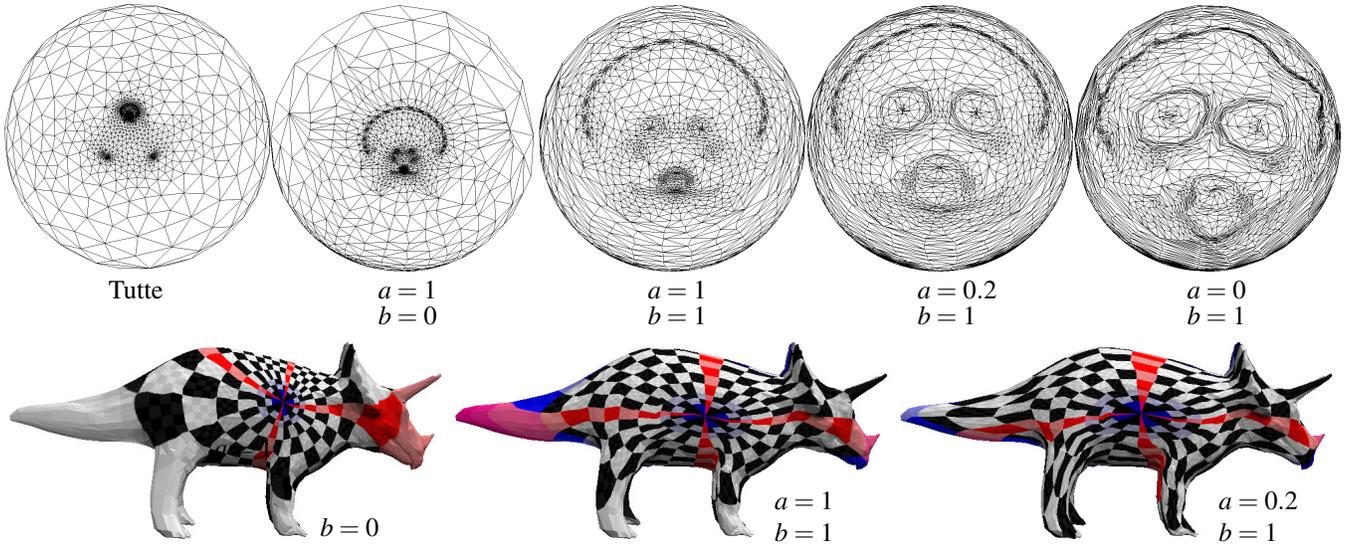


Figure 2: The top row shows the domain of parameterizing the triceratops model using the Tutte coefficients and the results of minimizing combined energies with weights as indicated. The lower row shows a texture map of the weighted parameterizations. Trade-offs are achieved between angle preservation (left) and the sampling density as can be observed on the extremities (right).

written as a combination of Dirichlet and area energies. We show that this is also true for the stretch energy (still using the notations defined in the previous paragraph):

$$\begin{aligned}
 E_{Stretch}(f) &= \mathbf{trace}(Df^{-T}Df^{-1}) \cdot \mathbf{area}(R) = \frac{\mathbf{trace}(Df^T Df)}{\mathbf{det}(Df^T Df)} \cdot \mathbf{area}(R) \\
 &= \frac{\mathbf{area}(R)^2}{\mathbf{area}(S)^2} \cdot \mathbf{trace}(Df^T Df) \cdot \mathbf{area}(R) = \frac{\mathbf{area}(R)^2}{\mathbf{area}(S)^2} \cdot E_{Dirichlet}(f).
 \end{aligned}$$

This expression explains how the stretch based parameterization strikes a balance between area and angle distortion. More importantly, we can apply our approach to obtain the upper-bound energy simply by substituting $d_{min}^{-2} \cdot E_{Dirichlet}$ for $E_{Dirichlet}$. Thus one can avoid the mesh refinement step used in [Praun and Hoppe 2003] and minimize the energy *without* requiring additional degrees of freedom.

For simplicity we decided to experiment with simple weighted averages of the energies we mentioned in Section 1, of the form:

$$E_{combined} = a \cdot E_{Dirichlet} + b \cdot E_{Area}.$$

Examples of balance between angle and area preservation are shown in Figure 2.

3 Results

Our approach based on upper bound approximation eliminates the need for constraints or reprojection during the minimization, which means that there is no need to write a custom solver: instead, one can directly use well-established numerical algorithms libraries. In our implementation, we used TAO [Benson et al. 2004], an existing solver that provides implementations of standard Newton and trust-region methods. We wrote the various parameterization energies in Maple, where they can be automatically differentiated and translated into C++ code. Note that to avoid negative cotangent weights in the case of the Dirichlet energy, we simply clamped all existing angles into the range of 5 to 85 degrees. (Alternatively one could use the intrinsic Laplace-Beltrami operator as described in [Fisher et al. 2006].) Our initial guess is taken as the central projection around the center of gravity of the input

mesh onto the unit sphere. Even with this simple, highly-distorted initialization, minimizing our energy approximations is efficient and the results is not prone to artificial distortion. In particular, while most of these initializations show multiple folds, all of our results are *fold-free* (without any attempt or need to enforce this). Run times for our implementation are typically within a few minutes; a small mesh ($< 10k$ triangles) requires less than a minute, a larger mesh like the *Igea* model (70k triangles) requires about 5 minutes.

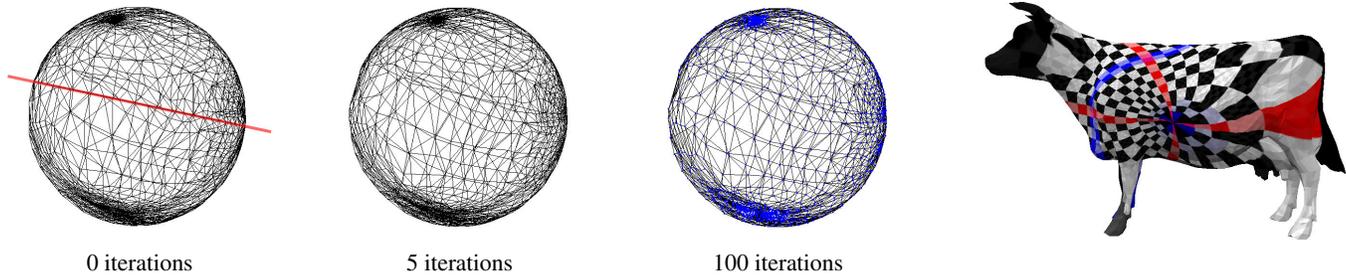


Figure 3: In Saba et al. [Saba et al. 2005] two planar parameterizations are stitched together and individual vertices moved according to “locally valid” Newton steps. The system is very stiff and the state after 5 iterations (center sphere) and 100 iterations (right sphere) is virtually identical. The general distribution of vertices (blue) does not change and remains clustered around the poles. Applying a texture to the model we observe that the resulting parameterization is locally reasonably angle preserving, but lacks this preservation across larger distances: the blue circle in the texture appears stretched and distorted.

Comparisons We performed several tests to compare our approach to existing methods. The only practical, direct approach to spherical parameterization of meshes for which code is available is that of [Saba et al. 2005]. In their approach the input mesh is cut into two “balanced” pieces of approximately equal size. Each of the two pieces is then independently parameterized into the plane, mapped back to the sphere, and stitched along the red line shown in Figure 3. A nonlinear Gauss-Seidel like sequence of “locally valid” Newton steps finally repositions individual vertices within their respective neighborhood. We observed that after about five such diffusion sequences the progress stops, and even 100 additional iterations do not visually change the vertex distribution. However, the vertex density in the vicinity of the (arbitrarily chosen) cut remains sparse, while the clusters around the poles (which are artifacts from the constrained planar parameterization) remain. Given that this behavior was consistently observed in all our tests, we believe that the optimization method presented in [Saba et al. 2005] fails to get close to a global energy minimum, even on small examples; the stiffness of the system to solve may well be the reason behind the numerical issues of their method, explaining the resulting large distortion shown in Figures 3 and 4. Note that our ability to derive energy functionals with a mix of area and angle conservation (Figure 2) also distinguishes our approach from previous work, as indicated in Figure 4.

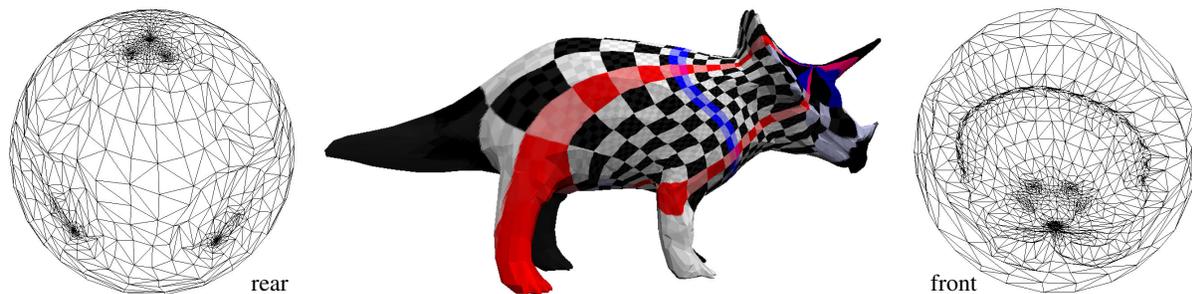


Figure 4: The method described by Saba et al. [Saba et al. 2005] offers no explicit distortion control. The parameterization of the triceratops head appears proportionate in size, even though slightly asymmetric in shape. The hind legs and the tail are mapped to a tiny area on the rear facing hemisphere (top three spots, the lower spots are the front legs). The texture mapped model shows clearly how this results in undersampling of these areas. Our method offers a continuous choice between area and angle preservation as shown in Figure 2.

Acknowledgments. The Lion-vase model is courtesy of SensAble Technologies, Inc.; all other models Cyberware, Inc.; Viewpoint and H. Hoppe. This work was partially supported by the NSF (ACI-0204932 & 0219979, DMS-0453145, CCF-

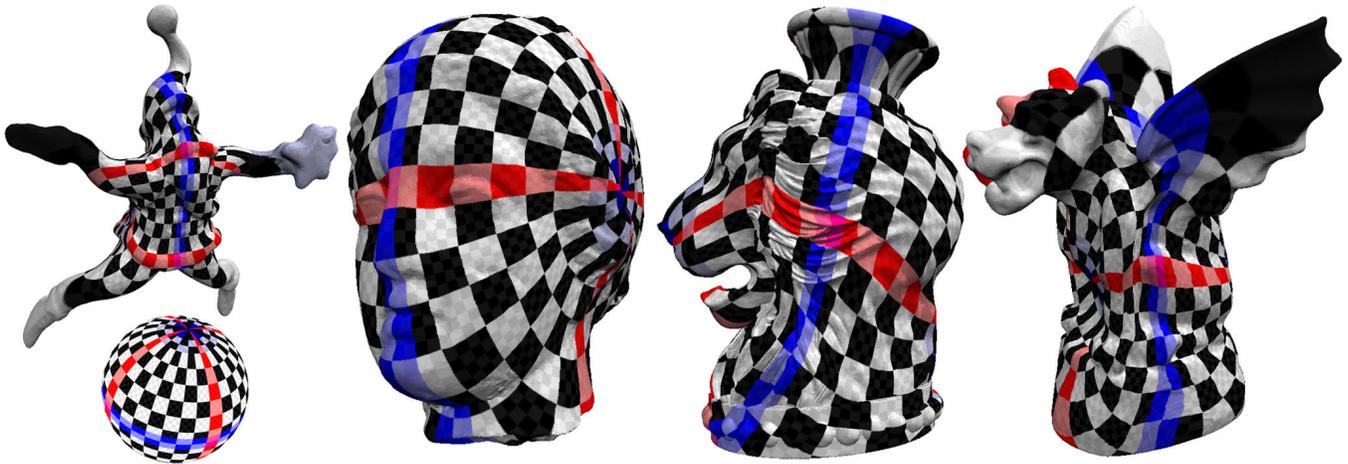


Figure 5: For our proposed method we computed parameterizations for several large (from 70k to 400k triangles) meshes using the combined energy $E_{combined}$ with weightings (1,1). No conditions were enforced during the solve, nevertheless the parameterizations are fold-free.

0503786 & 0528101), the DOE (DE-FG02-04ER25657, W-7405-ENG-48/B341492), the Okawa Foundation, Center for the Mathematics of Information, Autodesk, and Pixar Animation Studios. We would like to thank Shadi Saba for his library and the reviewers for their feedback.

References

- BENSON, S. J., MCINNES, L. C., MORÉ, J., AND SARICH, J. 2004. [TAO User Manual \(Revision 1.7\)](#). Tech. Rep. ANL/MCS-TM-242, Mathematics and Computer Science Division, Argonne National Laboratory. Available at <http://www-unix.mcs.anl.gov/tao>.
- FISHER, M., SPRINGBORN, B., BOBENKO, A. I., AND SCHRÖDER, P. 2006. [An Algorithm for the Construction of Intrinsic Delaunay Triangulations with Applications to Digital Geometry Processing](#). In *Discrete Differential Geometry*, E. Grinspun, M. Desbrun, and P. Schröder, Eds., Course Notes. ACM SIGGRAPH.
- FLOATER, M. S., AND HORMANN, K. 2005. [Surface Parameterization: a Tutorial and Survey](#). In *Advances on Multiresolution in Geometric Modelling*, N. Dodgson, M. S. Floater, and M. Sabin, Eds. Springer Verlag, Heidelberg.
- GOTSMAN, C., GU, X., AND SHEFFER, A. 2003. Fundamentals of spherical parameterization for 3d meshes. *ACM Transactions on Graphics* 22, 3 (July), 358–363.
- GU, X., WANG, Y., CHAN, T. F., THOMPSON, P. M., AND YAU, S.-T. 2004. Genus zero surface conformal mapping and its application to brain surface mapping. *IEEE Transaction on Medical Imaging* 23, 7, 949–958.
- HAKER, S., ANGENENT, S., TANNENBAUM, A., KIKINIS, R., SAPIRO, G., AND HALLE, M. 2000. Conformal surface parameterization for texture mapping. *IEEE Trans. on Vis. and Comp. Graphics* 6, 2, 181–189.
- HORMANN, K., AND GREINER, G. 2000. MIPS: An efficient global parametrization method. In *Curve and Surface Design: Saint-Malo 1999*, P.-J. Laurent, P. Sablonnière, and L. L. Schumaker, Eds., Innovations in Applied Mathematics. Vanderbilt University Press, Nashville, 153–162.
- PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces. *Experimental Mathematics* 2,1, 15–36.
- PRAUN, E., AND HOPPE, H. 2003. Spherical parametrization and remeshing. *ACM Trans. Graph.* 22, 3, 340–349.

- SABA, S., YAVNEH, I., GOTSMAN, C., AND SHEFFER, A. 2005. Practical Spherical Embedding of Manifold Triangle Meshes. *Proc. SMI 00*, 258–267.
- SANDER, P. V., SNYDER, J., GORTLER, S. J., AND HOPPE, H. 2001. Texture mapping progressive meshes. In *SIGGRAPH*, ACM Press, 409–416.
- SHEFFER, A., GOTSMAN, C., AND DYN, N. 2004. Robust spherical parameterization of triangular meshes. *Computing* 72, 1–2, 185–193.
- SNYDER, J. P., AND VOXLAND, P. M. 1989. *An Album of Map Projections*. U.S. Geological Survey.

Ilja Friedel, NVIDIA, 2701 San Thomas Expressway, Santa Clara, CA 95050

(ifriedel@nvidia.com)

Peter Schröder, Caltech, MS 256-80, Pasadena, CA 91125

(ps@cs.caltech.edu)

Mathieu Desbrun, Caltech, MS 256-80, Pasadena, CA 91125

(mathieu@cs.caltech.edu)