

Fundamentals of Spherical Parameterization for 3D Meshes

Craig Gotsman
gotsman@cs.technion.ac.il
Technion – Israel Inst. of Tech.

Xianfeng Gu
gu@eecs.harvard.edu
Harvard University

Alla Sheffer
sheffa@cs.technion.ac.il
Technion – Israel Inst. of Tech.

Abstract

Parameterization of 3D mesh data is important for many graphics applications, in particular for texture mapping, remeshing and morphing. Closed manifold genus-0 meshes are topologically equivalent to a sphere, hence this is the natural parameter domain for them. Parameterizing a triangle mesh onto the sphere means assigning a 3D position on the unit sphere to each of the mesh vertices, such that the spherical triangles induced by the mesh connectivity are not too distorted and do not overlap. Satisfying the non-overlapping requirement is the most difficult and critical component of this process. We describe a generalization of the method of barycentric coordinates for planar parameterization which solves the spherical parameterization problem, prove its correctness by establishing a connection to spectral graph theory and show how to compute these parameterizations.

CR Categories: I.3.5 [Computer Graphics] *Computational Geometry and Object Modeling – Curve, Surface and Solid and Object Representations.*

Keywords: Triangle mesh, parameterization, embedding.

1. Introduction

Parameterization of 3D mesh data is important for many graphics applications, in particular for texture mapping, remeshing and morphing. To date, mostly planar parameterizations have been considered. The main challenge is to produce a planar triangulation that best matches the geometry of the 3D mesh, minimizing some measure of *distortion*, yet is still valid. In this context, valid means that the individual planar triangles *do not overlap*. Most of the recent works on the subject of parameterization (e.g. [Desbrun et al. 2002; Floater 1997; Levy et al. 2002; Sander et al. 2001; Sheffer and de Sturler 2001]) have focused on defining the distortion, and showing how to minimize it.

While parameterizing to the plane is the most natural way to perform texture-mapping, this is less natural for other mesh processing operations which also require a parameterization. For applications such as morphing [Alexa 2000; Kanai et al. 2000; Shapiro and Tal 1998] and remeshing [Gu et al. 2002; Kobbelt 1999] it is best to parameterize the mesh over a domain which is topologically equivalent to it. This significantly reduces the distortion introduced by the parameterization without resorting to methods which introduce other artifacts, such as cutting seams.

If the mesh has the topology of a sphere, it is best to use a spherical parameter domain. Parameterizing a 3D triangle mesh over the sphere is equivalent to *embedding* its connectivity graph on the sphere, such that the resulting *spherical triangles* partition the sphere (their union is the sphere, and they are disjoint). A classical result due to Steinitz is that a graph may be embedded on the sphere if and only if it is planar and 3-connected. Thus a closed manifold genus-0 triangulation can always be mapped to a *spherical triangulation*.

The simplest way to map a closed triangle mesh to the sphere is to reduce the problem to the planar case. First cut out one triangle to serve as a boundary. Then parameterize the resulting open mesh over the unit triangle using any planar parameterization method, and finally use the inverse stereo projection to map the plane to the sphere [Haker et al. 2000]. See Fig. 1(b). The main problem with this method is severe distortion, and although the inverse stereo projection is conformal, namely, preserves angles in the continuous case, it does not preserve angles (or any other geometric properties) in the discrete case. The projection also does not guarantee that the result will be a spherical triangulation.

Another straightforward method to parameterize to the sphere is to cut the mesh into two pieces, each topologically equivalent to a disk, parameterize each over a planar disk with a common boundary, and then map each disk to a hemisphere (by adding an appropriate z component to each vertex). The common boundary guarantees that the two hemispheres fit together at the equator. See Fig. 1(c). Since this boundary will presumably contain more than just three vertices, each of the two disk parameterizations will be less distorted than the one obtained by using a single triangle as the boundary, so the spherical result will also be less distorted. However, the result will depend strongly on the specific cut used to obtain the two disks.

It is more natural to parameterize a mesh directly on the sphere without going back and forth to the plane. Several methods for direct parameterization on the sphere exist. The only one to date that seems to guarantee a valid spherical triangulation (i.e. with no triangle foldovers) is that of Shapiro and Tal [1998], similar to that of Das and Goodrich [1997]. This method works by simplifying the mesh by vertex removal until only a tetrahedron remains.

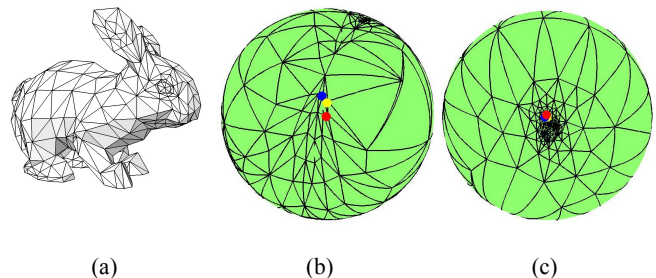


Figure 1: Parameterizing the (a) rabbit by (b) inverse stereo mapping (c) two hemispheres. Colored dots mark corresponding vertices.

Permission to make digital/hard copy of part of all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2003 ACM 0730-0301/03/0700-0358 \$5.00

The tetrahedron is easily embedded on the sphere, and then the vertices are inserted back one by one, so that the validity of the triangulation is preserved throughout the process. While this is quite an efficient process, it is difficult to optimize the parameterization, due to its greedy nature, and impossible to steer it to have any desirable mathematical properties. Other direct parameterization methods were proposed by Kobbelt et al [1999] and Alexa [2000]. These are heuristic iterative procedures, attempting to converge to a valid parameterization by applying local improvement (relaxation) rules. These work well in many cases, but there is no guarantee that they will terminate, and, even if they do, that the resulting embedding will be valid, or have any desirable mathematical properties. A method which guarantees a valid embedding was recently proposed by Sheffer et al. [2003]. This is a highly non-linear optimization procedure, working with the angles of the spherical triangulation (as opposed to the vertex positions), inspired by the angle-based method of Sheffer and de Sturler [2001] for planar parameterizations. So far it lacks an efficient numerical computation procedure, so it is not very practical.

1.1 Our contribution

The problem of mesh parameterization is that of mapping a piecewise linear surface with a discrete representation onto a continuous spherical surface. The theory of mappings between various Riemann surfaces is well understood in the continuous case using classical differential geometry [Do Carmo 1976]. Probably the most notable example of this is the so-called conformal mapping theory which shows how to map any continuous Riemann surface to another such that angles are preserved. However, the discrete case of meshes is much less understood. In the limit it obviously converges to the continuous case, but in practical applications the meshes involved may be far from this limit. Hence there is a need to treat the discrete case separately in a combinatorial manner, albeit inspired by the classical theory.

This paper introduces a precise mathematical characterization of all possible spherical parameterizations of a closed manifold genus-0 triangle mesh. We show that it is a natural non-linear extension of the linear theory of barycentric coordinates used in the planar case. The correctness of this methodology is proved by establishing a link to the so-called Colin de Verdiere matrices associated with planar graphs. We also describe a computational method for generating and controlling these parameterizations. These contributions are concentrated in Section 4, after establishing the theory in Section 3.

2. The Method of Barycentric Coordinates

2.1 The planar case

Floater [1997] described a generic method to embed a manifold 3D mesh with a boundary in the plane without foldovers. Floater's method is a generalization of the basic procedure originally proposed by Tutte [1963] for a planar graph, which can be traced back as far as I. Fary in 1948 and J.C. Maxwell in 1864. This method makes use of so-called *barycentric coordinates* (or *convex combinations*) and proceeds as follows:

1. To each interior (directed) edge $e = (i, j)$, assign a positive weight w_{ij} , such that

$$\sum_{j \in N(i)} w_{ij} = 1$$

where $N(i)$ is the list of vertices neighboring the i 'th vertex.

2. To all other entries (i, j) , assign $w_{ij} = 0$.
3. Embed the boundary vertices in the plane such that they form a closed convex polygon.
4. Solve the following two linear systems for the x and y coordinates of the n interior vertices: $(I - W)x = b_x$, $(I - W)y = b_y$, where W is a nxn matrix containing w_{ij} , and b_x and b_y vectors with non-zero entries corresponding to vertices adjacent to the boundary.

The cornerstone of this theory is the following theorem, first proven by Tutte [1963], and reproven over the years in different ways (e.g. [Floater 2003b, Richter-Gebert 1996, Chap 3]):

Theorem 1: Given a planar 3-connected graph with a boundary fixed to a convex shape in \mathbb{R}^2 , the positions of the interior vertices form a planar triangulation (i.e. none of the triangles overlap) if and only if each vertex position is some convex combination of its neighbor's positions.

Theorem 1 implies that the method of barycentric coordinates generates all possible valid embeddings of the graph in the plane, given the (convex) positions of the boundary. Tutte proposed using $w_{ij} = 1/\text{deg}(i)$ for all edges (i, j) , in effect placing each interior vertex at the centroid of its neighbors ($\text{deg}(i)$ is the degree, or valence of the i 'th vertex). This choice of W does not take into account the geometry of the mesh, just its connectivity. When the mesh is given with 3D geometry, a number of recipes for W have been proposed, each aiming for some effect related to reflecting the geometry of the mesh in the parameterization, namely, minimizing its metric distortion when flattened to the plane. The most popular methods seem to be the *shape-preserving* method [Floater 1997], the *conformal*, or *harmonic* method [Haker et al. 2000; Levy et al. 2002; Pinkall and Polthier 1993] and the *mean-value* method [Floater 2003a]. These methods all have the desirable *2D reproduction property* [Floater 1997], namely that when applied to a 2D triangulation, the embedding procedure will produce an output identical to the input. However, some of the methods, most notably the conformal method, do not always result in positive weights, hence cannot guarantee a valid embedding.

In general, the method of barycentric coordinates may be formulated as the solution to the 2D vector Laplace equation on the interior vertices:

$$(1) \quad L_W x = b$$

with boundary conditions derived from the convex boundary vertex positions, which prevents the trivial zero solution. This is a *linear* system. $L_W = I - W$ is the *general normalized Laplacian* operator, general because the weights of W are arbitrary positive values, and normalized because the rows of W all sum to unity. The special case of $w_{ij} = 1/\text{deg}(i)$ proposed by Tutte will be called the normalized *Tutte Laplacian*.

A simple numerical procedure to solve (1) is a relaxation procedure, where the boundary vertices are placed on a convex boundary, and the interior vertices are repeatedly updated to be at the weighted average of their neighboring vertex positions, as dictated by W . Since L_W is diagonally dominant, this Gauss-Seidel procedure is guaranteed to converge.

2.2 The spherical case

The Laplacian L_W has a unit diagonal, negative entries for each mesh edge, and vanishes otherwise. Also, all rows sum to zero, hence L_W is singular. L_W is, however, not symmetric. In what follows, we restrict the discussion to the class of *symmetric* Laplacians, which corresponds to sets of barycentric coordinates which are edge-symmetric up to normalization of each row. These are:

$$L_W(i, j) = \begin{cases} \text{negative number} & (i, j) \in E \\ -\sum_{k \neq i} L_W(i, k) & i = j \\ 0 & (i, j) \notin E \end{cases}$$

Note that the symmetric Tutte Laplacian has -1's at entries corresponding to edges, and the vertex degrees along the diagonal. Symmetric systems such as these can be given the physical interpretation of a mass-spring system at rest, where the vertices are point masses joined by springs along the edges. In this case, the Laplace equations are just the normal equations for minimizing the quadratic spring energy. Tutte and conformal barycentric coordinates have the symmetry property, but mean value coordinates unfortunately do not.

Generalizing the barycentric coordinates theory to spherical embeddings is not straightforward. Being non-planar, it will be impossible in general to express a vertex on the sphere as a convex combination of its neighbors (e.g. if a vertex's neighbors are all co-planar, this will imply that the vertex should also be on the same plane).

Inspired by classical differential geometry operator theory, Gu and Yau [2002] proposed to embed on a curved 3D surface using the generalization of the Laplacian to the Laplace-Beltrami operator. Intuitively, this is just the tangential component of the Laplacian at that point of the surface, and implies the following *non-linear* system:

$$(2) \quad L_W^\parallel x = 0 \quad \text{s.t.} \quad \|x_i\| = 1, \quad i = 1, \dots, n$$

where L^\parallel is the tangential component of L . The right hand side of (2) is zero, as opposed to the non-zero b in (1), because there is no boundary.

While Gu and Yau showed that solving (2) results in a bijective embedding of a continuous Riemann surface on the sphere, they did not show that this also holds for the discrete case of a piecewise-linear mesh, in the sense that the result is a valid spherical triangulation. In the next section we show how some recent deep results in spectral graph theory may be applied to establish this.

3. Connection to Spectral Graph Theory

We would like to prove the following analog of Theorem 1:

Theorem 2: *Given a planar 3-connected graph embedded in \mathbb{R}^3 , the positions of the vertices form a spherical triangulation (i.e. none of the spherical triangles overlap) if and only if each vertex position is some convex combination of the positions of its neighbors, which is then projected on the sphere.*

Theorem 2 means that the barycentric coordinate theory holds also on the sphere up to a radial residual, consistent with (2). In Section 4 we will prove this theorem. We start with some theory.

3.1 The Colin de Verdiere number

In 1990, Colin de Verdiere [1990] established an algebraic invariant for certain families of graphs. Given a n -vertex graph $G = \langle V, E \rangle$, consider the class $M(G)$ of symmetric matrices with element M_{ij} such that:

$$M_{ij} = \begin{cases} \text{negative number} & (i, j) \in E \\ \text{anything} & i = j \\ 0 & (i, j) \notin E \end{cases}$$

Note that $M(G)$ is a superset of the symmetric Laplacians for G , allowing the diagonal entries to assume arbitrary values (so that the rows do not necessarily sum to zero). Denote by $\lambda(M) = \{\lambda_0, \dots, \lambda_{n-1}\}$ the spectrum of M with corresponding eigenvectors $\{\xi_0, \dots, \xi_{n-1}\}$. Let $r = r(G)$ be the maximal integer such that $\lambda_1 = \lambda_2 = \dots = \lambda_r$ over all matrices in $M(G)$. Let M be a matrix which attains this maximum. This $r(G)$ is called the *Colin de Verdiere (CdV) number* of G , the matrix M a CdV matrix for G , its r identical eigenvalues CdV eigenvalues, and the corresponding eigenvectors CdV eigenvectors. Colin de Verdiere showed that:

G is a 3-connected planar graph if and only if $r(G) \leq 3$.

3.2 Nullspace embedding

An important extension of the results of Colin de Verdiere was obtained by Lovasz and Schrijver [1999], who showed that CdV eigenvectors of a graph G may be used to embed G in \mathbb{R}^r . For the special case $r(G) = 3$, this translates to:

G describes the edges of a convex polyhedron in \mathbb{R}^3 containing the origin if the three eigenvectors ξ_1, ξ_2 and ξ_3 of a CdV matrix of G are used as coordinate vectors for its vertices.

The fact that the polyhedron is convex and contains the origin is a key fact, since convexity implies star-shapedness. This in turn implies that by normalizing the vertices, the polyhedron may be projected onto the unit sphere to form a (valid) spherical triangulation.

Since the spectrum of a matrix may be shifted arbitrarily by adding an appropriate constant to the diagonal entries, we may assume, without loss of generality, that the three CdV eigenvalues are zero. Thus the corresponding CdV matrix has just one negative eigenvalue and co-rank 3. In this case the three CdV eigenvectors are independent non-trivial solutions to what looks like Laplace equations: $Mx = 0$, or a basis of the *nullspace* of M . Hence x is called the *nullspace embedding* of G .

Using eigenvectors of matrices as coordinate vectors for embedding graphs is not new. The traditional way of doing this is taking the eigenvectors corresponding to the smallest positive eigenvalues of the Tutte Laplacian (the smallest eigenvalue is zero, due to this matrix being singular). This dates back to Fiedler [1975] and Hall [1970]. See [Koren 2002] for a discussion of the different ways of using the Laplacian for "drawing" graphs. Eigenvalues and eigenvectors of the Tutte Laplacian of a graph are the cornerstone of spectral graph theory [Chung 1997], and have also been used for coding 3D mesh geometry [Karni and Gotsman 2000]. However, the embeddings resulting from eigenvectors of this Laplacian do not have very appealing geometric properties, and, specifically, the triangles overlap. The Colin de Verdiere theory reveals that more powerful generalizations of the Laplacian must

be used, yielding eigenvectors which are more “symmetric”, since they correspond to identical eigenvalues. On an intuitive level, this symmetry is what guarantees the validity of the embedding.

4. Generating Spherical Nullspace Embeddings

It is difficult to use the Colin de Verdiere theory directly to embed on a sphere, since, given a 3-connected planar graph G , neither Colin de Verdiere nor Lovasz and Schrijver provided any recipe to generate a CdV matrix for G .

In the planar case, once the boundary of the triangulation has been fixed and the barycentric coordinates chosen, the positions of the interior vertices are uniquely determined by solving a linear system. This is not the case for the spherical scenario. However, we propose to use a symmetric Laplacian as the starting point for constructing a CdV matrix. The off-diagonal values will not change, but the diagonal of the matrix is still lacking, and must be corrected. Only then can the embedding (the nullspace of the CdV matrix) be obtained.

The key observation of this paper is that we may solve for the diagonal of the CdV matrix and its nullspace *simultaneously*. We also force the resulting nullspace vectors to lie vertex-wise on the 3D sphere. This may be posed as the following set of $4n$ quadratic equations on the $3n$ positions of the vertices (x_i, y_i, z_i) and the n auxiliary variables α_i :

$$(3) \quad \begin{aligned} x_i^2 + y_i^2 + z_i^2 &= 1 & i &= 1, \dots, n \\ \alpha_i x_i - L_W[i]x &= 0 & i &= 1, \dots, n \\ \alpha_i y_i - L_W[i]y &= 0 & i &= 1, \dots, n \\ \alpha_i z_i - L_W[i]z &= 0 & i &= 1, \dots, n \end{aligned}$$

$L_W[i]$ denotes the i 'th row of the matrix L_W , and $x, y,$ and z are column vectors. The number of vertices in the mesh is n .

Fig. 2 illustrates the geometric interpretation of (3): the vector difference between the i 'th vertex and the weighted average (as dictated by L_W) of its neighbors is collinear with the vector difference between the vertex and the sphere's center.

To prove that this procedure is correct, assume that L is a symmetric Laplacian for G and that (3) has been solved for column n -vectors x, y, z and α . This means that the i 'th row $L[i]$ of L satisfies $L[i](x, y, z) = \alpha_i (x_i, y_i, z_i)$. Define the matrix M as:

$$M_{ij} = \begin{cases} L_{ij} & i \neq j \\ L_{ii} - \alpha_i & i = j \end{cases},$$

obtaining $M(x, y, z) = 0$. Hence M is a CdV matrix for G with vanishing CdV eigenvalue and nullspace spanned by x, y and z , implying that x, y, z form a valid spherical triangulation when used as coordinate vectors.

A similar argument shows the converse: If the three vectors (x, y, z) are a spherical triangulation of a 3-connected planar graph G , then these three vectors span the nullspace of a suitable CdV matrix M , which may easily be translated into L_W and α satisfying (3). This completes the proof of Theorem 2.

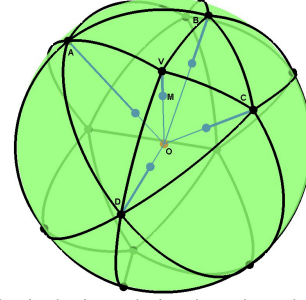


Figure 2: A spherical triangulation based on the Tutte Laplacian: The average of the neighbors of vertex V (A, B, C and D) is the point M , which is collinear with the sphere center O and the vertex V : $M-O = c(V-O)$. A similar relationship between a vertex and its neighbors holds for all mesh vertices.

5. Implementation Details

Steinitz's theorem guarantees that any planar 3-connected mesh admits a valid spherical triangulation. The theory of Colin de Verdiere guarantees that if a valid spherical triangulation exists, then it can be found by solving a system of the form (3) for *some* symmetric Laplacian L_W . However, there is no guarantee that an arbitrary symmetric Laplacian, when used in (3), will result in a non-degenerate triangulation.

Degenerate solutions always exist. The case $\alpha \equiv 0$ is the trivial solution when all vertices collapse to one point on the sphere. Gu and Yau [2002] tried to prevent this by requiring that the vertices average to zero. However, this is damaging and can prevent other solutions. Another degenerate solution can occur when the vertices are mapped to two antipodal points on the sphere. In this case, the vertices are partitioned into two sets such that the weighted average of the neighbors of a vertex in each set is still in the same hemisphere as the vertex. A more interesting situation can occur if the connectivity graph contains a Hamiltonian cycle. The cycle of vertices may then be mapped to the equator.

Beyond the degenerate solutions, a solution to (3) is not unique, since any spherical triangulation is certainly invariant to the rotation group $SO(3)$, which immediately gives three degrees of freedom. However, there are more. This can be seen by examining the simple special case of a tetrahedral mesh connectivity combined with the Tutte Laplacian. Observe that any rectangular parallelepiped circumscribed by the unit sphere defines an *equifaced* tetrahedron (i.e. a tetrahedron whose four faces are identical), by taking its edges to be the opposite diagonals of the six faces of the parallelepiped. It is easy to see that any vertex of the tetrahedron is collinear with the origin and the centroid of the three opposite vertices. This means that the vertex locations on the sphere satisfy (3) with $L_W =$ the Tutte Laplacian and $\alpha_i = 4$ for $1 \leq i \leq 4$. The CdV matrix is a 4×4 matrix whose entries are all -1's. There are two degrees of freedom in constructing such a rectangular parallelepiped, so all together there are five degrees of freedom.

To solve the nonlinear set of equations (3) for very large meshes, it is important to have a stable and efficient numerical procedure. A relaxation procedure where each vertex is updated to be the weighted average of its neighbors, and then projected onto the sphere, will not converge to the desired result, rather collapse to some degenerate configuration.

We use the *fsolve* procedure of MATLAB, a subspace trust region procedure [Coleman and Li 1996]. To better condition the system, it is useful to anchor an arbitrary vertex to a fixed position on the sphere. This eliminates only two degrees of freedom from $SO(3)$, so is not damaging. It appears that anchoring two arbitrary vertices limits the solution. Anchoring three vertices is damaging, since in the case of the equifaced tetrahedron there does not always exist a fourth vertex on the sphere forming the tetrahedron.

6. Experimental Results

We have used our methodology to generate a variety of embeddings of closed manifold genus-0 meshes on the unit sphere. The characteristics of the embedding may be controlled by the weights in the Laplacian matrix, similarly to the planar case. For example, uniform (Tutte) weights should be used if the spherical triangulation is required to be equi-angular, cotangential weights for a

conformal angle-preserving mapping (although these can be negative), inverse edge lengths for a triangulation preserving edge lengths, and the prescription of Desbrun et al. [2002] for an area-preserving triangulation.

Fig. 3 shows some spherical embeddings generated by our procedure on three sample meshes, and compares them to those generated by the procedure of Alexa [2000], and those based on reduction to the planar case by inverse stereo projection. The Tutte Laplacian produces an equi-angular triangulation, which is similar in some cases to that generated by Alexa's algorithm. Alexa's algorithm, though, sometimes tends to spread the vertices out over the sphere (as in the Triceratops model). The conformal embedding preserves many of the features of the 3D geometry, so the eyes and ears of the Rabbit are still noticeable in the result. The stereo embedding tends to lose most of the geometric structure.

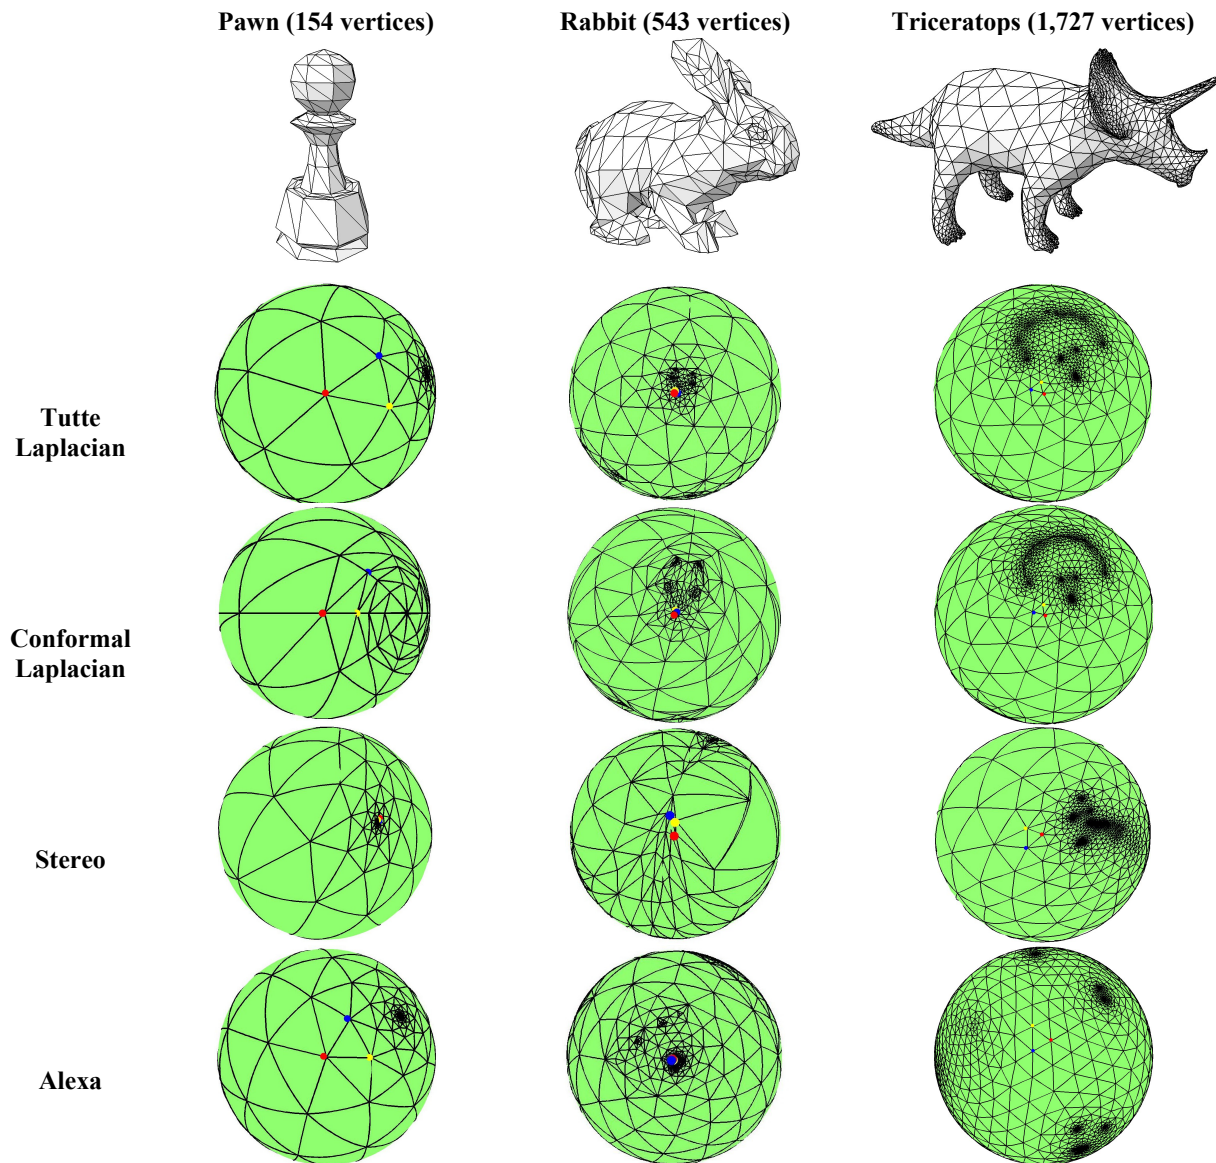


Figure 3: Some 3D models and their spherical parameterizations. Colored dots mark corresponding vertices.

In terms of runtimes, solving (2) using the MATLAB procedure required anywhere between a few seconds for the Pawn model of 154 vertices, and a few minutes for the Triceratops model (1,727 vertices) on a Xeon 2.8 GHz PC with 1GB RDRAM.

7. Conclusion

Parameterizing a closed manifold genus-0 mesh to the sphere is of paramount importance in digital geometry processing. It is a fundamental operation required for remeshing, morphing, filtering and texture mapping. While mapping between Riemann surfaces is well understood in the continuous case, the discrete case has not received satisfactory treatment to date. This paper has closed this gap, providing precise characterizations of discrete spherical parameterizations and methods to compute them.

We have shown that there is a natural extension of the barycentric coordinate theory from planar triangulations to spherical triangulations, corresponding to the extension of the Laplacian operator to the Laplace-Beltrami operator in differential geometry. Unfortunately, the extension involves a transition from a linear theory to a non-linear theory, so is much more difficult to analyze and compute.

The extension of the theory of continuous Riemann surfaces to a combinatorial treatment of discrete triangle meshes has been made based on recent results in algebraic and spectral graph theory. This translates to a set of equations which may be solved with not too much difficulty.

A few questions remain open, most notably on the existence of non-degenerate solutions and the analysis of the degrees of freedom in the various spherical embeddings, and how to control (or eliminate) them. Extensions to higher genus is also interesting.

Acknowledgements

Thanks to Zachi Karni for helpful discussions on CdV matrices. This research was partially funded by European grant HPRN-CT-1999-00117 (MINGLE), German-Israel Fund grant I-627-45.6/1999 and Israel Ministry of Science grant 01-01-01509.

References

- ALEXA, M. 2000. Merging Polyhedral Shapes with Scattered Features. *The Visual Computer* 16, 1, 26-37.
- CHUNG, F.R.K. 1997. Spectral Graph Theory. *CBMS 92*, AMS.
- COLEMAN, T.F., AND LI, Y. 1996. An Interior Trust Region Approach for Nonlinear Minimization Subject to Bounds. *SIAM Journal on Optimization*, 6, 418-445.
- COLIN DE VERDIERE, Y. 1990. Sur un Nouvel Invariant des Graphes et un Critere de Planarite. *Journal of Combinatorial Theory B* 50, 11-21. [English translation: On a New Graph Invariant and a Criterion for Planarity. In Graph Structure Theory. 1993. (N. Robertson, P. Seymour, Eds.) *Contemporary Mathematics*, AMS, 137-147.]
- DAS, G., AND GOODRICH, M.T. 1997. On the Complexity of Optimization Problems for 3-Dimensional Convex Polyhedra and Decision Trees. *Computational Geometry*, 8, 123-137.
- DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic Parameterizations of Surface Meshes. *Computer Graphics Forum*, 21, 3, 210-218.
- DO CARMO, M.P. 1976. *Differential Geometry of Curves and Surfaces*. Prentice-Hall.
- FIEDLER, M. 1975. A Property of Eigenvectors of Nonnegative Symmetric Matrices and Its Application to Graph Theory. *Czechoslovak Math. Journal*, 25, 619-633.
- FLOATER, M.S. 1997. Parameterization and Smooth Approximation of Surface Triangulations. *Computer Aided Geometric Design*, 14, 231-250.
- FLOATER, M.S. 2003. Mean-value Coordinates. *Computer Aided Geometric Design*, 20, 19-27.
- FLOATER, M.S. 2003. One-to-one Piecewise Linear Mappings Over Triangulations. *Mathematics of Computation* 2, 685-696.
- GU, X., AND YAU, S.-T. 2002. Computing Conformal Structures of Surfaces. *Communications in Information and Systems*, 2, 2, 121-146.
- GU, X., GORTLER, S., AND HOPPE, H. 2002. Geometry Images. *ACM Transactions on Graphics*, 21, 3, 355-361.
- GUSKOV, I., VIDIMCE, K., SWELDENS, W., AND SCHROEDER, P. 2000. Normal Meshes. In *Proceedings of ACM SIGGRAPH 2000*, ACM Press/ ACM SIGGRAPH, New York, K. Akeley, Ed., Computer Graphics Proceedings, Annual Conferences Series, ACM, 95-102.
- HAKER, S., ANGENENT, S., TANNENBAUM, A., KIKINIS, R., AND SAPIRO, G. 2000. Conformal Surface Parameterization for Texture Mapping. *IEEE Transactions on Visualization and Computer Graphics*, 6, 2, 1-9.
- HALL, K.M. 1970. An r-dimensional Quadratic Placement Algorithm. *Management Science*, 17, 219-229.
- KANAI, T., SUZUKI, H., AND KIMURA, F. 2000. Metamorphosis of Arbitrary Triangular Meshes. *IEEE Computer Graphics and Applications*, 20, 2, 62-75.
- KARNI, Z., AND GOTSMAN, C. Spectral Compression of Mesh Geometry. In *Proceedings of ACM SIGGRAPH 2000*, ACM Press / ACM SIGGRAPH, New York, K. Akeley, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM, 279-286.
- KOBBELT, L.P., VORSATZ, J., LABISK, U., AND SEIDEL, H.-P. 1999. A Shrink-wrapping Approach to Remeshing Polygonal Surfaces. *Computer Graphics Forum*, 18, 3, 119-129.
- KOREN, Y. 2001. On Spectral Graph Drawing. Preprint, Weizmann Institute of Science.
- LEVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. 2002. Least Squares Conformal Maps for Automatic Texture Atlas Generation. *ACM Transactions on Graphics*, 21, 3, 362-371.
- LOVASZ, L., AND SCHRIJVER, A. 1999. On the Nullspace of a Colin de Verdiere Matrix. *Annales de l'Institut Fourier* 49, 1017-1026.
- PINKALL, U., AND POLTHIER, K. 1993. Computing Discrete Minimal Surfaces and Their Conjugates. *Experimental Mathematics*, 2, 15-36.
- RICHTER-GEBERT, J. 1996. Realization Spaces of Polytopes. *Lecture Notes in Math #1643*, Springer.
- SANDER, P.V., SNYDER, J., GORTLER S.J., AND HOPPE, H. 2001. Texture Mapping Progressive Meshes. In *Proceedings of ACM SIGGRAPH 2001*, ACM Press/ ACM SIGGRAPH, New York, E. Fiume, Ed., Computer Graphics Proceedings, Annual Conferences Series, ACM, 409-416.
- SHAPIRO A., AND TAL, A. 1998. Polygon Realization for Shape Transformation. *The Visual Computer*, 14, 8-9, 429-444.
- SHEFFER, A., GOTSMAN C., AND DYN, N. 2003. Robust Spherical Parameterization of Triangular Meshes. In *Proceedings of 4th Israel-Korea Binational Workshop on Computer Graphics and Geometric Modeling*, Tel Aviv, 94-99.
- SHEFFER, A. AND DE STURLER, E. 2001. Parameterization of Faceted Surfaces for Meshing Using Angle Based Flattening. *Engineering with Computers*, 17, 3, 326-337.
- TUTTE, W.T. 1963. How to Draw a Graph. *Proc. London Math. Soc.* 13, 3, 743-768.