

Hexagonal Global Parameterization of Arbitrary Surfaces

Matthias Nieser, Jonathan Palacios, Konrad Polthier, and Eugene Zhang, *Senior Member, IEEE*

Abstract—We introduce *hexagonal global parameterization*, a new type of surface parameterization in which parameter lines respect sixfold rotational symmetries (6-RoSy). Such parameterizations enable the tiling of surfaces with nearly regular hexagonal or triangular patterns, and can be used for triangular remeshing. Our framework to construct a hexagonal parameterization, referred to as HEXCOVER, extends the QUADCOVER algorithm and formulates necessary conditions for hexagonal parameterization. We also provide an algorithm to automatically generate a 6-RoSy field that respects directional and singularity features in the surface. We demonstrate the usefulness of our geometry-aware global parameterization with applications such as surface tiling with nearly regular textures and geometry patterns, as well as triangular and hexagonal remeshing.

Index Terms—Surface parameterization, rotational symmetry, hexagonal global parameterization, triangular remeshing, pattern synthesis on surfaces, texture synthesis, geometry synthesis, regular patterns.



1 INTRODUCTION

IN this paper, we introduce *hexagonal global parameterization*, a new type of global parameterization that maps a surface onto the plane so that hexagonal or triangular patterns in this plane map seamlessly back onto the surface at all but a finite number of singular points. Such parameterizations facilitate regular pattern synthesis on surfaces and triangular remeshing.

1.1 Pattern Synthesis

Regular hexagonal patterns are one of the three regular patterns that can seamlessly tile a plane. They provide an optimal approximation to circle packings [1] which have been linked to the wide appearance of hexagonal patterns in nature, such as honeycombs, insect eyes, fish eggs, and snow and water crystals, as well as in man-made objects such as floor tiling, carpet patterns, and architectural decorations (Fig. 1).

Tiling a surface with regular texture and geometry patterns is an important yet challenging problem in pattern synthesis [2], [1]. Methods based on some local parameterization of the surface often lead to visible breakup of the patterns along seams, i.e., where the surface is cut open during parameterization. Global parameterizations can alleviate this problem when the translational and rotational discontinuity in the parameterization is compatible with the

tiling pattern in the input texture and geometry. For example, a quadrangular global parameterization is designed to be compatible with square patterns (Fig. 2a) and is incompatible with hexagonal patterns (Fig. 2b). In contrast, a hexagonal global parameterization is compatible with hexagonal or triangular patterns (Fig. 2c).

1.2 Remeshing

Another motivation of our work is triangular remeshing, which refers to generating a triangular mesh from an input triangular mesh to improve its quality. (Note that triangular and hexagonal meshes are dual to each other, and triangular remeshing can also be used to perform hexagonal remeshing.) In triangular remeshing, it is often desirable to have all the triangles in the mesh being nearly equilateral and of uniform sizes, and the edges following the curvature and feature directions in the surface. In addition, special treatment is needed for irregular vertices (whose valence is not equal to six) since they impact the overall appearance and quality of the mesh.

A hexagonal parameterization transforms these challenges to that of control over the singularities in the parameterization as well as the spacing and direction of parameter lines. Smooth parameter lines and a reduced number of singularities lead to highly regular meshes. Such meshes are desirable for subdivision surface applications [3].

1.3 Rotational Symmetry

Inspired by recent development in constructing a quadrangular global parameterization [4], [5], [6], we compute a global parameterization given a six-way rotational symmetry field, or a 6-RoSy field, an abbreviation introduced in [7]. An N -RoSy refers to a set of N vectors with evenly spaced angles, and a 1-, 2-, and 4-RoSy can represent a vector, a line segment, or a cross, respectively. While a 1-, 2-, and 4-RoSy field can each be used to compute a quadrangular parameterization, a 4-RoSy field provides the

• M. Nieser and K. Polthier are with the Institute of Mathematics, AG Mathematical Geometry Processing, Freie Universität Berlin, Arnimallee 6, Berlin D-14195, Germany.

E-mail: {matthias.nieser, konrad.polthier}@fu-berlin.de.
 • J. Palacios and E. Zhang are with the School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331.
 E-mail: {palacijo, zhang}@eecs.oregonstate.edu.

Manuscript received 2 Dec. 2010; revised 3 June 2011; accepted 6 June 2011; published online 16 June 2011.

Recommended for acceptance by J. Stam.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCG-2010-12-0280. Digital Object Identifier no. 10.1109/TVCG.2011.118.




Fig. 1. Hexagonal patterns in nature: (a) honeycombs, (b) insect eyes, and (c) snowflakes. Appearance in design: (d) star of David, (e) Islamic pattern, and (f) floor tiling.

most flexibility in terms of modeling branch points, and thus the types of irregular vertices in a quad mesh. Specifically, a 1- or 2-RoSy field can always be converted into a 4-RoSy field with the unfortunate constraint that a first-order singularity in the 1- or 2-RoSy field becomes a higher order singularity in the resulting 4-RoSy field. Consequently, when performing quadrangular remeshing with a 1- or 2-RoSy field, it is in general impossible to obtain a valence three or five vertex.

Similarly, while 1-, 2-, 3-, and 6-RoSy fields can all be used for triangular remeshing, only 6-RoSy fields can be used to model irregular vertices that have a valence of either five or seven which are desirable in many cases.

1.4 Parameterization

Automatic generation of a hexagonal parameterization from an input surface poses a number of challenges. First, unlike quadrangular parameterization whose parameter lines are parallel to either the major or the minor principal curvature directions, in hexagonal parameterization only one of the two directions can be used at each point on the surface. One must decide which direction to choose, and how to propagate such choices from a relatively small set of points to the whole surface to maintain the smoothness of the resulting parameterization. Second, existing techniques to explicitly control the singularities in a parameterization are user driven, and it is not an easy task to provide automatic control over the number and location of such singularities. Third, the continuity conditions developed for quadrangular parameterization along seams in the parameterization are not appropriate for hexagonal parameterization (Fig. 2).

1.5 Pipeline

To address these challenges, we present a two-step pipeline to generate a geometry-aware hexagonal global parameterization. First, we automatically select the most appropriate principal direction with which we align our 6-RoSy field. Moreover, we introduce an automatic singularity clustering algorithm that allows nearby singularities to be either




Fig. 2. A quadrangular parameterization ensures that the discontinuity along the cut is invisible (a). The same parameterization is incompatible with a hexagonal pattern (b), which leads to seams (yellow). In this case, a hexagonal parameterization is needed (c).

anceled or merged into a higher order singularity, thus reducing the total number of singularities in the field. Note that merging two higher order singularities with opposite signs can lead to a lower order (e.g., first order) singularity.

In the second step of the pipeline, we generate a global parameterization which is aligned to the 6-RoSy field as much as possible. The QUADCOVER algorithm [5] is adapted for handling the symmetries of a hexagonal parameterization. We also formulate a quadratic energy which measures the L^2 distance of the parameter lines to the field. During minimization, some variables are constrained to an integer grid. We point out that in the hexagonal parameterization this grid is the set of *Eisenstein integers*, which is different from the *Gauss integers* used in the quadrangular case. This leads to a parameterization method that we refer to as HEXCOVER. The resulting parameters can then be used to generate triangular meshes free of T-junctions as well as to seamlessly tile a surface with any regular hexagonal pattern.

1.6 Contributions

In summary, our contributions in this paper are as follows:

1. We introduce hexagonal global parameterization and demonstrate its uses with applications such as triangular remeshing and pattern synthesis on surfaces. For remeshing, we point out the need for a geometry-aware 6-RoSy field when generating a hexagonal global parameterization.
2. We present the first technique to construct a hexagonal global parameterization given an input surface with a guiding 6-RoSy field. We formulate the energy term as well as the continuity condition for hexagonal global parameterization.
3. We propose an automated pipeline for generating geometry-aware 6-RoSy fields. As part of the pipeline, we point out how to align the field with principal curvature directions as well as develop a way of automatically clustering singularities.

The remainder of this paper is organized as follows: We first cover work in relevant research areas in Section 2. Next, we describe our pipeline for generating a geometry-aware 6-RoSy field given an input surface in Section 3, and our parameterization technique in Section 4. In Section 5, we demonstrate the usefulness of our techniques with applications in triangular remeshing and surface tiling with regular texture and geometry patterns. We conclude in Section 6 with future work.

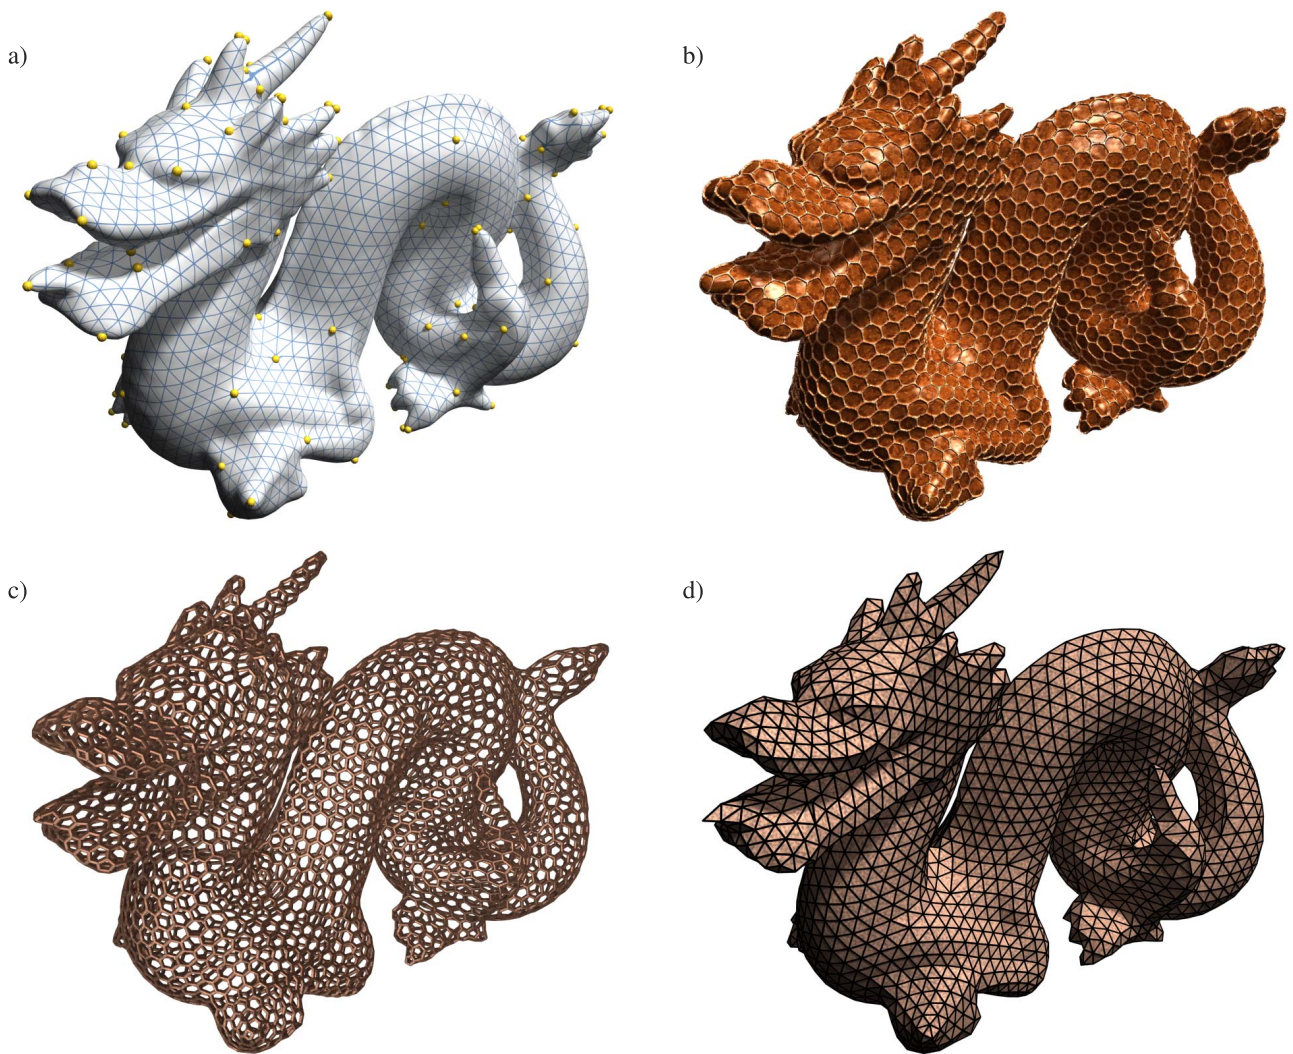


Fig. 3. Hexagonal global parameterization (a), used for regular texture and geometry synthesis with hexagonal patterns ((b) and (c)) and for geometry-aware triangular remeshing (d).

2 RELATED WORK

2.1 Surface Parameterization

Surface parameterization is a well-explored research area. We will not attempt a complete review of the literature but instead refer the reader to surveys by Floater and Hormann [8] and Hormann et al. [9].

Early global parameterization methods focus on conformal parameterization [10], [11], [12], which is aimed at angle preservation at the cost of length distortion. To reduce length distortion, Kharevych et al. [13] use cone singularities, which relax the constraint of a flat domain at few isolated points. Singularities have proven essential for high-quality parameterization and have been used in other parameterization schemes as well [14], [15].

Dong et al. [16] perform quadrangulation based on harmonics functions. Later, Dong et al. [17] use a similar idea for parameterization but create the quadrilateral meta layout automatically from the Morse-Smale complex of the eigenfunctions of the mesh Laplacian.

Tong et al. [18] use singularities at the vertices of a handpicked quadrilateral meta layout on a given surface. The patches of the meta layout are then parameterized by

solving for a global harmonic one form. Ray et al. [4] parameterize surfaces of arbitrary genus with periodic potential functions guided by two orthogonal input vector fields, or a 4-RoSy field. This leads to a continuous parameterization except in the vicinity of singularities on the surface. These singular regions are detected and reparameterized afterward.

The QUADCOVER algorithm [5] builds upon this idea by using the input 4-RoSy field to generate a global parameterization, based on a quadratic energy formulation. Also, the notion of covering spaces is used to describe a 4-RoSy field as a vector field and to provide a clear theoretical setting. Our algorithm to generate a parameterization from a 6-RoSy field is an adaptation of the QUADCOVER method.

Bommes et al. [6] propose a method similar to the aforementioned techniques based on the same energy formulation as in [5], but provide several advancements. Besides a robust generation of 4-RoSy fields, they propose to use a mixed-integer-solver for improving the rounding of the integer variables. They also add constraints that force parameter lines to capture sharp edges.

2.2 Field Processing

Much work has been done on the subject of vector (1-RoSy) and tensor (2-RoSy) field analysis. Note that a line field is equivalent to a symmetric tensor field with uniform magnitude [19]. To review all of this work is beyond the scope of this paper; here, we refer to only the most relevant work. Helman and Hesselink [20] propose a method of vector field visualization based on topological analysis and provide methods of extracting vector field singularities and separatrices. Topological analysis techniques for symmetric second-order tensor fields are later introduced in [21]. Numerous systems have been developed for the purpose of vector field design, most of which have been for specific graphics applications such as texture synthesis [22], [23], [24], fluid simulation [25], and vector field visualization [26]. Fisher et al. [27] propose a vector field design system based on discrete one forms. Note that the above systems do not employ any methods of topological analysis, and do not extract singularities and separatrices. Systems providing topological analysis include [28], [29], and [30]. The last has also been extended to design tensor fields [19], [31]. In contrast, relatively little work has been done on N -RoSy fields when $N > 2$. Hertzmann and Zorin [32] utilize *cross* or 4-RoSy fields in their work on nonphotorealistic pen-and-ink sketching, and provide a method for smoothing such fields. Ray et al. [33] extend the surface vector field representation proposed in [29] into a design system for N -RoSy fields of arbitrary N . Palacios and Zhang [7] propose an N -RoSy design system that allows initialization using design elements as well as topological editing of existing fields. They also provide analysis techniques for the purpose of locating both singularities and separatrices, and a visualization technique in [34]. Lai et al. [35] propose a design method based on a Riemannian metric that gives the user control over the number and locations of singularities. Their system also allows for mixed N -RoSy fields, with different values of N in different regions of the mesh. However, this method is based on user design while we focus on automatic and geometry-aware generation. Bommers et al. [6] offer a method of producing a smooth 4-RoSy field from sparse constraints, formulated as a mixed-integer problem. Zhang et al. introduce a quadrangulation method based on the notion of waves. Their method can also be used to generate 4-RoSy fields [36]. Crane et al. [37] handle cone singularities by using the notion of *trivial connection* in the surface. These singularities include those seen in 6-RoSy fields.

Ray et al. [38] propose a framework to generate an N -RoSy field that follows the natural directions in the surface and has a reduced number of singularities which tend to fall into natural locations. In this paper, we make use of this framework but automatically generate the input constraints, which relieves the user from labor-intensive manual design. Furthermore, we introduce to our knowledge the first automatic singularity clustering algorithm that reduces the number of singularities in the field.

3 GEOMETRY-AWARE 6-ROSY FIELD GENERATION

In this section, we describe our pipeline for generating a geometry-aware 6-RoSy field F given an input surface S .

This field will then be used to guide the parameterization stage of our algorithm (Section 4).

We first review some relevant properties of 6-RoSy fields [7], [33]. An N -RoSy field F has a set of N directions at each point \mathbf{p} in the domain of the field: $F(\mathbf{p}) = \{R_N^i v(\mathbf{p}) | 0 \leq i \leq N - 1\}$, where the vector $v(\mathbf{p}) = \rho(\mathbf{p})(\cos \theta(\mathbf{p}), \sin \theta(\mathbf{p}))^T$ is one of the N directions, and R_N^i is the linear operator that rotates a given vector by $\frac{2i\pi}{N}$ in the corresponding tangent plane. A *singularity* is a point \mathbf{p}_0 such that $\rho(\mathbf{p}_0) = 0$ and $\theta(\mathbf{p}_0)$ is undefined; \mathbf{p}_0 is *isolated* if the value of $\rho \neq 0$ for all points in a sufficiently small neighborhood of \mathbf{p}_0 , except at \mathbf{p}_0 . An isolated N -RoSy singularity can be measured by its *index*, which is defined in terms of the Gauss map [7] and has an index of $\frac{I}{N}$, where $I \in \mathbb{Z}$. A singularity \mathbf{p}_0 is of *first order* if $I = \pm 1$. When $|I| > 1$, \mathbf{p}_0 is referred to as a *higher order* singularity. A higher order singularity with an index of $\frac{I}{N}$ can be realized by merging I first-order singularities of the same sign.

3.1 Requirements and Pipeline

There are a number of goals that we wish to achieve with our automatic field generation.

First, we wish to control the number, location, and type of singularities in the field. When performing quadrangular and triangular remeshing, the singularities in the guiding 4- or 6-RoSy field correspond to irregular vertices in the mesh. Such singularities can also lead to the breakup of texture and geometry patterns during pattern synthesis on surfaces. Consequently, the ability to control the number, location, and type of singularities in the field can improve quality of remeshes and surface tilings.

Second, the field needs to be smooth, or distortion can occur in the resulting parameterization that has undesirable effects for triangular remeshing and surface tiling.

Third, we need the parameter lines in the parameterization to be aligned with the feature lines on the surfaces, such as ridge and valley lines (see Fig. 4). In addition, it has been documented that having texture directions aligned with the feature lines in the mesh can improve the visual perception of texture [39].

Note that these requirements may conflict with each other. For example, excessive reduction of singularities can lead to high distortion in the field, and an overly smoothed field may deviate from feature lines. To deal with this, we adopt the framework of Ray et al. [38]. In their framework, a set of user-specified constraints and a modified Gaussian curvature \bar{K} defined at the vertices are used to generate a sparse linear system whose solution (after several iterations) is the RoSy field that matches the constraints and \bar{K} in the least square sense. Each constraint represents a desired N -RoSy value, i.e., N directions, at a given point. In our case, we wish to have our field aligned with principal curvature directions. The user-specified \bar{K} is a vertex-based function defined on the mesh, whose value at a vertex represents the desired discrete Gauss curvature at this vertex to be reflected by resulting field curvature. The integral of \bar{K} over S must be equal to $2\pi\chi(S)$ where $\chi(S)$ is the *Euler characteristic* of the surface S . It allows the user to specify the location and type of singularities in the field. For example, a vertex with a \bar{K} value of $\frac{2k\pi}{N}$ should have a




Fig. 4. For remeshing, edges should follow principal curvature directions (right). Edges ignoring surface features (left) cause “twisting” artifacts (on the ears).

singularity of index $\frac{k}{N}$ in the resulting field. We would like to note that other field generation systems that allow directional constraints and the specification of singularities of index greater than $\frac{1}{N}$ can also be used (such as the one described in [33] and [37]). We use the geometry-aware method of Ray et al. because it gives additional control over the initial number singularities if desired.

Given a surface with complex geometry and topology, it can be labor intensive to provide all necessary constraints through a lengthy trial-and-error process. Consequently, we automatically generate the directional constraints as well as \bar{K} , which is at the core of our algorithm for field generation. Our algorithm consists of two stages. First, we identify a set of directional constraints based on the curvature and solve for an initial 6-RoSy field using these constraints only. Second, we extract all the singularities in the initial field and perform iterative singularity pair clustering until the distance between any singularity pair is above a given threshold. The remaining singularities will be used to generate new values for the vertex function \bar{K} , which will be used to generate the final RoSy field with reduced singularities. We describe each of these stages in more detail next.

3.2 Automatic Constraint Identification

To automatically identify directional constraints, we need to answer the questions of where to place constraints and what direction is assigned to each constraint.

Recall that we wish to align the parameter lines with feature lines such as ridges and valleys, i.e., the principal direction in which the least bending occurs. Note that the directions in the 6-RoSy field are the gradients of the parametrization (Section 4). Consequently, we will choose the principal direction that has the most bending, i.e., maximum absolute principal curvature, as one of the directions in the 6-RoSy. We estimate the curvature tensor of the mesh using the method of Meyer et al. [40].

Principal curvature directions are most meaningful in cylindrical and hyperbolic regions due to the strong anisotropy there. However, while purely hyperbolic regions possess strong anisotropy, the absolute principal curvatures are nearly indistinguishable, thus making both principal

curvature directions candidates. Moreover, the two bisectors between the major and minor principal curvature directions can also provide viable choices for the edge directions in hyperbolic regions. Due to the excessive choice of directions in hyperbolic regions and insufficient choice of directions in planar and spherical regions, we only generate directional constraints in cylindrical regions. Note that using the asymptotic directions could result in neighboring triangles being constrained with directions that differ by rotations of $\frac{\pi}{2}$. While this causes no problems in 4-RoSy field generation, such constraints conflict in the case of 6-RoSy field generation.

We make use of a representation of the curvature tensor that readily exposes where on this spectrum of classification any point on a given surface falls. Using the trace-and-deviator decomposition similar to those employed in [41], the curvature tensor T at a point $\mathbf{p} \in S$ can be rewritten as

$$T = \left(\frac{\kappa_1 - \kappa_2}{2} \begin{pmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{pmatrix} + \frac{\kappa_1 + \kappa_2}{2} \text{Id} \right) \\ = \frac{\rho}{\sqrt{2}} \left(\cos \phi \begin{pmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{pmatrix} + \sin \phi \cdot \text{Id} \right), \quad (1)$$

where κ_1 and κ_2 are the principal curvatures at \mathbf{p} , $\rho = \sqrt{\kappa_1^2 + \kappa_2^2}$, $\phi \in [-\pi/2, \pi/2] = \arctan\left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}\right)$, $\theta \in [0, \pi)$ is the angular component of the maximum principal direction measured in the local frame at \mathbf{p} , and Id denotes the identity matrix. Note that the first component in the sum is traceless and symmetric, while the second is a multiple of the identity matrix. $T(\mathbf{p})$ can now be classified using $(\rho(\mathbf{p}), \phi(\mathbf{p}))$, which spans a half plane. There are six special configurations on this half plane, the first satisfying $\rho(\mathbf{p}) = 0$, i.e., the local geometry near \mathbf{p} is planar. For the remaining five configurations, we have $\rho(\mathbf{p}) > 0$. Respectively, they correspond to $\phi(\mathbf{p}) = \frac{\pi}{2}$ (spherical), $\phi(\mathbf{p}) = \frac{\pi}{4}$ (cylindrical), $\phi(\mathbf{p}) = 0$ (purely hyperbolic), $\phi(\mathbf{p}) = -\frac{\pi}{4}$ (inverted cylindrical), and $\phi(\mathbf{p}) = -\frac{\pi}{2}$ (inverted spherical). With this representation, we can classify any point \mathbf{p} as being planar if $\rho(\mathbf{p})$ is smaller than a given threshold δ , elliptical if $\rho(\mathbf{p}) \geq \delta$ and $|\phi(\mathbf{p})| > \frac{3\pi}{8}$, hyperbolic if $\rho(\mathbf{p}) \geq \delta$ and $|\phi(\mathbf{p})| < \frac{\pi}{8}$, and cylindrical otherwise,




Fig. 5. Surface classification scheme to determine directional constraints. $\phi \in [-\pi/2, \pi/2]$ is color mapped to the [blue, red] arc in HSV color space: Left top: continuous mapping. Bottom: binned classification. The legend (right) shows surfaces patches which are locally similar to points with given values.

i.e., $\rho(\mathbf{p}) \geq \delta$ and $\frac{\pi}{8} \leq |\phi(\mathbf{p})| \leq \frac{3\pi}{8}$. We wish to point out that the tensor-based decomposition is equivalent to the concept of *shape index* [42].

Given the classification, we propagate the directions in the cylindrical regions into noncylindrical regions (planar, spherical, hyperbolic) using energy minimization, an approach taken in [6]. To accomplish this, we pick the points where ρ (the tensor magnitude) is above a certain threshold t_ρ , and label these points as having “strong” curvature (in all of our examples, we have chosen t_ρ so that 35 percent of the area of S is so labeled). From this set of points, we use only the directions of the cylindrical points as constraints, that is, the points for which $\phi \in [-3\pi/8, -\pi/8] \cup [\pi/8, 3\pi/8]$ (Fig. 6). Finally, we select the maximum direction θ as the constraint direction at points where $\phi > 0$ and the minimum direction $\theta + \pi/2$ where $\phi < 0$. Recall that the directions in the output field specify the gradients in our resulting parameterization, and we wish one of the isolines of the parameters to be orthogonal to the direction in which the surface is bending the most. Clearly, the above directions satisfy this requirement (see the shapes on the right side of the right image in Fig. 5). Finally, the constraints are used to set up a linear system [38] whose solution gives rise to our initial RoSy field.

For our solver, we use the geometry-aware N -RoSy field generation technique proposed by Ray et al. [38], as it allows us to control the level of geometric detail that is reflected by singularities, and also plays a role in the implementation of our singularity clustering technique. This system, based on discrete exterior calculus (DEC) [43], filters (locally averages) the Gauss curvature K of S to produce \bar{K} and then computes a target field curvature C^t using the difference between K and \bar{K} . C^t is then used to modify the angles by which directions rotate when parallel transported along mesh edges. This compensates for the actual curvature of S , and direction fields computed on S under these conditions behave as though S has a Gauss curvature of \bar{K} . Since \bar{K} is smoother than K , such fields have reduced topological noise, which makes them more suitable for our parameterization algorithm.



Fig. 6. Selection of constraints. Left: color mapping of ρ . Middle: highest 35 percent of values; colors are based on ϕ as in Fig. 5. We use maximum curvature directions where $\phi > 0$ (yellow) and minimum directions where $\phi < 0$ (cyan) as being orthogonal to the direction in which the surface is bending the most (see closeup, right). Notice that chosen directions in nearby yellow and cyan regions agree as they would not if we had selected only one of the curvature directions everywhere.

3.3 Automatic Singularity Clustering

Our initial field was obtained from directional constraints only. Consequently, it typically consists of only first-order singularities. Given a surface with complex geometry and topology, the number of singularities can be rather large. Furthermore, while the location of the singularities tend to be appropriate (in high-curvature regions), many of them form dense clusters. Having singularities in closer proximity can lead to difficulties in the resulting parameterization. This is because the singularities will be constrained to be on a regular lattice in the parameter space as typically required by most global parameterization methods [5], [6]. Consequently, the smallest distance between any singularity pair will be mapped to a unit in the parameter space. If the smallest distance is too small, the two involved singularities may be mapped to the same point on the lattice, leading to a locally infinite stretching in the parameterization. Fig. 15 illustrates this.

To address this, many field generation techniques constrain the number of singularities to be as few as possible [33], but this represents another extreme, where the field directions can become highly distorted in some regions. Furthermore, many of the aforementioned approaches require much user interaction [7], [33], [38], which can be time consuming for models with complex geometry and topology.

Our goal is to automatically reduce the number of singularities in the field while retaining the locations of the remaining singularities inside high-curvature regions. To achieve this, we employ the following process.

First, we extract the singularities in the initial RoSy field (using the method described in [38]) which we use to build a graph embedded in the surface. The nodes of this graph are the singularities in the field, and the edges representing proximity information between singularity pairs. We refer to this graph as the *singularity graph* G . To construct G , we compute a Voronoi diagram with the singularities as sites. The dual graph gives rise to the singularity graph [44].

Second, we iteratively perform edge collapses [55] on this graph, which is equivalent to performing singularity pair clustering (merging or cancellation), until the minimal surface distance between any singularity pair is above a




Fig. 7. Clustering pipeline: (a) Initial field. (b) Singularity graph G . (c) Reduced graph obtained by performing edge collapses. The region R is shown in green. (d) Reduced field generated by resolving in R with singular constraints at the nodes of G and directional constraints at the boundary of R .

given threshold. Every time a singularity pair is clustered, we compute the sum of the singularity indexes and place a singular constraint with the sum as its desired index. Note that we do this even if the sum is zero, i.e., singularity pair cancelation. The singularity constraint is placed on the path between the two original singularities \mathbf{p}_0 and \mathbf{p}_1 , closer to the one with the Gaussian curvature of highest magnitude. This is an attempt to keep singularities near the features that caused them to originally appear during initialization and is accomplished by interpolating along the geodesic from \mathbf{p}_0 to \mathbf{p}_1 using the value $|K(\mathbf{p}_1)|/(|K(\mathbf{p}_0)| + |K(\mathbf{p}_1)|)$, where $K(\mathbf{p})$ is the Gaussian curvature at $\mathbf{p} \in S$. We continue to collapse edges in the order of increasing edge length on G until no edge of length less than d_{sing} remains. At the end of this process, we will have generated a set of singularity constraints, i.e., the remaining vertices in the graph, which is then used to update the field in the vicinity of these singularities. In the case of fields generated for remeshing, d_{sing} can be selected based on the edge length of the output mesh. We choose d_{sing} to be $0.1B$ where B is the size of the bounding box for the model. For a visual summary of the algorithm, see Fig. 7.

Third, we modify \bar{K} based on the singularity constraints. Recall that the \bar{K} is simply a smoothed version of the discrete Gauss curvature during the generation of the initial field. The singularity constraints, produced in the previous step, consist of a set of vertices in the mesh and a desired singularity index $t(\mathbf{p})$ for each such singularity constraint \mathbf{p} . We modify \bar{K} such that it is zero everywhere on the surface except at singularity constraints where the value of \bar{K} is $\frac{2\pi}{N}t(\mathbf{p})$. Notice that such assignment satisfies the constraint that the integral of \bar{K} over S is equal to $2\pi\chi(S)$. We now modify the 6-RoSy field by solving the same system used to generate the initial field, with one difference: we do not update the field everywhere on the surface. Instead, we generate a region $R = \{\mathbf{p} | \text{dist}(\mathbf{p}, V_{collapse}) < d_{sing}\}$, where $V_{collapse}$ is the set of vertices that were members of collapsed edges in G , and update the field only in R . That is, the field values are fixed in the complement of R and the values on the boundary of R will serve as the boundary conditions when updating the field in R ; the original directional constraints are ignored in this step. In this way, we largely preserve the results of the field generated from the directional constraints, but force the merging and cancelation of singularities in the regions where large clusters had appeared before. The field values for triangles inside R are then updated. We have found this to be efficient in controlling the singularities.

We wish to point out that our automatic field generation method can be applied to N -RoSy field generation for any N that is even, in particular 4-RoSy fields. Fig. 8 shows an example generated using our method. The only change in the whole field generation pipeline occurs during automatic identification of directional constraints. Instead of choosing θ or $\theta + \frac{\pi}{2}$ as one of the six directions for constraints, we choose both for the case of 4-RoSy fields.

4 HEXCOVER PARAMETERIZATION

In this section, we describe the second stage of our pipeline, which constructs a hexagonal global parameterization given an input triangular mesh surface along with a 6-RoSy field defined on it. We will first introduce the notion of hexagonal parameterization before describing our HEXCOVER parameterization technique which is an extension of the QUADCOVER method for quad remeshing.

4.1 Hexagonal Parameterization and Energy.

Given a triangular mesh surface S with $|T|$ triangles, a *global parameterization* $\varphi : S \mapsto \mathbb{R}^2$ respecting an N -RoSy symmetry is a collection of linear maps $\{\varphi_i | 1 \leq i \leq |T|\}$ where each $\varphi_i : t_i \rightarrow \mathbb{R}^2$ maps triangle $t_i \in S$ onto \mathbb{R}^2 with the following property. For any adjacent triangles t_i and t_j , we have




Fig. 8. Geometry-aware 4-RoSy field and corresponding texture tiling.




Fig. 9. Left: texture with hexagonal rotational symmetries. Right: Eisenstein integer lattice G_6 .

$$\varphi_j(\mathbf{p}) = R_N^{r_{ij}} \varphi_i(\mathbf{p}) + w_{ij}, \quad \forall \mathbf{p} \in t_i \cap t_j, \quad (2)$$

where $r_{ij} \in \{0, 1, \dots, N-1\}$ and $w_{ij} \in \mathbb{R}^2$ are the rotational and translational discontinuities, respectively. Recall that R_N^k is the linear operator that rotates a vector by $\frac{2k\pi}{N}$ in its tangent plane (Section 3). The maps φ_i are restricted to be linear on each triangle. They are defined by their values at vertices, while r_{ij} and w_{ij} are defined on edges.

In quadrangular case where $N = 4$, parameter lines can be visualized by treating φ^{-1} as the map that textures the surface with a 2D regular unit grid. To ensure continuity in parameter lines, translational discontinuities w_{ij} are required to be on the set of *Gauss integers* $G_4 := \{(a, b)^T \mid a, b \in \mathbb{Z}\}$.

Hexagonal parameterization ($N = 6$) is similar, except that in this case the texture image needs to respect hexagonal rotational symmetries. A canonical choice is a hexagonal or triangular pattern as shown in Fig. 9 (left). The texture image has an aspect ratio of $1 : \sqrt{3}$ and tiles the plane seamlessly. It is furthermore invariant under rotations of $\frac{\pi}{3}$ around the center of each hexagon. The set of these center points is known as the *Eisenstein integer lattice*, shown in Fig. 9 (right)

$$G_6 := \left\{ a \begin{pmatrix} 1 \\ 0 \end{pmatrix} + b \begin{pmatrix} 1/2 \\ \sqrt{3}/2 \end{pmatrix} \mid a, b \in \mathbb{Z} \right\}. \quad (3)$$

Besides the rotational invariance, the hexagonal grid also remains invariant under translations by any vector in G_6 . While a hexagonal parameterization is a discontinuous map, the discontinuities are not visible if all w_{ij} are in G_6 because of the repeating structure of the texture image (Fig. 2).

A hexagonal parameterization can be generated from a guiding 6-RoSy field F . Given a point \mathbf{p} , the edges of the hexagons are aligned with the six vectors of F in \mathbf{p} . This is achieved by optimizing the alignment in L^2 sense. Specially, we minimize the quadratic energy

$$E(u, v) := \int_S (\|\nabla u - F_u\|^2 + \|\nabla v - F_v\|^2) dA, \quad (4)$$

where (u, v) is the parameterization, $F_u(\mathbf{p})$ is one of the six vectors of F at $\mathbf{p} \in S$, and $F_v(\mathbf{p}) := R_4^1 F_u(\mathbf{p})$ is perpendicular to it. We further define $u_i = u|_{t_i}$ and $v_i = v|_{t_i}$.

The parameterization must fulfill the integer constraints in (2), whereas r_{ij} 's encode which of the 6-RoSy vectors in adjacent triangles t_i and t_j are paired, i.e., F_u in t_i is paired with $R_6^{r_{ij}} F_u$ in t_j . r_{ij} 's are held fixed during energy minimization, whereas u_i 's, v_i 's, and w_{ij} 's are optimized.

Notice that the energy is independent of the choice of F_u (there are six choices per triangle) due to the rotational symmetries of φ from (2). A different choice of F_u in one triangle will result in the same change in the r_{ij} 's along all

adjacent edges. The resulting minimizer of the energy (4) is then locally rotated by a multiple of $\frac{\pi}{3}$ in this triangle, resulting in the same pattern.

A key observation in QUADCOVER [5] is that the optimization can be divided into two subproblems and solved independently:

1. **Local step.** Minimize the energy (4) for $u_i, v_i, w_{ij} \in \mathbb{R}$, ignoring the integer constraint on w_{ij} . In QUADCOVER, the minimizer is computed by removing the curl of F , making it locally integrable, and defining u_i, v_i as its potential. This leads to a local parameterization φ' .
2. **Global step.** Convert φ' into a global parameterization by incorporating the aforementioned integer constraints.

4.2 HEXCOVER and Covering Spaces

Minimizing (4) directly presents some challenges due to the fact that F_u and F_v are both multivalued (there are six values per triangle). Here, we make use the notion of *covering space*, which transforms the problem of computing a global parameterization on S under a guiding 6-RoSy field F to generating a global parameterization on an N -fold cover S' of S under a guiding vector field F' . The benefit of doing this is that we can use standard vector field calculus without having to deal with an N -RoSy field.

In fact, the covering is just used as theoretical foundation and is not explicitly computed in either QUADCOVER or HEXCOVER. The covering is implicitly represented by the values r_{ij} resulting in additional constraints (2) during optimization. Note that covering spaces are used implicitly by other approaches optimizing a piecewise-linear global parameterization [18], [6].

In the hexagonal case, F can be lifted to F' on a sixfold covering surface S' of S , which is defined as follows: every triangle t_i in S has six corresponding triangles in S' : $t_{i,0}, \dots, t_{i,5}$. The vector field F' distributes the six vectors of F onto the six copies, i.e., $F'(t_i) = R_6^i F_0(t_i)$ where $F_0(t_i)$ is one of the six directions of F in t_i . For adjacent triangles t_i, t_j in S , the corresponding copies are combinatorially connected, depending on the rotational discontinuity r_{ij} . The triangles $t_{i,k}, k \in \{0, \dots, 5\}$ are thereby connected with $t_{j, k+r_{ij} \bmod 6}$. Note that S' is a Riemann surface with branch points at those positions where the original 6-RoSy field has singularities. All six copies of a triangle are geometrically identical, so there is not necessarily an embedding without self-intersections. This does not present any difficulty for us, however, since the algorithm does not rely on an explicit embedding of S' .

The problem now turns into minimizing the energy in (4) on the covering space S' , using $F'_u := F'$, $F'_v := R_4^1 F'$ (see Fig. 10). Due to the symmetry of the covering surface and the symmetric behavior of the algorithm, the resulting texture images on different copies of each triangle are congruent and their projection onto the domain S is a global parameterization which satisfies (2). Again, the use of coverings is only a theoretical view. The algorithm will not compute the covering, but represents it implicitly by storing the values r_{ij} .




Fig. 10. Left: triangle t_i with a 6-RoSy field. Right: sixfold covering of t_i with vector fields F'_u, F'_v .

4.3 Local Step

In the local step, Energy (4) is minimized for values of the parameterization $u_i(p_j), v_i(p_j)$ at each vertex p_j in all incident triangles t_i , and for the translational discontinuities $w_{ij} \in \mathbb{R}^2$. Due to the high number of variables and additional constraints (2), QUADCOVER proposes to solve an alternative energy providing the same result but with a much smaller system of equations and no constraints. We use a similar simplification for HEXCOVER.

Let $\varphi = (u, v)^T$ be the minimizer of Energy (4). A key observation is derived from the discrete Hodge-Helmholtz decomposition of vector fields [45]: the field $(F_u - \nabla u, F_v - \nabla v)$ is exactly a *cogradient field* $(R_4^1 \nabla u^*, R_4^1 \nabla v^*)$ which minimizes the energy

$$E^*(u^*, v^*) := \int_S (\|R_4^1 \nabla u^* - F_u\|^2 + \|R_4^1 \nabla v^* - F_v\|^2) dA. \quad (5)$$

Here, u^* and v^* are scalar nonconforming finite element functions, which are linear in each triangle and defined by values on edge midpoints. At boundary edges, u^* and v^* are fixed to 0. The constraints (2) simplify to

$$\begin{pmatrix} u^*_{|t_i} \\ v^*_{|t_i} \end{pmatrix} = R_6^{r_{ij}} \begin{pmatrix} u^*_{|t_j} \\ v^*_{|t_j} \end{pmatrix}, \quad (6)$$

for adjacent triangles t_i, t_j . Notice that the translational discontinuities w_{ij} do not appear in this formulation.

Equation (6) directly relates the values of u^* and v^* in both adjacent triangles of each edge; therefore, only one free u^* variable and one free v^* variable remain left per edge. We build a system of linear equations by setting all partial derivatives of Energy (5) for the free variables to 0. The matrix of this system has dimension $2|E| \times 2|E|$, where $|E|$ is the number of edges in the mesh. We solve this system and obtain (u^*, v^*) from which we compute $(\nabla u, \nabla v)$.

The parameterization (u, v) is computed by first cutting the mesh open to a simply connected disk and then directly integrating the gradients. We cut the surface along the shortest homotopy generators similar to [46]. The result is a graph G on edges, such that the complement $S \setminus G$ is simply connected. We also need to connect all singularities with the cut graph, since they can be seen as infinitesimally small holes. For this purpose, the method was adapted to include the surface boundary and singularities in [47].

The gradient fields $(\nabla u, \nabla v)$ are integrated by setting $(u, v) = (0, 0)$ at an arbitrary root vertex v_0 in triangle t_0 and directly integrating the piecewise constant vectors in t_0 and adjacent triangles until the whole surface is covered. When

crossing an edge, the values of (u, v) must be rotated according to (2). Note that the translational discontinuities are set to 0 in the interior of $S \setminus G$. The solution is consistent and does not depend on the traversal of the triangles, as long as the edges in the cut graph G is not involved in this propagation.

4.4 Global Step

While the parameterization (u, v) is a minimizer of (4), it may be discontinuous along the edges of G . For a global hexagonal parameterization, such discontinuities lead to seams in the parameter lines if the w_{ij} 's are not in the set of G_6 (the Eisenstein integer lattice). However, when performing local integration in the previous step, we only require that $w_{ij} \in \mathbb{R}$. In this section, we discuss how to modify the initial parameterization to enforce the integer constraints.

The graph G can be considered as union of paths γ_i , each of which is either a closed loop or a segment starting and ending at a singularity. An important property of the solution of Energy (4) is that the translational discontinuity w_{ij} is constant for all edges on the same path γ_i . Let w_i be the constant for path γ_i , which can be computed from the coordinates of (u, v) at both sides of an edge of γ_i . Note that the translational discontinuities can add up if two paths partially overlap.

To enforce the integer constraints, we modify the translational discontinuity w_{ij} for every edge in G by rounding them to the nearest integer in G_6 . Then, Energy (4) is minimized, holding all discontinuities w_{ij} fixed.

The coordinates of a singularity are uniquely determined by the w_i of all of its incident paths γ_i . For a singular vertex with valence ν , there are ν constraints (2) that relate the ν coordinate vectors of the vertex in its adjacent triangles. Thus, rounding the values w_i is similar to prescribing the coordinates of singularities.

For each regular vertex of valence ν , one of the ν relations (2) is redundant since the total discontinuity adds up to zero, reflecting a zero Poincaré index. Therefore, its coordinates are determined by the coordinates in one of its incident triangles. Consequently, we obtain one free variable for u and one for v per vertex. Energy (4) is minimized by setting all partial derivatives to 0 resulting in a sparse linear system. The matrix has dimension $2|\bar{V}| \times 2|\bar{V}|$ with $|\bar{V}|$ being the number of regular vertices.

Fig. 11 shows the hexagonal parameterization of two minimal surfaces using our technique.

4.5 Rounding Technique

The presented rounding technique for the w_i is just a heuristic for the problem of finding an optimal parameterization yielding the integer conditions. In general, this problem is NP-hard, since it is equivalent to minimizing a quadratic function on a given lattice (also called the *closest vector problem*).

The rounding technique used in QUADCOVER [5] where all integer variables are rounded at once can be contrasted with that from Mixed Integer Quadrangulation (MIQ) [6], which iterates between rounding integer variables and solving the system with the new boundary condition. In QUADCOVER, the translational discontinuities w_{ij} are used as integer variables, whereas MIQ uses the coordinates of singularities. Since the coordinates of singularities are

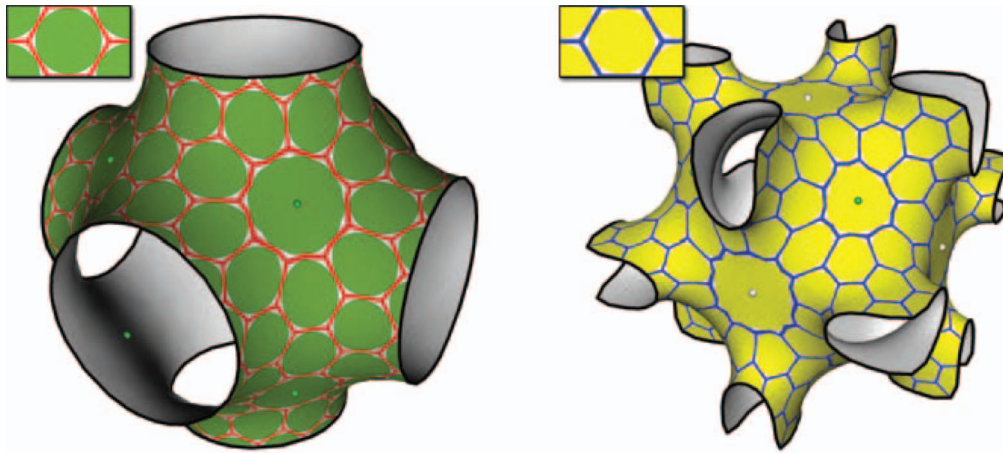


Fig. 11. Minimal surfaces. Left: Schwarz surface with eight singularities of index $-1/2$. Right: Neovius surface with eight index $-1/2$ and six index -1 singularities.

uniquely determined by the w_{ij} (up to global translation), both approaches consider a similar space but use a different basis for representation.

In all our tests, both rounding techniques (direct and mixed integer rounding) give similar results. We conjecture that the reason behind this is our use of the shortest cut graph G . It appears that shorter paths give the constants w_i a more local influence, hence directly rounding integer variables becomes more accurate.

In this work, we have opted to use the direct rounding, although one can easily replace it with a mixed integer solver.

5 RESULTS AND APPLICATIONS

Here, we apply hexagonal parameterization to two graphics applications: pattern synthesis, and triangular remeshing.

5.1 Pattern Synthesis on Surfaces

Example-based texture and geometry synthesis on surfaces has received much attention from the graphics community in recent years. We refer to [48] for a complete survey. Here, we will refer to the most relevant work.

Wei and Levoy [24] are the first to point out that N -RoSy fields of $N > 1$ are suitable for specification of special symmetries in textures. Liu et al. [49] propose techniques for the analysis, manipulation, and synthesis of *near-regular* textures (i.e., very structured textures with repeating patterns) in the plane. Kaplan and Salesin [2] address the design of Islamic star patterns in the plane. There has been some recent work in constructing circle patterns from a triangular mesh for architectural models [1].

Generating regular patterns on a surface can be greatly facilitated given an appropriate global parameterization. Given a regular hexagonal texture or geometry pattern, it is simply tiled in the parameter space of the mesh and the texture should stitch (relatively) seamlessly everywhere (Fig. 12). For example, to achieve circle packing for architectural patterns, our hexagonal parameterization allows nice hexagonal patterns to be generated from a surface, which can be used as input to such algorithms as shown in Fig. 12 (right). Our method provides necessary smoothness and feature alignment, thus leading to a high-quality model,

even in the case of relatively high geometric and topological complexity. Fig. 3b and 3c provides some additional examples in which regular hexagonal texture and geometry patterns are placed on the dragon.

We also comment that our field generation algorithm can also automatically generate geometry-aware 4-RoSy fields, which lead to coherent synthesized patterns that align with surface features (Fig. 8).

5.2 Triangular Remeshing

There has been much work in triangular remeshing. To review all past work is beyond the scope of this paper. We refer the reader to [50] for a complete survey of triangular remeshing literature, and review only the most relevant work here. Common methods of mesh triangulation are typically based on either a parameterization [51], [52], [53], [54], local optimization methods [55], [56], [57], or Delaunay triangulations and centroidal Voronoi tessellations [58], [59].

The focus of triangular remeshing is on shape preservation, good triangle aspect ratio, feature-aware triangle sizing, and control of irregular vertices (valence not equal to six). These objectives often conflict with one another, and the output mesh is a result of a compromise among these factors. For example, many parameterization-based methods suffer from artifacts in the triangulation at the locations of the chart boundaries (though this problem can be alleviated by using a global parameterization as in [54]). Direct and local optimization methods suffer from a lack of global control over the structure of the triangulation such as the location and number of irregular vertices.

In this paper, we perform triangular remeshing using a hexagonal global parameterization derived from a shape-aware 6-RoSy field. There are a number of benefits to this. First, such an approach can lead to overall better aspect ratio for triangles in the remesh (equilateral). Second, the number of irregular vertices can be reduced and their locations can be controlled as these vertices correspond exactly to the set of singularities in the 6-RoSy field. Third, we have incorporated the ability to match the orientations of the RoSy field based on natural anisotropy on the surfaces. Fourth, the size of the triangles can be controlled through a scalar sizing function. The frames are just scaled by the corresponding sizing value. A smaller scaling results




Fig. 12. Seamless tiling of hexagonal textures (left, middle) and geometry patterns (right).

in bigger triangles whereas a high value generates a finer triangle mesh (Fig. 13).

We can influence the number of singularities in the mesh by singularity clustering as described in Section 3. Fig. 14 shows that the distance between singularities impacts the smoothness of the parameterization, with more singularities reproducing more feature details of the surface. However, metric distortion also increases when more singularities are used as can be represented with the actual mesh resolution (see Fig. 15). Choosing the number of singularities can be considered as a trade-off between the smoothness of mesh elements and feature preservation. In Table 1, we compare the statistics for the three bunny remeshing results. Notice that the Hausdorff error and the standard deviation in angles of the triangles in the remesh are the lowest for the case when there are 65 singularities, corresponding to the parameter values that we used to generate all our models. The other two models have 23 and 151 singularities, respectively. They were the results of more and less aggressive singularity clustering. Fig. 14 compares the three models visually. Notice that features such as ridges along the ears are usually less preserved when there are too few singularities.

Fig. 16 compares the results of the foot and Venus models using our method with that of [52] and [59]. Table 1 provides the quality statistics of all tested models and the comparison. Notice that our method has better overall




Fig. 13. Adaptive sizing of triangles. Left: linear scaling along the y -axis. Right: scaling by the absolute maximal principle curvature value.

triangle aspect ratios (larger minimum angle, smaller maximum angle, and smaller standard deviation of angles) than [52]. All three methods capture the underlying geometry well (comparable Hausdorff distances to the original input mesh) but our method tends to have the fewest irregular vertices among all three methods. This is a direct result of automatic singularity clustering in the field generation step (Section 3) while achieving good triangle aspect ratios is due to the nature of the hexagonal parameterization. In addition, our method tends to produce edge directions that better align with the features in the mesh (such as along Venus' nose ridge) than [52]. Additional remeshing results can be found in Fig. 3.

5.3 Performance

The amount of time to automatically generate a geometry-aware 6-RoS field is on average 40 seconds for a model of 40K triangles, measured on a PC with a dual-core CPU of 2.8 GHz CPU and 4 GB RAM. The time to generate the parameterization is approximately 120 seconds per model, measured on a PC with a 2.13 GHz four-core CPU with 8 GB RAM. The running time of both stages is impacted by the mesh size as well as the number of singularities in the RoSy field. The computation time for both the field generation and parameterization stages is dominated by solving linear systems whose size is $O(|E|)$ where $|E|$ is the number of edges in the mesh. We solve these systems using a biconjugate gradient solver, whose complexity is sub-quadratic.




Fig. 14. Remeshing with 23, 65, and 151 singularities.




Fig. 15. Singularities which are closer than the grid size may force the parameterization to degenerate locally (left). This artifact can be avoided by either choosing a finer grid size (middle) or by merging nearby singularities with our clustering approach (right).

6 FUTURE WORK

There are a number of possible future research directions. First, we plan to add the capability to have parameter lines passing through sharp edges in the model, as considered in the quadrangulation case [6]. Second, we wish to study objects that are close to N -RoSy, which we refer to as near-regular RoSy's. In these objects, the N member vectors do not have identical magnitude nor even angular spacings. Such objects can allow more flexibility in both quadrangular and triangular remeshing. Third, pentagonal symmetry appears in many natural objects such as flowers. We wish to pursue graphics applications that deal with pentagonal symmetry. While an N -gon can tile a plane only if $N = 3, 4,$ and $6,$ it can tile a hyperbolic surface for any $N > 2$. Consequently, pentagonal patterns have the potential of being used to tile hyperbolic regions in a surface or for a hyperbolic parameterization. Notice our parameterization technique can actually handle a parameterization based on an N -RoSy field for any $N \geq 2$. In another direction, we plan to investigate

TABLE 1

Quality of Meshes: Hausdorff Distance (Percent of Bounding Box); Minimum, Maximum, and Standard Deviation (SD) of Angles, and Number of Irregular Vertices

model name	Hausdorff distance	min angle	max angle	SD angle	irregular vertices
Foot [52]	0.3373	2.82	173.88	11.92	146
Foot [59]	0.0094	26.92	115.85	7.40	3287
Foot	0.0129	22.65	125.09	5.11	13
Venus [52]	0.1005	0.42	178.99	17.48	38
Venus [59]	0.0439	19.89	121.13	10.37	1449
Venus	0.0543	24.87	114.80	6.84	36
Max Planck Fig.12	0.00263	12.44	145.79	5.00	44
Bunny Fig.14, left	0.6581	22.43	128.08	8.23	23
Bunny Fig.14, middle	0.0198	18.03	138.54	7.54	65
Bunny Fig.14, right	0.0309	16.99	133.8	8.50	151
Feline Fig.12	0.02695	5.75	167.34	11.09	121
Dragon Fig.3	0.00762	4.80	151.88	9.10	181
Blade Fig.13	0.84233	0.67	178.18	26.41	55

appropriate mathematical representations that handle other types of wallpaper textures which may contain reflections and gliding reflections. Surface tiling with at least two different types of rotational symmetries is another potential future direction. Such patterns have applications in cyclic weaving over surfaces [60] and remeshing [35].

ACKNOWLEDGMENTS

The authors wish to thank Felix Kälberer and Ulrich Reitebuch for fruitful discussions on parameterization and help on remeshing. Craig Anderson helped with the video production. Many thanks to all reviewers for their helpful




Fig. 16. Comparison of our method (right) to those from [52] (left) and [59] (middle). The histograms show occurring inner angles (on the X -axis from 0 to $\pi/3$). For each model, the scale on the Y -axis is the same.

comments which have led to significant improvements of the paper. The 3D models used in this paper are courtesy of Marc Levoy and the Stanford graphics group, and the AIMshape repository. The work is partially sponsored by the DFG research center MATHEON and the US National Science Foundation (NSF) grants IIS-0546881, CCF-0830808, and IIS-0917308.

REFERENCES

- [1] A. Schiftner, M. Höbinger, J. Wallner, and H. Pottmann, "Packing Circles and Spheres on Surfaces," *Proc. SIGGRAPH Asia*, pp. 1-8, 2009.
- [2] C.S. Kaplan and D.H. Salesin, "Islamic Star Patterns in Absolute Geometry," *ACM Trans. Graphics*, vol. 23, no. 2, pp. 97-119, 2004.
- [3] Z.J. Wood, P. Schröder, D. Breen, and M. Desbrun, "Semi-Regular Mesh Extraction from Volumes," *Proc. Conf. Visualization (VIS)*, pp. 275-282, 2000.
- [4] N. Ray, W.C. Li, B. Lévy, A. Sheffer, and P. Alliez, "Periodic Global Parameterization," *ACM Trans. Graphics*, vol. 25, no. 4, pp. 1460-1485, 2006.
- [5] F. Kälberer, M. Nieser, and K. Polthier, "QUADCOVER - Surface Parameterization Using Branched Coverings," *Computer Graphics Forum*, vol. 26, no. 3, pp. 375-384, 2007.
- [6] D. Bommes, H. Zimmer, and L. Kobbelt, "Mixed-Integer Quadrangulation," *ACM Trans. Graphics*, vol. 28, no. 3, pp. 1-10, 2009.
- [7] J. Palacios and E. Zhang, "Rotational Symmetry Field Design on Surfaces," *Proc. SIGGRAPH*, p. 55, 2007.
- [8] M.S. Floater and K. Hormann, "Surface Parameterization: A Tutorial and Survey," *Advances in Multiresolution for Geometric Modelling*, pp. 157-186, Springer Verlag, 2005.
- [9] K. Hormann, K. Polthier, and A. Sheffer, "Mesh Parameterization: Theory and Practice," *Proc. SIGGRAPH Asia*, no. 11, 2008.
- [10] S. Haker, S. Angenent, A. Tannenbaum, R. Kikinis, G. Sapiro, and M. Halle, "Conformal Surface Parameterization for Texture Mapping," *IEEE Trans. Visualization and Computer Graphics*, vol. 6, no. 2, pp. 181-189, Apr.-June 2000.
- [11] X. Gu and S.-T. Yau, "Global Conformal Surface Parameterization," *Proc. ACM SIGGRAPH Symp. Geometry Processing (SGP)*, pp. 127-137, 2003.
- [12] M. Jin, Y. Wang, S.-T. Yau, and X. Gu, "Optimal Global Conformal Surface Parameterization," *Proc. IEEE Visualization*, pp. 267-274, 2004.
- [13] L. Kharevych, B. Springborn, and P. Schröder, "Discrete Conformal Mappings Via Circle Patterns," *ACM Trans. Graphics*, vol. 25, no. 2, pp. 412-438, 2006.
- [14] M. Ben-Chen, C. Gotsman, and G. Bunin, "Conformal Flattening by Curvature Prescription and Metric Scaling," *Computer Graphics Forum*, vol. 27, no. 2, pp. 449-458, 2008.
- [15] B. Springborn, P. Schröder, and U. Pinkall, "Conformal Equivalence of Triangle Meshes," *ACM Trans. Graphics*, vol. 27, pp. 77:1-77:11, Aug. 2008.
- [16] S. Dong, S. Kircher, and M. Garland, "Harmonic Functions for Quadrilateral Remeshing of Arbitrary Manifolds," *Computer Aided Geometric Design*, vol. 22, pp. 392-423, 2005.
- [17] S. Dong, P.-T. Bremer, M. Garland, V. Pascucci, and J.C. Hart, "Spectral Surface Quadrangulation," *Proc. SIGGRAPH '06*, pp. 1057-1066, 2006.
- [18] Y. Tong, P. Alliez, D. Cohen-Steiner, and M. Desbrun, "Designing Quadrangulations with Discrete Harmonic Forms," *Proc. Fourth Eurographics Symp. Geometry Processing (SGP)*, 2006.
- [19] E. Zhang, J. Hays, and G. Turk, "Interactive Tensor Field Design and Visualization on Surfaces," *IEEE Trans. Visualization Computer Graphics*, vol. 13, no. 1, pp. 94-107, Jan./Feb. 2007.
- [20] J.L. Helman and L. Hesselink, "Visualizing Vector Field Topology in Fluid Flows," *IEEE Computer Graphics and Applications*, vol. 11, no. 3, pp. 36-46, May 1991.
- [21] T. Delmarcelle and L. Hesselink, "The Topology of Symmetric, Second-Order Tensor Fields," *IEEE Conf. Visualization*, pp. 140-147, 1994.
- [22] E. Praun, A. Finkelstein, and H. Hoppe, "Lapped Textures," *Proc. SIGGRAPH*, pp. 465-470, 2000.
- [23] G. Turk, "Texture Synthesis on Surfaces," *Proc. SIGGRAPH*, 2001.
- [24] L.Y. Wei and M. Levoy, "Texture Synthesis over Arbitrary Manifold Surfaces," *Proc. SIGGRAPH*, pp. 355-360, 2001.
- [25] J. Stam, "Flows on Surfaces of Arbitrary Topology," *Proc. SIGGRAPH*, 2003.
- [26] J.J. van Wijk, "Image Based Flow Visualization for Curved Surfaces," *Proc. IEEE Visualization*, pp. 123-130, 2003.
- [27] M. Fisher, P. Schröder, M. Desbrun, and H. Hoppe, "Design of Tangent Vector Fields," *Proc. SIGGRAPH*, p. 56, 2007.
- [28] H. Theisel, "Designing 2D Vector Fields of Arbitrary Topology," *Proc. Eurographics Conf.*, vol. 21, pp. 595-604, 2002.
- [29] W.-C. Li, B. Vallet, N. Ray, and B. Lévy, "Representing Higher-Order Singularities in Vector Fields on Piecewise Linear Surfaces," *IEEE Trans. Visualization and Computer Graphics*, vol. 12, no. 5, pp. 1315-1322, Sept./Oct. 2006.
- [30] E. Zhang, K. Mischaikow, and G. Turk, "Vector Field Design on Surfaces," *ACM Trans. Graphics*, vol. 25, no. 4, pp. 1294-1326, 2006.
- [31] G. Chen, G. Esch, P. Wonka, P. Müller, and E. Zhang, "Interactive Procedural Street Modeling," *ACM Trans. Graphics*, vol. 27, no. 3, pp. 1-10, 2008.
- [32] A. Hertzmann and D. Zorin, "Illustrating Smooth Surfaces," *Proc. SIGGRAPH*, pp. 517-526, 2000.
- [33] N. Ray, B. Vallet, W.C. Li, and B. Lévy, "N-Symmetry Direction Field Design," *ACM Trans. Graphics*, vol. 27, no. 2, pp. 1-13, 2008.
- [34] J. Palacios and E. Zhang, "Interactive Visualization of Rotational Symmetry Fields on Surfaces," *IEEE Trans. Visualization and Computer Graphics*, vol. 17, no. 7, pp. 947-955, July 2011.
- [35] Y.-K. Lai, M. Jin, X. Xie, Y. He, J. Palacios, E. Zhang, S.-M. Hu, and X. Gu, "Metric-Driven Rosy Field Design and Remeshing," *IEEE Trans. Visualization and Computer Graphics*, vol. 16, no. 1, pp. 95-108, Jan./Feb. 2010.
- [36] M. Zhang, J. Huang, X. Liu, and H. Bao, "A Wave-Based Anisotropic Quadrangulation Method," *ACM Trans. Graphics*, vol. 29, pp. 1-8, 2010.
- [37] K. Crane, M. Desbrun, and P. Schröder, "Trivial Connections on Discrete Surfaces," *Computer Graphic Forum*, vol. 29, no. 5, pp. 1525-1533, 2010.
- [38] N. Ray, B. Vallet, L. Alonso, and B. Lévy, "Geometry Aware Direction Field Design," *ACM Trans. Graphics*, vol. 29, no. 1, 2009.
- [39] K. Xu, D. Cohen-Or, T. Ju, L. Liu, H. Zhang, S. Zhou, and Y. Xiong, "Feature-Aligned Shape Texturing," *ACM Trans. Graphics*, vol. 28, no. 5, pp. 1-7, 2009.
- [40] M. Meyer, M. Desbrun, P. Schröder, and A.H. Barr, "Discrete Differential Geometry Operators for Triangulated 2-Manifolds," *Visualization and Mathematics III*, Springer, 2002.
- [41] E. Zhang, H. Yeh, Z. Lin, and R.S. Laramée, "Asymmetric Tensor Analysis for Flow Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 15, no. 1, pp. 106-122, Jan./Feb. 2009.
- [42] J.J. Koenderink and A.J. van Doorn, "Surface Shape and Curvature Scales," *Image and Vision Computing*, vol. 10, pp. 557-565, Oct. 1992.
- [43] A.N. Hirani, "Discrete Exterior Calculus," PhD dissertation, 2003.
- [44] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle, "Multiresolution Analysis of Arbitrary Meshes," *Proc. SIGGRAPH*, pp. 173-182, 1995.
- [45] K. Polthier and E. Preuss, "Identifying Vector Field Singularities Using a Discrete Hodge Decomposition," *Visualization and Mathematics III*, pp. 113-134, Springer Verlag 2003.
- [46] J. Erickson and K. Whittlesey, "Greedy Optimal Homotopy and Homology Generators," *Proc. ACM-SIAM Symp. Discrete Algorithms*, pp. 1038-1046, 2005.
- [47] F. Kälberer, M. Nieser, and K. Polthier, "Stripe Parameterization of Tubular Surfaces," *Topological Methods in Data Analysis and Visualization: Theory, Algorithms and Applications*, Springer Verlag, 2009.
- [48] L.-Y. Wei, S. Lefebvre, V. Kwatra, and G. Turk, "State of the Art in Example-Based Texture Synthesis," *Proc. Eurographics*, 2009.
- [49] Y. Liu, W.-C. Lin, and J. Hays, "Near-Regular Texture Analysis and Manipulation," *Proc. SIGGRAPH*, pp. 368-376, 2004.
- [50] P. Alliez, G. Ucelli, C. Gotsman, and M. Attene, "Recent Advances in Remeshing of Surfaces," *Shape Analysis and Structuring*, Springer, 2005.
- [51] X. Gu, S.J. Gortler, and H. Hoppe, "Geometry Images," *ACM Trans. Graphics*, vol. 21, no. 3, pp. 355-361, 2002.
- [52] P. Alliez, M. Meyer, and M. Desbrun, "Interactive Geometry Remeshing," *Proc. SIGGRAPH*, pp. 347-354, 2002.
- [53] P. Alliez, E.C. de Verdiere, O. Devillers, and M. Isenburg, "Isotropic Surface Remeshing," *Proc. Shape Modeling Int'l (SMI)*, p. 49, 2003.

- [54] A. Khodakovskiy, N. Litke, and P. Schröder, "Globally Smooth Parameterizations with Low Distortion," *Proc. SIGGRAPH '03*, 2003.
- [55] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Mesh Optimization," *Proc. SIGGRAPH*, pp. 19-26, 1993.
- [56] P. Lindstrom and G. Turk, "Image-Driven Simplification," *ACM Trans. Graphics*, vol. 19, no. 3, pp. 204-241, 2000.
- [57] V. Surazhsky and C. Gotsman, "Explicit Surface Remeshing," *Proc. ACM SIGGRAPH Symp. Geometry Processing (SGP)*, 2003.
- [58] G. Turk, "Re-Tiling Polygonal Surfaces," *Proc. SIGGRAPH*, 1992.
- [59] D.-M. Yan, B. Lévy, Y. Liu, F. Sun, and W. Wang, "Isotropic Remeshing with Fast and Exact Computation of Restricted Voronoi Diagram," *Proc. Symp. Geometry Processing (SGP)*, pp. 1445-1454, 2009.
- [60] E. Aklaman, J. Chen, Q. Xing, and J.L. Gross, "Cyclic Plain-Weaving on Polygonal Mesh Surfaces with Graph Rotation Systems," *ACM Trans. Graphics*, vol. 28, no. 3, pp. 1-8, 2009.



Matthias Nieser is currently working toward the PhD degree at the Department of Mathematics at Freie Universität Berlin, studying under Konrad Polthier. His current research focuses on discrete differential geometry, in particular the parameterization and structuring of surfaces and volumes. He is also a member of the DFG research center MATHEON.



Jonathan Palacios is currently working toward the PhD degree at the Department of Electrical Engineering and Computer Science at Oregon State University, studying under Dr. Eugene Zhang. His primary research areas are computer graphics, geometric modeling, symmetry, and higher order tensor field visualization and analysis. He is a US National Science Foundation (NSF) IGERT fellow, and a member of the ACM.



Konrad Polthier received the PhD degree from the University of Bonn in 1994, and headed research groups at Technische Universität Berlin and Zuse-Institute Berlin. He is a professor of mathematics at Freie Universität Berlin and DFG research center MATHEON, and chair of the Berlin Mathematical School. His current research focuses on discrete differential geometry and geometry processing. He co-edited several books on mathematical visualization, and coproduced mathematical video films. His recent video MESH (www.mesh-film.de, joint with Beau Janzen) has received international awards including "Best Animation" at the New York International Independent Film Festival. He served as paper or event cochair on international conferences including Symposium on Geometry Processing in 2006 and 2009.



Eugene Zhang received the PhD degree in computer science from Georgia Institute of Technology in 2004. He is currently an associate professor at Oregon State University, where he is a member of the School of Electrical Engineering and Computer Science. His research interests include computer graphics, scientific visualization, geometric modeling, and computational topology. He received the US National Science Foundation (NSF) CAREER award in 2006. He is a senior member of the IEEE and the ACM.

► **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**