

Spectral Quadrangulation with Orientation and Alignment Control

Jin Huang[†] Muyang Zhang Jin Ma Xinguo Liu[†] Leif Kobbelt* Hujun Bao[†]
State Key Lab. of CAD&CG, Zhejiang University *RWTH Aachen University




Figure 1: Quadrangulation on Rockarm model. Figure (a) shows the quasi-dual Morse-Smale complex of an unconstrained eigenfunction. Taking the direction (arrows), alignment (lines) and sizing (color) fields of figure (b) into account, our scheme computes a scalar function with the quasi-dual Morse-Smale complex shown in (c). In (d) we depict our final quadrangulation result.

Abstract

This paper presents a new quadrangulation algorithm, extending the spectral surface quadrangulation approach where the coarse quadrangular structure is derived from the Morse-Smale complex of an eigenfunction of the Laplacian operator on the input mesh. In contrast to the original scheme, we provide flexible explicit controls of the shape, size, orientation and feature alignment of the quadrangular faces. We achieve this by proper selection of the optimal eigenvalue (shape), by adaption of the area term in the Laplacian operator (size), and by adding special constraints to the Laplace eigenproblem (orientation and alignment). By solving a generalized eigenproblem we can generate a scalar field on the mesh whose Morse-Smale complex is of high quality and satisfies all the user requirements. The final quadrilateral mesh is generated from the Morse-Smale complex by computing a globally smooth parameterization. Here we additionally introduce edge constraints to preserve user specified feature lines accurately.

Keywords: quadrangular remeshing, Laplacian eigenfunctions, constrained optimization

1 Introduction

Although triangular meshes have been widely used in computer graphics community, quadrilateral meshes are more desired in some applications, such as texturing, simulation with finite elements and B-spline surface reconstruction, etc. Many approaches have been proposed to remesh triangular meshes into quadrilateral meshes. However, it is a challenging problem to generate a satisfactory

quadrilateral mesh with few irregular vertices and faces. Furthermore, the size, edge orientation and feature alignment of quadrilateral mesh should be under control in many applications, which add more difficulties.

In this paper, we propose a controllable spectral method to remesh triangular meshes into quadrilateral meshes. We focus on the control of the Morse-Smale complex according to the edge length of the coarse quadrilateral, direction field guided edge orientation and the alignment of feature lines. Users can use weighting to control the distance between nodes, assign a direction field over the mesh to make final quadrangulation result coincide with it and specify the feature lines to be preserved when remeshing. Most of the above constraints can be automatically derived from principal curvature directions and feature detections. To ensure better feature alignment, we also augment the global smooth parameterization by introducing edge constraints for feature lines.

2 Related Work

The problem about quadrangulation has been a focus of graphics research in many years. A detailed survey for remeshing can be found in [Alliez et al. 2005]. We only review the most related works here.

Direction field is an important guidance for the quadrangulation, and is widely used. The method proposed by [Alliez et al. 2003] integrates principal curvature lines in the parameter domain of the input mesh and constructs a quadrilateral dominant mesh through the intersections of the curvature lines. Marinov and Kobbelt [Boier-Martin et al. 2004] enhanced this method by directly integrating the curvature lines on the surface. Dong et al. [2005] create a harmonic function, then tracking the iso-lines and gradient lines to generate the quadrilateral mesh. But it is a challenging problem to create a harmonic function whose gradient aligns with the given direction field.

Global parameterization is another useful tool for quadrangulation. Periodic parameterization technique, proposed in [Ray et al. 2006], can automatically parameterize the input model under the parameter domain, whose coordinate axes coincide with the principle directions. It generates quadrilateral meshes with controlled orientation of mesh faces. To generate good results, non-linear optimization is needed together with an integrable direction field. Removing

[†]Corresponding authors: {hj,xgliu,bao}@cad.zju.edu.cn

the curl of the direction field will decrease the number of singular regions, but may stretch the parameterization greatly. As a consequence, it's hard to ensure good shape of quadrilateral faces while keeping a small number of irregular vertices. In [Kälberer et al. 2007], by converting a given frame field into a single vector field on a branched covering space, they can achieve high quality quadrangulation results with much fewer irregular vertices. This method does not require the irregular vertices located at the corners of some coarse meta mesh. All irregular vertices in both methods have to be known a priori from the input direction field. For noisy input direction field, they cannot optimize the irregular vertex positions for good quadrangulation.

Another way to control the quadrangulation result is to decompose a mesh into coarse patches which satisfy user's intention, and then convert these patches into quadrilaterals. For the CAD and CAE models, Marinov and Kobbelt [2006] propose a patch based quadrangulation method. They divide the model into a few polygonal patches according to the variational shape approximation, then remesh the patches individually into quadrilaterals. With given coarse meta mesh, [Tong et al. 2006] presents a linear solution with the consideration of Gaussian curvature, and achieves low distortion parameterization results. With the assistance of the curvature analysis proposed in [Tong et al. 2006], user specifies a quadrangular base domain as well as the types of singular continuity at the edges for further global parameterization. When the quadrangular base domain is not good enough, the result may contain fold and significant stretch.

[Dong et al. 2006] can automatically create coarse quadrangular domain by the Morse-Smale complex which comes from a eigenfunction of the mesh Laplacian. Then they track the iso-line of the global multi-chart parameterization to generate the final quadrangulation results. By virtue of the property of eigenfunctions, the coarse quadrangular complex contains only a few number of irregular vertices without any non-quad polygon. But the method is not good enough for controllable quadrangulation because of lacking controls over size, orientation and feature alignment.

3 Finding the Optimal Morse-Smale Complex

Similar to [Dong et al. 2006] we derive a coarse global quadrangular structure by computing the Morse-Smale Complex (MSC) of a scalar function f defined at mesh vertices. Our algorithm refines the idea of spectral surface quadrangulation in the sense that we provide flexible user controls over the shape, size, orientation and alignment of the resulting quadrangular structure. In particular we want to achieve the following goals:

- 1) In order to generate quadrangular patches with rectangular shape, the stationary points should be distributed periodically in two (locally) orthogonal directions across the mesh.
- 2) The size of these patches should follow a user specified sizing field.
- 3) The orientation of the two (locally) orthogonal directions should be consistent with the user specified tangent direction field.
- 4) The boundaries of some quadrangular patches should properly align with some feature lines on the mesh.

Notice that throughout this paper we make a distinction between *orientation* and *alignment* of a quadrilateral mesh. The term *orientation* refers to the rotational degree of freedom that is used to take the user specified direction field into account. The term *alignment* refers to the phase-shift degree of freedom that controls the actual location of the edges. With mere orientation control, we can only make the edges of a quadrilateral mesh parallel to the feature lines.

The addition of alignment control enables a parallel shift of mesh edges such that they lie exactly on the feature lines (see Figure 6).

As empirically observed in [Dong et al. 2006], the eigenfunctions of the Laplacian operator do have the potential to satisfy requirements (1) and (2) but lacks analysis of under which circumstances it actually happens. As we will show, these requirements can be achieved by selecting a proper eigenvalue λ . For the requirements (3) and (4) we have to add weighted constraints to the original eigenproblem.

By taking these additional constraints into consideration, our solution is no longer an exact eigenfunction to the Laplacian operator but rather a close enough approximation that inherits most of the nice analytical properties of the true eigenfunction while at the same time satisfying all the above requirements.

3.1 Setup

Let $\mathcal{M} = (V, T)$ be the input mesh with V the set of vertices and T the connectivity information. Each vertex $\mathbf{v}_i \in V$ is equipped with a position \mathbf{p}_i and a tangent vector \mathbf{d}_i which provides a 4-symmetry direction field on the mesh for the orientation of the final quadrilateral mesh.

A scalar function f on \mathcal{M} is defined by assigning a scalar value f_i to each vertex \mathbf{v}_i . In order to transfer concepts from differential geometry to this discrete setting, we compute a local quadratic approximation to f for the one-ring neighborhood of each vertex.

We first assign local parameter values (u, v) to \mathbf{v}_i and each of its adjacent neighbor vertices $\mathbf{v}_j \in \mathcal{N}(i)$ by exponential map. The rotational degree of freedom is used to orientate the u -direction in parameter space to the projected tangent direction \mathbf{d}_i associated with \mathbf{v}_i (we will exploit this orientation of the parametrization later in Section 3.4). Then we find the best fitting quadratic polynomial

$$q_i(u, v) = c_0^i + (c_u^i, c_v^i) \begin{pmatrix} u \\ v \end{pmatrix} + \frac{1}{2} (u, v) \begin{pmatrix} c_{uu}^i & c_{uv}^i \\ c_{uv}^i & c_{vv}^i \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix}$$

by minimizing the error functional

$$E(q_i) = \sum_{j \in \{i\} \cup \mathcal{N}(i)} D_j (q_i(u_j, v_j) - f_j)^2 \quad (1)$$

where the weight factor D_j takes the relative surface area associated with each vertex into account

$$D_j = \frac{1}{3} \sum_{t \in \mathcal{N}(j)} |t|. \quad (2)$$

Since the least squares minimizer of (1) is found by solving the corresponding normal equations, we can derive the linear operators Q_*^i for $*$ $\in \{0, u, v, uu, uv, vv\}$ such that

$$c_*^i = \langle Q_*^i, f \rangle.$$

Notice that these Taylor-operators Q_*^i depend only on the local parametrization of the input mesh \mathcal{M} and the direction field spanned by the vectors \mathbf{d}_i . Hence we can precompute these operators in advance and reuse them for any function f defined on \mathcal{M} . The local operator Q_*^i is the i -th row of the matrix operator \mathbf{Q}_* which maps the vector of all function values f to the vector of corresponding Taylor-coefficients \mathbf{c}_* . This operator will be used later in the optimization of f .

In order to optimally capture the geometry of the input mesh \mathcal{M} , we use the well-established cotangent form Laplacian operator \mathbf{L} [Pinkall and Polthier 1993]:

$$L(f_i) = \sum_{j \in \mathcal{N}(i)} (\cot(\alpha_{ij}) + \cot(\beta_{ij})) (f_j - f_i)$$

with α_{ij}, β_{ij} being the two angles opposite to the edge $[i, j]$ in the mesh \mathcal{M} . In order to keep the Laplacian operator \mathbf{L} symmetric, we do not normalize by the reciprocal surface area $1/D_i$ (cf. equation 2). Instead we formulate a generalized eigenproblem

$$\mathbf{L}f = -\lambda \mathbf{D}f \Leftrightarrow \begin{cases} \min_{\|g\|=1} \left\| (\sqrt{\mathbf{D}}^{-1} \mathbf{L} \sqrt{\mathbf{D}}^{-1} + \lambda I) g \right\|^2 \\ \text{where } f = \sqrt{\mathbf{D}}^{-1} g \end{cases} \quad (3)$$

where the diagonal matrix \mathbf{D} with area elements D_i appears explicitly. We will use this formulation later in Section 3.3 to adjust the local cell size of the MSC.




Figure 2: Example of a Laplacian eigenfunction $f(x,y) = \cos(\pi x) \cos(\pi y)$. The corresponding Morse-Smale complex is depicted by black arrows.

3.2 Square Patches

Let f be an eigenfunction of the discrete Laplacian operator, i.e.,

$$\mathbf{L}f = -\lambda f. \quad (4)$$

The Laplacian operator is tightly related to discrete cosine transform [Strang 1999]. By intuition, we extend the relationship to locally developable surface: there exist two intrinsic parameter directions x and y such that f is a discrete analogon to a function of the form

$$f(x,y) = A \cos(\alpha x + \phi_x) \cos(\beta y + \phi_y) \quad (5)$$

where α, β define the frequencies in x and y directions respectively, A is the amplitude, (ϕ_x, ϕ_y) defines the phase shift, and the eigenvalue is $\lambda = \alpha^2 + \beta^2$. The corresponding MSC is an affine grid consisting of two sets of lines with a slope of (β, α) and $(-\beta, \alpha)$ respectively. This affine grid turns out to be an orthogonal one if and only if $\alpha = \beta$. See Figure 2 for a depiction of this situation.

We use the ‘‘Multiresolution Spectral Analysis’’ technique [Dong et al. 2006] to solve the generalized eigenproblem on a large mesh, and then pick the eigenfunction for producing the best results. To find an eigenvalue λ whose eigenfunction f yields an as orthogonal as possible MSC, let

$$E_{\text{square}} = \frac{\alpha^2}{\beta^2} + \frac{\beta^2}{\alpha^2} = \frac{\alpha^4 + \beta^4}{\alpha^2 \beta^2} \quad (6)$$

be a measure for the ‘‘non-orthogonality’’ of the MSC grid. Then we compute a small number of eigenfunctions $f[\lambda_k]$ for several eigenvalues λ_k around the estimated λ which is selected according to the number of critical points in the complex (c.f. [Dong et al.

2006]). For each $f[\lambda_k]$ we check the non-orthogonality measure (6) at all the extrema (in the simplified MSC). This will produce orthogonal candidates to select from. Figure 3 shows an example of how this procedure allows us to find out the most orthogonal MSC.

To evaluate E_{square} , we consider the Hessian of (5)

$$H[f(x,y)] = A \begin{pmatrix} -\alpha^2 cc & \alpha\beta ss \\ \alpha\beta ss & -\beta^2 cc \end{pmatrix} \quad (7)$$

where we set $cc = \cos(\alpha x + \phi_x) \cos(\beta y + \phi_y)$ and $ss = \sin(\alpha x + \phi_x) \sin(\beta y + \phi_y)$. At an extremum, we have $cc = \pm 1$ and $ss = 0$. Hence

$$H[f(\frac{\pi i - \phi_x}{\alpha}, \frac{\pi j - \phi_y}{\beta})] = \begin{pmatrix} \mp A \alpha^2 & 0 \\ 0 & \mp A \beta^2 \end{pmatrix}.$$

This means that E_{square} can be expressed by the ratio of the eigenvalues of the Hessian and hence it does not depend on the amplitude A (which may vary at different extrema).

Since for a given eigenfunction f we cannot easily derive the intrinsic directions x and y , it is quite difficult to exploit (5) directly. However, we can replace $f(x,y)$ by the local quadratic fit $q(u,v)$ whose Hessian

$$H[q(u,v)] = \begin{pmatrix} c_{uu} & c_{uv} \\ c_{uv} & c_{vv} \end{pmatrix}$$

has approximately the same eigenvalues as $H[f(x,y)]$. This is obvious because the rotation between the (u,v) and the (x,y) coordinate systems only affect the eigenvectors, not the eigenvalues. From

$$\begin{aligned} \text{trace } H[q(u,v)] &= c_{uu} + c_{vv} \approx -A(\alpha^2 + \beta^2) \\ \det H[q(u,v)] &= c_{uu} c_{vv} - c_{uv}^2 \approx A^2 \alpha^2 \beta^2 \end{aligned}$$

at the extrema we conclude that

$$E_{\text{square}} = \frac{\alpha^4 + 2\alpha^2 \beta^2 + \beta^4}{\alpha^2 \beta^2} - 2 \approx \frac{(c_{uu} + c_{vv})^2}{c_{uu} c_{vv} - c_{uv}^2} - 2.$$

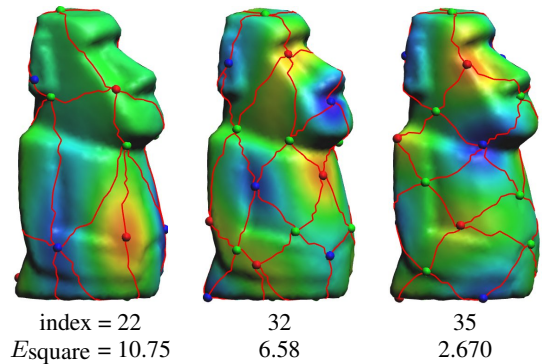


Figure 3: Finding eigenfunctions with approximately orthogonal MSC. By testing eigenfunctions $f[\lambda_k]$ for eigenvalues λ_k around the estimated λ (center column), we can find the most orthogonal eigenfunction (right column) among them.

3.3 Adaptive Size Control

In order to preserve fine details in some surface regions without introducing redundant over-tessellation in other regions, we would

like to vary the distance between stationary points locally in a controlled fashion. Therefore, we exploit the physical meaning of the Laplace eigenproblem.

The generalized eigenproblem in (3) can be viewed as a vibration analysis on the elastic membrane given by the input mesh. As discussed in [Vallet and Lévy 2007], we can interpret the Laplacian matrix \mathbf{L} as a stiffness matrix, and \mathbf{D} as a lumped mass matrix. The eigenvalue λ is related to the vibration frequency by $\omega = \sqrt{\lambda}$ and the eigenfunction f represents the corresponding stationary wave.

From the theory of wave propagation in an elastic body [Alford et al. 1974], we have the following relationship:

$$\text{wave speed} = \text{frequency} \times \text{wavelength} = k_1 \sqrt{\frac{\text{stiffness}}{\text{mass density}}}$$

and for the distance l between two adjacent extrema we obtain

$$l = k_2 \frac{1}{\sqrt{\lambda \mathbf{D}}} \quad (8)$$

where k_1, k_2 are constant coefficients.

Due to the surface area weights D_i in (3), the underlying elastic body is assumed to have uniform material properties. Hence, the wave speed is nearly constant over the mesh and thus for a given eigenvalue (frequency), the stationary points of the eigenfunction are uniformly distributed. From (8), we can also understand why the stationary points move closer to each other when the eigenvalue (frequency) is increasing.

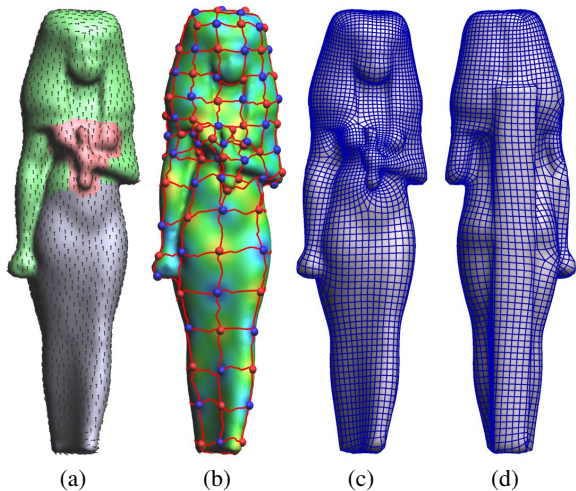


Figure 4: Creating a quadrangulation with adaptive density by changing the surface area weights \mathbf{D} . (a) Local surface areas are scaled by 1, 3, and 6 in the gray, green, and red regions respectively. (b) The quasi-dual Morse-Smale complex constructed from the generalized eigenproblem. (c)(d) The front and back view of the quadrangulation result with smaller quads in the green and red regions.

To control the cell size of the MSC, i.e., the density distribution of the eigenfunction’s stationary points, we can simply scale the corresponding entries in the lumped mass matrix \mathbf{D} up (to increase the density) or down (to decrease it). Thus we can effectively adapt the resulting quadrilateral mesh density for better feature preserving. See Figure 4 for an example.

In [Dong et al. 2006], the eigenproblem is solved without considering the mass matrix \mathbf{D} and consequently the distance between

stationary nodes in the MSC strongly depends on the tessellation of the input mesh. Our adaptive size control also provides a physical interpretation of the partial shifting technique used in [Dong et al. 2006] which reduces the diagonal entries of the Laplacian matrix for some vertices. From the relation:

$$(\mathbf{L} - \mathbf{S})f = \lambda f \quad \Leftrightarrow \quad \mathbf{L}f = \lambda(\mathbf{I} + \mathbf{S})f$$

with \mathbf{S} being a diagonal matrix, we find that partial shifting actually corresponds to adding offsets to the corresponding entries of the mass matrix (while we apply a scaling factor). In the region with partial shifting, the wavelength becomes much shorter, so extrema more likely appear.

3.4 Orientation Control

In section 3.2 we already exploited the Hessian (of the continuous analogon) of the eigenfunction f to identify the function which leads to a MSC with an orthogonal structure. Here we take this idea even further by using the coefficients of the Hessian for orientation control.

For eigenfunctions f with orthogonal MSC, we can assume the directional frequencies α and β to be equal and hence the Hessian of the continuous analogon (5) simplifies to

$$H[f(x, y)] = A\alpha^2 \begin{pmatrix} -cc & ss \\ ss & -cc \end{pmatrix}. \quad (9)$$

The eigenvectors $\mathbf{e}_+ = \begin{pmatrix} 1 \\ 1 \end{pmatrix}$ and $\mathbf{e}_- = \begin{pmatrix} -1 \\ 1 \end{pmatrix}$ of this matrix, i.e., the principal directions of f do not change with (x, y) . They coincide with the two orthogonal directions of the corresponding MSC.

The basic idea of orientating the MSC is to take an eigenfunction f which has an approximately orthogonal MSC and to “twist” it towards a function with the property that the principal directions (i.e., the eigenvectors of the Hessian) coincide with the prescribed direction field. Our approach is to add an energy term to the Laplace eigenproblem which penalizes the deviation of the principal directions.

Again, instead of using the Hessian of (5) directly we use the Hessian of the local quadratic fit $q(u, v)$. Since we have oriented the (u, v) -coordinate system such that its u -direction coincides with the prescribed direction field \mathbf{d}_i , we simply have to make sure that the principal directions of $q(u, v)$ coincide with the main axes in (u, v) -parameter space. This is guaranteed if the Hessian of $q(u, v)$ is a diagonal matrix. Hence, our local penalizing energy is

$$(\mathbf{c}_{uv}^i)^2 = \langle Q_{uv}^i, f \rangle^2.$$

The total penalizing energy is then obtained by integration over the entire mesh, i.e.,

$$E_{\text{orient}} = \sum_{\mathbf{v}_i \in V} D_i \langle Q_{uv}^i, f \rangle^2 = \|\mathbf{Q}_{\text{orient}} f\|^2. \quad (10)$$

As mentioned in [Dong et al. 2006], the angle between the edge direction of quasi-dual and primal complex is $\pi/4$, we can use the $\pi/4$ rotated direction field to make the orientation of a quasi-dual complex coincide with this direction field. See Figure 5 for orientation examples.

3.5 Alignment Control

When the input mesh has sharp or non-sharp feature lines, they should be represented by a sequence of polygon edges in the output.




Figure 5: Orientation control on the torus model. (a) make the edges of the quasi-dual Morse-Smale complex coincide with the principal curvatures (b) rotate the direction field in (a) by 20 degree. (c) rotate the direction field in (a) by 45 degree.

Hence, besides *orientation* control we have to introduce control on feature *alignment*. Many feature detection techniques have been proposed, for example [Yoshizawa et al. 2005]. In this paper we assume that the relevant feature lines on the input mesh are given.

In the local DCT approximation (5) of the eigenfunction f , we can observe that the edges of the primal MSC lie on the diagonals: $(\alpha x + \phi_x) \pm (\beta y + \phi_y) = k\pi$, and the edges of dual MSC are on: $\alpha x + \phi_x = k\pi$ and $\beta y + \phi_y = k\pi$. These lines represent the set of symmetry axes of f . Hence in order to make the MSC of the function f align to a certain feature line l , we enforce f to be symmetric with respect to l .

Let (\hat{u}, \hat{v}) be a local parametrization such that the \hat{u} -axis is aligned to the feature line l . Then the symmetry condition on f written in terms of the local quadratic approximation $q(\hat{u}, \hat{v})$ is

$$q(\hat{u}, -\hat{v}) = q(\hat{u}, \hat{v})$$

which is equivalent to

$$c_{\hat{u}\hat{v}} = 0 \quad \text{and} \quad c_{\hat{v}} = 0.$$

Thus, the penalizing energy for feature alignment over the whole mesh becomes:

$$E_{\text{align}} = \sum_{\mathbf{v}_i \in l} D_i (\langle Q_{\hat{u}\hat{v}}^i, f \rangle^2 + \langle Q_{\hat{v}}^i, f \rangle^2) = \|\mathbf{Q}_{\text{align}} f\|^2. \quad (11)$$

In the primal MSC case, we can expect that the feature is oriented along the prescribed tangent direction field \mathbf{d} . In this case the parametrization (\hat{u}, \hat{v}) coincides with (u, v) and the first condition $c_{\hat{u}\hat{v}} = 0$ simply repeats the orientation condition. Consequently the penalizing energy for misalignment to the feature line l is simplified as

$$E_{\text{align}} = \sum_{\mathbf{v}_i \in l} D_i \langle Q_{\hat{v}}^i, f \rangle^2. \quad (12)$$

3.6 Meshes with Boundaries

The original framework proposed in [Dong et al. 2006] cannot handle meshes with boundaries which are very common in practice. By treating boundaries of the mesh as feature lines, we can easily exploit the orientation and alignment control to handle them.

As shown in Figure 7 and Figure 8, we create a direction field which is perpendicular or parallel to the boundaries in the nearby region, and use the orientation control to ensure that some of the edges of the MSC are parallel with the boundaries. The alignment control finally snaps edges precisely to the boundaries. In Figure 8, we also add size controls to make the wavelength compatible with the alignment configuration (which will be discussed in section 5).




Figure 6: Using alignment control to snap some edges of the quasi-dual MSC in (b) to the specified feature lines in (a). (c) shows the quasi-dual MSC for the unconstrained eigenproblem.




Figure 7: Handle meshes with boundaries by treating boundaries as feature lines. (a) the direction field which is perpendicular to the boundaries (b) the unconstrained MSC which is not aligned with the boundaries. (c) some edges of the MSC are aligned with the boundaries by orientation and alignment control. (d) the final quadrangulation result.

Since there is no appropriate definition of the Laplacian operator on boundaries, we use the definition proposed in [Vallet and Lévy 2008]: if the edge $[i, j]$ is on the border, the term of $\cot(\beta_{ij})$ in ω_{ij} vanishes. This definition matches the FEM formulation with Neumann boundary conditions.

3.7 Constrained Eigenproblem

To eventually find a scalar function f on \mathcal{M} which satisfies all the requirements, we first select a proper λ and \mathbf{D} with the methods presented in Sections (3.2) and (3.3), then combine Equations (3), (10) and (11) to obtain:

$$\begin{aligned} \min_f \left(\left\| (\sqrt{\mathbf{D}}^{-1} L \sqrt{\mathbf{D}}^{-1} + \lambda \mathbf{I}) \sqrt{\mathbf{D}} f \right\|^2 \right. \\ \left. + w_1 \|\mathbf{Q}_{\text{orient}} f\|^2 + w_2 \|\mathbf{Q}_{\text{align}} f\|^2 \right) \\ \text{s.t. } \|\sqrt{\mathbf{D}} f\|^2 = 1 \end{aligned} \quad (13)$$




Figure 8: Quadrangulation on the car model. (a) the direction field. Different weights \mathbf{D} are indicated by different color. (b) the quasi-dual Morse-Smale complex, (c) the final quadrilateral mesh.

where w_1 and w_2 are the user specified weights to balance the orientation energy, alignment energy and the quality of MSC (see Figure 9). We experimentally find that $w_1 = 0.5$ and $w_2 = 1.0$ works well in all our results.

The above constrained optimization problem are iteratively solved by the following equation:

$$\begin{pmatrix} A & J_k^T \\ J_k & 0 \end{pmatrix} \begin{pmatrix} f_{k+1} \\ \nu \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (14)$$

where ν is the Lagrange multiplier, and

$$\begin{aligned} A &= \hat{L}^T \hat{L} + w_1 \mathbf{Q}_{\text{orient}}^T \mathbf{Q}_{\text{orient}} + w_2 \mathbf{Q}_{\text{align}}^T \mathbf{Q}_{\text{align}} \\ \hat{L} &= (\sqrt{\mathbf{D}}^{-1} L \sqrt{\mathbf{D}}^{-1} + \lambda \mathbf{I}) \sqrt{\mathbf{D}} \\ J_k &= f_k^T D. \end{aligned}$$

In (14) we take f_k to calculate the approximate Hessian matrix at iteration $k + 1$. By exploiting the block structure in (14), it can be solved efficiently through the following two equations with pre-factorized the matrix A :

$$\begin{aligned} (J_k A^{-1} J_k^T) \nu &= -1 \\ A f_{k+1} &= -\nu J_k^T. \end{aligned} \quad (15)$$

The iteration stops when the $\|\sqrt{\mathbf{D}}(f_{k+1} - f_k)\| < 10^{-7}$. In our experiments, about 20 iterations are enough.

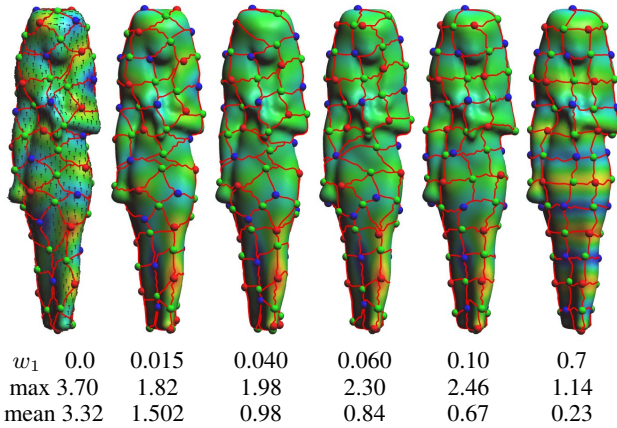


Figure 9: Visualization of how the eigenfunction f and the MSC change as more and more weight w_1 is put on the orientation. The other two rows show the maximum and average value of c_{uv}^2 .

After removing the topological noise using the algorithm proposed in [Dong et al. 2006] which includes cancellation (removing redundant saddle-extremum pairs) and anti-cancellation (the inverse of cancellation), we have constructed a quadrangular base complex over the mesh which satisfies all the input requirements. Then we build a parameterization over this complex and generate a regular $d \times d$ grid of quadrilaterals in this parametric domain with a user-specified density d .

4 Implementation Details and Results

Alike [Dong et al. 2006], we construct the MSC by the algorithm proposed in [Edelsbrunner et al. 2003; timo Bremer et al. 2004]. After classifying the critical points (maximum/minimum/saddle) of f , we construct

the steepest ascending/descending lines starting from each saddle to a maximum/minimum, and use them to partition the mesh into quadrangular regions. We also use the *cancellations* [Edelsbrunner et al. 2003; timo Bremer et al. 2004] which eliminates a connected saddle-extremum pair each time to simplify the MSC. The priority of cancellation is ranked by their *persistence* [Edelsbrunner et al. 2002]. In practice, we only perform the *anti-cancellation* for the very high valency (larger than 7) extremum of MSC, and split it along the longest edge. This procedure is actually rarely used, because high valency extrema rarely occur in our experiments. If the quasi-dual MSC is required, we connect the maximum-minimum diagonal within each quadrangular region on the simplified primal MSC.

The parameterization method proposed in [Dong et al. 2006] will relocate the edges of the MSC when swapping vertices across boundaries to adjust patches. As a result, after the iterative relaxation procedure, the parameterization often distorts the edge away from the feature lines. To accurately align the feature lines, we augment the original algorithm by *edge constraints*. The parameterization coordinates of vertices on the feature lines are interpolated from the corresponding nodes of the complex. Then we put these known values as hard constraints into the parameterization equations.

The analysis of the performance, some additional results and the quality of the final quadrilateral meshes will be demonstrated in the following.

	$ \mathcal{M} $	$ \mathcal{M}_q $	MSC	Parameterization
Car	2818	3360	0.65s	2.28s
Rockarm	9405	9301	2.32s	19.08s
Elephant	18074	18173	4.57s	51.19s
Pegaso	23930	16693	5.84s	83.57s

Table 1: Performance of our system.

Table 1 summarizes the size of some models and the performance of our system on these models. $|\mathcal{M}|$ and $|\mathcal{M}_q|$ are the number of vertices of the input triangle mesh and output quadrangulation mesh respectively. The column named as MSC shows the time (in seconds) for solving the constrained eigenproblem. And the column Parameterization lists the time used in the parameterization. Running times are measured on a 1.8G Intel DualCore 2 CPU with 2GB memory. We solve all the sparse linear equations by UMF-PACK [Davis 2004].

Our method is insensitive to the noise in the direction field, and can even apply to noisy geometry. In Figure 10, we directly use the direction field which comes from the curvature tensor and weight the orientation energy by the curvature tensor magnitude. Although the hair of the David head model is very noisy, the orientation control still works well. Even in the hair region, some singularities automatically appear to make the quadrangulation result consistent with some of the significant features in the direction field. Comparing with [Ray et al. 2006; Kälberer et al. 2007], it's an advantage that our method can optimize irregular vertex positions for such noisy direction field.

Finally, we measure the quality of our results. The number of irregular vertices, the distribution of the angle and edge length can be found in Figure 11. Both models are constrained by direction field, and size control is added on the face of pegaso horse to capture more details. By virtue of the period property of the eigenfunction, our result has only a few irregular vertices.

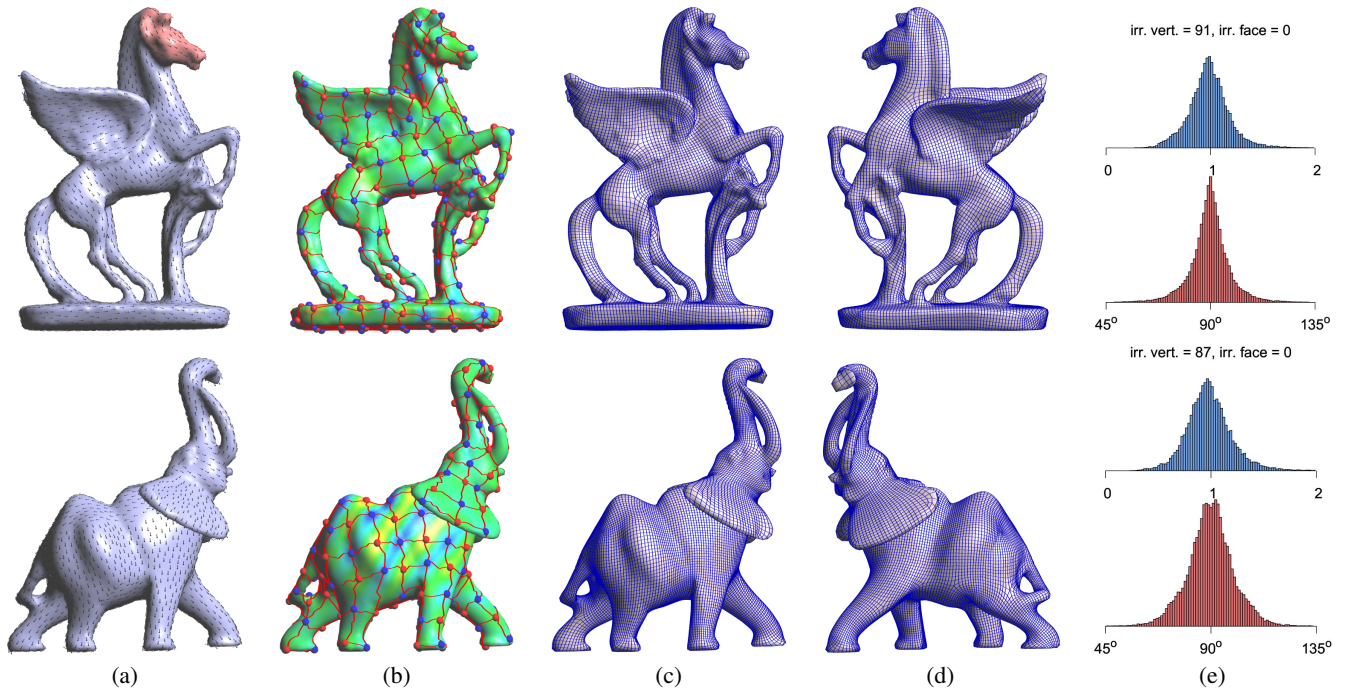


Figure 11: Quadrangulation on the pegaso horse and elephant model. (a) the direction field. Different weights \mathbf{D} are indicated by different color on the pegaso horse. (b) the quasi-dual Morse-Smale complex, (c)(d) the final quadrilateral mesh from the left and right view, (e) the plot of quality measurement.




Figure 10: Our method can directly apply to noisy geometry. (a) the direction field which comes from the curvature tensor. (b) the quasi-dual Morse-Smale complex. (c)(d) the final quadrilateral mesh from two views.

5 Discussion and Future Work

In this paper, we have presented a method to control the topological structure of the quadrangulation result by finding the optimal Morse-Smale complex. The scalar function from which the MSC is derived is approximated locally by DCT basis function and a quadratic polynomial. We look into the relationships between the MSC and the DCT basis function, then using the quadratic polynomial to represent these relationships. By solving a constrained least square problem which is dominated by a *generalized eigenproblem*, we find the best scalar function according to the requirements. After getting a good topological structure derived from the scalar function, we parameterize it over the surface. For better feature preserving, edge constraints are introduced into the global smooth parameterization framework.

Our method can generate the quadrangulation result with low distortion and only a few number of the irregular vertices. But in some cases, the result may not follow user’s intention, especially

when the applied constraints conflict with the period property of the eigenfunction. As shown in Figure 12, the Fandisk model has many corners, and the distribution of them is very irregular. It’s a tough job to set proper size, orientation and alignment controls to get a scalar function which has stationary points on every corner.




Figure 12: Using the size, orientation and alignment controls can get an approximately feature aligned result. But a perfect result requires more controls except for the ones proposed in this paper.

The alignment control cannot work well on some complex feature line configurations. It is hard to simultaneously align the edges of MSC with all feature lines if the distance between them is contradict with wavelength, e.g. the distance between two parallel feature lines is quite different from the integer times of the wavelength/4.

We also tried a constrain to ensure the corners to be nodes of the MSC (see Figure 13), which is similar to the alignment control but without the requirement of orientation control. The nodes of the MSC are saddles or extrema of f and hence at the positions where nodes appear, the following equation must hold:

$$\begin{aligned} 0 &= \frac{\partial f(u,v)}{\partial u} = c_u = \langle Q_u, f \rangle \\ 0 &= \frac{\partial f(u,v)}{\partial v} = c_v = \langle Q_v, f \rangle. \end{aligned} \quad (16)$$

By including above linear equations as hard constraints into (13), we can impose node position constraints. However, applying po-




Figure 13: We can make nodes lie on the corners of the gargoyle model by position control. (a) without node position control. (b) the node positions are specified by the user on some vertices of the model (indicated by yellow spheres), (c) the Morse-Smale Complex with node position control.

sition constraints at sparse locations on the mesh does not guarantee that the arcs of the MSC properly align to the feature lines in between and hence alignment control along feature lines remains necessary.

It would be another interesting control to constrain both the position and valance of irregular vertex simultaneously. As pointed out in [Tong et al. 2006], the optimal quadrangulation should match the curvature tensor. The position and valance of a MSC node should reflect the local average of Gaussian curvature, which is heavily related to the Hessian of $f(x, y)$ and $q(u, v)$.

6 Acknowledgments

We would like to thank the reviewers for their valuable comments. This work is supported in partial by NSFC (No.60703039), the 973 Program of China (No.2002CB312100), the Program for New Century Excellent Talents in University of China (No. NCET-06-0516), Fok Ying Tung Education Fund, UMIC Research Cluster and AICES Graduate School (both funded by DFG).

References

ALFORD, R. M., KELLY, K. R., AND BOORE, D. M. 1974. Accuracy of finite-difference modeling of the acoustic wave equation. *Geophysics* 39, 6 (December), 834–842.

ALLIEZ, P., COHEN-STEINER, D., DEVILLERS, O., LÉVY, B., AND DESBRUN, M. 2003. Anisotropic polygonal remeshing. *ACM Trans. Graph.* 22, 3, 485–493.

ALLIEZ, P., UCELLI, G., GOTSMAN, C., AND ATTENE, M. 2005. Recent advances in remeshing of surfaces.

BOIER-MARTIN, I., RUSHMEIER, H., AND JIN, J. 2004. Parameterization of triangle meshes over quadrilateral domains. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, ACM, 193–203.

BOTSCH, M., AND PAULY, M. 2007. Course 23: Geometric modeling based on polygonal meshes. In *SIGGRAPH courses*.

DAVIS, T. A. 2004. A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. *ACM Trans. Math. Softw.* 30, 2, 165–195.

DONG, S., KIRCHER, S., AND GARLAND, M. 2005. Harmonic functions for quadrilateral remeshing of arbitrary manifolds. *Comput. Aided Geom. Des.* 22, 5, 392–423.

DONG, S., BREMER, P.-T., GARLAND, M., PASCUCI, V., AND HART, J. C. 2006. Spectral surface quadrangulation. *ACM Trans. Graph.* 25, 3, 1057–1066.

EDELSBRUNNER, H., LETSCHER, D., AND ZOMORODIAN, A. 2002. Topological persistence and simplification. *Discrete Computer Geometry* 28, 511–533.

EDELSBRUNNER, H., HARER, J., AND ZOMORODIAN, A. 2003. Hierarchical morse-smale complexes for piecewise linear 2-manifolds. *Discrete and Computational Geometry* 30, 1, 87–107.

FISHER, M., SCHRÖDER, P., DESBRUN, M., AND HOPPE, H. 2007. Design of tangent vector fields. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, ACM, 56.

GU, X., AND YAU, S.-T. 2003. Global conformal surface parameterization. In *SGP '03: Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, Eurographics Association, 127–137.

IGARASHI, T., MOSCOVICH, T., AND HUGHES, J. F. 2005. As-rigid-as-possible shape manipulation. *ACM Trans. Graph.* 24, 3, 1134–1141.

KÄLBERER, F., NIESER, M., AND POLTHIER, K. 2007. Quadcover - surface parameterization using branched coverings. *Computer Graphics Forum* 26, 3 (September), 375–384.

KHODAKOVSKY, A., LITKE, N., AND SCHRÖDER, P. 2003. Globally smooth parameterizations with low distortion. In *SIGGRAPH '03: ACM SIGGRAPH 2003 Papers*, ACM, 350–357.

LI, W.-C., RAY, N., AND LÉVY, B. 2006. Automatic and interactive mesh to t-spline conversion. In *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing*, Eurographics Association, 191–200.

MARINOV, M., AND KOBELT, L. 2004. Direct anisotropic quad-dominant remeshing. In *PG '04: Proceedings of the Computer Graphics and Applications, 12th Pacific Conference on (PG'04)*, IEEE Computer Society, 207–216.

MARINOV, M., AND KOBELT, L. 2006. A robust two-step procedure for quad-dominant remeshing. *Computer Graphics Forum* 25, 3 (Sept.), 537–546.

MEYER, M., DESBRUN, M., SCHRÖDER, P., AND BARR, A. H. Discrete differential-geometry operators for triangulated 2-manifolds.

PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 1, 15–36.

RAY, N., LI, W. C., LÉVY, B., SHEFFER, A., AND ALLIEZ, P. 2006. Periodic global parameterization. *ACM Trans. Graph.* 25, 4, 1460–1485.

SHI, L., AND YU, Y. 2004. Inviscid and incompressible fluid simulation on triangle meshes: Research articles. *Comput. Animat. Virtual Worlds* 15, 3-4, 173–181.

STRANG, G. 1999. The discrete cosine transform. *SIAM Review* 41, 1, 135–147.

TIMO BREMER, P., EDELSBRUNNER, H., HAMANN, B., AND PASCUCI, V. 2004. A topological hierarchy for functions on triangulated surfaces. *IEEE Transactions on Visualization and Computer Graphics* 10, 385–396.

TONG, Y., ALLIEZ, P., COHEN-STEINER, D., AND DESBRUN, M. 2006. Designing quadrangulations with discrete harmonic forms. In *SGP '06: Proceedings of the fourth Eurographics symposium on Geometry processing*, Eurographics Association, 201–210.

VALLET, B., AND LÉVY, B. 2007. Spectral geometry processing with manifold harmonics. *Technical Report* (April).

VALLET, B., AND LÉVY, B. 2008. Spectral geometry processing with manifold harmonics. *Computer Graphics Forum* 27, 2 (Apr.), 251–260.

YOSHIZAWA, S., BELYAEV, A., AND SEIDEL, H.-P. 2005. Fast and robust detection of crest lines on meshes. In *SPM '05: Proceedings of the 2005 ACM symposium on Solid and physical modeling*, ACM, 227–232.