

Spectral Surface Quadrangulation

Shen Dong*

Peer-Timo Bremer*


Michael Garland*

Valerio Pascucci†

John C. Hart*

*University of Illinois at Urbana-Champaign

†Lawrence Livermore National Laboratory



(a) Laplacian eigenfunction

(b) Morse-Smale complex

(c) Optimized complex

(d) Semi-regular remeshing

Figure 1: We quadrangulate a given triangle mesh by extracting the Morse-Smale complex of a selected eigenvector of the mesh Laplacian matrix. After optimizing the geometry of the base complex, we remesh the surface with a semi-regular grid of quadrilaterals.

Abstract

Resampling raw surface meshes is one of the most fundamental operations used by nearly all digital geometry processing systems. The vast majority of this work has focused on triangular remeshing, yet quadrilateral meshes are preferred for many surface PDE problems, especially fluid dynamics, and are best suited for defining Catmull-Clark subdivision surfaces. We describe a fundamentally new approach to the quadrangulation of manifold polygon meshes using Laplacian eigenfunctions, the natural harmonics of the surface. These surface functions distribute their extrema evenly across a mesh, which connect via gradient flow into a quadrangular base mesh. An iterative relaxation algorithm simultaneously refines this initial complex to produce a globally smooth parameterization of the surface. From this, we can construct a well-shaped quadrilateral mesh with very few extraordinary vertices. The quality of this mesh relies on the initial choice of eigenfunction, for which we describe algorithms and heuristics to efficiently and effectively select the harmonic most appropriate for the intended application.

Keywords: quadrangular remeshing, spectral mesh decomposition, Laplacian eigenvectors, Morse theory, Morse-Smale complex

1 Introduction

Meshes generated from laser scanning, isosurface extraction and other methods often suffer from irregular element and sampling ar-

tifacts of the process. Because these problems arise so easily and can hinder the accuracy and efficiency of subsequent operations, the ability to remesh surfaces with well-shaped well-spaced elements is an important tool for mesh processing.

Much of the remeshing work in the graphics literature focuses on triangle meshes, though many graphics and scientific applications benefit from good quadrilateral meshes. Such meshes should have as few extraordinary vertices as possible and their elements should have internal angles near 90° . Quadrilaterals are the preferred primitive in several simulation domains, including computational fluid dynamics, where extraordinary points can lead to numerical instability [Stam 2003]. Catmull-Clark subdivision of a poor mesh can yield wrinkles [Halstead et al. 1993], and the tensor-product NURBS patches still used in CAD/CAM production software work best on a mesh composed exclusively of quadrilaterals. Furthermore, decomposing a surface into well-shaped quadrangles simplifies the construction of a texture atlas.

We have developed a new approach for building a quadrangular base complex over a triangulated manifold of arbitrary genus. This approach is based on the Morse theorem that for almost all real functions, the Morse-Smale complex (reviewed in §4), consisting of the ridge lines that extend from its saddles to its extrema, forms quadrangular regions. To space these regions evenly over the surface, we choose as our real function a shape harmonic of the appropriate frequency, computed in §3 as an eigenvector of the Laplacian matrix of the input mesh. A new iterative relaxation algorithm described in §5 simultaneously improves this base mesh layout while computing a globally smooth parameterization used to generate the final semi-regular grid of well-shaped quadrilaterals.

The complete spectrum of the mesh defines two families of complexes: the primal Morse-Smale and their quasi-dual complexes, a construction we propose in §4.3. The quality of the final mesh is intimately tied to the choice of complex, a choice we make based on parametric distortion. Section 3 provides a detailed analysis of the Laplacian spectrum, using spectral shifts to efficiently limit computation only to the eigenvectors around a desired frequency.

The resulting method produces fully conforming semi-regular

*{shendong, ptbremer, garland, jch}@uiuc.edu †pascucci@llnl.gov

Copyright © 2006 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.

© 2006 ACM 0730-0301/06/0700-1057 \$5.00

quad-only meshes that §6 shows contain fewer extraordinary points, have better element quality and competitive geometric fidelity when compared to meshes produced by existing quadrangulation methods. Though this method is designed for uniform surface sampling, §3.2 explores its adaptation to follow large-scale sharp features using a selective feature-based shifting of the Laplacian matrix.

2 Related Work

Meshed surface patching and resampling touch on a number of closely interrelated and extensively studied areas. We review only the most relevant results here, leaving the details to survey articles on parameterization [Floater and Hormann 2004], remeshing [Alliez et al. 2005], surface simplification [Garland 1999], and mesh generation [Bern and Eppstein 1995].

Semi-Regular Triangle Remeshing. Semi-regular schemes map the input surface onto a triangulated base domain and then regularly sample each triangular patch by recursive subdivision. Eck *et al.* [1995] formed a triangular base mesh from the dual of a quasi-Voronoi surface decomposition whereas the MAPS system [Lee et al. 1998] used simplification. Normal meshes [Guskov et al. 2000; Friedel et al. 2004] improve the encoding efficiency using a multiresolution hierarchy, globally smooth parameterization [Khodakovsky et al. 2003] increases parametric smoothness across triangular base-domain boundaries, and non-linear parameterization methods [Schreiner et al. 2004] can further reduce distortion.

Meshing with Quadrilaterals. Geometry images [Gu et al. 2002] produce a fully regular quadrangulation by cutting the mesh into a single component mapped onto a square domain, whereas multi-chart geometry images [Sander et al. 2003] cut the mesh into regularly sampled patches, improving distortion at the expense of patch continuity. Hormann and Greiner [2000] present a most-isometric parameterization of individual patches. Boier-Martin *et al.* [2004] produce a fully conforming quadrilateral mesh by quadrangling and grid sampling a general patch decomposition.

Eck and Hoppe [1996] build a quadrangular base complex by constructing a maximal pairing over a triangular base complex, though such methods cannot guarantee a purely quadrangular result. Shimada *et al.* [1998] further explored this approach for planar finite element mesh generation, whereas the Q-Morph algorithm [Owen et al. 1999] enhanced it with an advancing front traversal.

Recent methods have proposed a rather different approach to generating quad-dominant meshes. Alliez *et al.* [2003] numerically compute integral lines of the two principal direction fields of the surface in a conformal parametric domain. The spacing of these lines is controlled by the local surface curvature, and vertices are created where two orthogonal lines intersect. The stability of this method hinges on carefully smoothing the curvature tensor field from which the principal directions are derived. Marinov and Kobbelt [2004] developed a non-parametric variant. Dong *et al.* [2005] similarly trace curves, though using a harmonic scalar field over the surface. By instead fitting parametric functions to the principal direction fields, Ray *et al.* [2005] developed a quad-dominant method that produces noticeably cleaner meshes than these tracing methods.

Spectral Methods. Spectral graph theory is a well-developed branch of mathematics and has produced many fascinating results [Chung 1997]. As the spectral decomposition of a mesh exposes a great deal of its structure, it has been successfully applied in many diverse ways. It defines a natural frequency domain over the mesh, providing an attractive formalism for surface smoothing

and mesh signal processing [Taubin 2000]. Retaining and quantizing only the most important frequency bands provides a very effective means for compressing the surface geometry [Karni and Gotsman 2000]. Eigenvectors corresponding to the first d non-zero eigenvalues can also automatically and aesthetically embed graphs in \mathbb{R}^d [Koren et al. 2002].

Morse Theory. Discrete Morse theory [Banchoff 1967; Edelsbrunner et al. 2003] and its related data structures provide the underpinnings for a variety of topological graphics applications for implicit surfaces [Stander and Hart 1997], volumetric isosurfaces [van Kreveld et al. 1997; Pascucci and Cole-McLaughlin 2002; Weber et al. 2002] and shape database searches [Hilaga et al. 2001]. Our approach builds on the Morse-Smale complex used for surface function simplification [Bremer et al. 2004] and the minimal Morse base domain described by Ni *et al.* [2004].

3 Laplacian Eigenfunctions

We seek, for an input oriented manifold mesh M of any genus, a well-defined quadrangulation of well-shaped evenly distributed all-quad elements with few extraordinary vertices. Our approach is built upon the property that the Morse-Smale complex (described later in §4) connecting the saddles and extrema of almost any real surface function $f : M \rightarrow \mathbb{R}$ quadrangulates the surface. Though this complex is well-defined for any non-degenerate f , the *quality* of the quadrangulation is intimately tied to the choice of f . For arbitrary choices of f , the resulting complex can very poorly quadrangulate the surface. This paper explores the key insight that shape harmonics evenly distribute their extrema and so serve as ideal functions from which to generate a quadrangled base domain. This section shows how to efficiently compute these harmonics of the input mesh as *eigenfunctions* of its Laplacian matrix.

3.1 Spectral Surface Analysis

The discrete Laplacian operator on piecewise linear functions over triangulated manifolds is

$$\Delta f_i = \sum_{j \in N_i} w_{ij}(f_j - f_i), \quad (1)$$

where N_i is the set of vertices adjacent to vertex i and w_{ij} is a scalar weight assigned to the directed edge (i, j) . For graphs, it is customary to use the combinatorial weights $w_{ij} = 1/\deg(i)$ in defining this operator. However, for 2-manifold surfaces the appropriate choice are the discrete harmonic weights

$$w_{ij} = \frac{1}{2}(\cot \alpha_{ij} + \cot \beta_{ij}). \quad (2)$$

Here α_{ij} and β_{ij} are the angles opposite the edge (i, j) . Pinkall and Polthier [1993] provide details on the derivation of these weights.

This formulation of the Laplacian is clearly a linear operator. We represent the function f by the column vector of its per-vertex values \mathbf{f} , rewriting Laplace's equation as

$$\Delta \mathbf{f} = -\mathbf{L}\mathbf{f}, \quad \text{where } L_{ij} = \begin{cases} \sum_k w_{ik} & \text{if } i = j, \\ -w_{ij} & \text{if } \text{edge}(i, j) \in M, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The eigenvalues $\lambda_1 = 0 \leq \lambda_2 \leq \dots \leq \lambda_n$ of \mathbf{L} form the *spectrum* of the mesh M and the corresponding eigenvectors $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$




Figure 2: The first 8 non-constant eigenfunctions over a 15×15 planar grid, plotted as heightfields.

of \mathbf{L} define piecewise linear functions over M of progressively higher frequencies [Taubin 2000]. These functions are the Laplacian *eigenfunctions* of the mesh.

Laplacian eigenfunctions represent the natural harmonics of a shape; in the physical domain they are the vibrational modes of the surface. For a planar grid, the eigenvectors of \mathbf{L} are the basis functions of the discrete cosine transform (see Figure 2). Similarly for a sphere or torus they are, respectively, the discrete spherical and toroidal harmonics. The eigenvalues identify the squared frequency of the corresponding eigenfunction.

For our goal of producing a well-shaped quadrangulation of M , Laplacian eigenfunctions have several crucial properties. Their critical points are well-spaced over the surface. Minima and maxima are interleaved in such a way that high valence nodes are extremely rare, except in cases where they are geometrically desirable. Multi-saddles almost never arise, thus practically guaranteeing that extraordinary points can only occur at extrema. Without these properties the Morse-Smale complex can produce an extremely poor quadrangulation of the surface.

It is also important to note that the eigenfunctions occur in order of increasing frequency, and hence in (roughly) increasing order of critical point count. Specifically, it is known that the number of nodal domains of the eigenfunction with eigenvalue λ_k is at most k [Courant and Hilbert 1953], although this is not a sharp upper bound. Thus it is fairly easy to select an eigenfunction of a desired complexity. This is also important from an efficiency standpoint, as it means we only need to compute the first k low frequency eigenvectors of the matrix. For this, we use the ARPACK sparse eigensystem solver, which implements an efficient iterative Arnoldi method.

In practice, we are generally interested in producing relatively sparse quadrangulations. We find that for most surfaces eigenfunctions in the range 40–80 will produce the most desirable complexes. Simple surfaces (sphere, torus) work well with lower harmonics whereas higher genus surfaces require higher harmonics. To capture features such as those of the polytope shown in Figure 3, for example, one needs more nodes and hence a higher frequency.




Figure 3: Surfaces whose “ideal” complex contains more nodes require higher frequency eigenfunctions. Here we see the 10th, 46th, and 108th harmonic of a sphere, dodecahedron, and corner-cut icosahedron, respectively.

3.2 Multiresolution Spectral Analysis

Solving for a substantial number of eigenfunctions on a large mesh can be quite costly, as the running time of the eigensolver will be super-linear in the number of eigenvectors being computed. Fortunately, these eigenfunctions are an intrinsic property of the *shape*, allowing multiresolution techniques to overcome this performance bottleneck.

As a general rule, surface simplification methods remove high frequency detail while preserving low frequency shape, and so should preserve the low end of the mesh spectrum. Figure 4 demonstrates this phenomenon in practice. We produced multiple approximations of the kitten model shown in Figure 1 using the QSLim algorithm [Garland and Heckbert 1997]. The spectral plots of mass-adjusted eigenvalues, which are simply scaled by the fraction of vertices remaining, are extremely similar, even after aggressive simplification. The number of critical points in the complex for each frequency—after noise removal; see Section 4.2—are also quite consistent across the various resolutions.




Figure 4: Mass-adjusted spectra and critical point count for the first 80 eigenfunctions of a progressively simplified model.

When remeshing a surface, we typically want to find the small range of eigenfunctions that will produce complexes with a given number of critical points. Our observation above leads to a very efficient way of finding this narrow frequency band. We coarsen the model to a small size, while preserving the topological type of the surface. On this coarse mesh, we compute the first k eigenfunctions, their eigenvalues, and Morse-Smale complexes. We can now select a target eigenvalue λ based on the number of critical points in the corresponding complex.

Given this target eigenvalue, we can compute a small number (e.g., 20) of eigenfunctions on the original mesh with eigenvalues close to λ . We do this by using a *spectral shift*, a feature supported by the ARPACK solver. If λ is an eigenvalue of \mathbf{L} with eigenvector \mathbf{e} , then \mathbf{e} is an eigenvector of $(\mathbf{L} - \sigma\mathbf{I})$ with eigenvalue $(\lambda - \sigma)$.

$$(\mathbf{L} - \sigma\mathbf{I})\mathbf{e} = (\lambda - \sigma)\mathbf{e} \quad (4)$$

Therefore, we can shift \mathbf{L} by the ideal frequency λ and the k' eigenvalues with smallest absolute value are those closest to the ideal frequency. Overall, this process takes less than 2 minutes for a model with 130,000 vertices.

Spectral shifting can also align eigenfunctions to user-defined features, such as sharp corners in CAD models. By applying an additional shift only to vertices on feature lines, the gradients of the eigenfunctions start to align with feature edges. Such a partial shift can be interpreted as giving feature vertices a smaller weight [Koren et al. 2002], making them more likely to become eigenfunction extrema than the Morse-Smale complex samples at a higher rate. Although, such a feature driven approach is not our focus we have included a small example in Figure 17 to demonstrate the versatility of this approach.

4 Building a Quadrangular Base Complex

Given a Laplacian eigenfunction, this section describes how to construct its Morse-Smale complex to coarsely quadrangulate the surface. It begins by reviewing the formal definition of the Morse-Smale complex and an algorithm for topological noise removal. It also describes the construction of a second family of quasi-dual complexes to increase the diversity of the pool of potential quadrangular base meshes. The geometric embedding of the complex is further improved in Section 5.

4.1 The Morse-Smale Complex

The Morse-Smale complex is a cellular decomposition of a scalar (Morse-Smale) function over a manifold, defined formally as the refinement of its ascending manifolds by its descending manifolds, but computed more practically by tracing lines of steepest ascent/descent. Given a function defined on the vertices of a triangulated manifold M , a vertex is labeled a *maximum/minimum* if its function value is higher/lower than those of its neighbors, *regular* if its lower neighbors form a connected chain, and a *saddle* otherwise. We compute steepest ascending/descending lines starting in each chain of higher/lower neighbors of all saddles until we reach a maximum/minimum. In general, these lines segment M into four-sided regions with two opposing saddles, a maximum, and a minimum as corners. In practice, a number of special cases exist and care must be taken to correctly handle degeneracies in the function [Edelsbrunner et al. 2003; Bremer et al. 2004; Ni et al. 2004].

4.2 Topological Noise Removal

In principle, our eigenfunctions are smooth and we encounter increasing numbers of critical points at progressively higher frequencies. However, numerical issues can result in the presence of high frequency noise, which manifests itself as clusters of superfluous critical points (see Figure 5a).




Figure 5: The 30th eigenfunction of a triple torus shows a typical noise pattern (a) that we filter out by topological simplification (b).




Figure 6: Morse-Smale complex before and after canceling u, v . One saddle, one extremum, four paths, and two cells are removed.

To remove the extraneous critical points we use *cancellations* that simplify the Morse-Smale complex [Edelsbrunner et al. 2003; Bremer et al. 2004]. Cancellations can be seen as a double edge contraction that removes a connected saddle-extremum pair, all paths incident to the saddle, and two 2-cells (see Figure 6). We rank

cancellations by their *persistence* [Edelsbrunner et al. 2002]—the difference in function value between the two critical points they remove—and greedily cancel critical points in order of increasing persistence, up to a noise threshold.

The amount of noise we have encountered in practice is minimal and well separated from the signal. To find a good threshold we first normalize all persistences with respect to the maximal spread in function value. This allows us to set coarse bounds on the minimal and maximal allowed noise threshold valid for all models. For all examples in the paper we have used a range of 0.05–0.5% persistence. Within this range we optimize the complex based on mesh quality criteria. In particular, we have found that the only problematic configurations are valence-2 extrema like those shown in Figure 5. We therefore, pick the threshold as the minimal persistence within the given range that removes the maximal number of valence-2 extrema.

In general, not every valence-2 extremum can be removed by choosing an appropriate noise threshold. It might have high persistence and be vital for a balanced base mesh. In this case, we perform an *anti-cancellation* (the inverse of cancellation) of the neighboring extremum with highest valence to raise the valence of the problem extremum to three. A similar refinement mechanism is described in §5.2.3, but driven by parametric distortion rather than valence.

4.3 Quasi-Dual Complexes

The set of all eigenfunctions defines an entire family of complexes over the surface, which we refer to as the *primal* complexes. From each primal complex, we can also derive a *quasi-dual* complex. Since the vertices of a Morse-Smale complex are two-colorable (extrema being one color and saddles the other) and saddles have valence four, computing the minimum-maximum diagonal within each Morse-Smale region creates another purely quadrangular complex. This quasi-dual complex contains about half the vertices (only the extrema) of the original complex and one patch per original saddle. In the (extremely rare) case of a multi-saddle, computing diagonals yields a polygonal patch that can be trivially quadrangulated. Figure 7 illustrates the process for a small example.




Figure 7: A primal complex (a) is replaced with min-max diagonals (b) to produce the quasi-dual complex (c).

These quasi-dual complexes serve to expand the pool of possible base meshes. This is advantageous, since it expands the symmetries of the object that can be captured. For instance, it is the quasi-dual rather than the primal family that yields optimal complexes on the torus (see Figure 8). The quasi-dual transformation effectively reverses one step of 4-8 subdivision [Velho and Zorin 2001], and rotates the initial complex by $\pi/4$. Quasi-dual complexes are also more compact than primal complexes. Roughly half the critical points of a Morse-Smale complex will be Morse saddles, and hence valence four nodes in the complex. Quasi-dual complexes do not contain such additional saddle points.

We construct quasi-dual complexes after denoising the Morse-Smale complex. During iterative relaxation (§5.2), we may need to perform anti-cancellations on the quasi-dual complex. These

are induced directly by the anti-cancellation of the corresponding Morse-Smale complex. In particular, each anti-cancellation on a quasi-dual complex adds one extremum and one patch and is akin to a traditional vertex split.




Figure 8: Complexes of torus eigenfunctions 8 and 16.

5 Parameterization and Remeshing

At this point, we have constructed a quadrangular base complex over the surface. We must now build a parameterization over this complex. We must also correct and optimize the geometry of the complex’s embedding on the surface. For all but the simplest shapes, some paths connecting extrema will be less than satisfactory. As they follow the gradient field of the eigenfunction, they may not follow the surface shape in a natural way. It is also possible for multiple paths to merge and follow the same edge chain, thus producing degenerate patches. These kinds of artifacts are clearly evident in the complex shown in Figure 11a, for example.

We resolve both problems simultaneously using a globally smooth parameterization algorithm inspired by the work of Khodakovsky *et al.* [2003]. We construct a parameterization over the initial quadrangular complex, and then use this parameterization to correct the patch shapes. We iterate this relaxation procedure until convergence, at which point all patches will be valid and well-shaped.

5.1 Globally Smooth Parameterization

For each quadrangular patch P_α , we seek to construct a parameterization ϕ_α mapping the patch onto the unit square $D_\alpha = [0, 1] \times [0, 1]$. We also define the *transition function* $\phi_{\alpha\beta} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ mapping the coordinates of a point in P_α to the corresponding coordinates w.r.t. P_β (see Figure 9). The transition functions between two arbitrary patches can be determined by pair-wise composition along a path in the dual graph of the complex. Note, however, that the resulting functions will only be path-invariant in the absence of extraordinary points.

To begin constructing coordinate charts for each patch, we arbitrarily mark one corner as the origin of its coordinate system. We orient all charts to form a right-handed coordinate system with the surface




Figure 9: We define a coordinate chart ϕ_α for patch P_α , as well as transition functions $\phi_{\alpha\beta}$ to the chart of each adjacent patch P_β .




Figure 10: A single relaxation step: (a) parameterization, (b) boundary adjustment, and (c) node relocation.

normal. Every vertex that is not a node of the complex is contained within exactly one patch P_α and assigned parametric coordinates $\mathbf{u}_i^\alpha = (s_i^\alpha, t_i^\alpha)$. The parametric coordinates for nodes of the complex are always constrained to lie at the corners of D , although they will have distinct parametric coordinates in each of their incident patches.

For the particular configuration of patch origins marked in Figure 9, the two coordinate transition functions are

$$\phi_{\alpha\beta}(s^\alpha, t^\alpha) = (s^\alpha - 1, t^\alpha), \quad \phi_{\alpha\gamma}(s^\alpha, t^\alpha) = (s^\alpha, t^\alpha + 1). \quad (5)$$

All other markings produce simply the inverses of these functions and/or compositions with rotations of $k\pi/2$.

We solve for the parameterization using a linear system with the usual structure [Floater and Hormann 2004]. For each vertex i , which is not a node of the complex, we have

$$\sum_{(j,\beta) \in N_i} \bar{w}_{ij} \left((\phi_{\beta\alpha} \mathbf{u}_j^\beta) - \mathbf{u}_i^\alpha \right) = 0, \quad (6)$$

where \bar{w}_{ij} are the discrete harmonic weights (Eq. 2) normalized to sum to 1. Other choices are possible, notably mean value coordinates [Floater 2003], but we find that quasi-conformality of discrete harmonic weights yields better elements in the final mesh.

Notice that some of the \mathbf{u}_j coordinates may be known quantities, if they are vertices that are nodes of the complex. Furthermore, they do not have unique coordinates, since they have different coordinates in each of their incident patches. We eliminate these variables from the system by rewriting the parameterization system as

$$\mathbf{u}_i^\alpha - \sum_{(j,\beta) \in A_i} \bar{w}_{ij} \left(\phi_{\beta\alpha} \mathbf{u}_j^\beta \right) = \sum_{j \in C_i} \mathbf{u}_j^\alpha, \quad (7)$$

where C_i are the members of N_i that are corners and A_i are the rest. This equation gives us 2 rows of a $2n \times 2n$ linear system.

5.2 Iterative Relaxation

At this point, we have an initial parameterization of the complex. As noted earlier, the complex itself is likely to be deficient in that multiple paths may merge, thus causing degeneracies in the patches. Furthermore, this is likely to cause the parameterization itself to be non-bijective, and hence invalid.

We have developed an iterative relaxation procedure to resolve these degeneracies that

1. builds a parameterization by solving linear system Eq. 7,
2. swaps vertices across boundaries to adjust patches, and
3. relocates nodes of the complex to better positions.

This process repeats until it computes a parameterization (1) that is not improved by vertex swapping (2) or node relocation (3).

Our basic approach to the problem is inspired by that taken by Khodakovsky *et al.* [2003]; however, two distinctive features of our problem lead to a fairly different algorithm. First, we are working with quadrangular rather than triangular complexes. Second, we encounter, and must successfully repair, far more significant defects in the initial complex. Whereas Khodakovsky *et al.* were able to craft a decimation technique that considered patch quality and parametric distortion in the construction of the base domain, we begin with a complex built by a purely topological method that is unaware of such geometric issues. Consequently, we must perform far more substantial node relocations and may even need to refine the complex during relaxation.

5.2.1 Adjusting Patch Boundaries

The essential problem with the parameterization computed over a degenerate complex is illustrated in Figure 10a. For each vertex, there is some patch P_α such that the parametric coordinates of the vertex under ϕ_α fall within D . We refer to this vertex as *in-range* with respect to this patch. As we clearly see in Figure 10a, where vertices are color-coded according to the patch for which they are in-range, the in-range regions can differ substantially from the patches themselves. The fundamental goal of relaxation is to conform the patches to their in-range vertex sets. In this first step of relaxation, we adjust the boundaries by iteratively swapping vertices between patches.

A vertex i may be swapped across an edge (i, j) where $i \in P_\alpha$ and $j \in P_\beta$ if it is out-of-range in P_α but in-range in P_β (i.e., $\mathbf{u}_i^\alpha \notin D$ and $\mathbf{u}_i^\beta \in D$). It is important to note that since multiple patch boundaries may cross the edge (i, j) , the patch into which i moves may not be adjacent to its current patch in the complex. We continue swapping vertices until no more swappable vertices remain.

5.2.2 Relocating Nodes of the Complex

After swapping, the overall shape of patches in the complex is generally improved. However, there will still be patch boundaries that merge and follow the same edge sequence. Indeed, the extent of such merging is likely to increase in areas where patch corners do not coincide with the natural “corners” of the in-range vertex set. We can see an example of this behavior in Figure 10b. Having performed one round of boundary adjustment, the patch shapes are markedly improved. However, poorly placed nodes, most notably the one on the nose, produce path merging. Consequently, in this second stage of relaxation we seek to reposition nodes of the complex wherever necessary.

For each node with incident merging paths, we construct the set of branch points at which the merged paths diverge (see Figure 10). We consider each branch point as a potential target to which we might relocate the current node. The node’s current position is also a potential target. A relocation candidate is valid if the parameterization is locally bijective in at least one incident patch sector and




Figure 11: Iterative relaxation on the initial complex produces well-shaped final patches equipped with a parameterization.

the corresponding boundary of D crosses the same mesh edge as the patch boundary in question. To select the best candidate position, we wish to choose the one whose parametric coordinates are closest to a corner of D . Each non-node candidate is already assigned coordinates by our parameterization. For the node itself, we compute coordinates using a mapping of its incident patches into the complex plane by the power transform $z^{4/k}$ [Khodakovsky *et al.* 2003; Ying and Zorin 2004].

For some number of nodes with incident merging paths, their current location will prove preferable to the alternatives provided by their branch points. There are also potentially many nodes whose incident paths do not merge. To locally optimize the complex, we allow all such nodes to relocate to any immediately adjacent vertex if that will reduce the parametric distortion.

5.2.3 Refining the Complex

In some circumstances, relaxation will converge while merged paths still remain. This is essentially the result of either (1) nodes of unnecessarily high degree or (2) patches of such size or complexity that the linear parameterization method is unable to find a valid parameterization. We address this problem by locally refining the complex. We refine by topological anti-cancellation, the inverse of the cancellation operation discussed in Section 4. Note that the refinement operations on primal and quasi-dual complexes are slightly different.

As long as the original complex was not unreasonably coarse, we can expect relatively few refinements to be necessary. Therefore, we perform only a single refinement at a time. We consider all branch points that were considered as targets for relocation, and refine the one furthest from its associated node. Having performed this refinement, we resume iterative relaxation.

Figure 11 shows a sample of the results produced by our iterative relaxation algorithm. We begin with an initial complex that may have poorly-shaped patches, such as the patch near the bunny’s shoulder. After relaxation, all patches are well-shaped and we have also constructed a globally smooth parameterization over the surface.

5.3 Mesh Generation

Once a valid parameterization has been built and degeneracies in the complex have been removed, we can produce a final semi-regular mesh. For each path in the complex, we must trace out the corresponding parametric boundary over the mesh. This gives us surface patches, each of which is equipped with a parametric mapping onto the unit square. Given a user-specified density d , we




Figure 12: Remeshing the torus with progressively higher harmonics, with both primal (above) and quasi-dual (below) complexes.

construct a regular $d \times d$ grid of quadrilaterals in this parametric domain and map their corners back onto the surface, thus producing our output mesh. By sampling at a fixed rate, we can trivially guarantee that the final mesh is fully conforming. We can also guarantee that extraordinary points can only occur at nodes of the complex. In fact, we expect extraordinary points only at extrema of the eigenfunction, as Morse saddles will have valence four.

5.4 Selecting a Base Complex

This process of parameterization, relaxation, and remeshing can be applied to any of the primal or quasi-dual complexes defined over the spectrum of the surface. However, we are interested in picking the eigenvector that will produce the best result. We have already discussed in Section 3 how we select a small frequency band according to a target number of critical points in the complex. Within this band, we normally select the complex with the lowest parametric distortion. Low distortion leads to well-shaped elements and we find that it is generally well correlated with final RMS error. This is most accurately done after iterative relaxation is complete; however, using the distortion after a single relaxation step produces substantially the same ordering of complexes at a fairly low cost.

In certain cases, we may specifically desire a complex whose edges are aligned with a predominant direction, as in Figure 13, or with ridge lines, as in Figure 17. It is a straightforward process to select the eigenfunction whose gradient field most closely follows any such user-specified orientation.

6 Results

We begin our analysis of the performance of our method with the torus, shown in Figure 12. This is a simple surface whose eigenfunctions, as we have mentioned earlier, are discretizations of continuous toroidal harmonics. The spectrum of the torus is in fact highly structured, and the eigenfunctions and complexes we extract exhibit near-perfect symmetry (see Figure 8). This regularity and symmetry are apparent in the remeshing output as well. We generated meshes using, from left to right, the 8th, 16th, and 32nd eigenfunctions. Each quadrangular patch was resampled with an 8×8 grid of quadrilaterals. Each of these meshes consists exclusively of valence-4 vertices. Higher harmonics lead to more patches in the base complex and more quadrilaterals in the final mesh. We can also clearly see that the quasi-dual complexes capture different




Figure 13: A Moai statue is remeshed and then rendered as a Catmull-Clark surface.

symmetries of the surface, and in this case provide a nearly ideal mesh of the torus.

Figure 13 shows a quadrangulation of a scanned Moai statue. The raw surface data (on which the raw complex is superimposed) is fairly noisy and the mesh is moderately irregular. Nevertheless, our quadrangulation algorithm is quite stable. The final mesh is highly regular, containing only 12 extraordinary vertices and no vertices with valence higher than 5. The individual elements are also generally very well-shaped, with an average internal angle of 89.73° .




Figure 14: Remeshing a genus-1 statue of a dancer.

Further remeshing results can be seen on the kitten in Figure 1 and the dancer in Figure 14. These surfaces each exhibit fairly complex geometry, which our quadrangulation is able to capture and pre-

serve during resampling. The dancer is particularly challenging as it has significant elongated features protruding from the main body of the surface. If these are not captured well by the base complex, we would experience very high parametric distortions, and hence very poor meshes, in these areas. However, we see that the surface is covered at a reasonably uniform rate, indicating that the base complex has done a good job of identifying and capturing these features. Similarly, both examples are genus-1 surfaces, having large topological handles that can degenerate significantly if the base complex does not adequately capture this topological detail. Again, we see that our base complex and the derived remeshing result samples this feature quite well.




Figure 15: Remeshing a genus-22 “heptoroid”.


In Figure 15 we see a challenging high-genus model. The genus of this “heptoroid” surface is 22, and it has many hyperbolic regions. Despite the topological complexity of this surface, we are able to produce a well-shaped mesh with a relatively modest 175 extraordinary vertices and a maximum vertex degree of 7.

	Vertices				Time (s)		
	In	Out	Ext.	Freq.	Eigen.	Complex	Relax
torus	1600	1024	0	32d	1.36	3.28	0.332
moai	10,092	5834	12	33d	6.97	1.59+3.29	8.67
kitten	10,000	6600	15	51p	0.59+1.17	1.58+3.47	28.08
dancer	24,998	16,272	33	64p	0.57+2.67	1.62+8.06	441.43
bunny	72,023	10,370	26	46p	0.59+10.2	1.55+24.0	1259.15

Table 1: Performance data for our system on results shown.

Table 1 summarizes the performance of our system, reporting for the examples shown the input vs. output complexities, the number of extraordinary vertices (which we note is universally low), and the spectral index of the eigenfunction used to generate the complex. Running times were measured on a 1.8 GHz PowerPC G5 processor with 2 GB of memory. For each model, we extracted 80 eigenvectors. For the kitten, dancer, and bunny the table separates the times to produce these 80 eigenvectors on a 1000 vertex approximation and using spectral shifts to solve for the 20 eigenvectors of the full mesh nearest the eigenvalue whose complex on the coarse mesh best met the target node count, as described in Section 3. The final column reports the time for iterative relaxation and remeshing including the substantial time spent solving the GSP linear system, for which we use UMFPACK.

The eigenvectors of the combinatorial graph Laplacian are dependent only on the connectivity of the graph. However, this is *not* true of the Laplacian eigenfunctions that we use, which are indeed functions of the shape. To demonstrate this fact, consider the three poses of a cat shown in Figure 16a. Each mesh has identical connectivity, but substantially different shape. We show the complex induced by the 61st eigenfunction on each of them. Note how, in



(a) Original models (b) Projected complexes

Figure 16: Our eigenfunctions and complexes are shape-dependent.

each case, the complex follows the overall shape of the surface. In Figure 16b we see complexes for the two sitting poses mapped onto the reference pose. Clearly, the eigenfunctions themselves are quite different and the complexes still reflect the shapes of their original poses. In particular, notice that the reference complex exhibits near perfect left-right symmetry while the others do not. And in those parts of the model that do not change, the complexes remain essentially the same. For an example, see the right front paw of the sitting cat, whose shape and complex is identical (up to sign) to the reference pose.




Figure 17: Feature-sensitive remeshing using a quasi-dual complex.

As discussed in Section 3.2, partial shifts of \mathbf{L} can align its eigenfunctions with features. Figure 17 shows an example of a sphere with a toroidal cut-out where all vertices along creases have been shifted by $\sigma = 1$. The resulting complexes are quite different from those of an unmodified sphere and many are perfectly aligned with the features. The quasi-dual complex shown is part of a large symmetry class of complexes—which we discover by examining their gradient directions along the feature curves—all showing essentially the same connectivity at different frequencies.

	Vertices		Angles		Edges		Efficiency	
	Total	Ext.	Mean	σ	Mean	σ	L_2	$L_{2,1}$
PQD	12,738	175	89.88°	12.71°	0.0019	0.00093	0.628	0.328
PGP	6,355	314	89.44°	9.63°	0.0030	0.00074	0.997	0.411
SSQ	10,370	26	89.79°	6.87°	0.0023	0.00070	0.968	0.413

Table 2: Quality metrics for meshes shown in Figure 18.

Finally, we examine the quality of our results in comparison to existing techniques. We obtained results for the algorithms of Boier-Martin *et al.* [2004], denoted by the acronym PQD, and Ray *et al.* [2005], denoted by PGP, on the bunny model¹. Like our own, the PQD method produces a pure quadrilateral mesh, whereas the PGP method produces a quad-dominant mesh. The remeshed surfaces are shown in Figure 18, and our measurements of mesh quality are summarized in Table 2 together with histograms of edge length and angle distributions in Figure 19.

¹This mesh differs from the standard bunny; it is denser and has no holes.




Figure 18: Comparison of meshes generated by PQD [Boier-Martin et al. 2004], PGP [Ray et al. 2005], and our SSQ algorithm.




Figure 19: Edge and angle distributions for meshes in Figure 18.

Examining these results, it is clear that our method produces a mesh with far fewer extraordinary points. As we would expect, this is particularly true with respect to the PGP method, which produces quad-dominant rather than quadrilateral meshes. In fact, the number of extraordinary points generated by these other two methods on the genus-0 bunny is no less than the number of extraordinary points our method generates on the genus-22 heptoroid (Fig. 15). All three methods generate a large number of well-shaped elements, with mean internal angles very near the ideal 90° . However, our method was able to produce a rather more narrow angle distribution with a standard deviation of only 6.87° . Our meshes also exhibit a greater uniformity in edge length. In order to quantify geometric fidelity for these meshes with variable numbers of vertices, we use a measure of *error efficiency*, which we define to be $-\log E / \log m$ where E is the error—either L_2 (geometric) or $L_{2,1}$ (normal)—and m is the number of output vertices. The mesh produced by our method has a noticeably higher efficiency than the PQD mesh, which appears to include a number of elements that do not significantly reduce the error. On the other hand, our result is comparable to that from PGP, being somewhat lower in geometric (L_2) efficiency but somewhat higher in normal ($L_{2,1}$) efficiency. In summary, our experiments demonstrate that our method produces well-shaped meshes with far fewer critical points and greater uniformity than other state-of-the-art methods, while generating surfaces of comparable, and in some cases superior, geometric fidelity.

7 Conclusions

In this paper, we have outlined a new theoretical framework for quadrangulating polygonal manifolds. By using Morse theory to analyze the structure of the Laplacian eigenfunctions of the surface, we are able to produce appealing quadrangulations that arise directly from the intrinsic shape of the manifold. Our use of the Morse-Smale complex is topologically robust and guarantees that the base complex is always quadrangular. We have proposed a new globally smooth parameterization method for quadrangular complexes that can successfully optimize the shape of even highly degenerate complexes. Finally, we have also demonstrated that our method produces semi-regular pure quadrilateral meshes that have far fewer extraordinary points than comparable methods while maintaining high geometric fidelity.

The results we have presented in this paper open a new line of research aimed at describing and understanding shapes and geometry. We have chosen semi-regular quadrangular remeshing as an example which demonstrates that the Morse-Smale structure of Laplacian eigenfunctions encodes fundamental information about the shape of a piecewise linear manifold. But more broadly, these results are based on several intriguing properties of the Laplace matrix and its eigenfunctions that we have only begun to explore.

There is much that could be learned from a more thorough theoretical understanding of the structure of the Laplacian spectrum. The majority of results in spectral graph theory tie the Laplacian eigenvalues to various properties of the graph. The structure of the eigenvectors is relatively unexplored. A clearer understanding of the spectral structure should enable us to prove stronger results about the quality of the final quadrangulation.

We have outlined a basic technique for producing “feature-sensitive” eigenfunctions. This is clearly an avenue in which significant new contributions could be made, as there are certainly applications where some measure of user control over the flow of the mesh is quite desirable.

Acknowledgements. The scanned models used in this paper were provided by the AIM@SHAPE project and the Stanford, Princeton, Berkeley, and MPI graphics groups. Ioana Boier-Martin and Bruno Lévy generously provided the sample results analyzed in Fig. 18. This work was funded in part by the National Science Foundation (grants SGER #0432257 and ITR #0121288), and was performed in part under the auspices of the U.S. Department of Energy by the University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

References

- ALLIEZ, P., COHEN-STEINER, D., DEVILLERS, O., LÉVY, B., AND DESBRUN, M. 2003. Anisotropic polygonal remeshing. *TOG* 22, 3, 485–493. (Proc. SIGGRAPH).
- ALLIEZ, P., UCELLI, G., GOTSMAN, C., AND ATTENE, M. 2005. Recent advances in remeshing of surfaces. ftp://ftp-sop.inria.fr/geometrica/alliez/survey_remeshing.pdf.
- BANCHOFF, T. F. 1967. Critical points and curvature for embedded polyhedral surfaces. *Differential Geometry* 3, 1, 257–268.
- BERN, M. W., AND EPPSTEIN, D. 1995. Mesh generation and optimal triangulation. In *Computing in Euclidean Geometry, Lecture Notes on Computing #4*. World Scientific, 47–123.

- BOIER-MARTIN, I., RUSHMEIER, H., AND JIN, J. 2004. Parameterization of triangle meshes over quadrilateral domains. In *Proc. Eurographics Symposium on Geometry Processing*, 197–207.
- BREMER, P.-T., EDELSBRUNNER, H., HAMANN, B., AND PASCUCCI, V. 2004. A topological hierarchy for functions on triangulated surfaces. *TVCG 10*, 4, 385–396.
- CHUNG, F. R. K. 1997. *Spectral Graph Theory*. American Mathematical Society.
- COURANT, R., AND HILBERT, D. 1953. *Methods of Mathematical Physics*, vol. 1. Interscience Publishers, New York.
- DONG, S., KIRCHER, S., AND GARLAND, M. 2005. Harmonic functions for quadrilateral remeshing of arbitrary manifolds. *CAGD 22*, 5, 392–423.
- ECK, M., AND HOPPE, H. 1996. Automatic reconstruction of B-spline surfaces of arbitrary topological type. In *Proc. SIGGRAPH*, 325–334.
- ECK, M., DEROSE, T. D., DUCHAMP, T., HOPPE, H., LOUNSBERRY, M., AND STUETZLE, W. 1995. Multiresolution analysis of arbitrary meshes. In *Proc. SIGGRAPH*, 173–182.
- EDELSBRUNNER, H., LETSCHER, D., AND ZOMORODIAN, A. 2002. Topological persistence and simplification. *Discrete Comput. Geom.* 28, 511–533.
- EDELSBRUNNER, H., HARER, J., AND ZOMORODIAN, A. 2003. Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds. *Discrete Comput. Geom.* 30, 87–107.
- FLOATER, M. S., AND HORMANN, K. 2004. Surface parameterization: A tutorial and survey. In *Multiresolution in Geometric Modelling*.
- FLOATER, M. S. 2003. Mean value coordinates. *Computer Aided Geometric Design 20*, 1 (Mar.), 19–27.
- FRIEDEL, I., SCHRÖDER, P., AND KHODAKOVSKY, A. 2004. Variational normal meshes. *TOG 23*, 4, 1061–1073.
- GARLAND, M., AND HECKBERT, P. S. 1997. Surface simplification using quadric error metrics. In *Proc. SIGGRAPH*, 209–216.
- GARLAND, M. 1999. Multiresolution modeling: Survey & future opportunities. In *State of the Art Report*, Eurographics, 111–131.
- GU, X., GORTLER, S. J., AND HOPPE, H. 2002. Geometry images. *TOG 21*, 3, 355–361. (Proc. SIGGRAPH).
- GUSKOV, I., VIDIMCE, K., SWELDENS, W., AND SCHRÖDER, P. 2000. Normal meshes. In *Proc. SIGGRAPH*, 95–102.
- HALSTEAD, M., KASS, M., AND DEROSE, T. 1993. Efficient, fair interpolation using Catmull-Clark surfaces. In *Proc. SIGGRAPH*, 35–44.
- HILAGA, M., SHINAGAWA, Y., KOHMURA, T., AND KUNII, T. L. 2001. Topology matching for fully automatic similarity estimation of 3D shapes. In *Proc. SIGGRAPH*, 203–212.
- HORMANN, K., AND GREINER, G. 2000. Quadrilateral remeshing. In *Proc. Vision Modeling and Visualization*, 153–162.
- KARNI, Z., AND GOTSMAN, C. 2000. Spectral compression of mesh geometry. In *Proc. SIGGRAPH*, 279–286.
- KHODAKOVSKY, A., LITKE, N., AND SCHRÖDER, P. 2003. Globally smooth parameterizations with low distortion. *TOG 22*, 3, 350–357. (Proc. SIGGRAPH).
- KOREN, Y., CARMEL, L., AND HAREL, D. 2002. ACE: A fast multiscale eigenvector computation for drawing huge graphs. In *Proc. InfoVis '02*, 137–144.
- LEE, A. W. F., SWELDENS, W., SCHRÖDER, P., COWSAR, L., AND DOBKIN, D. 1998. MAPS: Multiresolution adaptive parameterization of surfaces. In *Proc. SIGGRAPH*, 95–104.
- MARINOV, M., AND KOBBELT, L. 2004. Direct anisotropic quad-dominant remeshing. In *Proc. Pacific Graphics*.
- NI, X., GARLAND, M., AND HART, J. C. 2004. Fair Morse functions for extracting the topological structure of a surface mesh. *TOG 23*, 3, 613–622. (Proc. SIGGRAPH).
- OWEN, S., STATEN, M. L., CANANN, S. A., AND SAIGAL, S. 1999. Q-Morph: An indirect approach to advancing front quad meshing. *Intl. J. Num. Methods in Engineering 9*, 1317–1340.
- PASCUCCI, V., AND COLE-MCLAUGHLIN, K. 2002. Efficient computation of the topology of level sets. In *Proc. Visualization*, 187–194.
- PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Exp. Math.* 2, 1, 15–36.
- RAY, N., LI, W. C., LEVY, B., SHEFFER, A., AND ALLIEZ, P. 2005. Periodic global parameterization. *TOG*. (Accepted, pending revision).
- SANDER, P. V., WOOD, Z. J., GORTLER, S. J., SNYDER, J., AND HOPPE, H. 2003. Multi-chart geometry images. In *Proc. Eurographics Symposium on Geometry Processing*, 146–155.
- SCHREINER, J., ASIRVATHAM, A., PRAUN, E., AND HOPPE, H. 2004. Inter-surface mapping. *TOG 23*, 3, 870–877. (Proc. SIGGRAPH).
- SHIMADA, K., LIAO, J.-H., AND ITOH, T. 1998. Quadrilateral meshing with directionality control through the packing of square cells. In *Seventh Int'l Meshing Roundtable*, 61–75.
- STAM, J. 2003. Flows on surfaces of arbitrary topology. *TOG 22*, 3, 724–731. (Proc. SIGGRAPH).
- STANDER, B. T., AND HART, J. C. 1997. Guaranteeing the topology of implicit surface polygonization for interactive modeling. In *Proc. SIGGRAPH*, 279–286.
- TAUBIN, G. 2000. Geometric signal processing on polygonal meshes. In *State of the Art Report*, Eurographics, 81–96.
- VAN KREVELD, M. J., VAN OOSTRUM, R., BAJAJ, C. L., PASCUCCI, V., AND SCHIKORE, D. 1997. Contour trees and small seed sets for isosurface traversal. In *Sym. Comp. Geo.*, 212–220.
- VELHO, L., AND ZORIN, D. 2001. 4-8 subdivision. *CAGD 18*, 5, 397–427. Spec. Issue on Subdiv. Techniques.
- WEBER, G., SCHEUERMANN, G., HAGEN, H., AND HAMANN, B. 2002. Exploring scalar fields using critical isovalues. In *Proc. Visualization*, 171–178.
- YING, L., AND ZORIN, D. 2004. A simple manifold-based construction of surfaces of arbitrary smoothness. *TOG 23*, 3, 271–275. (Proc. SIGGRAPH).