

Multiresolution Analysis of Arbitrary Meshes

Matthias Eck* Tony DeRose* Tom Duchamp*
Hughes Hoppe† Michael Lounsbery‡ Werner Stuetzle*

*University of Washington, Seattle, WA †Microsoft Research, Redmond, WA ‡Alias Research, Seattle, WA

Abstract

In computer graphics and geometric modeling, shapes are often represented by triangular meshes. With the advent of laser scanning systems, meshes of extreme complexity are rapidly becoming commonplace. Such meshes are notoriously expensive to store, transmit, render, and are awkward to edit. Multiresolution analysis offers a simple, unified, and theoretically sound approach to dealing with these problems. Lounsbery *et al.* have recently developed a technique for creating multiresolution representations for a restricted class of meshes with *subdivision connectivity*. Unfortunately, meshes encountered in practice typically do not meet this requirement. In this paper we present a method for overcoming the subdivision connectivity restriction, meaning that completely arbitrary meshes can now be converted to multiresolution form. The method is based on the approximation of an arbitrary initial mesh M by a mesh M' that has subdivision connectivity and is guaranteed to be within a specified tolerance.

The key ingredient of our algorithm is the construction of a parametrization of M over a simple domain. We expect this parametrization to be of use in other contexts, such as texture mapping or the approximation of complex meshes by NURBS patches.

CR Categories and Subject Descriptors: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling. - surfaces and object representations; J.6 [Computer-Aided Engineering]: Computer-Aided Design (CAD); G.1.2 [Approximation]: Spline Approximation.

Additional Keywords: Geometric modeling, subdivision surfaces, wavelets.

1 Introduction

In computer graphics and geometric modeling, shapes are often represented by triangular meshes. With the advent of laser scanning systems, meshes of extreme complexity are rapidly becoming commonplace. The objects shown in Color Plates 1(k) and 2(g) for instance, consist of 69,473 and 103,713 triangles, respectively. Such meshes are notoriously expensive to store, transmit, and render. They are also awkward to edit, as many vertices typically must be moved to make a change of substantial spatial extent.

Multiresolution analysis offers a promising new approach for addressing these difficulties in a simple, unified, and theoretically sound way. A multiresolution representation of a mesh, as recently developed by Lounsbery *et al.* [11], consists of a simple base mesh (Color Plate 1(e)) together with a sequence of local correction terms,

called *wavelet coefficients*, capturing the detail present in the object at various resolutions. Color Plates 1(g)–(h) show a sequence of intermediate resolution models incorporating an increasing number of wavelets.

Multiresolution mesh representations are particularly convenient for a number of applications, including:

- *Compression/simplification:* A multiresolution mesh can be compressed by removing small wavelet coefficients. Moreover, the threshold for removal can be chosen such that the resulting approximation is guaranteed to be within a specified error tolerance of the original mesh. A number of examples are shown in the color plates.
- *Progressive display and transmission:* An attractive method for displaying a complex object is to begin with a low resolution version that can be quickly rendered, and then progressively improve the display as more detail is obtained from disk or over a network. Using a multiresolution representation, this is simply achieved by first displaying the base mesh, and then progressively adding the contributions of wavelet coefficients in order of decreasing magnitude.
- *Level-of-detail control:* High performance rendering systems often use a level-of-detail hierarchy, that is, a sequence of approximations at various levels-of-detail. The crudest approximations are used when the viewer is far from the object, while higher detail versions are substituted as the viewer approaches. Multiresolution representations naturally support this type of display by adding successively smaller wavelet coefficients as the viewer approaches the object, and by removing them as the viewer recedes. Moreover, the coefficients can be added smoothly, thereby avoiding the visual discontinuities encountered when switching between approximations of different resolution. This use of multiresolution representations is illustrated in Color Plates 1(k) and 1(l).
- *Multiresolution editing:* Editing at various scales can proceed along the lines developed by Finkelstein and Salesin [4] by ordering coefficients according to their support, that is, by the spatial extent of their influence. Color Plates 2(e) and 2(f) show edits of a mesh at low and high levels of detail.

Although the multiresolution analysis of Lounsbery *et al.* [11] can be applied to meshes of arbitrary topological type, it has a serious shortcoming: it is restricted to meshes with *subdivision connectivity*, that is, to meshes obtained from a simple base mesh by recursive 4-to-1 splitting (see Figure 1). Figure 6(a) shows an example of a mesh with subdivision connectivity — it results from recursively splitting the faces of an octahedron four times. Unfortunately, few of the meshes encountered in practice have this restricted structure.

In this paper we present a method for overcoming the subdivision connectivity restriction, meaning that completely arbitrary meshes can now be converted to multiresolution form. Our approach is to develop an algorithm for approximating an arbitrary mesh M (as in Color Plate 1(a)), which might not have subdivision connectivity, by a mesh M^J that does (as in Color Plate 1(f)), and is guaranteed to be within a prescribed tolerance ϵ_1 of M . We refer to this process as *remeshing*, and we call M^J the *remesh*.

Multiresolution analysis of an arbitrary mesh M thus proceeds in two steps: we first use remeshing to approximate M by a mesh M^J with subdivision connectivity, and then use the method of Lounsbery *et al.* to convert M^J to multiresolution representation. (Although we cannot reproduce here all the results of Lounsbery *et al.* [11], we have included a brief summary in Appendix A.)

The key ingredient of the remeshing procedure — and the principal technical contribution of the paper — is the construction of a parametrization of M over a *base complex* K^0 possessing a small number of faces. We then sample the parametrization to produce the remesh. Considerable care is taken to create a parametrization and a sampling pattern so that the resulting remesh can be well approximated with relatively few wavelet coefficients.

The construction of parametrizations for complex shapes over simple domains is a fundamental problem that occurs in numerous applications, including texture mapping, and the approximation of meshes by NURBS patches. We therefore expect that our parametrization algorithm will have uses outside of remeshing.

The remainder of the paper is organized as follows. In Section 2, we describe the relationship between our work and previously published methods. In Section 3, we give a high level overview of the major steps of the remeshing algorithm. The details of the algorithm are presented in Sections 4–7. In Section 8, we apply our method to meshes of varying complexity, and give examples of compression, level-of-detail control, and editing. We close with conclusions and future work in Section 9.

2 Related Work

The difficulty of dealing with complicated shapes is evidenced by the extensive recent research on the topic.

The problems of compression/simplification and level-of-detail control have been addressed by Turk [20], Schroeder *et al.* [19], Hoppe *et al.* [8], Rossignac and Borrel [16], and Varsney [22]. Our approach differs from these methods in three principal respects. First, it provides guaranteed error bounds, whereas the approaches of Turk, Schroeder *et al.*, and Hoppe *et al.* do not. Second, it produces a single compact representation from which a continuous family of lower resolution approximations can be quickly and easily constructed, whereas the previous methods generate a discrete set of models of varying complexity. (We should note, however, that Turk, and Rossignac/Borrel, and Varsney present methods for interpolating between models.) Third, our representation can be simply and conveniently edited at multiple scales, whereas it is hard to imagine how one would achieve similar results using the previous approaches.

The editing of complex shapes was a central motivation for the introduction of hierarchical B-splines by Forsey and Bartels [6]. Forsey and Bartels [5] and Forsey and Wang [7] have subsequently developed methods for fitting hierarchical B-splines to meshes topologically equivalent to a disk. Finkelstein and Salesin [4] have demonstrated how wavelet representations of B-spline curves and tensor product surfaces can be used to achieve similar benefits. The main advantage of our method is its ability to deal with shapes of arbitrary topological type.

The problem of parametrizing meshes has recently been con-

sidered by Maillot *et al.* [12] in the context of texture mapping. However, the parametrizations they construct are not useful for our purpose: their surface tiles are not triangular, and their local parametrizations do not fit together continuously. Additionally, our local parametrizations, based on the well-established theory of harmonic maps, are simpler to compute than the ones used by Maillot *et al.*, and seem to produce parametrizations of comparable quality (see Section 4).

Finally, the technique of Schröder and Sweldens [18] could be used in place of Lounsbery *et al.* for multiresolution analysis of the remesh.

3 Overview of Remeshing

The basic idea of remeshing is to construct a parametrization of M over a suitably determined domain mesh K^0 . This parametrization is then resampled to produce a mesh M^J that has subdivision connectivity and is of the same topological type as M .

Our remeshing algorithm consists of three steps, as illustrated in Color Plates 1(a)–1(h):

1. *Partitioning*: Partition M into a number of triangular regions T_1, \dots, T_r , as shown in Color Plate 1(d). We want the number r of regions to be small, because the lowest complexity approximation we can construct has r faces, as shown in Color Plate 1(e). Basic tools used in partitioning are harmonic maps, maps that preserve as much of the metric structure (lengths, angles, etc.) of M as possible. Harmonic maps are described in Section 4. A detailed description of our partitioning algorithm is given in Section 5.

Identifying each of the m vertices or *nodes* of the triangulation T_1, \dots, T_r with a canonical basis vector of \mathbf{R}^m defines a mesh in \mathbf{R}^m , called the *base complex*, with a face corresponding to each of the r triangular regions. This mesh serves as the domain of the parametrization constructed in the next step.

2. *Parametrization*: For each region T_i of M construct a (local) parametrization $\rho_i : F_i \rightarrow T_i$ over the corresponding face F_i of the base complex K^0 . The local parametrizations are made to fit together continuously, meaning that collectively they define a globally continuous parametrization $\rho : K^0 \rightarrow M$. We want the coordinate functions of the parametrization to vary as little as possible since such functions have multiresolution approximations with few significant wavelet coefficients, leading to high compression ratios. Harmonic maps in a sense minimize distortion and therefore are particularly well suited for this purpose. A description of the parametrization step is presented in Section 6.

3. *Resampling*: Perform J recursive 4-to-1 splits on each of the faces of K^0 (see Figure 1). This results in a triangulation K^J of K^0 with subdivision connectivity. The remesh M^J , as shown in Color Plate 1(f), is obtained by mapping the vertices of K^J into \mathbf{R}^3 using the parametrization ρ , and constructing an interpolating mesh in the obvious way; M^J therefore has vertices lying on M , and has subdivision connectivity.

The resampling step is described more fully in Section 7, and it is shown that J can be determined so that M^J and M differ by no more than a specified remeshing tolerance ϵ_1 .

4 Harmonic maps

A crucial building block of our remeshing algorithm is a method for constructing a parametrization of a (topological) disk $D \subset M$ over a convex polygonal region $P \subset \mathbf{R}^2$. This method is used in two places: in the construction of the triangulation T_1, \dots, T_r of




Figure 1: 4-to-1 splitting of a triangular face: (a) the initial face; (b) after one 4-to-1 split; (c) after two 4-to-1 splits.

M (see Section 5), and in the parametrization of M over the base complex K^0 (see Section 6). We want this parametrization to have small distortion; for example, if D is (close to) planar, we want the parametrization to be (close to) linear. Because the region may be geometrically complex (see, for example, Figure 2), some distortion is usually inevitable.

While it is not clear in general how to find a parametrization ρ with small distortion, there is a closely related and well-studied problem that has a unique solution: Fix a homeomorphism g between the boundary of D and the boundary of the polygonal region P ; then there is a unique *harmonic map* $h : D \rightarrow P$ that agrees with g on the boundary of D and minimizes *metric dispersion* (see Eells and Sampson [3], pages 114–115, and the survey article by Eells and Lemaire[2]). Metric dispersion is a measure of the extent to which a map stretches regions of small diameter in D . It is thus a measure of metric distortion.

In addition to minimizing metric distortion, the harmonic map h has a number of important properties: (i) It is infinitely differentiable on each face of D ; (ii) it is an embedding [17]; and (iii) it is independent of the triangulation of D . Because $h : D \rightarrow P$ is an embedding, the inverse h^{-1} is a parametrization of D over P . We will return below to the issues of choosing the boundary map g and of computing approximations to h .

The dispersion minimizing property of harmonic maps is illustrated in Figure 2, which shows a piecewise linear approximation of a harmonic map from a geometrically complex region onto a polygon. The relatively dense regions of the polygon correspond to the ears and nose of the cat. Notice that the aspect ratios of triangles tend to be preserved. Notice also that the map introduces a certain amount of area compression. This is inevitable because the region has a large area relative to its circumference, and consequently any embedding must introduce some distortion in edge lengths. The harmonic map tends to minimize such distortion while maintaining the embedding property and attempting to preserve aspect ratios of triangles.

Harmonic maps can be visualized as follows. Imagine D to be composed of elastic, triangular rubber sheets sewn together along their edges. Stretch the boundary of D over the boundary of the polygon P according to the map g . The harmonic map minimizes the total energy $E_{\text{harm}}[h]$ of this configuration of rubber sheets.

Rather than constructing the harmonic map directly, we compute a piecewise linear approximation. Assume that n vertices v_1, \dots, v_n , called *corners*, have been selected on the boundary ∂D of D (see Figure 2), and (for technical reasons) assume that the degree of each of the remaining boundary vertices is at least 3.

We choose the polygon P by mapping the corners of D onto the vertices of an n -gon in \mathbf{R}^2 . The vertices of the n -gon are positioned on a circle such that the sides subtend angles proportional to the arc lengths of the boundary segments of D joining the corresponding corners. We then define g to be the piecewise linear map that sends the corners of ∂D to the vertices of P , and is a homothety (i.e. an isometry up to a constant factor) between each boundary segment of D and the corresponding side of P (Figure 2).




Figure 2: The harmonic map for the head of a cat. The neck of the cat is mapped onto the boundary of the polygon. The “corner” vertices (those sent to vertices of the polygon) are indicated by small balls.

Now suppose that h is any piecewise linear map that agrees with g on the boundary. It is therefore uniquely determined by its values $h(i)$ at the vertices of D . By explicitly integrating the functional E_{harm} over each face, one finds that E_{harm} can be reinterpreted as the energy of a configuration of springs with one spring placed along each edge of D :

$$E_{\text{harm}}[h] = 1/2 \sum_{\{i,j\} \in \text{Edges}(D)} \kappa_{i,j} \|h(i) - h(j)\|^2, \quad (1)$$

where the spring constants $\kappa_{i,j}$ are computed as follows: For each edge $\{i, j\}$, let $L_{i,j}$ denote its length as measured in the initial mesh D , and for each face $\{i, j, k\}$, let $\text{Area}_{i,j,k}$ denote its area, again as measured in D . Each interior edge $\{i, j\}$ is incident to two faces, say $\{i, j, k_1\}$ and $\{i, j, k_2\}$. Then

$$\kappa_{i,j} = (L_{i,k_1}^2 + L_{j,k_1}^2 - L_{i,j}^2) / \text{Area}_{i,j,k_1} + (L_{i,k_2}^2 + L_{j,k_2}^2 - L_{i,j}^2) / \text{Area}_{i,j,k_2}$$

The formula for spring constants associated to boundary edges has only one term.

Although the spring constants $\kappa_{i,j}$ can assume negative values, the function (1) is positive definite, and its unique minimum can be found by solving a sparse linear least-squares problem for the values $h(i)$. In contrast to the harmonic map itself, its piecewise linear approximation is not always an embedding. In our experience, this problem occurs extremely rarely (3 times in the roughly 1000 harmonic maps we computed). In these cases we use uniform spring constants.

For the remainder of this paper we refer to the unique piecewise linear function minimizing (1) as a harmonic map, although strictly speaking it is only an approximation.

Others have developed similar approaches to embedding disk-like regions. One such approach, described by Kent *et al.* [9], is also based on minimizing the energy of a network of springs. They choose spring constants to be either all equal or inversely proportional to edge lengths. Maillot *et al.* [12] introduced another functional, also based on elasticity theory.

Figure 3 illustrates the behavior of the various embedding schemes in a simple example where the region D (see Figure 3(a)) is a triangulation of a planar polygon P and $g : \partial D \rightarrow P$ is the identity. The harmonic map (Figure 3(b)) is the identity map and therefore has no metric distortion. The method of Kent *et al.* with either choice of spring constants produces considerable metric distortion (Figure 3(c) and (d)).




Figure 3: Comparison of various “spring embeddings”. From left to right: (a) Original mesh; (b) Harmonic map and embedding of Maillot *et al.* with $\alpha = 1$; (c) $\kappa_{i,j} = 1$; (d) $\kappa_{i,j} = 1/L_{i,j}$; (e) Embedding of Maillot *et al.* with $\alpha = 1/2$.

The mathematical properties of the functional proposed by Maillot *et al.* are not entirely clear. In particular, the smooth theory to which it is an approximation does not yield planar embeddings of geometrically complex regions. This led them to introduce a user-specified tuning parameter α . In the example of Figure 3, the choice $\alpha = 1$ also produces the identity map, whereas the choice $\alpha = 1/2$ leads to small distortion (see Figure 3(e)). The method of Maillot *et al.* appears produce results whose quality is comparable to ours (for appropriately chosen α). However, their method requires non-linear optimization, whereas our method requires only the solution of a sparse linear least-squares problem.

5 Partitioning

Our partitioning scheme is based on the concepts of Voronoi diagrams and Delaunay triangulations. Let us first see how these concepts could be used to partition a dense triangulation of a planar region into a small number of large triangles. We could begin by selecting a set of relatively uniformly distributed vertices of the dense triangulation, and then compute the Delaunay triangulation for the selected vertices. One method for computing the Delaunay triangulation for a set of sites in the plane is to first construct the Voronoi diagram. Its polyhedral dual is the Delaunay triangulation if Voronoi tiles meet three at a corner.

By analogy, our approach is to first partition the faces of M into a set of Voronoi-like tiles τ_i using a discrete approximation of the Voronoi diagram as described in Section 5.1. Unlike typical uses of Voronoi diagrams, we do not know the sites *a priori* — they are determined dynamically as the Voronoi diagram is constructed.

We then construct the dual to the Voronoi diagram, resulting in a Delaunay-like partition of M into triangular regions T_i , as described in Section 5.2.

5.1 Constructing the Voronoi diagram

As mentioned above, we use a discrete version of the Voronoi diagram to partition M into a set of Voronoi-like tiles. An efficient algorithm for constructing true Voronoi diagrams on the surface of a mesh has been developed by Mount [15], but it is rather difficult to implement, and is unnecessary for our purposes.

We first describe an algorithm for computing tiles τ_1, \dots, τ_s given a set of sites logically positioned at the centroids of the *site faces* $S = \{f_1, \dots, f_s\}$. We then present an algorithm for selecting a set S of site faces for which the induced Voronoi diagram is dual to a triangulation. The results of applying the Voronoi algorithm is shown in Color Plate 1(b).

Constructing the Voronoi diagram for a given set of site faces

A Voronoi tile τ_i consists of all faces for which the closest site face is f_i . Our measure of distance between faces is an approximation of geodesic distance over the surface. It is defined by constructing a dual graph to the mesh: the nodes of the graph correspond to faces of M , and the edges of the graph connect nodes corresponding to

adjacent faces. We set the cost of edges in this dual graph to be the distance between centroids of the corresponding faces. The distance between two faces is defined as length of the shortest path in this dual graph.

Constructing the Voronoi diagram is a multi-source shortest path problem in the dual graph, which we solve using a variant of Dijkstra’s algorithm [1]. The algorithm simultaneously grows the Voronoi tiles from their site faces until they cover M .

Selecting the site faces In this section we describe an algorithm for selecting a set S of site faces such that the induced Voronoi diagram, computed as above, is dual to a triangulation. Although our algorithm for selecting such site faces can be applied to any mesh M , let us assume for the moment that M does not possess boundaries. With this assumption, the Voronoi diagram must satisfy the following conditions to be dual to a triangulation:

1. tiles must be homeomorphic to disks;
2. no pair of tiles may share more than one *cut* (a cut is a contiguous set of edges of M along which a pair of tiles touch);
3. no more than three tiles can meet at any vertex.

The algorithm begins by initializing S with a single randomly chosen site face. In the outer loop we then incrementally add faces to S until the induced tiling satisfies conditions (1) through (3) above.

In the inner loop (tile growth), tiles associated with the faces in S are grown until either they cover M , in which case tile growth terminates, or until condition (1) is violated. Violation of condition (1) can be detected by examining only the neighborhood of the most recently added face. If condition (1) is violated, this face is added to S and tile growth is resumed.

When tile growth is complete, conditions (2) and (3) are checked. If condition (2) is violated, a face along one of the offending shared cuts is selected as a new site face. If condition (3) is violated, one of the faces adjacent to the offending vertex is selected as a site. If all adjacent faces already are sites, the Voronoi algorithm fails. This has never happened in any of the examples we have run. If it were to happen, we would simply use the original mesh as the base mesh.

To accommodate boundaries, we introduce a single *fictitious* Voronoi tile, logically outside of M , that touches each of the boundaries of M . Conditions (1) through (3) can then be applied without change. To ensure that the Delaunay-like triangulation covers M , we require that boundary tiles (those adjacent to the fictitious tile) have sites on the boundary of M . This issue is addressed again in the next section. To achieve this requirement, the algorithm adds a new boundary site face whenever an interior tile touches a boundary. As before, when tile growth stops, conditions (2) and (3) are checked, and if violated, appropriate new sites are added.

It sometimes happens that tiles have adjacent short cuts, a situation that leads to Delaunay-like triangles with poor aspect ratios, and hence to poor compression rates. We therefore add to the list of conditions one that disallows such tiles. Adjacent cuts of a tile are deemed short if the sum of their lengths is less than 10% of the length of the boundary of the tile. When an offending pair of cuts is found, one of the faces they share is added as a new site.

Properties of the Voronoi algorithm The time complexity of the Voronoi algorithm depends on the number s of sites that are needed, for which there is no general formula. For a fixed set of s sites, the Voronoi tiles can be constructed using an s -source version of Dijkstra’s algorithm. Like the ordinary single source Dijkstra algorithm, the s -source version can be implemented efficiently ($O(n \log n)$ time) using a priority queue, where the priority of a face is the distance to the nearest site.




Figure 4: Construction of initial Delaunay paths on M .

Naively rerunning Dijkstra’s algorithm from scratch each time a new site face is added would require $O(sn \log n)$ time. However, the algorithm can be sped up significantly by incrementally updating the priority queue as new sites are added to S .

Finally, because our site selection algorithm uses a greedy search, it cannot be expected to produce a minimal set of sites.

5.2 Constructing the Delaunay triangulation

The partition of M into Voronoi tiles obtained in the previous section has the property that its dual graph consists of 3-sided faces. However, mapping these 3-sided faces onto the surface is a non-trivial problem. The obvious approach of connecting pairs of Voronoi sites by the shortest paths on the surface — as is done in constructing the Delaunay triangulation in the plane — is not guaranteed to produce a valid triangulation for arbitrary manifolds since the resulting paths can cross. Moreover, finding the shortest paths between two points on a mesh is itself a difficult problem [14]. Our alternative uses harmonic maps twice: once to produce an initial Delaunay triangulation, and then again to improve the triangulation by straightening its edges.

Constructing an initial Delaunay triangulation. The first step is to compute the harmonic map h_i that carries each Voronoi tile τ_i into an appropriate planar polygon P_i , as described in Section 4. The inverse of h_i provides a parametrization of τ_i over P_i which we use to construct paths lying on M .

Let τ_i and τ_j denote two adjacent interior Voronoi tiles as illustrated in Figure 4. The path of the initial Delaunay triangulation joining these tiles is constructed as follows: the cut shared by the tiles is mapped to an edge $e_{i,j}$ of P_i by the harmonic map h_i ; similarly, the cut is mapped to an edge $e_{j,i}$ of P_j by h_j (see Figure 4). We construct a line $L_{i,j}$ from the centroid of P_i to the midpoint of $e_{i,j}$, and a line $L_{j,i}$ from the centroid of P_j to the midpoint of $e_{j,i}$. The path is formed by mapping these lines onto M using the inverse harmonic maps. That is, the path is obtained by joining $h_i^{-1}(L_{i,j})$ and $h_j^{-1}(L_{j,i})$.¹

The construction of a path between an interior tile τ_i and a boundary tile τ_k is slightly different, as indicated in Figure 4. In order for the Delaunay triangulation to cover M , it is necessary to construct paths from the boundary. (The site selection algorithm of Section 5.1 was designed with this goal in mind in that it guarantees

¹Note that this path does not connect the site faces as one might expect. We have found that the method described here produces more uniform triangulations than were obtained by connecting site faces.

that boundary tiles have site faces on the boundary.) We therefore select a boundary vertex v_k of the site face f_k , and construct $L_{k,i}$ as the line from $h_k(v_k)$ to the midpoint of $e_{k,i}$. The line $L_{i,k}$ is constructed as before from the centroid of P_i to the midpoint of $e_{i,k}$.

Finally, two adjacent boundary tiles τ_k and τ_ℓ are connected by the path along the boundary between v_k and v_ℓ .

The edges of the paths thus constructed are generally not edges of M . For convenience in constructing the parametrizations of Section 6, we refine M to include the path edges.

Straightening the Delaunay edges. The edges of the initial Delaunay triangles constructed in the first step can have kinks where they cross the border between two Voronoi tiles (see Color Plate 1(c)). To straighten a Delaunay edge, we construct a second harmonic map from the union of the two Delaunay triangles adjacent to the edge into a planar quadrilateral, as described in Section 4. We then replace the edge by the image of the corresponding diagonal of the quadrilateral under the inverse harmonic map. This straightening step is applied to all Delaunay edges in an arbitrary order, resulting in a final triangulation T_1, \dots, T_r . Color Plate 1(d) shows the result of straightening the edges in Color Plate 1(c).

6 Parametrization

Identifying each of the m vertices or *nodes* of the triangulation T_1, \dots, T_r with a canonical basis vector of \mathbf{R}^m defines the base complex $K^0 \subset \mathbf{R}^m$, with a face corresponding to each of the r triangular regions. The goal of this section is to construct a continuous parametrization $\rho : K^0 \rightarrow M$ of the initial mesh over K^0 . We map each triangle T_i onto a triangular region of the plane, again using harmonic maps described in Section 4. We then affinely map the triangular region onto the corresponding face F_i of the base complex. The composition of the two maps is an embedding, and therefore its inverse ρ_i defines a parametrization of T_i over F_i . By construction, the maps ρ_i agree on shared boundaries, and thus the ρ_i collectively define a continuous parametrization ρ of M over K^0 .

7 Resampling

In this section, we describe a method for producing a mesh M^J with subdivision connectivity from the parametrization $\rho : K^0 \rightarrow M$ constructed in Section 6. We also show how to determine the subdivision level J so that M^J and M differ by no more than a specified remeshing tolerance ϵ_1 .

For a given value of J , we first produce a triangulation K^J of K^0 by performing J recursive 4-to-1 splits of the faces of K^0 . We then approximate ρ by a function ρ^J defined as the piecewise linear interpolant to ρ on K^J ; that is, ρ^J is such that $\rho^J(\mathbf{x}_i^J) = \rho(\mathbf{x}_i^J)$, where the points \mathbf{x}_i^J (called *knots*) denote the vertices of K^J .

The simplest strategy for performing a 4-to-1 split of a face is to position the split points at midpoints of edges, as illustrated in Figure 1. We refer to this process as *parametrically uniform resampling* since the faces of K^J are of equal size. Alternatively, we could attempt to place the knots so that the images of triangles of K^J , that is, the triangles of the remesh M^J , are of equal size. We refer to this as *geometrically uniform resampling*.

As one of our fundamental objectives is high compression rate, we evaluate the performance of a resampling strategy by the number of wavelet coefficients needed for a given compression tolerance ϵ_2 . This number is governed by at least two competing factors:

1. As mentioned in Section 3, the coordinate functions of ρ should be as slowly varying as possible; this is largely achieved by the distortion minimizing property of the harmonic map parametrizations.

ϵ_2	Geom. Uniform	Hybrid	Param. Uniform
0.5%	(2679) [5422]	(1768) [3562]	(2224) [4502]
1.0%	(1100) [2180]	(795) [1591]	(1044) [2079]
2.0%	(416) [809]	(385) [758]	(455) [881]
5.0%	(112) [223]	(130) [245]	(143) [302]

Table 1: Performance of the three sampling strategies on the cat model. Parentheses denote the number of significant wavelet coefficients; square brackets denote the number of triangles. All examples were run using $\epsilon_1 = 1.0\%$. Errors are measured as a percentage of the object’s diameter.

- The triangles of M^j should be of roughly uniform size. Lounsbury *et al.* define wavelets so that the magnitude of a wavelet coefficient is a measure of the “unweighted” least-squares error that would be incurred if the coefficient were set to zero. By unweighted we mean that deviations on large triangles of M^j are counted no more heavily than deviations on small triangles. If M^j has triangles of roughly uniform size, magnitudes of wavelet coefficients are better measures of geometric error.

The strategy that has performed best in our experiments is a hybrid strategy using geometrically uniform sampling in the first few splitting steps (the first three steps in all our examples), and parametrically uniform sampling in subsequent steps. Intuitively, this strategy does a reasonable job of uniformly distributing the triangles on a coarse scale, while still remaining faithful to the harmonic parametrization on smaller scales.

This intuition is supported by numerical results. Our tests have shown that hybrid resampling typically results in wavelet expansions with fewer significant coefficients than either parametrically uniform or geometrically uniform resampling. Moreover, the number of subdivisions J necessary to satisfy a remeshing tolerance ϵ_1 is often smaller and hence the remesh is often faster to compute and requires less storage. Table 1 presents the results of an experiment for the cat mesh (shown in Color Plate 2(d)) for various wavelet compression tolerances ϵ_2 . Notice that hybrid resampling is particularly advantageous for small tolerances.

7.1 Geometrically uniform resampling

The task of determining new knots $\mathbf{x}_i^j \in K^0$ so that the triangles generated are roughly uniform in size is an optimization problem whose solution we approximate using the following recursive greedy algorithm.

In the parametrically uniform resampling process the knot \mathbf{x}_i^j at level j is simply computed as midpoint of the edge of the two (neighboring) knots $\mathbf{x}_{n_1(i)}^{j-1}$ and $\mathbf{x}_{n_2(i)}^{j-1}$ at level $j - 1$. Instead of performing uniform subdivision, we define

$$\mathbf{x}_i^j = (1 - \lambda_i^j) \cdot \mathbf{x}_{n_1(i)}^{j-1} + \lambda_i^j \cdot \mathbf{x}_{n_2(i)}^{j-1} \quad \text{with } \lambda_i^j \in (0, 1),$$

where the splitting parameter λ_i^j is determined as follows: Split the two faces adjacent to the edge from $\mathbf{x}_{n_1(i)}^{j-1}$ to $\mathbf{x}_{n_2(i)}^{j-1}$, as illustrated in Figure 5. Our goal is to find λ_i^j so that the regions $R_i^j = \Delta_{i,1}^j \cup \Delta_{i,3}^j$ and $S_i^j = \Delta_{i,2}^j \cup \Delta_{i,4}^j$ map to regions of equal area on M .

In the current implementation we have simplified the area computations by using a discrete approximation: We scatter a roughly uniform collection of points on the faces of M , then map these sample points back to K^0 using ρ^{-1} (ρ^{-1} is the harmonic map, so it is already known). We then use binary search to compute the parameter λ_i^j so that the number of sampled points in the regions R_i^j and S_i^j are nearly equal.




Figure 5: Computing the new knot \mathbf{x}_i^j

7.2 Bounding the remeshing error

In this section we describe how to determine J such that the remesh M^j and the initial mesh M deviate by no more than a remeshing tolerance ϵ_1 in an L^∞ sense. That is, we seek to find the smallest J such that

$$\max_{\mathbf{x} \in K^0} \|\rho(\mathbf{x}) - \rho^j(\mathbf{x})\| \leq \epsilon_1.$$

Our strategy for determining J will be to perform successive steps of 4-to-1 splitting until the error bound is satisfied.

To bound the error for a given value of J , let $E(\mathbf{x}) := \rho(\mathbf{x}) - \rho^j(\mathbf{x})$ denote the (vector-valued) error function. First, note that the preimages of the triangles of M under ρ form a partition π of K^0 , and that ρ is a linear function on each triangle of π . Next, recall that ρ^j is linear within each of the triangles of the partition K^j of K^0 . Thus, $E(\mathbf{x})$, the difference between the two, is linear within each cell of the union partition $\pi^j = \pi \cup K^j$. The squared norm of $E(\mathbf{x})$ is therefore quadratic and convex up over each cell of π^j , and so must achieve its maximum value at a vertex of π^j .

The L^∞ error for a given value of J can therefore easily be determined by evaluating $E(\mathbf{x})$ at the vertices of π^j . Using a local marching technique such as the one in Kent *et al.* [9], these vertices can be found in time proportional to the total number of vertices in π and K^j .

8 Results

Color Plates 1 and 2 illustrate the steps of the algorithm and present examples of its applications.

Color Plate 1(a)-1(h) demonstrate the complete process of multiresolution analysis for a mesh of genus 3. We first partition the original mesh of Color Plate 1(a) into Voronoi-like tiles shown in Color Plate 1(b). We then construct the initial Delaunay-like triangulation (Color Plate 1(c)), and straighten its edges (Color Plate 1(d)). The Delaunay triangles define a simple base complex that serves as the domain for the parametrization of the mesh. Resampling this parametrization using the hybrid strategy described in Section 7 with a remeshing tolerance of $\epsilon_1 = 0.75\%$ required $J = 4$ subdivision steps and produced the remesh shown in Color Plate 1(f) consisting of 17,920 triangles. The lowest resolution approximation, shown in Color Plate 1(e), is a piecewise linear embedding of the base complex. Color Plates 1(g) and 1(h) show more detailed approximations using, respectively, 366 and 2,614 faces.

Table 2 summarizes the remeshing process for a variety of other meshes. (All computing times were measured on a SGI Onyx Reality Engine 2 with 256MB of memory.) Note that the number of Voronoi tiles is influenced more by the geometry of the model than by the number of faces of M .

Approximating the dinosaur and the phone with low tolerances would require high subdivision levels. This is due to the presence of jagged boundaries which can only be well approximated using a large number of subdivisions.

Computing times are strongly dependent on the ratio of the num-

object	# faces of M	# Voronoi tiles τ_i	# Delaunay triangles T_i	remesh. tol. ϵ_1	subdiv. level J	time mins
holes3	11,776	31	70	0.5 %	4	4.6
bunny	69,473	88	162	0.5 %	5	33.5
cat	698	7	9	1.0 %	6	0.8
dino	103,713	117	229	1.0 %	5	39.3
phone	165,896	69	132	2.5 %	5	346.6

Table 2: Summary of results of the remeshing algorithm

Color Plate	object	# faces of M'	compr. tol. ϵ_2	# wavelet coeff.	# faces	time mins
1(g)	holes3	17,920	4.0 %	179	366	0.9
1(h)	holes3	17,920	0.5 %	1,298	2,614	1.0
2(a)	bunny	165,888	0.07 %	18,636	37,598	4.5
2(b)	bunny	165,888	0.7 %	2,268	4,639	3.7
2(c)	bunny	165,888	1.5 %	952	1,921	3.5
2(i)	dino	234,496	0.5 %	2,329	4,725	5.0
2(l)	phone	135,168	0.1 %	7,920	16,451	3.3

Table 3: Summary of results of the algorithm of Lounsbery *et al.*

ber of faces of M to the number of Delaunay triangles, since the bottleneck of the algorithm, the harmonic map computation, requires solving sparse least-squares problems whose time complexity is proportional to the square of the number of vertices in the triangles.

Table 3 summarizes the results of wavelet compression applied to remeshed models. Each line of the table gives the number of faces of the remesh, the compression tolerance ϵ_2 used in the wavelet compression method described in Appendix A.2, the number of wavelet coefficients, the number of faces of the resulting approximation, and the time required for filterbank analysis and synthesis. The total deviation between the compressed model and the original is bounded by $\epsilon = \epsilon_1 + \epsilon_2$, the sum of the remeshing tolerance and the compression tolerance. Note that for storage and transmission purposes, the relevant performance measure is the number of wavelet coefficients rather than the number of faces, since only the wavelet coefficients (and their indices) have to be stored or transmitted.

Color Plates 1(k)-1(l) illustrate level-of-detail control. The original model (Color Plates 1(f) and 1(k)) was created from laser range data using the mesh zipping algorithm of Turk and Levoy [21]. Color Plates 1(k) and 1(l) show views of the original model and of lower resolution approximations from three different distances. Color Plates 2(a)-2(c) are close-ups of the approximations in Color Plate 2(l). Note the enormous reduction in the number of triangles when the multiresolution approximations are viewed from afar.

Color Plates 2(d)-2(f) illustrate *multiresolution editing*. Color Plate 2(e) shows a large-scale modification caused by changing a wavelet coefficient at the coarsest level, whereas Color Plate 2(f) corresponds to changing two coefficients at an intermediate level-of-detail.

Color Plates 2(g)-2(l) show the application of remeshing and multiresolution analysis to two additional meshes.

9 Conclusion

We have described an algorithm for solving the remeshing problem, that is, the problem of approximating an arbitrary mesh by a mesh with subdivision connectivity. Combined with the previous work of Lounsbery *et al.*, our remeshing algorithm allows multiresolution analysis to be applied to arbitrary meshes. Multiresolution representations support efficient storage, rendering, transmission, and editing of complex meshes in a simple, unified, and theoretically

sound way.

We have applied our remeshing algorithm and multiresolution analysis to complicated meshes consisting of more than 100,000 triangles. Examples of compression, level-of-detail rendering, and editing are shown in the Color Plates.

The key ingredient of our remeshing procedure — and the principal technical contribution of the paper — is the construction of a continuous parametrization of an arbitrary mesh over a simple domain mesh. Parametrizing complex shapes over simple domains is a fundamental problem in numerous applications, including texture mapping and the approximation of meshes by NURBS patches. We therefore expect that our parametrization algorithm will have uses outside of multiresolution analysis. We intend to explore these uses in future work.

Acknowledgments

This work was supported in part by a postdoctoral fellowship for the lead author (Eck) from the German Research Foundation (DFG), Alias Research Inc., Microsoft Corp., and the National Science Foundation under grants CCR-8957323, DMS-9103002, and DMS-9402734. We are grateful to Marc Levoy and his students at Stanford University for providing the bunny, dinosaur, and phone models.

References

- [1] A. Aho, J.E. Hopcroft, and J.D. Ullman. *Data structures and algorithms*. Addison-Wesley, Reading, Mass., 1983.
- [2] J. Eells and L. Lemaire. Another report on harmonic maps. *Bull. London Math. Soc.*, 20:385–524, 1988.
- [3] J. Eells and J.H. Sampson. Harmonic mappings of Riemannian manifolds. *Amer. J. Math.*, 86:109–160, 1964.
- [4] Adam Finkelstein and David Salesin. Multiresolution curves. *Computer Graphics (SIGGRAPH '94 Proceedings)*, 28(3):261–268, July 1994.
- [5] D. Forsey and R. Bartels. Hierarchical B-spline fitting. *ACM Transactions on Graphics*. To appear.
- [6] D. Forsey and R. Bartels. Hierarchical B-spline refinement. *Computer Graphics*, 22(4):205–212, 1988.
- [7] David Forsey and Lifeng Wang. Multi-resolution surface approximation for animation. In *Proceedings of Graphics Interface*, 1993.
- [8] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle. Mesh optimization. *Computer Graphics (SIGGRAPH '93 Proceedings)*, pages 19–26, August 1993.
- [9] James R. Kent, Wayne E. Carlson, and Richard E. Parent. Shape transformation for polyhedral objects. *Computer Graphics (SIGGRAPH '92 Proceedings)*, 26(2):47–54, July 1992.
- [10] J. Michael Lounsbery. *Multiresolution Analysis for Surfaces of Arbitrary Topological Type*. PhD thesis, Department of Computer Science and Engineering, University of Washington, September 1994. Available as ftp://cs.washington.edu/pub/graphics/LounsPhd.ps.Z.
- [11] Michael Lounsbery, Tony DeRose, and Joe Warren. Multiresolution analysis for surfaces of arbitrary topological type. Submitted for publication. Preliminary version available as Technical Report 93-10-05b, Department of Computer Science and Engineering, University of Washington, January, 1994. Also available as ftp://cs.washington.edu/pub/graphics/TR931005b.ps.Z.
- [12] J. Maillot, H. Yahia, and A. Verroust. Interactive texture mapping. *Computer Graphics (SIGGRAPH '93 Proceedings)*, 27(3):27–34, August 1993.
- [13] Stephane Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, July 1989.
- [14] J.S. Mitchell, D.M. Mount, and C.H. Papadimitriou. The discrete geodesic problem. *SIAM Journal of Computing*, 16(4):647–668, 1987.
- [15] David M. Mount. Voronoi diagrams on the surface of a polyhedron. Department of Computer Science CAR-TR-121, CS-TR-1496, University of Maryland, May 1985.




Figure 6: Decomposition of a mesh.

- [16] J. Rossignac and P. Borrel. Multi-resolution 3D approximations for rendering. In B. Falcidieno and T.L. Kunii, editors, *Modeling in Computer Graphics*, pages 455–465. Springer-Verlag, June-July 1993.
- [17] Richard Schoen and Shing-Tung Yau. Univalent harmonic maps between surfaces. *Inventiones math.*, 44:265–278, 1978.
- [18] P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. *Computer Graphics, (SIGGRAPH '95 Proceedings)*, 1995.
- [19] William Schroeder, Jonathan Zarge, and William Lorensen. Decimation of triangle meshes. *Computer Graphics (SIGGRAPH '92 Proceedings)*, 26(2):65–70, July 1992.
- [20] Greg Turk. Re-tiling polygonal surfaces. *Computer Graphics (SIGGRAPH '92 Proceedings)*, 26(2):55–64, July 1992.
- [21] Greg Turk and Marc Levoy. Zippered polygon meshes from range images. *Computer Graphics (SIGGRAPH '94 Proceedings)*, 28(3):311–318, July 1994.
- [22] Amitabh Varshney. *Hierarchical Geometric Approximations*. PhD thesis, Department of Computer Science, University of North Carolina at Chapel Hill, 1994.

A Multiresolution Analysis of Subdivision Meshes

As mentioned in Section 1, the main idea of multiresolution analysis is to decompose a function into a low resolution part and a set of correction or “detail” terms at increasing resolutions. Multiresolution analysis was first formalized by Mallat [13] for functions defined on \mathbf{R}^n . Lounsbery [10] and Lounsbery *et al.* [11] have recently extended the notion of multiresolution analysis to functions defined on base complexes of arbitrary topological type. Their results can be used to construct multiresolution representations of meshes with subdivision connectivity. The purpose of this appendix is to summarize their basic results and algorithms at a high level.

A.1 Background

The two basic ingredients of multiresolution analysis are a sequence of nested linear function spaces and an inner product. Lounsbery *et al.* use a sequence of spaces $V^0 \subset V^1 \subset \dots$ associated with the base complex. To describe meshes, the approximation spaces V^j consist of piecewise linear functions; specifically, V^j is the space of continuous piecewise linear functions over a partition K^j of K^0 created by performing j recursive steps of 4-to-1 splitting to the faces of K^0 , as shown in Figure 1. As j increases, the triangulation K^j becomes more dense, and so the functions in V^j are better able to model arbitrary continuous functions on K^0 . The inner product used by Lounsbery *et al.* is the standard inner product defined as

$$\langle f, g \rangle := \int_{\mathbf{x} \in K^0} f(\mathbf{x})g(\mathbf{x})d\mathbf{x}$$

where $d\mathbf{x}$ is the differential area of K^0 embedded in \mathbf{R}^m , so that all faces have unit area.

The inner product is used to define the following orthogonal complement spaces, also called *wavelet spaces*,

$$W^j := \{f \in V^{j+1} \mid \langle f, g \rangle = 0 \quad \forall g \in V^j\}.$$

Intuitively, W^j captures the detail that is missed when a function in V^{j+1} is approximated by a function in V^j .

Basis functions for V^j are called *scaling functions*. In the piecewise linear case, particularly simple scaling functions for V^j are the “hat functions” on K^j : the i -th hat function $\phi_i^j \in V^j$ is the unique function in V^j that is one at \mathbf{x}_i^j and zero at all other knots of K^j .

A *wavelet* $\psi_i^k(\mathbf{x})$ is a basis function for one of the wavelet spaces W^k . Lounsbery *et al.* [11] give constructions for wavelet bases on arbitrary base complexes K^0 . A *wavelet basis* for V^j consists of a basis for V^0 together with bases for the wavelet spaces W^0, \dots, W^{j-1} .

The parametrization $\rho^j \in V^j$ for V^j can be expanded in the hat function basis as

$$\rho^j(\mathbf{x}) = \sum_i v_i^j \phi_i^j, \quad \mathbf{x} \in K^0 \quad (2)$$

where v_i^j denote the vertex positions of M^j . A *multiresolution representation* of ρ^j refers to its expansion in a wavelet basis

$$\rho^j(\mathbf{x}) = \sum_i v_i^0 \phi_i^0(\mathbf{x}) + \sum_{j=0}^{j-1} \sum_i w_i^j \psi_i^j(\mathbf{x}), \quad \mathbf{x} \in K^0, \quad (3)$$

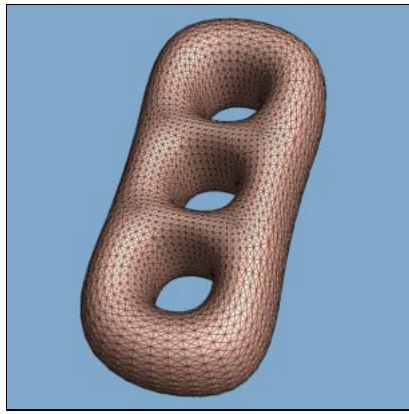
where w_i^j denote the wavelet coefficients.

An algorithm known as *filterbank analysis* can be used to convert between the hat function expansion and the multiresolution representation. The geometric interpretation of filterbank analysis is shown in Figure 6. The full detail model, described by $\rho^j(\mathbf{x})$ is successively decomposed into a lower resolution approximation together with a collection of coefficients that multiply the wavelets. The result is a simple base mesh together with wavelet coefficients at various levels of detail. The operators **A** and **B** in Figure 6 refer to sparse matrices whose entries are given by Lounsbery *et al.* The filterbank analysis has an inverse process called *filterbank synthesis* that recovers the full resolution model from its multiresolution representation.


A.2 L^∞ Wavelet compression

The L^∞ error caused by wavelet compression is the L^∞ norm of the difference function $\delta(\mathbf{x}) = \rho^j(\mathbf{x}) - \tilde{\rho}^j(\mathbf{x})$, where $\tilde{\rho}^j(\mathbf{x})$ denotes the compressed approximation to $\rho^j(\mathbf{x})$. This difference function is simply the sum of the wavelet terms that have been removed from $\rho^j(\mathbf{x})$. Since $\delta(\mathbf{x})$ is a piecewise linear function on K^j , its L^∞ norm can be determined as in Section 7.2 by recording its values at the vertices of K^j .

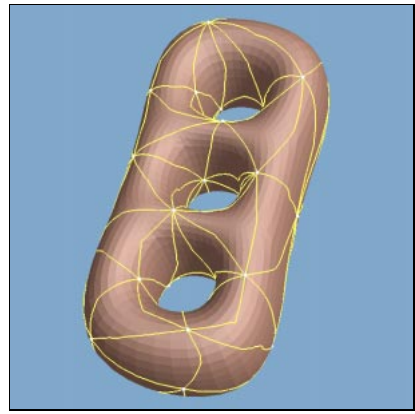
Compression in principle proceeds by considering the wavelet coefficients in order of increasing magnitude. A coefficient is removed if doing so does not cause the L^∞ norm of $\delta(\mathbf{x})$ to exceed ϵ_2 . If removal of a coefficient would violate the error tolerance, the coefficient is retained and the next coefficient is examined. The procedure terminates when all coefficients have been considered for removal. The examples presented in this paper have used a conservative approximation to this approach where a bound on the L^∞ norm of $\delta(\mathbf{x})$ is maintained, rather than maintaining $\delta(\mathbf{x})$ itself; we plan to implement the principled approach in the near future.




(a) Original mesh M (11,776 faces)




(b) Voronoi diagram (31 tiles)



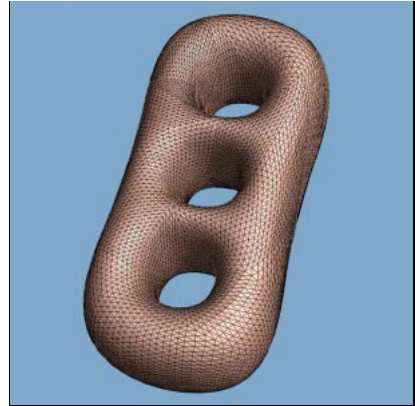
(c) Initial Delaunay triangulation (70 tri.)



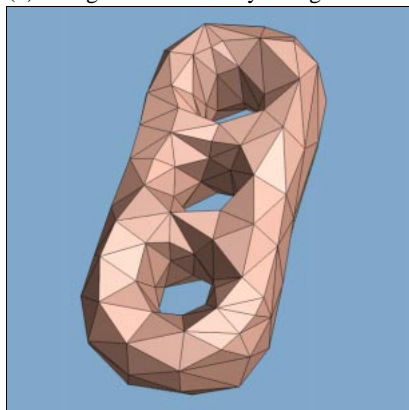
(d) Straightened Delaunay triangulation



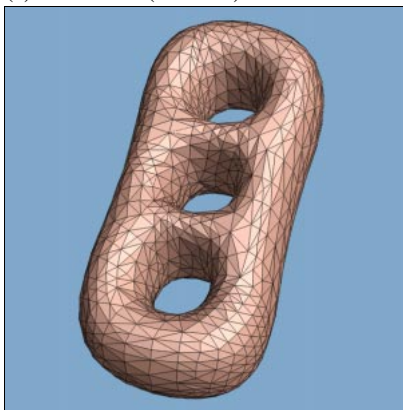
(e) Base mesh (70 faces)




(f) Remesh M' ($J = 4$; 17,920 faces)




(g) Approx. ($\epsilon = 4.5\%$; 366 faces)



(h) Approx. ($\epsilon = 1.0\%$; 2,614 faces)



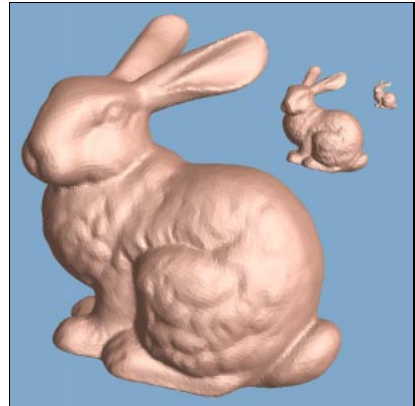
(i) Delaunay triangulation (162 tri.)



(j) Base mesh (162 faces)




(k) Original mesh (69,473 faces)




(l) LOD using multiresolution approx.

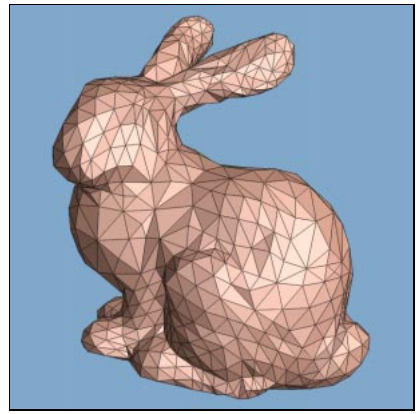
Color Plate 1: (a-g) Example of partition, parameterization, resampling, and approximation of a mesh using multiresolution analysis; (h-k) Level-of-detail approximations of a dense mesh.



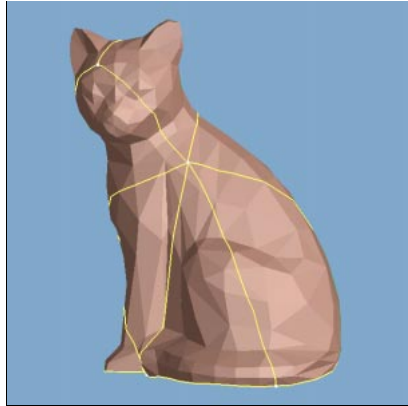
(a) Approx. ($\epsilon = 0.57\%$; 37,598 faces)




(b) Approx. ($\epsilon = 1.2\%$; 4,639 faces)




(c) Approx. ($\epsilon = 2.0\%$; 1,921 faces)




(d) Original mesh (698 faces)




(e) Surface editing at a coarse level




(f) Surface editing at a finer level




(g) Original mesh (103,713 faces)




(h) Base mesh (229 faces)



(i) Approx. ($\epsilon = 1.5\%$; 4,725 faces)



(j) Original mesh (165,896 faces)



(k) Base mesh (132 faces)



(l) Approx. ($\epsilon = 2.6\%$; 16,451 faces)

Color Plate 2: (a-c) Multiresolution approximations used in Color Plate 1(l); (d-f) Example of multiresolution surface editing; (g-l) More results of multiresolution surface approximation.