Remeshing triangulated surfaces with optimal parameterizations

Kai Hormann* Ulf Labsik† Günther Greiner‡

Computer Graphics Group, University of Erlangen

Abstract

The use of polygonal meshes, especially triangle meshes, is manifold, but a lot of algorithms require the mesh to be structured in a certain way and cannot be applied to an arbitrarily structured mesh. The process of replacing an arbitrarily structured mesh by a structured one is called *remeshing* and the most important class of structered meshes are triangle meshes with *subdivision connectivity*. In this paper we present an algorithm for remeshing triangle meshes with boundary that is based on parameterizing the mesh over a planar domain. We discuss what kind of parameterizations are optimal for the purpose of remeshing and show the advantages of our approach in a series of examples.

Keywords: Remeshing, Parameterization, Subdivision Surfaces, Meshes, Multiresolution

1 Introduction

Triangle meshes are of great importance as a standard surface representation, especially in computer graphics, due to their simplicity, flexibility and the fact that they are widely supported by the graphics hardware. But the use of 3D acquisition techniques like laser range scanning results in dense meshes with a large number of triangles, which can still be quite awkward to handle because of the sheer size of data that has to be processed. It is therefore necessary to use efficient algorithms for storing, transmitting and editing those meshes. Many algorithms in the context of multiresolution modeling need a special structure of the mesh, the so called *subdivision connectivity*. This special kind of connectivity is generated by iteratively subdividing a coarse base mesh $\mathcal{S}^0$ with a uniform refinement operator.

With most meshes not having this special kind of connectivity it is necessary to have algorithms at hand that enable the transformation of a mesh with arbitrary connectivity into a mesh with subdivision connectivity. This process is called *remeshing* and several approaches exist to solve this problem, e.g. [4, 11, 13]. There are many advantages resulting from this conversion. The different refinement levels automatically provide levels of detail which can be used by multiresolution algorithms like level-of-detail rendering [3], wavelets [16, 20], progressive transmission [12] or multiresolution editing [24].

In this paper we present an algorithm for remeshing triangulated, topologically disk-like surfaces which can be parameterized over a planar domain. We have studied the influence of the parameterization method on the quality of the resulting remesh and concluded to favor the most isometric parameterization strategy (MIPS) [8] that provides minimal distortion of the triangles in the mesh. The main advantage of this parameterization method is that the boundary vertices can be included in the optimization process, thereby avoiding the distortions that tend to occur near the boundary of the parameterizations obtained by other techniques.

*hormann@cs.fau.de
†labsik@cs.fau.de
‡greiner@cs.fau.de

The paper is organized as follows. In Section 2 a short overview of related work in the field of remeshing polygonal surfaces is given. The general concept of our algorithm is explained in Section 3. In Section 4 we show how to compute a most isometric parameterization of a given triangle mesh and in Section 5 the remeshing algorithm is explained in detail. In Section 6 we demonstrate the efficiency of our algorithm by some examples and close with a summary in Section 7.

2 Previous Work

During the last years a lot of effort has been spent on the task of remeshing polygonal surfaces. In this section we will give an overview of the most important work.

In [4] Eck et al. have presented a remeshing algorithm consisting of three steps. First, a partitioning of the given mesh $\mathcal{M}$ into a number of triangular regions has to be found. This is done by distributing vertices over $\mathcal{M}$ and growing Voronoi tiles around each vertex. With the help of the constructed Voronoi diagram a triangulation of the new vertices can be found. In the next step a parameterization of the given mesh within each triangle of the base mesh is computed. This is done with harmonic maps trying to minimize the local distortion. The remesh is achieved by uniformly subdividing each base triangle and mapping the vertices into 3-space using the parameterization.

A different approach is proposed by Lee et al. in [13]. Here the mesh $\mathcal{M}$ is coarsened by vertex removal leading to a base mesh that is used as a non-planar parameter domain for the remeshing. A parameterization is obtained by projecting the removed vertex onto the remaining mesh. This parameterization is also only locally smooth. To achieve global smoothness of the remesh the parameterization is not sampled at the dyadic points. A smoothed version of the dyadic points is computed instead by applying a variant of Loop's subdivision scheme [15] and mapping these vertices into 3-space.

The basic idea of the shrink wrapping approach [11] by Kobbelt et al. is to place a triangle mesh with subdivision connectivity around an arbitrary triangle mesh and shrink it onto the surface of that mesh. In this process two forces have to be simulated, an attracting force, moving the vertices of the subdivision in the direction of the original surface and a relaxing force, distributing the vertices over the surface. With this approach it is possible to remesh topologically sphere-like (genus-0) meshes.

3 General concept

The task of remeshing can be understood as the problem of finding an approximation operator $\mathfrak{A} : \mathbb{M} \to \mathbb{S}$ that maps from a given set $\mathbb{M}$ of meshes to the set $\mathbb{S}$ of all meshes with subdivision connectivity such that the resulting mesh is a good approximation of the original mesh, i.e. the distance of the two meshes is below a user-specified threshold. The resulting subdivision connectivity mesh is called a *remesh* of the original mesh.

In this paper we concentrate on the set of 2-manifold triangle meshes that have a boundary and no holes. In such a *spatial* mesh

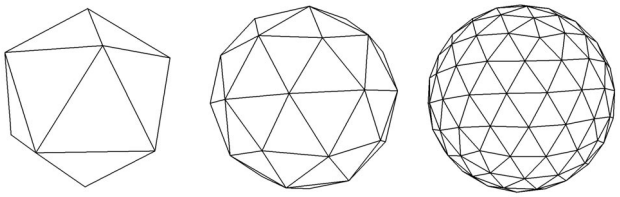

Figure 1: Starting with a coarse base mesh $\mathcal{S}^0$, a sequence of SCMs can be generated by iteratively applying the uniform subdivision operator that performs a 1-to-4 split on every triangle.

the intersection of two triangles may be either an edge, a vertex or empty. Additionally, an edge of the mesh belongs to either two or only one triangle and is a boundary edge in the latter case. The mesh $\mathcal{M}_\Psi = (P, T)$ is represented by a set of *vertices* $P = \{P_i\}$ in the *space* $\Psi \subset \mathbb{R}^3$, representing the geometry, and a set $T = \{T_j\}$ of triangles $T_j = \triangle(P_{j_0}, P_{j_1}, P_{j_2})$, defining the connectivity of the mesh.

A subdivision connectivity mesh (SCM) $\mathcal{S}_\Psi$ of level m has the further property of being a refinement of a coarse base mesh $\mathcal{S}^0$, i.e. it is the last element of a sequence of meshes $\mathcal{S}^0, ..., \mathcal{S}^m = \mathcal{S}_\Psi$ where each $\mathcal{S}^{l+1}$ emerges from $\mathcal{S}^l$ by uniformly subdividing each triangle of $\mathcal{S}^l$ into 4 subtriangles (cf. Fig. 1). If the vertices P^l of $\mathcal{S}^l$ are a subset of the vertices P^{l+1}, this operation is called *interpolatory subdivision*, and *non-interpolatory subdivision* otherwise. Note that all vertices of $\mathcal{S}_\Psi$ have valence 6, except for the boundary vertices and those who correspond to the vertices P^0 of the base mesh.

Two different aspects relate to the *quality* of a SCM. First, the number of triangles of the base mesh $\mathcal{S}^0$ shall be as small as possible, which results in a large number of hierarchy levels m of the remesh $\mathcal{S}_\Psi$. This assures a maximal utilization of the multiresolution techniques that can be applied to $\mathcal{S}_\Psi$. Another criterion for the quality of a SCM is the visual appearance, i.e. all triangles T_j of $\mathcal{S}_\Psi$ should have uniform size and aspect ratio.

Since we consider only meshes with boundary and no holes, these meshes are topologically disk-like and can therefore be parameterized over a simply connected *planar domain* $\Omega \subset \mathbb{R}^2$ by assigning each vertex P_i a *parameter value* $p_i \in \Omega$. By triangulating these parameter values in the same way as the original data points, i.e. by defining a set of triangles $t = \{t_j\}$, $t_j = \triangle(p_{j_0}, p_{j_1}, p_{j_2})$ we obtain a *planar* mesh $\mathcal{M}_\Omega = (p, t)$ that corresponds to $\mathcal{M}_\Psi$ in a natural way. We call the function $f : \Psi \to \Omega$ that linearly maps each spatial triangle to the corresponding planar triangle (i.e. $f(T_j) = t_j$, $f(P_i) = p_i$ and $f(\mathcal{M}_\Psi) = \mathcal{M}_\Omega$) the *projection* of $\mathcal{M}_\Psi$, whereas the inverse function $F = f^{-1}$ that maps back from $\mathbb{R}^2$ to $\mathbb{R}^3$ is called the *inflation*.

The general concept of our remeshing algorithm is to use the projection and the inflation in order to shift the approximation problem from the space Ψ to the planar domain Ω and consider the operator $\mathfrak{A}_\Omega$ instead of $\mathfrak{A}_\Psi$ (cf. Fig. 2). After determining a projection f with minimal distortion we approximate the planar mesh $\mathcal{M}_\Omega$ with a high quality SCM $\mathcal{S}_\Omega$ that is mapped back into space by the inflation F to yield a remesh $\mathcal{S}_\Psi = F(\mathcal{S}_\Omega)$ of the original mesh $\mathcal{M}_\Psi$. Note that this method guarantees the vertices of the remesh to lie on the surface of the given mesh.

4 Parameterization

The parameterization problem is vital for many applications in computer graphics such as surface fitting, texture mapping, and remeshing. The most important type of this problem is the parameterization of 3D data points which has been addressed in many

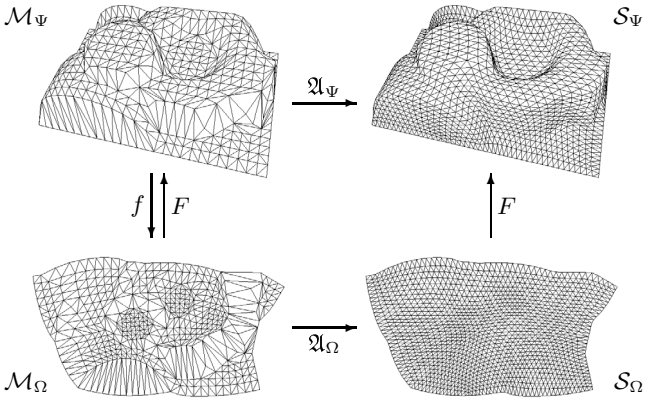

Figure 2: The general concept: $\mathfrak{A}_\Psi \approx F \circ \mathfrak{A}_\Omega \circ f$.

papers before [2, 4, 5, 6, 14, 17, 18, 19].

In general, a set of points $P_i \in \mathbb{R}^3$, a certain kind of neighborhood information that defines the topology of the point set, e.g. a triangulation, and a domain Ω that is of the same topological type are given. The task is now to assign every vertex P_i a parameter value $p_i \in \Omega$ such that the topology is preserved. Typical domains are the disk $D_2 = \{x \in \mathbb{R}^2 : \|x\| \leq 1\}$ and the sphere $S_2 = \{x \in \mathbb{R}^3 : \|x\| = 1\}$, but other objects like the torus are possible as well. In this section we will concentrate on the special case of parameterizing triangulated point sets that are topologically disk-like and a simply connected planar domain Ω, but the general idea of the presented parameterization technique is also applicable to other domain types.

Using the notation of the previous section, we are looking for a projection $f : \Psi \to \Omega$. Since the given mesh $\mathcal{M}_\Psi$ is normally geometrically complex, this function will inevitably cause some deformation to the shape of the triangles T_j. A projection without distortion is called *isometric* and can only be found for developable surfaces, e.g. planes, cylinders and conical surfaces. Such a projection would be optimal for our purposes, because the inflation $F = f^{-1}$, that is used to create the final remesh $\mathcal{S}_\Psi$ from the planar remesh $\mathcal{S}_\Omega$, would be isometric, too. Hence, $\mathcal{S}_\Psi$ would inherit the quality from $\mathcal{S}_\Omega$, since the shape of the triangles would not change. Therefore, we need to find a projection f that is "as isometric as possible".

The concept of MIPS [8] can be used to find such projections. This approach decomposes the piecewise linear projection f into *atomic linear maps* f_j that map the spatial triangle T_j to the corresponding planar triangle t_j. By introducing a local coordinate system at T_j with the third axis chosen to be perpendicular to T_j, these linear maps can be written as $f_j : \mathbb{R}^2 \to \mathbb{R}^2, x \mapsto A_j x + b_j$ and the deformation of the triangle T_j can be measured by the *2-norm condition* of the matrix A_j,

$$\kappa_2(A_j) = \|A_j\|_2 \, \|A_j^{-1}\|_2 = \sigma_1 \cdot \frac{1}{\sigma_2} = \frac{\sigma_1}{\sigma_2},$$

where $\sigma_1 \geq \sigma_2 \geq 0$ are the *singular values* of A_j. The singular values represent the lengths of the semi-axes of the ellipse $\{A_j x : \|x\|_2 = 1\}$ and their ratio therefore is a measure of the deformation that is applied to any planar region when being transformed by A_j. However, for practical reasons it is easier to use the *Frobenius norm condition*

$$\kappa_F(A_j) = \|A_j\|_F \, \|A_j^{-1}\|_F = \sqrt{\sigma_1^2 + \sigma_2^2}\sqrt{\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}},$$

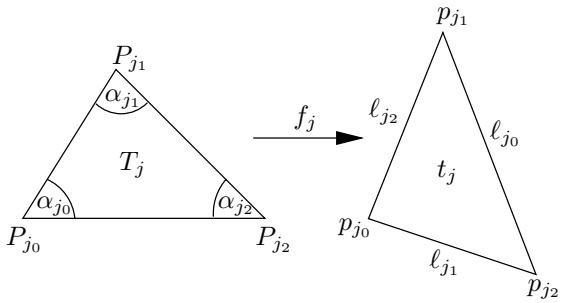

Figure 3: The atomic linear map f_j maps the spatial triangle T_j to the corresponding planar triangle t_j.

which is closely related to κ_2 and can be computed by the following formula (see [19] and [8] for details)

$$\begin{aligned}\kappa_F(A_j) &= \kappa_2(A_j) + \frac{1}{\kappa_2(A_j)} \\ &= \frac{\cot\alpha_{j_0}\|\ell_{j_0}\|^2 + \cot\alpha_{j_1}\|\ell_{j_1}\|^2 + \cot\alpha_{j_2}\|\ell_{j_2}\|^2}{2\,\mathrm{area}(t_j)} \end{aligned} \quad (1)$$

with the notations of Fig. 3, i.e.

$$\ell_{j_0} = p_{j_2} - p_{j_1}, \quad \ell_{j_1} = p_{j_2} - p_{j_0}, \quad \ell_{j_2} = p_{j_1} - p_{j_0},$$

and

$$2\,\mathrm{area}(t_j) = \|\ell_{j_1} \times \ell_{j_2}\| = \ell^x_{j_1}\ell^y_{j_2} - \ell^y_{j_1}\ell^x_{j_2}.$$

Minimization of $\kappa = \sum_j \kappa_F(A_j)$ will then lead to an optimal projection f in the sense of minimal deformation.

This is an optimization problem in the unknown parameter values p_i and according to Eq. (1) highly nonlinear. Instead of solving this problem globally, we use a Gauss-Seidel approach to minimize κ with a series of local optimization steps. Starting with an initial set of parameter values that can be obtained by one of the methods described in [4, 5, 6], we perform the following algorithm, illustrated in Fig. 4(a).

repeat

choose a vertex p_i by random

let $t_{i_1}, ..., t_{i_n}$ be the triangles that surround p_i

fix the positions of all parameter values except for p_i

minimize the local deformation energy $\kappa_i = \sum_{k=1}^{n} \kappa_F(A_{i_k})$ in order to get the optimal position $\tilde{p}_i$ of p_i

until numerical convergence

Each function $\kappa_F(A_{i_k})$ w.r.t. p_i is convex on both sides of the line that is defined by the edge in the corresponding triangle t_{i_k} opposite to p_i and tends to infinity along that line. As the sum of such functions, the local functional κ_i is convex and grows to infinity along the edges that surround the parameter value p_i (cf. Fig. 4(b)). Since Eq. (1) can be used to derive explicit formulas for the first and second derivates of κ_i, Newton's method is suitable for solving the local optimization problem. Usually the optimal $\tilde{p}_i$ can be found with two of three Newton steps. It can be proven that these local optimization steps converge globally and are stable in the sense of validity: if the initial configuration is valid (i.e. neighboring triangles do not overlap) this property is preserved by each optimization step.

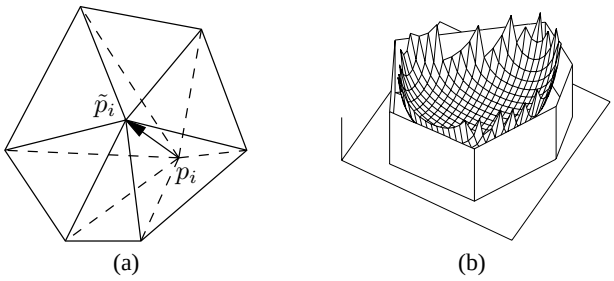

Figure 4: Each optimization step finds the optimal position $\tilde{p}_i$ of the parameter value p_i (a) by minimizing the convex local functional κ_i (b).

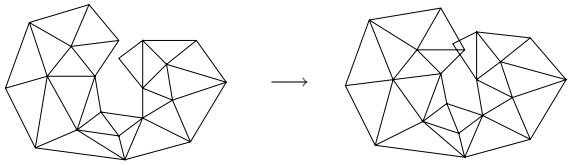

Figure 5: MIPS may generate overlapping parts in the parameter domain.

A major advantage of MIPS is that the optimization step can also be performed on the parameter values at the boundary of the triangulation and does not need the boundary to be fixed in advance. Of course this may lead to overlaps of different boundary regions (cf. Fig. 5) although we never experienced it in our examples. However, if we think of the different overlapping parts being assigned to different depth layers, it is still possible to use this configuration in order to "lift" the planar SCM $\mathcal{S}_\Omega$ up to the space Ψ as long as we add some depth layer information to the geometric information of the vertices p_i of $\mathcal{S}_\Omega$, e.g. by storing the number of the triangle of $\mathcal{M}_\Omega$ in which each p_i is located.

We close this section by briefly mentioning that the optimization process of MIPS can be significantly accelerated by a multiresolution approach as explained in [9]. The idea is to use the concept of progressive meshes [7] to build a hierarchy $\mathcal{M}_\Psi = \mathcal{M}^0, ..., \mathcal{M}^k$ of meshes and create the optimal projection for the coarsest mesh $\mathcal{M}^k$ first. Then the remaining vertices of $\mathcal{M}_\Psi$ are successively inserted and optimized until the final planar mesh $\mathcal{M}_\Omega$ is generated.

5 Remeshing

As described in Section 3, the task of generating a remesh $\mathcal{S}_\Psi$ with subdivision connectivity of a given mesh $\mathcal{M}_\Psi$ consists of three steps. We have already explained how to compute the parameterization of $\mathcal{M}_\Psi$. The second step is to find an appropriate remesh $\mathcal{S}_\Omega$ for the mesh $\mathcal{M}_\Omega = f(\mathcal{M}_\Psi)$ in the parameter domain Ω.

In order to find the remesh $\mathcal{S}_\Omega$ in the parameter domain we have to solve several problems which are discussed in this section. First, we have to determine an appropriate base mesh $\mathcal{S}^0_\Omega$, then we have to subdivide this base mesh until a prescribed error bound is reached. During the subdivision process we have to reconstruct the boundary polygon of $\mathcal{M}_\Omega$ with the remesh $\mathcal{S}^l_\Omega$ and the triangles $t_j \in \mathcal{S}^l_\Omega$ have to be scaled so that the corresponding triangles $T_j = F(t_j)$ are of equal size.

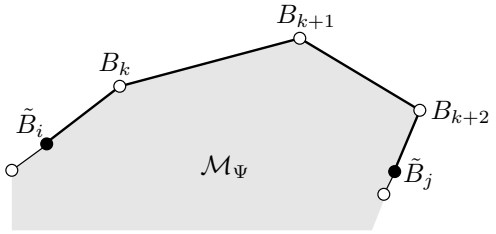

Figure 6: Computing the distance between two boundary points: $\text{dist}(\tilde{B}_i, \tilde{B}_j) = \|\tilde{B}_i - B_k\| + \|B_k - B_{k+1}\| + \|B_{k+1} - B_{k+2}\| + \|B_{k+2} - \tilde{B}_j\|$

5.1 Constructing the base mesh

For remeshing meshes with boundary it is possible to use just a single triangle as a base mesh. But in most cases this will not lead to a satisfying remesh of the given surface because of a high distortion of the triangles. It is therefore important to find a good initial mesh $\mathcal{S}^0_\Omega$ in order to avoid this problem. This base mesh should meet all quality criterions mentioned in Section 3, i.e. consisting of a few number of triangles that are as equilateral as possible. Note that the latter criterion refers to the triangles mapped into 3-space. Furthermore the base mesh should approximate the boundary of $\mathcal{M}_\Omega$.

One important aspect of the remeshing algorithm is the reconstruction of the boundary of $\mathcal{M}_\Psi$. Hence we start the construction of the base mesh by specifying extraordinary vertices on the boundary polygon of the given mesh $\mathcal{M}_\Psi$ which consists of a set of vertices $B = \{B_0, B_1, ..., B_n = B_0\}$. We compute the angle $\measuredangle(\overrightarrow{B_iB_{i-1}}, \overrightarrow{B_iB_{i+1}})$ of the boundary polygon at every boundary vertex B_i. If the value exceeds a given threshold α, the corresponding vertex $b_i = f(B_i)$ of the mesh $\mathcal{M}_\Omega$ will be a vertex $\tilde{b}_j$ of the base mesh $\mathcal{S}^0_\Omega$.

These vertices $\tilde{b} = \{\tilde{b}_i\}$ may be distributed irregularly along the boundary polygon. To avoid triangles of different size in the base mesh we now specify additional vertices on the boundary polygon so that the distance between all boundary vertices in the base mesh is approximately the same. The distance of two boundary vertices $\tilde{b}_i$ and $\tilde{b}_j$ of the base mesh $\mathcal{S}^0_\Omega$ is computed as follows. Using the inflation function F, we obtain the points $\tilde{B}_i = F(\tilde{b}_i)$ and $\tilde{B}_j = F(\tilde{b}_j)$ on the boundary polygon of $\mathcal{M}_\Psi$. We now compute the distance between these points by following the boundary line from one point to the other (cf. Fig. 6).

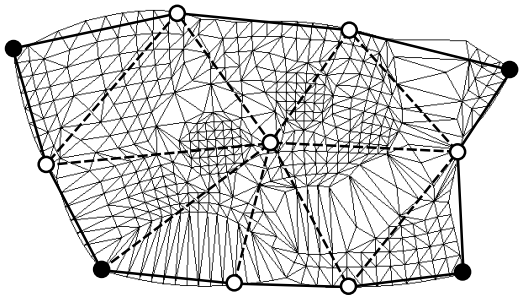

Figure 7: Constructing a base mesh. The filled vertices were found by searching for extraordinary vertices of the original mesh, the outlined boundary vertices were inserted to adapt the distances on the boundary polygon and the interior vertex was inserted to improve the mesh structure.

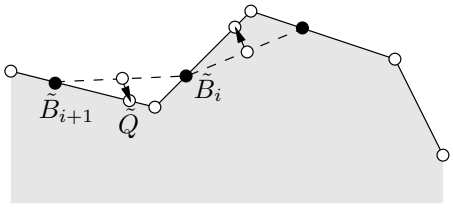

Figure 8: New boundary vertices of the remesh (dashed) have to be moved onto the boundary polygon of the parameterization.

The set of vertices we have chosen so far defines the boundary of the base mesh and has to be triangulated. Notice that the vertices also lie on the boundary polygon of $\mathcal{M}_\Omega$. Triangulating only the boundary vertices does not lead to a regular structure with almost equilateral triangles. Therefore the triangulation algorithm is allowed to insert new interior vertices, called *Steiner points*, to the base mesh [21]. There are two constraints for the position of these interior vertices. The angles of the triangles should be close to 60° and the area representing the triangles in the original mesh $\mathcal{M}_\Psi$ should be about the same.

This algorithm enables us to compute a base mesh $\mathcal{S}^0_\Omega$ which is the starting point of the remeshing process. In practice these base meshes will have only a small number of triangles (cf. Fig. 7).

5.2 Border adaption

While parameterizing the mesh $\mathcal{M}_\Psi$ using MIPS, the boundary is allowed to develop naturally. Hence, the parameterization will not be a rectangular area and not even convex. All boundary vertices of $\mathcal{S}^l_\Omega$ also belong to the boundary polygon of $\mathcal{M}_\Omega$. The subdivision scheme we use just computes the new vertices as the midpoints of the edges on the coarser level $\mathcal{S}^l_\Omega$ and retriangulates the mesh leading to $\mathcal{S}^{l+1}_\Omega$. The additional boundary vertices of $\mathcal{S}^{l+1}_\Omega$ are not located on the boundary of $\mathcal{M}_\Omega$ but somewhere inside or outside the mesh (cf. Fig. 8).

For these vertices a new position on the boundary polygon of the parameterization has to be found. The algorithm we use works as follows. We have to find a position for the vertex $\tilde{q}$, which has been inserted between the boundary vertices $\tilde{b}_i$ and $\tilde{b}_{i+1}$. Using the inflation function F, we obtain $\tilde{B}_i = F(\tilde{b}_i)$ and $\tilde{B}_{i+1} = F(\tilde{b}_{i+1})$ which are points on the boundary of $\mathcal{M}_\Psi$. We now compute the distance between these points along the boundary line of the mesh as described in the last section and search for the position $\tilde{Q}$ midway between these points. The new position of the vertex $\tilde{q}$ in the parameter domain is then computed by $\tilde{q} = f(\tilde{Q})$. If the distances of the boundary vertices in the mesh $\mathcal{S}^l_\Omega$ are equal then this strategy is optimal and will lead to equidistant boundary vertices in $\mathcal{S}^{l+1}_\Omega$.

5.3 Density weighted smoothing

The deformation functional κ that we use to obtain the parameterization is invariant to scalings (see Eq. (1)). This may lead to scalings of triangles between the meshes $\mathcal{M}_\Psi$ and $\mathcal{M}_\Omega$. Hence, to minimize the distortion and to get approximately the same size for each triangle T_j of $\mathcal{S}^m_\Psi$ it is essential for the remeshing algorithm to use a relaxing operator to optimize the interior vertices of the remesh. For this relaxing we use the umbrella operator $\mathfrak{U}$ [10]. The umbrella operator minimizes the membrane energy of a mesh, i.e. the surface area. The update rule for a vertex $p_i \in \mathcal{S}^l_\Omega$ is

$$\mathfrak{U} : p_i \mapsto \tilde{p}_i = (1-\omega)p_i + \frac{\omega}{\sum_j d_{i_j}} \sum_{j=1}^{n} d_{i_j} p_{i_j}$$

with n being the valence of vertex p_i and $p_{i_1}, ..., p_{i_n}$ its adjacent neighbors in $\mathcal{S}_\Omega^l$. With the help of the density coefficients d_i the distribution of the vertices can be controlled.

In our case we want a vertex p to be moved so that neighboring triangles transformed into 3-space have approximately the same size. Therefore we use the sum of the area of the adjacent triangles T_j as the density coefficient d_i for the vertex p_i. We apply the density weighted umbrella only to the interior vertices of the mesh. The boundary vertices are not affected by this relaxing operation.

The density weighted umbrella operator is applied iteratively. With the vertices moving in the parameter domain the size of the triangles has to be recomputed after every iteration of the operator $\mathfrak{U}$. The iterations stop when the change of the vertex positions falls below a certain threshold. Then the mesh $\mathcal{S}_\Omega^l$ can be refined to the next finer level $\mathcal{S}_\Omega^{l+1}$ and the process of border adaption and relaxing iterations starts again. When reaching a certain level m the subdivision process is stopped and the mesh is transformed into 3-space by evaluating the inflation function F. For every vertex $\tilde{p}$ of $\mathcal{S}_\Omega^m$ we determine the triangle $t_j = \triangle(p_{j_0}, p_{j_1}, p_{j_2}) \in \mathcal{M}_\Omega$ that surrounds $\tilde{p}$ and calculate the barycentric coordinates

$$\tilde{p} = \alpha p_{j_0} + \beta p_{j_1} + \gamma p_{j_2}, \quad \alpha + \beta + \gamma = 1.$$

The corresponding point $\tilde{P} = F(\tilde{p})$ in 3-space can now be found by applying the same barycentric combination to the vertices $P_{j_0} = F(p_{j_0})$, $P_{j_1} = F(p_{j_1})$, $P_{j_2} = F(p_{j_2})$ of the corresponding triangle $T_j \in \mathcal{M}_\Psi$

$$\tilde{P} = \alpha P_{j_0} + \beta P_{j_1} + \gamma P_{j_2}.$$

Since MIPS generates minimal distortion it is not necessary to use the smoothing and projecting technique in 3-space as proposed in [11]. So the risk of failing projections is excluded. Examples which are remeshed with our algorithm are shown in Section 6.

5.4 Adaptive remeshing

In order to reach the user specified error threshold for the distance of the two meshes $\mathcal{M}$ and $\mathcal{S}$ in regions of small local features, the mesh $\mathcal{S}$ has to be subdivided up to a very fine level. When using uniform subdivision the number of triangles will grow exponentially. As in [13] we use a straight forward adaptive subdivision approach to keep the number of triangles small and avoid the overhead of first subdividing the mesh to a fine level and afterwards decimating triangles by a wavelet threshold like in [3, 4].

Adaptive subdivision splits only those triangles which do not satisfy some prescribed criterion while other triangles remain coarse. There are a few restrictions in this approach. In order to keep the number of special configurations small we will only allow *balanced* meshes, i.e. the refinement level of two neighboring triangles may only differ by one. In order to avoid cracks in the mesh where two triangles from different levels meet, we have to use a special technique, the so-called *red-green triangulation* [1, 22, 23]. A normal 1-to-4 split is called *red split*. To fix cracks in the mesh, triangle bisection is used, which is called a *green split*. A green split is only temporary, i.e. if a green split triangle is to be further subdivided in a subsequent refinement step then the green split is undone first and a red split is applied to the original triangle.

As a subdivision criterion we compute the distance function

$$D(t_j) = \max_i \operatorname{dist}(F(p_i), F(t_j))$$

for every triangle t_j of the actual mesh $\mathcal{S}_\Omega^l$. In this equation p_i are those vertices of $\mathcal{M}_\Omega$ which are surrounded by the triangle t_j. The computation of the distance function is performed in 3-space Ψ with the help of the inflation mapping F. If the maximum distance between such a vertex $F(p_i)$ and the plane defined by the triangle $T_j = F(t_j)$ is greater than a given threshold ε the triangle t_j will be subdivided. With this method only a minimal number of triangles has to be refined.

6 Examples

In this section we show some results of our remeshing algorithm. We also illustrate the quality of the used parameterization method (MIPS) by showing an example with different parameterizations.

Fig. 9 shows a remesh of a head data set consisting of 21,680 triangles. We used a simple mesh of four triangles as a base mesh. By uniform subdivision and smoothing with the density umbrella operator we get triangles of almost equal size in the remesh. The results of both the planar ($\mathcal{S}_\Omega^i$) and the corresponding spatial remesh ($\mathcal{S}_\Psi^i$) can be seen at different levels $i = 3, 5, 7$ in Fig. 9. These remeshes consist of $4{\cdot}4^3 = 256$, $4{\cdot}4^5 = 4,096$ and $4{\cdot}4^7 = 65,536$ triangles, resp.

If we want to improve the quality of the remesh at regions with local detail, e.g. at the ear, we have to sample these areas with a higher density. Using adaptive subdivision with an error bound of $\varepsilon = 0.1\%$ results in the adaptively refined remesh with 64,398 triangles in Fig. 10 (right). Comparing this mesh with the uniformly refined one, it can be clearly seen that the adaptive approach reproduces more detail with the same number of triangles.

In Fig. 11 the result of the remeshing process of a triangulated data set with 7,938 triangles (top left) is illustrated. The remesh (top center and right) consists of 10,240 triangles and was generated by splitting a base mesh with 10 triangles 5 times uniformly. In the last row we can see how the choice of the parameterization method affects the remesh. Using a uniform parameterization for the remeshing process leads to a mesh with highly distorted triangles. The structure of the triangles is much better when the discrete harmonic parameterization is used and can further be improved at the boundary of the mesh by using MIPS.

However, the minimization of this non-linear functional requires 13 minutes, while the linear problem of finding the uniform or the discrete harmonic parameterization can be solved in only 34 seconds on a Silicon Graphics O_2 with a 195 Mhz R10000 processor.

7 Conclusion

In this paper we have presented an algorithm for converting triangle meshes with arbitrary connectivity into meshes with subdivision connectivity. The algorithm works on triangle meshes with boundary and no holes. For this class of meshes we can compute a planar parameterization with minimal distortion. Then we compute a remesh for this parameterization which is mapped into 3-space in order to obtain a remesh of the original mesh. By using an iterative relaxation process on every level of the remesh we adapt the remesh to the local scalings of the parameterization.

With this method all triangles of a uniformly refined remesh will have almost equal size. Since a higher sampling density of the original mesh is often needed in highly detailed regions, we use adaptive subdivision to compute remeshes with guaranteed error bounds.

By using the most isometric parameterization strategy on a sphere the results for remeshing polygonal surfaces of genus 0 with the shrink-wrapping approach [11] can be improved, too. There, a parameterization generated by a density weighted umbrella operator was used which fails to guarantee minimal distortion of the triangles in the parameter domain.

Future work will aim at remeshing not only disk-like objects without holes. To be able to handle meshes of arbitrary topology an appropriate segmentation of the whole object has to be found. Then the segments can be remeshed separately with the presented algorithm and merged afterwards. Therefore it is important for adjacent segments to share the vertices along the common boundary.

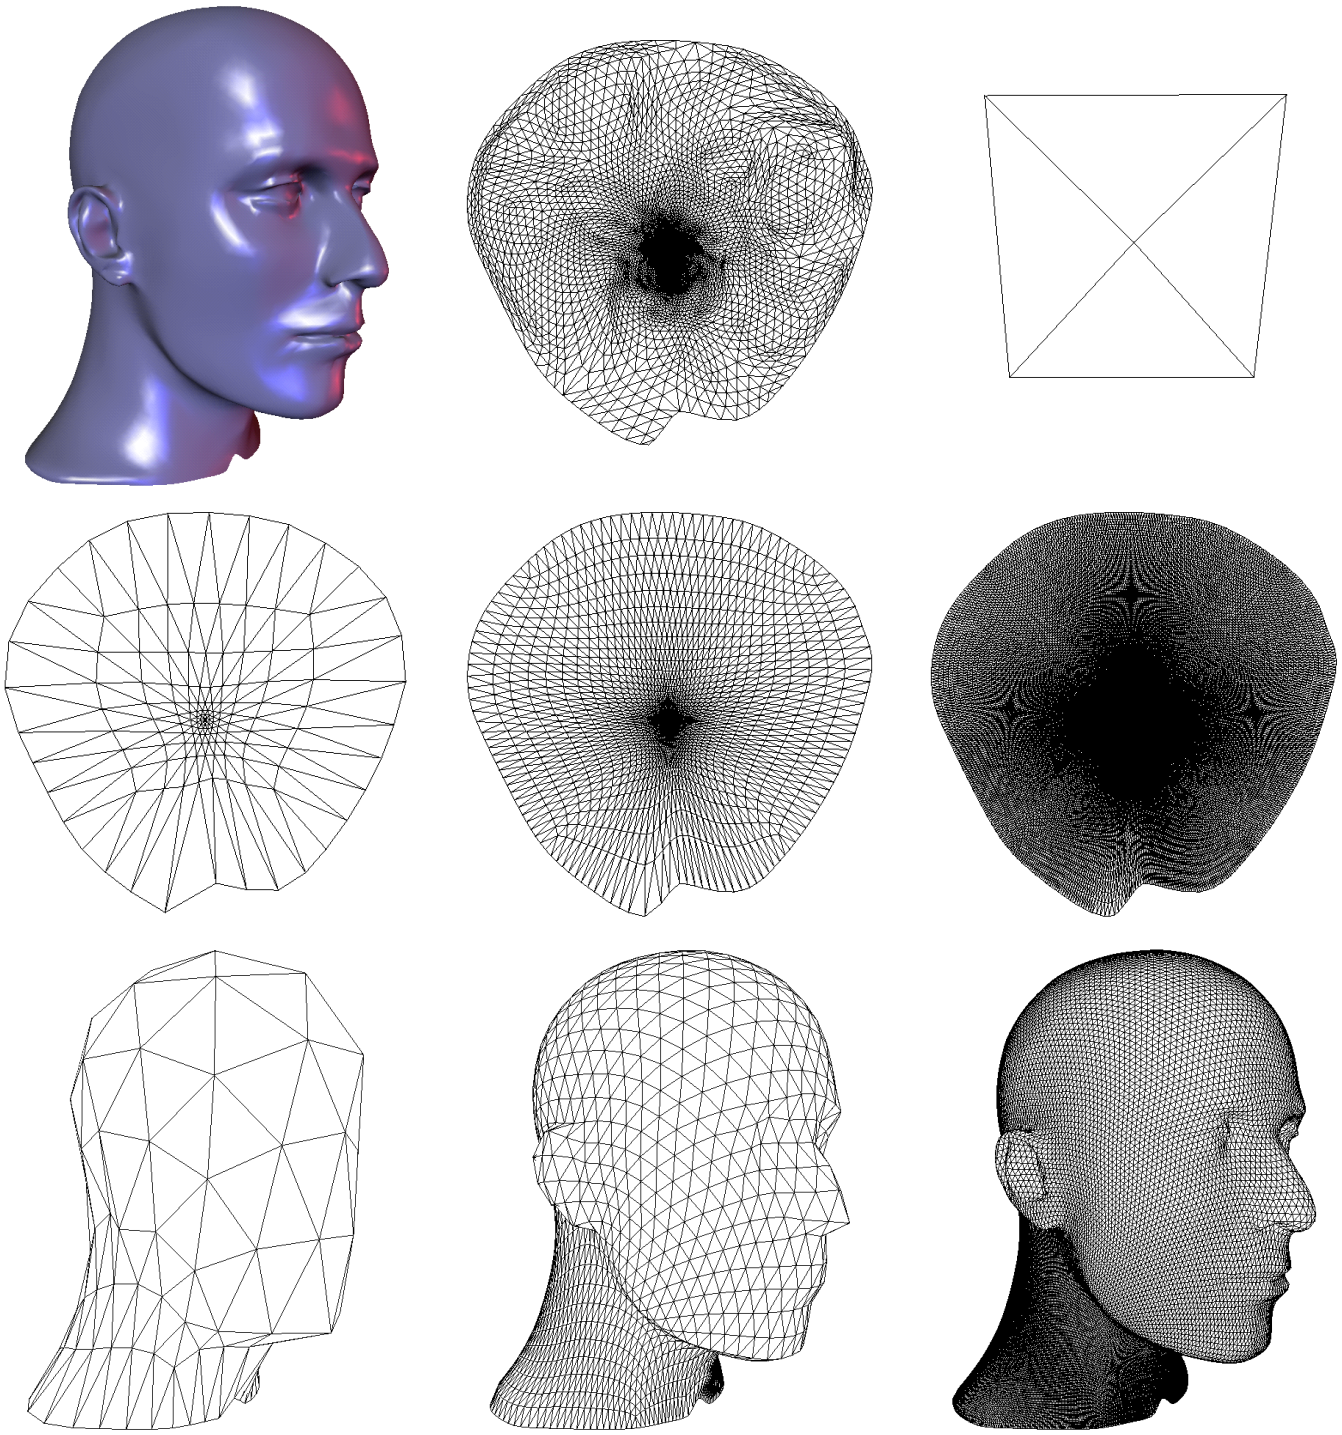

Figure 9: Remeshing the head data set. The original data set (top left) was parameterized with MIPS (top center). The base mesh for the remeshing procedure consists of four triangles (top right). The central and lower row show the planar and the corresponding spatial remesh at the 3rd, 5th and 7th refinement level.

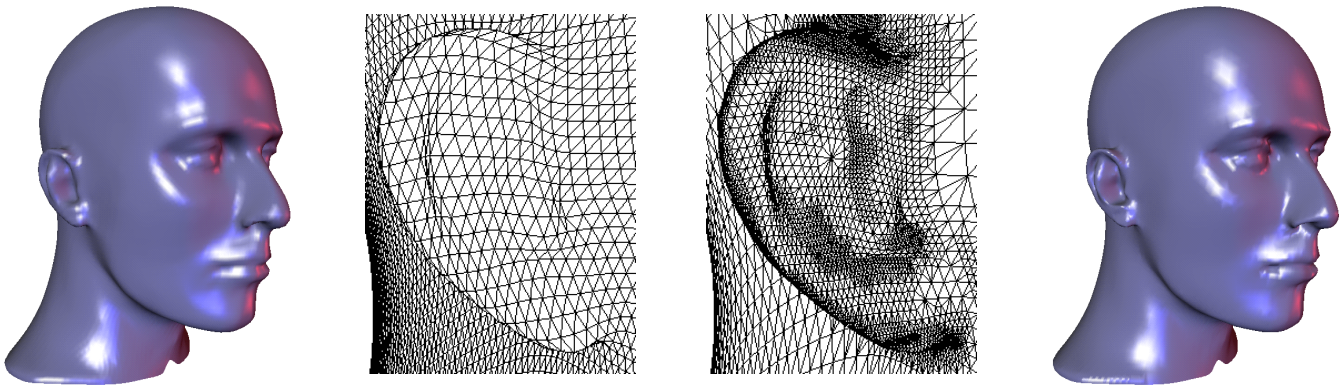

Figure 10: The left side shows the uniformly refined remesh at the 7th refinement level and a close up view of the right ear, while the adaptively refined remesh is shown on the right side.

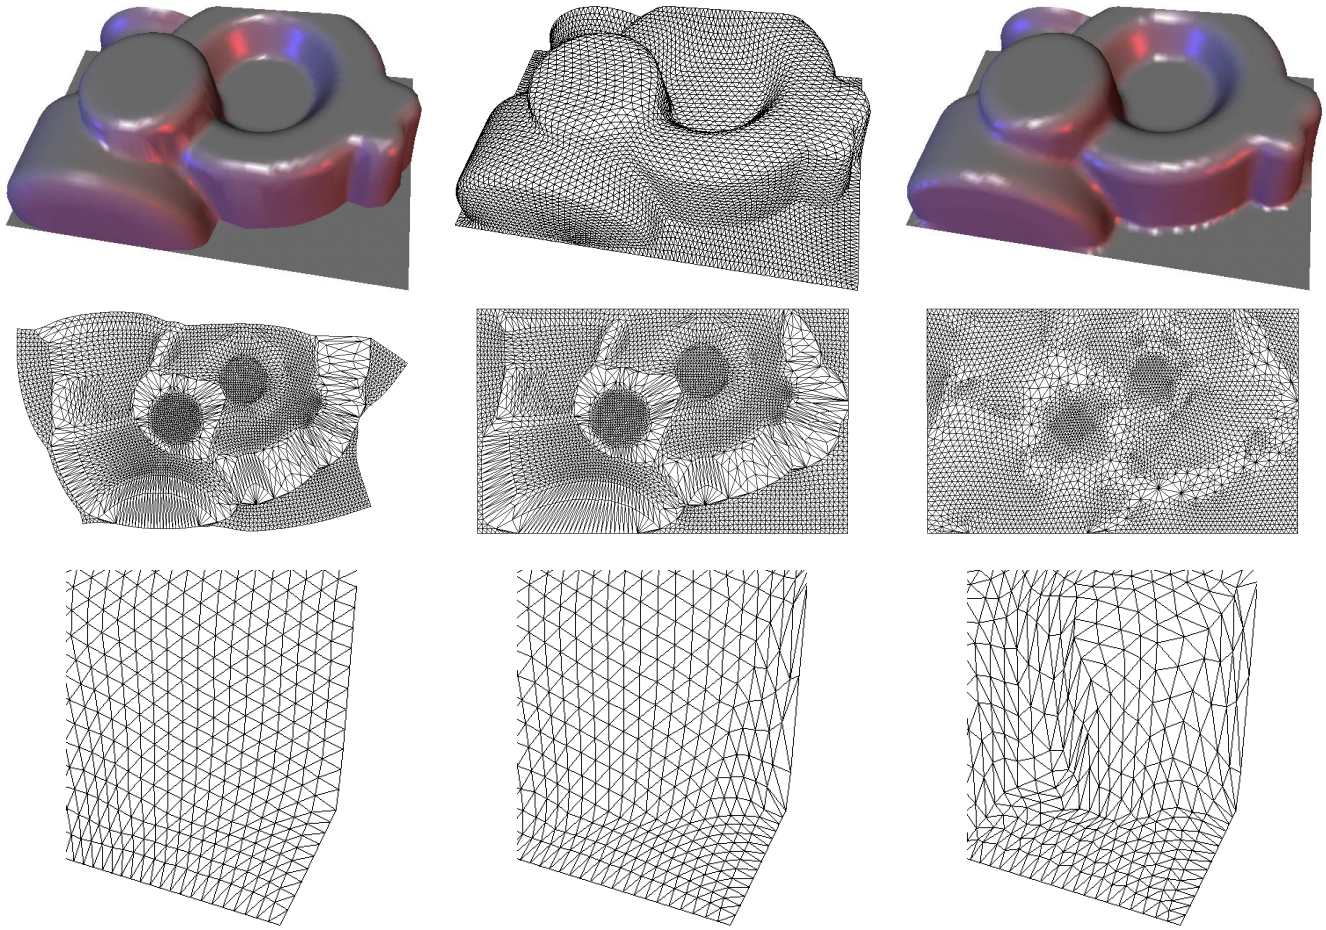

Figure 11: The top row shows a technical data set (left) and its remesh (center and right) using MIPS. The central row shows different parameterizations of the data set (from left to right: MIPS, discrete harmonic and uniform parameterization) and a close up view of the corresponding remeshes can be seen at the bottom.

Prominent candidates for the boundaries of such segments are feature or character lines. Segmenting the given mesh along such lines would also help to overcome the problem of oscillations that the current method tends to produce if the edges of the remesh are not aligned to them (cf. Fig. 11). Alternatively, this problem could be solved by making the edges of the coarse base mesh follow the feature lines as in [13].

References

[1] R. E. Bank, A. H. Sherman, and A. Weiser. Refinement algorithms and data structures for regular local mesh refinement. In R. Stepleman, editor, *Scientific Computing*, pages 3–17, Amsterdam, 1983. IMACS/North Holland.

[2] C. Bennis, J.-M. Vézien, and G. Iglésias. Piecewise surface flattening for non-distorted texture mapping. In *ACM Computer Graphics (SIGGRAPH '91 Proceedings)*, pages 237–246, 1991.

[3] A. Certain, J. Popović, T. DeRose, T. Duchamp, D. Salesin, and W. Stuetzle. Interactive multiresolution surface viewing. In *ACM Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 91–98, 1996.

[4] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In *ACM Computer Graphics (SIGGRAPH '95 Proceedings)*, pages 173–182, 1995.

[5] M. S. Floater. Parameterization and smooth approximation of surface triangulations. *Computer Aided Geometric Design*, 14:231–250, 1997.

[6] G. Greiner and K. Hormann. Interpolating and approximating scattered 3D data with hierarchical tensor product B-splines. In A. Méhauté, C. Rabut, and L. L. Schumaker, editors, *Surface Fitting and Multiresolution Methods*, pages 163–172. Vanderbilt University Press, 1997.

[7] H. Hoppe. Progressive meshes. In *ACM Computer Graphics (SIGGRAPH '96 Proceedings)*, pages 99–108, 1996.

[8] K. Hormann and G. Greiner. MIPS: an efficient global parametrization method. In P.-J. Laurent, P. Sablonnière, and L. L. Schumaker, editors, *Curve and Surface Design: Saint-Malo 1999*, pages 153–162. Vanderbilt University Press, 2000.

[9] K. Hormann, G. Greiner, and S. Campagna. Hierarchical parametrization of triangulated surfaces. In B. Girod, H. Niemann, and H.-P. Seidel, editors, *Vision, Modeling and Visualization '99*, pages 219–226. infix, 1999.

[10] L. Kobbelt, S. Campagna, J. Vorsatz, and H-P. Seidel. Interactive multi-resolution modeling on arbitrary meshes. In *ACM Computer Graphics (SIGGRAPH '98 Proceedings)*, pages 105–114, 1998.

[11] L. Kobbelt, J. Vorsatz, U. Labsik, and H.-P. Seidel. A shrink wrapping approach to remeshing polygonal surfaces. In *Computer Graphics Forum (EUROGRAPHICS '99 Proceedings)*, pages 119–130, 1999.

[12] U. Labsik, L. Kobbelt, R. Schneider, and H.-P. Seidel. Progressive transmission of subdivision surfaces. *Computational Geometry*, 15:25–39, 2000.

[13] A. Lee, W. Sweldens, P. Schröder, L. Coswar, and D. Dobkin. Multirelosution adaptive parametrization of surfaces. In *ACM Computer Graphics (SIGGRAPH '98 Proceedings)*, pages 95–104, 1998.

[14] B. Lévy and J. L. Mallet. Non-distorted texture mapping for sheared triangulated meshes. In *ACM Computer Graphics (SIGGRAPH '98 Proceedings)*, pages 343–352, 1998.

[15] C. Loop. Smooth sudivsion surfaces based on triangles. Master's thesis, Utah University, 1987.

[16] M. Lounsbery, T. DeRose, and J. Warren. Multiresolution analysis for surfaces of arbitrary topological type. Technical report 93-10-05, University of Washington, Department of Computer Science and Engineering, 1993.

[17] W. Ma and J. P. Kruth. Parameterization of randomly measured points for least squares fitting of B-spline curves and surfaces. *Computer-Aided Design*, 27:663–675, 1995.

[18] J. Maillot, H. Yahia, and A. Verroust. Interactive texture mapping. In *ACM Computer Graphics (SIGGRAPH '93 Proceedings)*, pages 27–34, 1993.

[19] U. Pinkall and K. Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2:15–36, 1993.

[20] P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. In *ACM Computer Graphics (SIGGRAPH '95 Proceedings)*, pages 161–172, 1995.

[21] J. R. Shewchuk. Triangle: Engineering a 2D Quality Mesh Generator and Delaunay Triangulator. In M. C. Lin and D. Manocha, editors, *Applied Computational Geometry: Towards Geometric Engineering*, volume 1148 of *Lecture Notes in Computer Science*, pages 203–222. Springer-Verlag, 1996.

[22] M. Vasilescu and D. Terzopoulus. Adaptive meshes and shells: Irregular triangulation, discontinuities, and hierarchical subdivision. In *Proceedings of Computer Vision and Pattern Recognition conference*, pages 829–832, 1992.

[23] R. Verfürth. *A review of a posteriori error estimation and adaptive mesh refinement techniques*. Wiley-Teubner, 1996.

[24] D. Zorin, P. Schröder, and W. Sweldens. Interactive multiresolution mesh editing. In *ACM Computer Graphics (SIGGRAPH '97 Proceedings)*, pages 259–268, 1997.