Isotropic Surface Remeshing

Pierre Alliez, Éric Colin de Verdière, Olivier Devillers, Martin Isenburg

► To cite this version:

Pierre Alliez, Éric Colin de Verdière, Olivier Devillers, Martin Isenburg. Isotropic Surface Remeshing. International Conference on Shape Modeling and applications., May 2003, Seoul, South Korea. inria-00413144

HAL Id: inria-00413144

<https://hal.inria.fr/inria-00413144>

Submitted on 3 Sep 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Isotropic Surface Remeshing

Pierre Alliez
INRIA Sophia-Antipolis

Éric Colin de Verdière
ENS Paris

Olivier Devillers
INRIA Sophia-Antipolis

Martin Isenburg
UNC at Chapel Hill

Abstract

This paper proposes a new method for isotropic remeshing of triangulated surface meshes. Given a triangulated surface mesh to be resampled and a user-specified density function defined over it, we first distribute the desired number of samples by generalizing error diffusion, commonly used in image halftoning, to work directly on mesh triangles and feature edges. We then use the resulting sampling as an initial configuration for building a weighted centroidal Voronoi tessellation in a conformal parameter space, where the specified density function is used for weighting. We finally create the mesh by lifting the corresponding constrained Delaunay triangulation from parameter space. A precise control over the sampling is obtained through a flexible design of the density function, the latter being possibly low-pass filtered to obtain a smoother gradation. We demonstrate the versatility of our approach through various remeshing examples.

Keywords: Surface sampling, error diffusion, centroidal Voronoi tessellation, constrained Delaunay triangulation, parameterization, optimal cutting, polygonal schema.

1 Introduction

Many applications in simulation, visualization, or computer graphics require a model of 3D surface geometry. The most commonly used representation is the triangle mesh. Such meshes can be the result of careful design using modeling software, or may come as an output of a scanning device, associated with reconstruction or computer vision algorithms.

In this paper, we make the assumption that geometric details are captured accurately in the given model. The original mesh can be seen as one particular instance of the surface geometry of interest. We aim at generating new instances (*e.g.*, new triangle meshes) of this surface geometry that better fit user-specified demands on complexity, sampling, regularity, connectivity, gradation, and quality. For example, a laser scanner often performs a uniform sample acquisition while sweeping a line or point-based beam, without any a priori knowledge of the surface content. This may lead to under-sampling or oversampling of certain regions, mainly depending on the angle between the laser beam and the surface normal. Such drawbacks are often alleviated by tuning the sampling rate sufficiently high. Note that the same problems arise for surface extraction from volume data using marching cubes algorithms, where choosing a sufficiently fine grid may lead to overly complex models and a vertex distribution not consistent with respect to sampling theory. In general, simplification techniques are then applied to reduce the complexity and better adapt the sampling to geometry. The lack of flexibility found in the (re-)sampling strategies of most simplification algorithms motivates this work.

Our goal is to provide a flexible technique to remesh triangulated surfaces so that the remeshed models are better suited to a subsequent process, *e.g.* visualization, finite element simulation, storage, transmission or any mesh processing technique. See Figure 1 for example results of the proposed remeshing technique.

1.1 Related Work

Mesh generation received much attention from various interest groups ranging from Computer Graphics over Numerical Analy-

Figure 1: *Top left: the Darmstadt benchmark model. Bottom: two examples of uniform and curvature-adapted remeshing of the Darmstadt benchmark model with 5k and 10k vertices respectively. Top right, two close-ups of the weighted centroidal Voronoi tessellation computed in a conformal parameter space to produce the sampling.*

sis to Computational Geometry. Finite element mesh generation usually amounts to find a partition of a given domain that is optimal according to some criteria related to shape of elements, angles, sizes or complexity (see [3, 4]). In most cases only the frontier of the domain has to be given, the goal being to discretize this domain in accordance with an importance function. Our problem is slightly different since it is dealing with surface *remeshing*. The domain to discretize is now given by an original surface mesh that has to be re-discretized so that the result best matches some user-specified properties.

Mesheres for numerical analysis Numerical analysts essentially focus on mesh quality, since it impacts the numerical accuracy of computations performed on the mesh elements [3, 4]. We distinguish between techniques that use a parameter space and techniques that act on an explicit mesh. The key idea of the first is to partition a parameter domain into sets of adjacent elements that have the same specified properties [14, 50, 23]. The key idea of the second is to progressively adapt an explicit mesh by performing elementary operations on its elements until it matches some specified properties [12, 44, 8, 7, 21, 46, 9].

Mesheres for graphics Computer graphists mainly focus on remeshing for efficient visualization or geometry processing [25, 49]. In an early work, Turk [52] proposed a re-tiling technique that resamples an input mesh by first applying a relaxation method to initially randomly place points, by then inserting those points into the mesh, and by finally removing the original vertices. We should also mention mesh simplification [22, 38, 28, 56] and refinement [5] methods, which also generate a new mesh starting from a given one. Such schemes primarily aim at adapting the complexity of the mesh to an acceptable level for graphics visualization hardware or simulation algorithms. For efficient mesh processing, most previous work have focused on semi-regular remeshing [36, 34, 26, 33, 29], the lat-

ter techniques often requiring a first simplification stage. Kobbelt *et al.* focus on feature sensitive remeshing techniques [54, 10, 11] to reduce the artifacts produced when converting a given geometry into a triangle mesh.

More recently, Gu *et al.* [24] proposed a technique for regular remeshing of surface meshes. Surfaces of arbitrary genus are first cut into patches, then parameterized using a signal-adapted technique [48], and finally represented as a set of images that store the geometry, the normals and any attributes used for visualization purpose. Such a storage is compact and drastically simplifies the rendering pipeline since all cache indirections found in usual irregular mesh rendering are simply removed. In a recent work, we have proposed an interactive sampling technique [2]. A mesh is decomposed into a set of disk-like patches, and each patch is parameterized. We then measure some geometric and differential quantities and generate a set of maps inserted in a pipeline of signal processing algorithms. The output of this pipeline is a density map, interactively resampled using an error diffusion technique commonly used for gray level image halftoning. Although the technique described in [2] offers a fairly good level of flexibility, its main drawback is the pixel grid layout constraint that turns out to be memory consuming for complex models with a high range of area distortion.

1.2 Goals and Contributions

Based on the above observations, the goal of this paper is to remove the constraint coming from the regular and uniform structure of images for both sampling and optimization. This led us to extend the concept of direct error diffusion onto triangle meshes for sampling. The motivation of this paper being also to formulate the issue of surface sampling with a larger set of Computational Geometry tools at hand, we demonstrate the relevance of building a weighted centroidal Voronoi tessellation for repartitioning a set of samples in accordance with a specified density function. Another motivation of this paper is to move from the *unit length* paradigm used for numerical analysis [23] to the *unit cell* tiling, well suited for targeted application, *i.e.* surface shape modeling. The first technique aim at generating meshes with unit edge length measured in a control space metric, while our algorithm tends to partition the surface with unit density *integrated* over the cells of a centroidal Voronoi tessellation. The latter property is indeed intricately related to the notion of *isotropic sampling*.

1.3 Overview of the algorithm

The first stage of our algorithm provides an initial geometry resampling by performing an *error diffusion* process directly over the original triangle mesh (see Section 3). The second stage computes a conformal parameterization of the original model over a planar domain, connects the samples using a constrained Delaunay triangulation built in parameter space, then optimizes the sampling by building a *weighted centroidal Voronoi tessellation* in parameter space (see Section 4). Figure 2 shows the tasks requiring a parameter space.

Figure 2: Remeshing pipeline.

2 Preliminaries

All meshes considered in this paper are presumed to be oriented manifolds of arbitrary genus and possibly with boundary. They are

seen as an approximation of a underlying piecewise smooth surface. The input of the algorithm is a set $(\mathcal{M}, \mathcal{F}, d_s, d_f)$, where $\mathcal{M}$ is a triangular mesh, $\mathcal{F}$ is a set of feature edges (defined below), $d_s : \mathcal{M} \rightarrow \mathbb{R}$ is an importance function that specifies the *ideal* sampling density for every surface point and $d_f : \mathcal{F} \rightarrow \mathbb{R}$ is an importance function that specifies the density for every point located on a feature edge.

We assume the density functions d_s and d_f to be specified by the user or deduced from geometric quantities measured on the given model. In our experiments, we use discrete differential geometry techniques [42] to approximate the curvature on every vertex or on the wedges surrounding a corner or a crease vertex. We also provide some options for the user to specify a transfer function (typ. a gamma function as described in [2]) and the amount of low-pass filtering over the density functions to control the final mesh gradation. Before diffusing the error using the resulting density function, we extract a set of *features edges* from the original mesh, which receive special treatment during the subsequent processing. They are of three types:

- ◊ the *sharp edges* represent the main features of the object; they are classified using dihedral angle thresholding, yet a more sophisticated approach could be used (see *e.g.* [30, 55] to cite a few);
- ◊ the *boundary edges* are incident to exactly one face;
- ◊ the *cut edges* are additional edges which we also need to consider separately for parameterization of closed or genus > 0 models (see Section 4.1).

Similarly to Botsch and Kobbelt [10], we then chain the set of incident feature edges as a set of backbones. The resulting feature skeleton will be resampled as a set of curves (see Figure 3).

Figure 3: Two models and their feature skeleton, made of both open and closed backbones.

3 Error Diffusion

Because of its simplicity and efficiency, the error diffusion algorithm received much attention [53, 43] since its introduction in 1976 by Floyd and Steinberg [20]. Before describing our algorithm, let us review the concept of error diffusion when used for grey level halftoning of images.

3.1 Concept

The core principle of error diffusion in image halftoning consists of processing each pixel of an input signal image according to a path: every pixel is binary quantized according to a given threshold. The signed quantization error is then distributed to its unprocessed neighbors according to error distribution coefficients. The art of error diffusion, among others, consists of finding the best *processing path* and choosing the best *distribution coefficients* so that the sampling spectrum exhibits a so-called blue noise profile (see [43, 31] for more details). One of the most appealing property of this concept for our application lies in the global conservation of density offered by the error diffusion. Such a property means that the average grey level of the discretized image is closely equal to the one of

the original image. We showed in [2] that this guarantees an exact vertex budget for any specified sampling.

In our context, the task of the sampler consists of distributing a set of samples on the original mesh triangles or feature edges so that they locally match the specified density. If one looks at the density as an input signal to discretize, this task is not far from the issue of image halftoning. It has been showed in [2] that recent error diffusion techniques [43] work fairly well for a certain class of models. In this paper we remove the discrete constraints of the image pixel grid layout by generalizing the concept of error diffusion directly over the original mesh triangles. This task is achieved by organizing a fluency over the mesh triangles for the smooth parts (*i.e.*, 2D error diffusion) and along the backbones of the feature skeleton (*i.e.*, 1D error diffusion).

Before performing error diffusion, the user specifies a number of vertices V to distribute on the original mesh. The total integral of the density functions, defined both on smooth parts and on the feature skeleton, represents the amount of density to equally distribute among the V samples. To fit the exact budget we first need to calibrate the sampler.

3.2 Sampling Calibration

Recall that we have a density function d_s and d_f over the surface and the features, respectively, which specify what would be the local ideal distribution of the samples. By summing the density function over the surface and the features, we obtain some quantities with different unit. This requires to define two distinct sampling rates: R_s , expressed in #samples per unit amount of surface density and R_f , in #samples per unit amount of feature density.

Uniform sampling For uniform sample distribution, the density functions d_s and d_f simplify to equally constants. Uniform sampling means that every surface sample ideally covers the same amount of area and that every feature sample covers the same length of feature curve. Since we seek isotropic sampling, the triangles can be considered as being ideally equilateral, allowing to deduce the local ideal edge length once each triangle area is known, and therefore link the sampling rates R_s and R_f .

Figure 4: Uniform plane tiling with triangles.

Let us consider a perfect uniform isotropic tiling of the plane with a triangle mesh (Figure 4). By Euler's formula, the number of triangles T per area unit is equal to twice the number of samples per area unit, namely $T = 2 \cdot R_s$. The ideal area of every triangle is therefore equal to $A = 1/(2R_s)$. From geometry of equilateral triangle we know:

$$A = \frac{a^2}{4}\sqrt{3}$$

where a is the triangle edge length. From the relation $R_f = 1/a$ we deduce:

$$R_s = 2 \frac{R_f^2}{\sqrt{3}}.$$

We now write the repartition of resources between smooth parts, features and corners to be preserved:

$$R_s \cdot \int_{\text{surface}} d_s(u, v) du dv + R_f \cdot \int_{\text{features}} d_f(u) du + C = V \quad (1)$$

where C denotes the total number of corners and V denotes the global vertex budget. Practically, the density functions d_s and d_f are linearly interpolated over the mesh triangles and the feature edges. The number of corners being known, as well as the triangle areas and the feature edges lengths, we deduce R_s and R_f required for the uniform sampling process. The same formula is performed for the non-uniform case after suited conversion.

3.3 Diffusion over triangles

We replace the pixels of a standard error diffusion technique by the mesh triangles, and organize upon them a fluency that goes through their edges (see Figure 5). The next task is to find a processing path over the triangles so that any residual error teleport is minimized and to build a set of distribution coefficients.

Figure 5: A fluency is organized over the mesh triangles. Each triangle is classified according to its number of input and output edges during the error diffusion process.

For genus 0 models with at most one boundary, the processing path is trivial. The algorithm picks an arbitrary seed face (on the boundary if there is) and processes by extending a region triangle by triangle without generating any split on its edge boundary (similarly to the *skip* regime described in [1]). The unprocessed region being topologically equivalent to a disk, it is considered as a triangulated, simple polygon. If this polygon contains no interior vertex, then it has at least two so-called ear-triangles [41] that we used to expand the processed region. If this polygon contains some interior vertices, then we search for ear-triangles and use one of those if exist. If there is no, we deduce that there exists at least one triangle with exactly one edge incident to the boundary and one interior vertex, used to expand the processed region. The diffusion algorithm then proceeds as follows for every current face f :

1. read the total amount of density on f ;
2. from the sampling rate described in Section 3.2, deduce the number of samples to distribute on the current face (the distribution inside a triangle is described later). This number is rounded to the nearest integer value, such a rounding generating a signed quantization error e , translated in amount of density, to diffuse on incident unprocessed faces;
3. the error e is diffused to the unprocessed faces incident to f through the corresponding edges and proportionally to their length. This heuristic mimics the notion of “geometric aperture” for diffusion;
4. flag f as processed;
5. pick the next face to proceed and restart from step 1 if any, stop otherwise.

Instead of choosing an arbitrary triangle among the valid candidates for expansion, we put them into a priority queue sorted by a region compacity score. The compacity ($\text{area}/\text{perimeter}$) of the expanding region is computed before and after simulating the expansion, and the face corresponding to the maximum signed change of compacity is chosen for diffusion. Note that several samples may be distributed on a face, the latter case occurring often in the case of over-sampling. In the uniform case they are randomly sampled in the current triangle using Turk's technique [51], combined with a low-discrepancy Sobol sequence for efficiency. The non uniform case is correctly handled using an adaptation of Rocchini *et al.* technique [47] after suited discretization of the density function.

Surfaces with handles and boundaries For models with genus > 0 , or with at least two boundaries, the traversal described above ends up with a *connected* set of unprocessed triangle strips, on which some error has been diffused during the region expansion process. We describe now a technique to sample these triangles without error teleport.

By dualization, we look at these unprocessed triangles as the vertices of a graph. First, we compute a spanning tree of this connected graph, rooted at an arbitrary triangle. Then we sample the triangles by diffusing the errors from the end leafs to the root. More precisely, assume that each unprocessed triangle t has received a signed error $e(t)$ in the previous part of the algorithm; to sample a given triangle t :

- ◊ recursively sample the children of t ;
- ◊ compute the ideal (float) number of samples to distribute on t , and round it. The ideal number of samples is computed by summing up the ideal density of t and the result of up to two errors distributed upon it at distinct moments in time during the error diffusion process;
- ◊ transmit the excess or loss of samples obtained with the rounding to the father f of t , by increasing or decreasing $e(f)$ accordingly.

We run this process from the root of the tree. All unprocessed triangles will be sampled recursively; at the end, the root is sampled and no error has been teleported nor dropped, matching thus the exact vertex budget. Once the smooth parts are sampled, the algorithm moves to the feature skeleton for sampling along its backbones.

3.4 Diffusion over feature edges

As for the triangles of surface parts, we design an error diffusion algorithm that distributes some samples along every backbone, the processing path being deduced from the ordering over its halfedges [32]. The diffusion process starts by picking the first edge of a feature backbone and do the following for every current edge :

1. read the total amount of density on the edge;
2. from the sampling rate described in Section 3.2, deduce the number of samples to distribute on the current edge. This number is rounded to the nearest integer value, such a rounding generating a signed quantization error e , translated in amount of density, to diffuse on the next unprocessed edge sitting on the current backbone. If the current sampled edge is the last of the processed backbone then we count the number of unprocessed feature edges connected to its end vertex (the latter may be a corner for an open backbone), and distribute the error to each of them with equal repartition. If there is no unprocessed edge connected to the end vertex, we teleport the error to the first edge of the next backbone picked for sampling;
3. pick the next edge in the current backbone to proceed if any, else move to the next unprocessed backbone while giving priority to any of a backbone connected to the end vertex of the current edge. When possible, the latter heuristic reduces the error teleport;
4. restart from step 1 if next feature edge exists.

The feature skeleton being possibly not connected, the error diffusion algorithm jumps from one feature backbone to the other and teleports the very last error on another. In a sense this is not consistent with the first goal of diffusing only on nearby areas, yet such a teleport allows to match the exact vertex budget and provides a way to repartition the samples between features, corners and surface parts of the original model to remesh since no error is dropped during the process.

Consistent boundary filtering Notice that such a diffusion technique offers one direct way to interleave boundary filtering with the sampling process. If one considers a smooth model with B boundaries, the feature skeleton is composed of B closed backbones. Each backbone corresponds to a certain amount of density, *i.e.*, a floating point number of samples once converted by the calibration process (note that this number may be negative once a negative error has been diffused from another backbone). After complete diffusion over the backbone, the total number of vertices sampled may be lower than 3, which leads to a closure of the boundary. In other terms, reducing the vertex budget may *filter* some boundaries in a consistent manner with respect to the sampling process. The latter is an appealing feature for applications similar in spirit to topology filtering as described in [27].

3.5 Discussion

The samples are sitting on the triangles of the original mesh for the smooth parts of the surface, and on the edges of the feature skeleton. The present generalization of direct error diffusion over a triangle mesh does not provide a sampling with a blue-noise profile spectrum such as *e.g.* [43]. This is due to the uniform sampling performed inside each triangle with no dependent probability nor any other sophistication (*i.e.*, white-noise profile), and to the heuristic chosen for distributing the error through the edges. Nevertheless, and thanks to the core principle of error diffusion the resulting sampling is consistent with respect to the specified density function since any error teleport is minimized. The error diffusion thus provides a simple way to *distribute* the samples between boundary loops, features, and connected components of smooth areas. It also serves as the initial guess for building a weighted centroidal Voronoi tessellation, which is the actual mechanism used to *place* the samples.

4 Centroidal Voronoi Tessellation

Before improving the sampling by building a (weighted) centroidal Voronoi tessellation, we first need to *mesh* the samples so that it generates an oriented manifold. The meshing is performed by parameterizing the original mesh onto a planar domain (possibly with holes), then applying 2D constrained Delaunay triangulation over the newly sampled vertices. If the mesh is closed or of genus > 0 , such a correspondence between the surface and a planar domain does not exist, and we need to cut the surface first.

4.1 Cutting

We discuss below several ways to cut the original surface into a genus 0 surface. Note that, although the cutting is necessary only at this point of the remeshing process, the edges which have to be cut are computed in the *first* stage of the algorithm, even before the error diffusion process, since they have to be marked as feature edges as described in Section 2. Since all cut edges are considered as features and need to be sampled separately, like the sharp and boundary edges, our goal is to minimize the total length of the cut and to snap corresponding cut graph onto the feature skeleton as much as possible. We first describe our current implementation, then explain how recent theoretical results may help us to improve our cutting algorithm.

4.1.1 Our current implementation

In [24], the authors use a simple algorithm to cut the surface into a genus zero surface; after that, they describe a way to improve the cut graph. We implemented a retraction technique similar in spirit to the first part of their method. Note that cutting a mesh along a seaming backbone generates two instances of it in parametric space and one or more instances of the branching nodes encountered in

the cut graph (Figure 6). To guarantee a perfect stitching, we temporarily split the feature, resp. branching samples, into two, resp. their multiplicity, during the error diffusion process and the univariate Lloyd relaxation along the seaming backbones.

Figure 6: *Left: the cut graph of the torus is composed of two pairs of twin backbones $\{a; a'\}$ and $\{b; b'\}$, intentionally separated for better visualization purpose. Middle: closeup on the branching node of the cut graph. Vertices are labeled according to their multiplicity. Right: the torus has been parameterized on a disk-like domain, 4 instances of the branching node have been generated and will later be merged during the stitching.*

4.1.2 Promising approaches

Some recent work in Computational Geometry express interest in the problem of cutting a surface along a set of curves to obtain a topological disk. The latter is called a *polygonal schema*, because the edges on its boundary can be pairwise identified to re-obtain the surface; see Figure 7. Such a polygonal schema is *reduced* if the curves are simple loops, meeting at a common vertex v_0 , and pairwise disjoint except at this vertex. Note that, for our cutting stage, we can simulate that we work on a surface without boundaries: fill its holes, cut this surface (without entering the filled holes) with any cutting algorithm, then re-open the holes.

Figure 7: *From an orientable surface to a polygonal schema: the torus.*

Erickson and Har-Peled [18] have focused on finding the shortest polygonal schema of a surface (possibly non-oriented and with boundary). They have proved that this problem is NP-hard, yet they describe a greedy algorithm that outputs a $O(\log^2 g)$ -approximation of the minimum schema. To our knowledge, no experiment has been done to see whether it yields a visually acceptable result, and it would be worth implementing it (though it is not so easy) for our purposes. Because it is desirable to cut the surface along sharp edges, we advise to simulate that these edges are considered shorter.

In contrast to the NP-hardness of the above problem, Colin de Verdière and Lazarus [13] give a polynomial algorithm which computes the shortest reduced polygonal schema homotopic to a given reduced polygonal schema (computed e.g., with the help of [35]). The algorithm described in [13] consists of iterating *elementary optimization steps*, each of them optimizing the current schema by shortening one loop while maintaining the other ones fixed; stability is reached after a finite number of steps. The loops drawn on the edges of the manifold can partly overlap, yet they are considered as disjoint if we imagine that they are spread apart with a thin space. We believe this algorithm is better suited to our purposes. In particular, in our context, it may be desirable to further generate overlappings of loops on edges, because this decreases the number of distinct cut edges. This can be easily achieved by artificially decreasing the lengths of the edges that contain loops at elementary

steps. It remains to describe how we deal with multiple loops on an edge during the parameterization and sampling processes. If several loops go along a given edge, this creates strips of infinitely small area, and no sample should go inside them. In order to achieve this goal, these strips are given a density of zero in parameter space. Moreover, in this case, a given edge of the original mesh can correspond to an arbitrary large (even) number of boundary edges in parameter space, all these edges being sampled consistently before stitching.

4.2 Parameterization

The goal of surface parameterization is to remove the embedding by flattening the original mesh on a plane. Our goal being to build an isotropic sampling technique, we advocate for using a conformal parameterization [45, 17] with free boundaries since it is known to be both *angle-preserving* and *locally isotropic* (see [37, 15] for more details on the technique). Those two properties are of crucial importance for the meshing and optimization stages since the latter operations are performed only from parameter space. This way every decision on parameter space makes sense with respect to the original mesh, at least for the two preserved properties. Intuitively, locally isotropic and angle-preserving means that a small circle mapped on the surface will be transformed into a circle in parameter space (see Figure 8). Hence a well-shaped triangle in parameter space will not be deformed too much once lifted back into $\mathbb{R}^3$, except for its size. One distortion remains: the area-stretching. This one being isotropic, it can easily be compensated by modifying the density function expressed in parameter space for the optimization stage, as explained in Section 4.4, contrary to a shearing deformation that would be much more difficult to compensate for.

4.3 Meshing

Every sample now lives in parameter space, be it on a corner, on a feature edge or on a face of the original mesh. We perform a 2D constrained Delaunay triangulation [19] in parameter space so that every segment joining two consecutive vertices sampled along a feature backbone is added to the list of constrained edges (see Figure 9).

Figure 9: *Constrained Delaunay triangulation.*

4.4 Construction

With a Delaunay triangulation, we can already deduce an initial Voronoi Tessellation [6]. From this, we aim at building a *weighted centroidal* Voronoi tessellation [16] to improve the initial sampling obtained by error diffusion.

Definition A weighted centroidal Voronoi tessellation is a Voronoi diagram such that the associated sites coincide with the center of mass of the corresponding Voronoi cells. In the current state of the algorithm, we know the density function and the position of the samples in parameter space. The goal is to determine

Figure 8: *Conformal parameterization of the David head model over a disk-like domain. The two textures mapped onto the model illustrates the two main properties of conformal parameterization: preservation of angles and local isotropy.*

the final locations of the samples so that they coincide with the centroids of the corresponding Voronoi regions; the centroids have to be evaluated over the density function expressed in *parameter space* rather than on the surface.

Density Function in Parameter Space Similarly in spirit to [2], we compute a stretching factor for every vertex as a weighted averaging over the stretching factors of its incident faces. The weight being equal to the area of every face measured in parameter space, the stretching factor s of a vertex v is equal to: $s(v) = \sum area(f_i) / \sum area_{uv}(f_i)$, where the sums iterate over its incident faces and any subscript uv denotes a quantity measured in parameter space. The corrected density expressed on a vertex in parameter space is now the product of its stretching factor by its density expressed on the surface: $d_{uv}(v) = s(v) \cdot d(v)$. Such a density function is linearly interpolated over the vertices in parameter space and used for weighting the centroidal Voronoi tessellation.

4.4.1 2D Lloyd relaxation

One way to build such a tessellation is to use Lloyd’s relaxation method [16]. The Lloyd algorithm is a deterministic, fixed point iteration [40]. Given a density function and an initial set of n sites, it consists of the following steps:

1. build the Voronoi tessellation corresponding to the n sites;
2. compute the centroids of the n Voronoi regions with respect to the density function expressed in parameter space, and move the n sites onto their respective centroid;
3. repeat steps 1 and 2 until satisfactory convergence is achieved.

Computing the centroid (*i.e.*, center of mass) of every cell is the most delicate part of the algorithm. Let us recall that the original mesh is now parameterized and that a density function lives in parameter space. The new mesh has been created using constrained Delaunay triangulation over the new samples obtained by error diffusion. The Voronoi regions are computed over the new samples and we aim at computing their centroid evaluated over the density function sitting over the *original* mesh. This requires to compute the intersection between each cell and the set of overlapped triangles, then compute the center of mass of the set of resulting polygons on which the corrected density function is piecewise linear (see Figure 10). Figure 11 illustrates several iterations of the Lloyd algorithm over the *mushroom* model uniformly sampled with 3k vertices. The variable density compensates for any area distortion due to parameterization, and the final tiling tends to make each cell cover the same amount of density. Notice how the main improvement is already achieved after 10 iterations of Lloyd algorithm, the sampling being then polished with additional iterations.

Cell Clipping Clipping the Voronoi cells with constrained edges [39] allows to disconnect two smooth regions separated by a backbone during the computation of the centroid. This leads to a nice quality of the sampling in the vicinity of the features. The latter is obtained through a non-symmetric behavior of the algorithm in the sense that the boundary and the feature backbones influence the surface samples but the inverse is not true. Intuitively, two samples incident in the Voronoi tessellation but separated by a feature are not influencing each other anymore, and the samples closed to a boundary or from a feature backbone are rather repulsed by the constraints (see Figure 12). Indeed, clipping a cell by the set of constraints may remove some regions from the computation of the centroid, making the Lloyd relaxation consistent with respect to the constraints.

4.4.2 1D Lloyd relaxation

The goal is to distribute on each feature backbone a number of samples so that each of them covers the same amount of density. To proceed, we parameterize each backbone independently on a segment without any length distortion, then apply univariate Lloyd relaxation on the density function sitting on every parameterized feature backbone. A special care is taken for seaming backbones, the twin samples being reflected on opposite halfedges to guarantee a perfect stitching during the lifting. For the sake of consistency with the asymmetric influence between feature and surface samples, the univariate relaxation described here is applied first, then the feature samples are not moved any more during the 2D Lloyd relaxation process previously described.

4.5 Lifting and Stitching

The lifting stage restores back the embedding by locating every vertex in its associated triangle in parameter space and computing its barycentric coordinates. We then project back every ordinary sample onto its corresponding triangle in $\mathbb{R}^3$ and every feature

Figure 10: *Left: the Voronoi cell surrounding a sample overlaps the original mesh. Right: closeup on the intersection computed between a Voronoi cell and the original mesh. The new centroid (filled) of the cell is computed as a weighted contribution of the centroids computed on the linearly interpolated polygons.*

Figure 11: *Top row: original mushroom model to resample and its boundary-free parameterization on a disk-like domain. Row 2: initial uniform sampling of 3k vertices generated by direct error diffusion (300 ms). Row 3: first iteration of Lloyd algorithm with a variable density coming from the area distortion due to parameterization (1.2 s). Row 4: 10th iteration of Lloyd algorithm (8 s). Row 5: 100th iteration of Lloyd algorithm (55 s).*

sample onto its corresponding feature edge or corner. For closed or genus > 0 objects, a stitching stage is required. It involves to merge each set of twin vertices with a multiplicity > 1 generated

for meshing in parameter space. This welds the model and wraps up the complete remeshing pipeline.

5 Results

Figure 13 illustrates a uniform remeshing of the Michelangelo David head model. The model considered for remeshing has a complexity of 25k vertices, and the new remeshed model has 50k vertices. After 100 iterations of Lloyd relaxation performed in parameter space, the weighted centroidal Voronoi tessellation provides a satisfactory uniform sampling after lifting. Running times are the following on a PIII 1GHz: 7 s for the mesh parameterization, 400 ms for differential geometry analysis, 850 ms for sampler calibration, 2.8 s for error diffusion, and 26 minutes for polishing the sampling using 100 iterations of Lloyd algorithm. Note that starting from iteration 20 the result is already very similar to the present one. As shown by Figure 11, our technique is in a sense progressive in terms of sampling quality since the Lloyd algorithm progressively improves the repartition of the density function among all Voronoi cells. Figure 14 illustrates three curvature-adapted remeshings of the David head using 30k vertices. The curvature-adapted behavior is obtained by specifying a variable $\gamma > 0$ to the transfer function x^γ applied over the initial density function, related to absolute mean discrete curvature. Increasing γ from 0 (Figure 13) to 2.5 on our examples allows us to vary continuously from uniform to curved-adapted remeshing, while matching the exact desired vertex budget.

Figure 15 illustrates uniform remeshing of the bottomless *fan-disk* model with 10k vertices. The Lloyd algorithm optimizes the sample positions both along backbones (univariate relaxation) and on the surface. The repartition of samples between corners, backbones and smooth parts exactly matches the vertex budget thanks to the sampler calibration described in Section 3.2. A few closeups nearby the features show the consistent behavior due to cell clipping with the constrained edges.

Figure 16 illustrates uniform remeshing of the genus-1 *rotor* model with 10k vertices. The initial mesh is cut and parameterized on a disk-like domain. The cut graph (highlighted) is composed of several seaming backbones on which curve sampling technique is applied, as done for ordinary features. Although the pairing of seaming backbones and samples allow a perfect stitching of the remeshed model, we believe that the seams generated by univariate sampling of the cut graph limit the proposed remeshing technique to the class of low-genus surfaces.

Figure 17, top illustrates uniform remeshing of the *pig* model with 15k vertices. This model has 7 boundaries and generates an

Figure 12: *Left: a Voronoi tessellation in parameter space with a feature skeleton. All the cells are drawn according to the circumcircle property. Computing the centroid without clipping by the constraints makes the sampling inconsistent, while the effect of clipping is to repulse the samples from the boundary or sharp edges, the centroid being computed on the truncated cell. A constrained edge separating two samples thus acts as a barrier [39] annihilating their mutual influence as expected.*

Figure 13: *Weighted centroidal Voronoi tessellation built with 100 iterations of Lloyd relaxation, and final remeshing with 50k vertices*

Figure 14: *Curvature-adapted remeshing using 30k vertices with $\gamma = 1$, $\gamma = 2$ and $\gamma = 2.5$ respectively.*

extreme range of area distortion in parameter space. Notice that the longest boundary has been automatically chosen for parameterization (foot boundary is highlighted). Due to cell clipping using constrained edges, the Lloyd relaxation is guaranteed to correctly handle models with several boundaries. The density function computed in parameter space correctly compensates the area distortion due to flattening. Figure 17, bottom illustrates four curvature-adapted remeshings of the *pig* model with exactly 30k vertices and $\gamma = 3$ set in the transfer function. The only difference between the models lies into the mesh gradation progressively smoothed with respectively 0, 3, 10, and 100 iterations of low-pass filtering over the density function. A few closeups of the ear shows the effect of smoothing the density function over the final mesh gradation. This result illustrates the key concept demonstrated in this paper, *i.e.*, the user can act on some parameters until the density function fits his desire in terms of sampling distribution and gradation on the surface and the features.

6 Conclusions and Future Work

We have presented a new technique for isotropic remeshing of triangle meshes. Our approach consists of specifying a density function over the surface to be repartitioned between a set of samples so that each of them covers the same amount of density on a neighborhood as compact as possible. This task is performed using two algorithms: the first one performs a fast mesh resampling by generalizing the core principle of error diffusion over the original mesh triangles. The second tool is a Lloyd relaxation used to build a weighted centroidal Voronoi tessellation in a conformal parameter space. The latter, commonly used for optimal clustering [16], tends to optimally distribute the density function between all the samples both on the feature skeleton and on smooth parts of the models. The weak part of the algorithm is the need for cutting closed or genus > 0 surfaces. Such a cut generates a seaming backbone on which *curve sampling* is applied instead of *surface sampling*, which is not consistent if the cut graph does not coincide with a set of fea-

Figure 15: Uniform remeshing of the fandisk model with 10k vertices. Top row: Initial sampling computed by direct error diffusion, and sampling obtained after 20 iterations of Lloyd relaxation. Bottom row: closeup over the Voronoi tessellation after Lloyd convergence, global view of the remeshed model, and several closeups nearby the features.

Figure 16: Uniform remeshing of the genus-1 rotor model with 10k vertices. From left to the right: cut graph (colored) and feature skeleton (light grey); mesh parameterization on a disk-like domain, and associated parameterized feature skeleton (boundary backbones with identical colors are associated pairwise). Far right, top to bottom: original, and remeshed model. Two closeups over the cut graph illustrate the univariate sampling performed along seaming backbones.

ture edges. Such an approximation becomes unacceptable for high genus surfaces, which would require a complex cut graph.

As future work we plan to remove the cutting and parameterization stages and compute a *geodesic* version of the weighted centroidal Voronoi tessellation to handle surfaces of high genus. As a separate issue we would like to investigate the conditions over the density function to guarantee the convergence of Lloyd’s algorithm. Being able to remesh volume meshes with error diffusion and centroidal Voronoi diagram is also of some interest.

Acknowledgements The work reported here was supported in part by the ECG project of the EU No IST-2000-26473, and by an ARC TéléGéo grant from INRIA. Special thanks to Mariette Yvinec, Sylvain Pion and Lutz Kettner for help with CGAL. The David’s head mesh is courtesy of Marc Levoy and the Digital Michelangelo Project, and Kai Hormann provided us with the Darmstadt model.

References

- [1] ALLIEZ, P., AND DESBRUN, M. Valence-driven connectivity encoding of 3d meshes. In *Eurographics 2001 Conference Proceedings* (2001), pp. 480–489.
- [2] ALLIEZ, P., MEYER, M., AND DESBRUN, M. Interactive geometry remeshing. *ACM Transactions on Graphics* 21(3) (2002), pp. 347–354. SIGGRAPH conference proceedings.
- [3] BERN, M., AND EPPSTEIN, D. Mesh Generation and Optimal Triangulation. *Computing in Euclidean Geometry*, edited by D.-Z. Du and F. K. Hwang (1992), pp. 23–90.

- [4] BERN, M., AND PLASSMANN, P. Mesh generation. In *Chapter 6 in Handbook of Comp. Geometry*, J. Sack and J. Urrutia, Eds. Elsevier Science, 1999.
- [5] BOISSONNAT, J.-D., AND CAZALS, F. Coarse-to-fine Surface Simplification with Geometric Guarantees. In *Eurographics proceedings* (2001), pp.490–499.
- [6] BOISSONNAT, J.-D., AND YVINEC, M. *Algorithmic Geometry*. Cambridge University Press, UK, 1998.
- [7] BOROUCHAKI, H. Geometric Surface Mesh. In *Int. Conf. on Integrated and Manufacturing in Mechanical Engineering* (1998), pp.343–350.
- [8] BOROUCHAKI, H., GEORGE, P. L., HECHT, F., LAUG, P., AND SALTET, E. Delaunay Mesh Generation Governed by Metric Specifications. *Finite Elements in Analysis and Design* 25 (1997), pp.61–83.
- [9] BOROUCHAKI, H., HECHT, F., AND FREY, P. J. Mesh Gradation Control. In *Proceedings of 6th International Meshing Roundtable* (1997), pp.131–141.
- [10] BOTSCH, M., AND KOBELT, L. Resampling Feature and Blend Regions in Polygonal Meshes for Surface Anti-Aliasing. In *Eurographics proceedings* (sep 2001), pp.402–410.
- [11] BOTSCH, M., RÖSSL, C., AND KOBELT, L. Feature Sensitive Sampling for Interactive Remeshing. In *Vision, Modeling and Visualization proceedings* (2000), pp.129–136.
- [12] CHEW, L. P. Guaranteed-quality mesh generation for curved surfaces. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.* (1993), pp.274–280.
- [13] COLIN DE VERDIÈRE, É., AND LAZARUS, F. Optimal system of loops on an orientable surface. In *Proceedings of 43rd Annu. IEEE Sympos. Foundations of Computer Science* (Nov. 2002).
- [14] DE COUGNY, H. L., AND SHEPHARD, M. S. Surface Meshing Using Vertex Insertion. In *Proceedings of 5th International Meshing Roundtable* (1996), pp.243–256.
- [15] DESBRUN, M., MEYER, M., AND ALLIEZ, P. Intrinsic parameterizations of surface meshes. In *Proceedings of Eurographics* (2002), pp.209–218.
- [16] DU, Q., FABER, V., AND GUNZBURGER, M. Centroidal Voronoi Tessellations: Applications and Algorithms. *SIAM review* 41(4) (1999), pp.637–676.
- [17] ECK, M., DEROSE, T., DUCHAMP, T., HOPPE, H., LOUNSBURY, M., AND STUETZLE, W. Multiresolution Analysis of Arbitrary Meshes. In *Proceedings of SIGGRAPH* (1995), pp.173–182.
- [18] ERICKSON, J., AND HAR-PELED, S. Optimally cutting a surface into a disk. In *Proceedings of the 18th Annual ACM Symposium on Computational Geometry* (2002), pp.244–253.
- [19] FABRI, A., GIEZEMAN, G.-J., KETTNER, L., SCHIRRA, S., AND SCHÖNHERR, S. On the design of CGAL a computational geometry algorithms library. *Softw. – Pract. Exp.* 30, 11 (2000), pp.1167–1202.
- [20] FLOYD, R., AND STEINBERG, L. An adaptive algorithm for spatial grey scale. In *Proc. of the Society for Information Display* (1976), vol. 17, pp.75–77.
- [21] FREY, P. J. About Surface Remeshing. In *Proceedings of the 9th Int. Meshing Roundtable* (2000), pp.123–136.
- [22] GARLAND, M., AND HECKBERT, P. Simplifying Surfaces with Color and Texture using Quadric Error Metrics. In *IEEE Visualization Conference Proceedings* (1998), pp.263–269.
- [23] GEORGE, P. L., AND BOROUCHAKI, H., Eds. *Delaunay Triangulation and Meshing Application to Finite Elements*. Hermes, Paris, 1998.
- [24] GU, X., GORTLER, S., AND HOPPE, H. Geometry Images. In *Proceedings of SIGGRAPH* (2002), pp.355–361.
- [25] GUSKOV, I., SWELDENS, W., AND SCHRÖDER, P. Multiresolution Signal Processing for Meshes. In *Proceedings of SIGGRAPH* (1999), pp.325–334.
- [26] GUSKOV, I., VIDIMCE, K., SWELDENS, W., AND SCHRÖDER, P. Normal Meshes. In *Proceedings of SIGGRAPH* (2000), pp.95–102.
- [27] GUSKOV, I., AND WOOD, Z. Topological noise removal. In *Proceedings of Graphics Interface* (2001), pp.19–26.
- [28] HOPPE, H. New quadric metric for simplifying meshes with appearance attributes. In *IEEE Visualization Proceedings* (1999), pp.59–66.
- [29] HORMANN, K., LABSIK, U., AND GREINER, G. Remeshing Triangulated Surfaces with Optimal Parameterizations. *Computer-Aided Design* 33 (2001), pp.779–788.
- [30] HUBELI, A., AND GROSS, M. Multiresolution feature extraction from unstructured meshes. In *IEEE Visualization conf. proceedings* (2001), pp.287–294.
- [31] JODOIN, P.-M., AND ÖSTRÖMUKHOV, V. Error-Diffusion with Blue-Noise Properties for Midtones. *SPIE 4663* (2002), pp.293–301.
- [32] KETTNER, L. Using generic programming for designing a data structure for polyhedral surfaces. *Comput. Geom. Theory Appl.* 13 (1999), pp.65–90.
- [33] KHODAKOVSKY, A., SCHRÖDER, P., AND SWELDENS, W. Progressive Geometry Compression. In *Proceedings of SIGGRAPH* (2000), pp.271–278.
- [34] KOBELT, L., VORSATZ, J., LABSIK, U., AND SEIDEL, H.-P. A Shrink Wrapping Approach to Remeshing Polygonal Surfaces. *Computer Graphics Forum, Eurographics '99 issue 18* (1999), pp.119–130.
- [35] LAZARUS, F., POCCHIOLA, M., VEGTER, G., AND VERRONST, A. Computing a Canonical Polygonal Schema of an Orientable Triangulated Surface. In *Proceedings of 17th Annu. ACM Sympos. Comput. Geom.* (2001), pp.80–89.
- [36] LEE, A. W. F., SWELDENS, W., SCHRÖDER, P., COWSAR, L., AND DOBKIN, D. MAPS: Multiresolution Adaptive Parameterization of Surfaces. In *Proceedings of SIGGRAPH* (1998), pp.95–104.
- [37] LÉVY, B., PETITJEAN, S., RAY, N., AND MAILLOT, J. Least Squares Conformal Maps for Automatic Texture Atlas Generation. In *Proceedings of SIGGRAPH* (2002), pp.362–371.
- [38] LINDSTROM, P., AND TURK, G. Fast and Memory Efficient Polygonal Simplification. In *IEEE Visualization Proceedings* (1998), pp.279–286.
- [39] LINGAS, A. Voronoi diagrams with barriers and the shortest diagonal problem. *Inform. Process. Lett.* 32 (1989), pp.191–198.
- [40] LLOYD, S. Least square quantization in PCM. *IEEE Trans. Inform. Theory* 28 (1982), pp.129–137.
- [41] MEISTERS, G. Polygons have ears. *American Mathematical Monthly* (jun/july 1975), pp.648–651.
- [42] MEYER, M., DESBRUN, M., SCHRÖDER, P., AND BARR, A. H. Discrete Differential-Geometry Operators for Triangulated 2-Manifolds, 2002. VisMath.
- [43] ÖSTRÖMUKHOV, V. A Simple and Efficient Error-Diffusion Algorithm. In *Proceedings of SIGGRAPH* (2001), pp.567–572.
- [44] P. VÉRON, J.-C. L. Static Polyhedron Simplification using Error Measurements. *Computer-Aided Design* 29(4) (1997), pp.287–298.
- [45] PINKALL, U., AND POLTHIER, K. Computing Discrete Minimal Surfaces and Conjugates. *Experimental Mathematics* 2(1) (1993), pp.15–36.
- [46] RASSINEUX, A., VILLON, P., SAVIGNAT, J.-M., AND STAB, O. Surface Remeshing by Local Hermite Diffuse Interpolation. *International Journal for numerical methods in Engineering* 49 (2000), pp.31–49.
- [47] ROCCHINI, C., AND CIGNONI, P. Generating random points in a tetrahedron. *Journal of Graphics tools, A.K. Peters Ltd* 5(4) (2000), pp.9–12.
- [48] SANDER, P., GORTLER, S., SNYDER, J., AND HOPPE, H. Signal-specialized parameterization. In *Eurographics Workshop on Rendering 2002* (2002).
- [49] SWELDENS, W., AND SCHRÖDER, P., Eds. *Digital Geometry Processing*. Course Notes. ACM SIGGRAPH, 2001.
- [50] TRISTANO, J. R., OWEN, S. J., AND CANANN, S. A. Advancing Front Surface Mesh Generation in Parametric Space Using a Riemannian Surface Definition. In *Proceedings of 7th International Meshing Roundtable* (1998), pp.429–445.
- [51] TURK, G. *Generating random points in triangles*. Graphics Gems (A.S. Glassner). Academic Press, 1990.
- [52] TURK, G. Re-Tiling Polygonal Surfaces. In *Proceedings of SIGGRAPH* (1992), pp.55–64.
- [53] ULICHNEY, R. A review of halftoning techniques. *SPIE* (2000), pp.378–391.
- [54] VORSATZ, J., RÖSSL, C., KOBELT, L., AND SEIDEL, H.-P. Feature Sensitive Remeshing. In *Eurographics proceedings* (sep 2001), pp.393–401.
- [55] WATANABE, K., AND BELYAEV, A. Detection of Salient Curvature Features on Polygonal Surfaces. *Computer Graphics Forum, Eurographics 2001* 20(3) (2001), pp.385–392.
- [56] WU, J., AND KOBELT, L. Fast Mesh Decimation by Multiple-Choice Techniques. In *Vision, Modeling and Visualization Proceedings* (2002). To appear.

Figure 17: Top: uniform remeshing of the pig model with 15k vertices. The model has been parameterized from the longest boundary and undergoes a high range of area distortion due to flattening. Bottom: curvature-adapted remeshing with 30k vertices. The gradation is more and more smoothed by specifying 0, 3, 10 and 100 iterations for low-pass filtering the density function.