Good Random Multi-Triangulation of Surfaces

Esdras Medeiros and Marcelo Siqueira

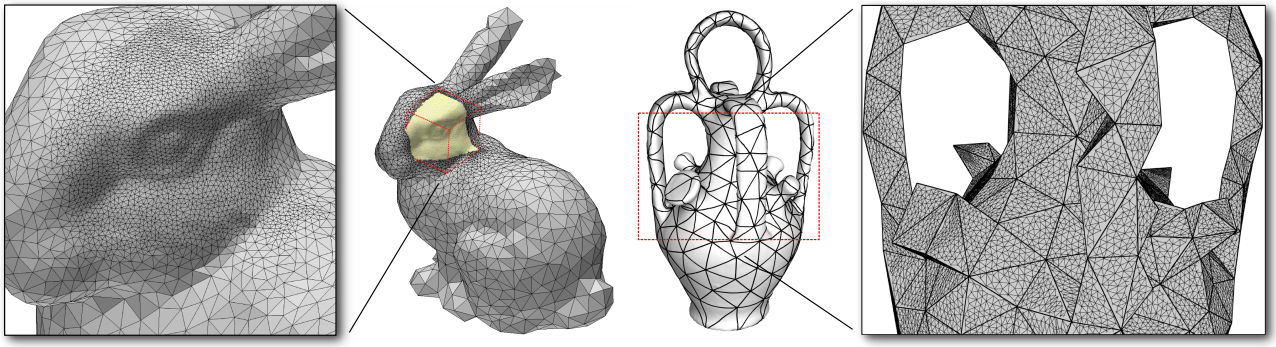

Fig. 1: HPDS-MT applications. (Left) Magic lens on a 3D object with good triangles. (Right) Base domain generation with good triangles.

Abstract—We introduce the Hierarchical Poisson Disk Sampling Multi-Triangulation (HPDS-MT) of surfaces, a novel structure that combines the power of multi-triangulation (MT) with the benefits of Hierarchical Poisson Disk Sampling (HPDS). MT is a general framework for representing surfaces through variable resolution triangle meshes, while HPDS is a well-spaced random distribution with blue noise characteristics. The distinguishing feature of the HPDS-MT is its ability to extract adaptive meshes whose triangles are guaranteed to have good shape quality. The key idea behind the HPDS-MT is a preprocessed hierarchy of points, which is used in the construction of a MT via incremental simplification. In addition to proving theoretical properties on the shape quality of the triangle meshes extracted by the HPDS-MT, we provide an implementation that computes the HPDS-MT with high accuracy. Our results confirm the theoretical guarantees and outperform similar methods. We also prove that the Hausdorff distance between the original surface and any (extracted) adaptive mesh is bounded by the sampling distribution of the radii of Poisson-disks over the surface. Finally, we illustrate the advantages of the HPDS-MT in some typical problems of geometry processing.

Index Terms—Multiresolution, Poisson Disk Sampling, Triangulation.

1 Introduction

Multiresolution modeling is an active research area that encompasses a number of techniques for handling complex geometric data at varying levels of detail. These techniques have been successfully applied to problems in graphics and geometric modeling. In particular, they have been employed to separate an input 3D polygonal surface – viewed as a surface signal – into high and low frequencies for shape editing [6], [28], and to adaptively construct surfaces that satisfy a view-dependent criteria error [25], [31].

The Multi-Triangulation (MT) [40] is a hierarchical structure for representing triangular decompositions of 2-dimensional domains. More specifically, MT is a collection of fragments of triangulations arranged as nodes in a directed acyclic graph (DAG). MT was devised to be a unifying framework for the entire class of multiresolution models. By defining a resolution filter and a focus condition [13] and by traversing the DAG, one can extract a mesh of the represented model at any specified level of detail.

MTs constructed via incremental simplification differ by the user-defined criteria. Typical criteria include geometric distortion (e.g., Hausdorff distance) and visual appearance (e.g., texture). Here, we restrict our attention to criteria that improve the shape quality of the triangles. The aspect-ratio and the min/max angle are two of the mostly used metrics to measure the shape quality of a triangle. For both metrics the optimal reference is the equilateral triangle. Hierarchical Poisson Disk Sampling (HPDS) is a method for computing PDS which enables us to more rapidly generate adaptive blue noise distributions over planar [33] and surface domains [34].

Well-spaced distributions of sample points are characterized by the minimum spacing between samples (conflicting distance) and by the maximum distance between a domain point and a closest sample point (covering distance). Poisson disk sampling (PDS) [29] is an example of well-spaced random distribution with blue noise characteristics. In particular, its conflicting and covering distances are equal. PDS yields provably good quality 2D and 3D Delaunay triangulations [36], which in turn lead to better conditioned stiffness matrices in the finite element method (FEM) [42].

- *Esdras Medeiros is with Department of Mathematics, Federal University of Ceará, Fortaleza, CE, Brazil. E-mail: esdras@mat.ufc.br.*
- *Marcelo Siqueira is with DIMAP, Federal University of Rio Grande do Norte, Natal, RN, Brazil. E-mail: mfsiqueira@dimap.ufrn.br.*

Manuscript submitted for peer review to IEEE Transactions on Visualization and Computer Graphics on May 10, 2017.

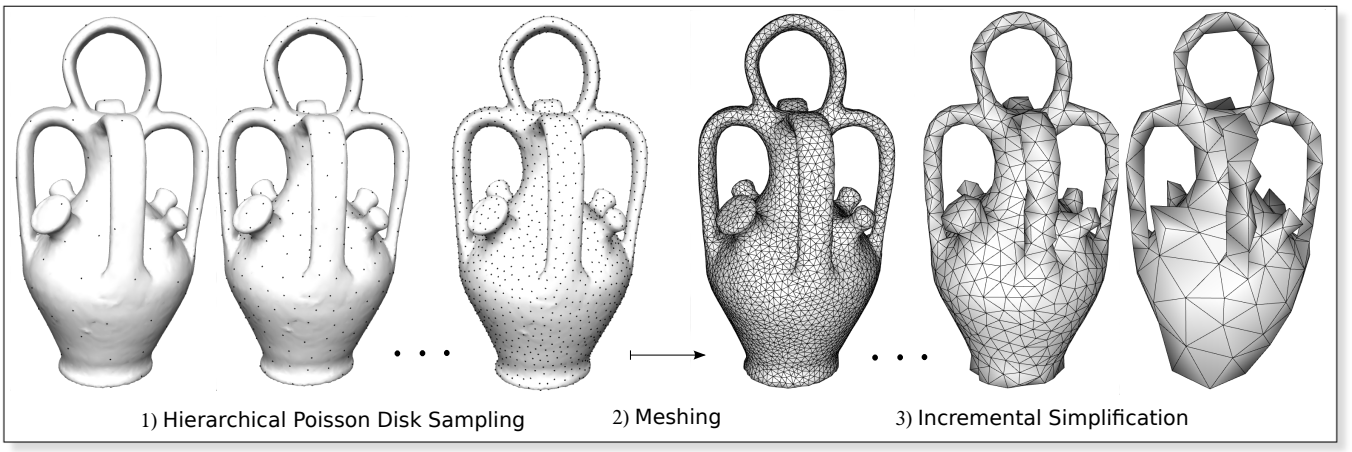

Fig. 2: Overview of the algorithm for constructing a Multi-Triangulation with good triangles, the HPDS-MT.

The random nature of the samples also plays an important role in preventing directional bias in the edge directions of a triangulation. This isotropic behavior can be highly desirable in FEM-based simulations involving fractures [17], for instance. Hierarchical Poisson Disk Sampling (HPDS) is a method for computing PDS which enables us to more rapidly generate adaptive blue noise distributions over planar [33] and surface domains [34].

In this paper, we introduce the *Hierarchical Poisson Disk Sampling MT (HPDS-MT)* of surfaces, which is a multiresolution structure that combines the power of MT with the benefits of HPDS. The pipeline for constructing the MT first generates a HPDS on the input surface, and then builds a MT of the surface. The latter is created by incremental simplification, following the hierarchy of samples in a natural way. The main feature of the HPDS-MT is its ability to extract *adaptive meshes* with *provably good quality shape triangles* from the hierarchy of samples (see Figure 1).

To construct the HPDS-MT we decimate incrementally a sequence of vertices from an inital mesh. At each step, we remove the triangles incident on the candidate vertex and re-triangulate the hole. While the HPDS continues to preserve the well-spaced property across the samples, a Localized Restricted Delaunay Triangulation (LRDT) [45] of the current set of vertices is maintained to re-triangulate the holes and to obtain triangles with provably good quality shape (i.e., triangles whose minimum angles are bounded from below). The LRDT is a more general version of the Restricted Delaunay Triangulation (RDT) [19]. In principle, one can adopt the RDT instead. However, our experiments indicated that the RDT decreased the expressive power[1] of the MT. Indeed, the RDT of sparse vertex sets is more likely to change the topology of the input surface, causing the simplication process to end too soon.

Our approach for constructing the HPDS-MT is top-down, and it starts with a triangle mesh topologically equivalent to the input surface. We ensure that the topology of the intermediate triangulations do not change during the entire simplification process. One could think of building the HPDS-MT in a bottom-up fashion, i.e., by incrementally adding new samples to the surface. To that end, one could rely on an algorithm for constructing provably good quality Delaunay triangulations of surfaces (e.g., [4], [9], [10]). However, since the coarse mesh may not be topologically equivalent to the input surface, the refinement process must handle topological changes. Unfortunately, the commonly used operators for constructing a MT cannot achieve changes of topological type, as they are topology-preserving operators.

1. The *expressive power* of a MT is the number of different triangulations that one can build with the DAG.

The remainder of this paper is organized as follows. Section 2 reviews previous work and states our contributions. Section 3 introduces basic notions. Section 4 defines the HPDS-MT. Section 5 provides an analysis of the shape quality of the triangles. Section 6 describes a computational method for constructing the HPDS-MT. Section 7 offers an experimental evaluation of our computational method as well as a comparison with other methods. Section 8 demonstrates the application of HPDS-MT to typical problems in computer graphics and geometric modeling. Section 9 discusses current limitations of our method, and Section 10 offers final remarks.

2 Related Work

Multiresolution representation of triangle meshes is currently considered a mature research area from the theoretical point of view. The reason is that any such representation can in principle be reduced to a MT. In addition, algorithms for efficiently constructing MTs and for extracting meshes with varying levels of detail have long been known by the graphics community [40], [41], [44]. However, to the best of our knowledge, none of the previous algorithms for constructing MTs can produce triangle meshes with theoretical guarantees on the quality shape of the triangles. Our proposed algorithm consists of three steps, namely (see Figure 2): 1) generation of a HPDS over the underlying space of the input mesh, 2) construction of a dense remeshed surface from the samples, and 3) construction of a DAG by incremental simplification. In what follows we review prior research efforts related to the three steps.

Computing PDS on surfaces. There is an extensive literature on PDS (see [29] and [46] for more details). The method more related to ours is the hierarchical dart-throwing algorithm [33] for planar domains. This algorithm was extended to surface domains in [34], allowing for real-time rendering of visually pleasant stippling effects. Behind the key idea of the latter algorithm is a preprocessed hierarchy of points, which enables us to more rapidly compute adaptive blue noise. Our proposed method also makes use of a hierarchy of points but the purpose is to produce a sorted sequence of points for decimating a dense remeshed mesh.

Surface remeshing using PDS. Delaunay triangulations of PDS's over planar (resp. volumetric) domains enjoy desirable properties, such as bounded radius-edge ratio, leading to good quality meshes [17], [35], [36]. The restricted Delaunay triangulation (RDT) [19] can be used to remesh adaptive PDS on surfaces [45], [47]. Experimental results from [45], [47] indicate that the triangles of the resulting meshes have good quality shape. However, the proposed algorithm is not accompanied by any theoretical properties to ensure bounds on values obtained from any shape quality metric. Moreover, although the algorithm is capable of generating adaptive meshes, the authors do not offer any hints on how to create multiresolution models.

Incremental simplification yielding good shape triangles. Incremental simplification methods that generate triangles with good shape quality are mostly included in the context of surface parametrization and remeshing. Global mesh parametrization methods are used for semi-regular remeshing by means of a coarsely triangulated parameter domain [37]. In this approach a (coarse) base mesh is produced as the output of a greedy algorithm that incrementally simplifies the input (dense) mesh. During the generation of the base mesh, optimization criteria are heuristically employed to improve the shape and/or restrict the size of the base mesh triangles [26], [38]. The heuristic nature of the simplification process makes it difficult to prove any bounds on the shape quality and size of the resulting triangles (even if this is possible).

2.1 Contributions

Our contributions are two-fold. First, we introduce the HPDS-MT, a multi-triangulation specialization over surfaces based on hierarchical Poisson-disk samplings. Second, we describe an implementation for constructing an approximate HPDS-MT of a given triangle mesh of a surface in 3D. The main feature of the HPDS-MT is its ability to extract adaptive meshes, whose triangles under a certain sampling condition have provably good quality shape, from the hierarchy. In addition to the definition of the HPDS-MT and its implementation, this paper also offers (a) a formal definition of a special case of well-spaced sampling over surfaces: the (α, β)-*Poisson disk sampling*; (b) A formal definition of *well-placed Poisson-disk sampling*, a sufficient sampling condition over the distribution of samples that guarantees good remeshing; (c) a proof that the Hausdorff distance between the original surface and any adaptive mesh is limited by the sampling distribution of the Poisson-disks radii over the surface; and (d) an experimental comparison with other methods for generating simplified meshes, which shows the practical benefits of the theoretical guarantees of our method.

3 PRELIMINARIES

Throughout the remainder of this paper we denote a connected, compact, orientable, 2-dimensional manifold without boundary in $\mathbb{R}^3$ by $\mathbf{S}$. Unless explicitly stated otherwise, we refer to $\mathbf{S}$ as *surface*. We shall also use the terms *localized* and *restricted* to refer to any connected component of a subset of $\mathbf{S}$, and to the intersection of a 3-ball in $\mathbb{R}^3$ with $\mathbf{S}$, respectively. In particular, these terms are used to extend the notions of Voronoi Diagrams and Delaunay Triangulations [39] in $\mathbb{R}^3$ to surfaces in $\mathbb{R}^3$. We refer the reader to [19] for more details on the notions presented in Sections 3.1-3.4.

3.1 Topological Concepts

A *simplex* σ^p of dimension p, or p-simplex (for short), is the convex hull of $p+1$ affinely independent points $\{v_i\}_{i=0}^{p} \subset \mathbb{R}^3$. Simplices of dimension 0, 1 and 2 are called *vertices*, *edges* and *triangles*, respectively. A *face* of σ is the simplex of some subset of vertices of σ. A *simplicial complex* $\mathcal{K}$ is a finite set of simplices together with all their faces such that if $\sigma, \tau \in \mathcal{K}$, then either σ and τ share a common face λ in $\mathcal{K}$ or $\sigma \cap \tau = \emptyset$. The *dimension* of a simplicial complex is the highest dimension of its simplices. The *underlying space* $|\mathcal{L}| \subset \mathbb{R}^3$ of a subset of simplices $\mathcal{L}$ of $\mathcal{K}$, corresponds to the union of the interior of the simplices of $\mathcal{L}$. The *star* $St(\sigma, \mathcal{K})$ *of* σ *in* $\mathcal{K}$ is the collection of all simplices of $\mathcal{K}$ of which σ is a face including σ itself. A subset $\mathcal{L} \subset \mathcal{K}$ is said to be a *subcomplex* of $\mathcal{K}$ iff $\mathcal{L}$ is a simplicial complex. The *closure* $\mathrm{Cl}(\mathcal{L}, \mathcal{K})$ of any subset $\mathcal{L} \subset \mathcal{K}$ is the smallest subcomplex of $\mathcal{K}$ that contains $\mathcal{L}$, i.e., the set of all τ in $\mathcal{K}$ such that τ is face of some simplex σ in $\mathcal{L}$.

A *covering* $\mathcal{C}$ of $\mathbf{S}$ is a collection $\{\mathcal{C}_i\}_{i=1}^{n}$ of subsets of $\mathbf{S}$ such that $\mathbf{S} = \cup_{i=1}^{n} \mathcal{C}_i$. The *nerve* $Nrv(\mathcal{C})$ of a covering $\mathcal{C}$ is the collection of subsets $\mathcal{F}$ of $\mathcal{C}$ such that the intersection of all elements of $\mathcal{F}$ is nonempty, i.e., $Nrv(\mathcal{C}) = \{\mathcal{F} \in 2^{\mathcal{C}} \mid \bigcap \mathcal{F} \neq \emptyset\}$. The *dual complex* of a covering $\mathcal{C}$ is a simplicial complex $\mathcal{K}$ together with a bijection ψ between $\mathcal{C}$ and the vertices of $\mathcal{K}$, so that $\mathcal{F}$ is in the nerve of $\mathcal{C}$ iff the simplex spanned by $\psi(\mathcal{F})$ is in $\mathcal{K}$.

A subset $A \subset \mathbf{S}$ is *regular closed* if $\overline{int(A)} = A$. We say that a covering $\{\mathcal{C}_i\}_{i=1}^{n}$ of $\mathbf{S}$ is a *regular partition* if $\mathcal{C}_i$ is connected, regular closed and $\mathcal{C}_i \cap \mathcal{C}_j \subset \partial\mathcal{C}_i \cap \partial\mathcal{C}_j$, for $i \neq j$. A regular partition is a *cell decomposition* iff it satisfies the *closed-ball property*: (B1) each $\mathcal{C}_i$ is a disk, (B2) $\mathcal{C}_i \cap \mathcal{C}_j$ is an interval, for $i \neq j$ and $\mathcal{C}_i \cap \mathcal{C}_j \neq \emptyset$, and (B3) $\mathcal{C}_i \cap \mathcal{C}_j \cap \mathcal{C}_k$ is a point, for i, j, k pairwise distinct and $\mathcal{C}_i \cap \mathcal{C}_j \cap \mathcal{C}_k \neq \emptyset$.

A *triangulation* $\mathcal{T}$ is a 2-dimensional simplicial complex in which $|St(\mathbf{v}, \mathcal{T})|$ is either an open disk (i.e., it satisfies the *topological disk property* [14]) or a half plane, for every vertex v. The *boundary* $\partial\mathcal{T}$ of a triangulation $\mathcal{T}$ is the set of all edges $e \in \mathcal{T}$ together with its vertices such that e is a face of at most one triangle. The *interior* $int(\mathcal{T})$ of $\mathcal{T}$ is the complement set $\mathcal{T} \setminus \partial\mathcal{T}$. A surface $\mathbf{S}$ is said to be *triangulable* iff there exists a triangulation $\mathcal{T}$, whose vertices belong to $\mathbf{S}$, such that $|\mathcal{T}|$ is homeomorphic to $\mathbf{S}$.

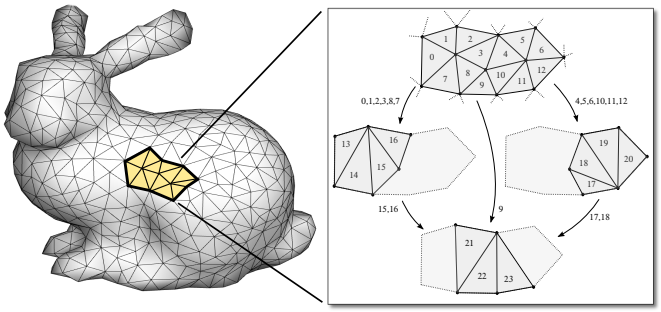

Fig. 3: Partial view of a lattice representation.

3.2 Multi-triangulation

We adapt the definition of MT, which was originally formulated for planar domains by Puppo in [40], to surfaces. Let $\mathcal{T}$ be a triangulation of $\mathbf{S}$. *A minimally compatible local modification* is a substitution of a triangulation $\widehat{\mathcal{T}_j} \subset \mathcal{T}$ by a new triangulation $\mathcal{T}_j$, such that

(LM1) $\partial\widehat{\mathcal{T}_j} = \partial\mathcal{T}_j$, and
(LM2) $\widehat{\mathcal{T}_j} - \mathcal{T}_j = int(\widehat{\mathcal{T}_j})$.

The above local modification yields a new triangulation, which is denoted by $\mathcal{T} \oplus \mathcal{T}_j$. Triangulation $\widehat{\mathcal{T}_j}$ is referred to as the *floor* of $\mathcal{T}_j$. Compatibility is enforced by condition (LM1), which ensures that $\partial\widehat{\mathcal{T}_j}$ is the same as $\partial\mathcal{T}_j$. In turn, minimality is enforced by condition (LM2), which ensures that $\widehat{\mathcal{T}_j}$ and $\mathcal{T}_j$ have no triangle in common. Given a triangulation $\mathcal{T}_0$, a sequence $\{\mathcal{T}_s\}_{s=1}^m$ of triangulations is said to be a *minimally compatible sequence of local modifications* iff $\mathcal{T}_j$ is minimally compatible with the result of $\bigoplus_{s=0}^{j-1}\mathcal{T}_s$. A minimally compatible sequence of local modifications $\{\mathcal{T}_s\}_{s=1}^m$ is *non-redundant* iff $\mathbf{t} \in \mathcal{T}_i$ implies $\mathbf{t} \notin \mathcal{T}_j$, for $i \neq j$, i.e., there are no duplicated triangles. Let $\{\mathcal{T}_s\}_{s=1}^m$ be a non-redundant, minimally compatible sequence of local modifications. A *multi-triangulation* (MT) with respect to $\{\mathcal{T}_s\}_{s=1}^m$ is a direct acyclic graph (DAG) $G = (V, E)$ such that

(MT1) there exists a bijection $f : V \rightarrow \{\mathcal{T}_1, \ldots, \mathcal{T}_m\}$ between the set V of nodes of G and the set $\{\mathcal{T}_s\}_{s=1}^m$ of all triangulations.
(MT2) for every $i \neq j$, there exists an arc in E from $f^{-1}(\mathcal{T}_i)$ to $f^{-1}(\mathcal{T}_j)$ iff there is a triangle in $\mathcal{T}_i$ that belongs to the floor of $\mathcal{T}_j$.

The *root* of the DAG is the initial triangulation, $\mathcal{T}_0$. A *lattice representation* of the MT is obtained as follows (see Figure 3). A *drain* node is added to the DAG. This node corresponds to the *combination* $\bigoplus\mathcal{T}_s$ of the whole sequence. An arc is added from a node $f^{-1}(\mathcal{T}_j)$ to the drain iff there exists a triangle in $\mathcal{T}_j$ that belongs to the drain. Finally, each arc is labeled by the set of triangles shared by the triangulations associated with its source and end nodes. A *cut* of a DAG is a set of arcs whose removal disconnects the DAG. A *front* $\mathcal{E}$ is a cut that contains exactly one arc for each path from the root to the drain. The following proposition establishes a relationship between fronts and triangulations extracted from a MT:

Proposition 1 ([40]). *Let $\mathcal{M}$ be a MT, and let $\mathcal{T}_\mathcal{E}$ be the subsequence of $\{\mathcal{T}_s\}_{s=0}^m$ formed by nodes in the upper part of a given front $\mathcal{E}$. Then $\mathcal{T}_\mathcal{E}$ is a non-redundant minimally compatible sequence, and its combination $\bigoplus\mathcal{T}_\mathcal{E}$ is formed by all triangles labeling the arcs of $\mathcal{E}$.*

3.3 Localized Restricted Voronoi Diagram

Let $\mathbf{p}$ be a point of $\mathbf{S}$. For any given scalar r, with $r > 0$, let $\mathcal{B}_r(p)$ denote the intersection of $\mathbf{S}$ with an open 3-ball of radius r centered at $\mathbf{p}$, i.e., $\mathcal{B}_r(p) = \{\mathbf{q} \in \mathbf{S} \mid ||\mathbf{p} - \mathbf{q}|| < r\}$, where $||\cdot||$ denotes the Euclidian norm. The *localized restricted r-neighborhood* $\mathcal{N}_r(p)$ of $\mathbf{p}$ is the connected component of $\mathcal{B}_r(p)$ that contains $\mathbf{p}$. For convenience, we refer to $\mathcal{N}_r$(p) as *r-neighborhood*. Let $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^n$ be a set of samples from $\mathbf{S}$. Then, the *localized restricted Voronoi diagram* (LRVD) is a regular partition of $\mathbf{S}$ by *Voronoi regions* $\{\mathcal{V}_i\}_{i=1}^n$ such that

(V1) $\mathbf{x}_i \in \mathcal{V}_i$, for $1 \leq i \leq n$, and
(V2) for $i \neq j$, if $\mathbf{q} \in \partial\mathcal{V}_i \cap \partial\mathcal{V}_j$ then $||\mathbf{x}_i - \mathbf{q}|| = ||\mathbf{x}_j - \mathbf{q}||$.

Whenever the parent sampling must be made explicit to avoid ambiguities, the Voronoi region $\mathcal{V}_i$ shall be denoted by the pair $(\mathcal{V}_i, \mathbf{X})$.

Inflating balls construction. Depending on the geometry and topology of the surface as well as on the positions of the samples, there can be more than one LRVD for the same surface and set of samples [45]. We now show how to precisely build a particular LRVD for a given $\mathbf{S}$ and a given $\mathbf{X}$ using front propagation. For each sample $\mathbf{x}_i$ in $\mathbf{X}$, set a self inflating 3-ball such that the radius increases at a constant rate. Let $\mathcal{N}_{r_t}^i$ be the r_t-neighborhood of $\mathbf{x}_i$ at instant time t. Initialize each region by setting $\mathcal{V}_i \leftarrow \mathcal{N}_{r_0}^i$. Here, r_0 is an initial radius that keeps all r_0-neighborhoods pairwise disjoint. As every region front propagates along the surface, each point of the surface is assigned to the first region that reaches it. Two or three region fronts can reach the same point at the same time, in which case the point belongs to the intersecting boundaries of the regions. This construction process ends when every point of the surface is assigned to at least one region. It might be the case that four or more Voronoi regions meet at a common point $\mathbf{q}$ of $\mathbf{S}$. This means that the points generating the four or more regions have the same distance from $\mathbf{q}$: they lie on a common sphere whose center is $\mathbf{q}$. To avoid ambiguity in defining the LRVD, we shall assume that the points are in general position, which is equivalent to assuming the absence of degeneracies. From now on, whenever we refer to the LRVD of a set of points sampled from $\mathbf{S}$, we mean the LRVD constructed by the above process and we will denote it by LRVD($\mathbf{X}$).

3.4 Localized Restricted Delaunay Triangulation

Given a set of sample points $\mathbf{X} = \{\mathbf{x_i}\}_{i=1}^n$ from $\mathbf{S}$, the dual complex of the corresponding LRVD($\mathbf{X}$) is a complex called *localized restricted Delaunay complex*, denoted by LRDC($\mathbf{X}$), for which $\psi(\mathcal{V}_i) = \mathbf{x}_i$ (see Section 3.1). The subcomplex formed by all triangles in LRDC($\mathbf{X}$), along with their edges and vertices, is called *localized restricted Delaunay triangulation* and denoted by LRDT($\mathbf{X}$).

By the *nerve theorem* [19] applied to regular partitions, it follows that if LRVD($\mathbf{X}$) is a cell decomposition of a

triangulable surface $\mathbf{S}$, then LRDC is a triangulation $\mathcal{T}$ such that $|\mathcal{T}|$ and $\mathbf{S}$ are topologically equivalent; that is, the nerve theorem gives us a sufficient condition, the *closed-ball property*, under which the LRDC is the LRDT.

If every vertex of the LRDC satisfies the topological disk property then LRDC is the LRDT. Moreover, if $\mathcal{X}(\mathbf{S}) = \mathcal{X}(\text{LRDT}(\mathbf{X}))$, where $\mathcal{X}$ is the Euler characteristic function, then $\mathbf{S}$ and $\text{LRDT}(\mathbf{X})$ are topologically equivalent by the *Classification Theorem for Compact Surfaces* [20].

3.4.1 LRDT versus RDT

Given a set of sample points $\mathbf{X} = \{\mathbf{x_i}\}_{i=1}^{n}$, the *Delaunay triangulation of* $\mathbf{X}$ *restricted to* $\mathbf{S}$, denoted by $\text{RDT}_{|\mathbf{S}}(\mathbf{X})$ is the subcomplex of the 3D Delaunay triangulation of $\mathbf{X}$ consisting of the triangles – and their edges and vertices – whose dual Voronoi edges intersect $\mathbf{S}$. If every ball circumscribing a triangle $\mathbf{t}$ of $\text{LRDT}(\mathbf{X})$ with center in $\mathbf{S}$ is empty, then $\mathbf{t}$ is also a triangle in $\text{RDT}(\mathbf{X})$. Also, if $\text{RDT}(\mathbf{X})$ is a triangulation of $\mathbf{S}$) then $\text{LRDT}(\mathbf{X}) = \text{RDT}_{|\mathbf{S}}(\mathbf{X})$ [45]. For low resolution samplings and highly-varying curvature surfaces the RDT is more unlikely to be a triangulation of $\mathbf{S}$. On the other hand, the construction of the LRDT is more resilient to low resolution samplings, and it can correctly represent the topology of the surface where the RDT fails [45]. The inset figure illustrates such a scenario in the bidimensional case.

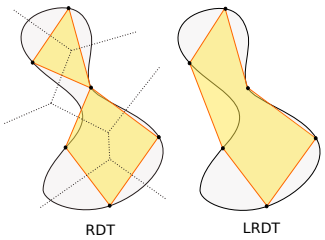

4 Defining the HPDS-MT

We first define (H)PDS using the Euclidian distance metric. Then, we describe an incremental simplification algorithm that receives a HPDS over a surface as input and outputs the corresponding HPDS-MT.

4.1 Localized Restricted Poisson Disk Sampling

Given a surface $\mathbf{S}$ and two fixed scalars $\beta \geq \alpha > 0$, the set of sampled points $\mathcal{P}_{\alpha\beta} = \{\mathbf{p}_i\}_{i=1}^{n}$ of $\mathbf{S}$ is called a *localized restricted* (α, β)*-Poisson disk sampling* (LR-PDS for short) iff:

(P1) $\{\mathcal{N}_\beta^i\}_{i=1}^{n}$ is a covering of $\mathbf{S}$.
(P2) Each $\mathcal{N}_\alpha^i$ does not contain any other sample of $\mathcal{P} - \{\mathbf{p_i}\}$.

In particular, whenever $\beta = \alpha = r$ we obtain a sampling $\mathcal{P}_r$ called *localized restricted Poisson disk sampling*. The following lemma states an important property of the LR-VD defined from a given LR-PDS:

Lemma 1. *Let* $\mathcal{P}_{\alpha\beta} = \{\mathbf{p}_i\}_{i=1}^{n}$ *be a LR-PDS generated over* $\mathbf{S}$*, and let* $\{\mathcal{V}_i\}_{i=1}^{n}$ *be its LRVD. If* $\mathbf{q}$ *is a point lying on the boundary of a Voronoi region* $\mathcal{V}_k$*, then* $||\mathbf{p}_k - \mathbf{q}|| < \beta$.

Proof. See supplemental meterial. □

We say that two sampled points $\mathbf{p}_i$, $\mathbf{p}_j$ are *adjacent* if their Voronoi regions, $\mathcal{V}_i$ and $\mathcal{V}_j$, intersect, i.e., $\mathcal{V}_i \cap \mathcal{V}_j \neq \emptyset$. For a LR-PDS $\mathcal{P}_{\alpha\beta}$, a pair of adjacent points is called *misplaced* iff their Euclidian distance is less than α. We say that $\mathcal{P}_{\alpha\beta}$ is *well-placed* iff $\mathcal{P}_{\alpha\beta}$ contains no misplaced pair.

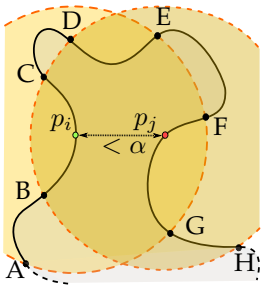

Fig. 4: Misplaced pair.

In Figure 4 we show an 1-dimensional example of misplaced pair. We have $\mathcal{N}_\beta^i = \text{arc}(\mathbf{AE})$, $\mathcal{N}_\beta^j = \text{arc}(\mathbf{DH})$ and $\mathcal{N}_\beta^i \cap \mathcal{N}_\beta^j = \text{arc}(\mathbf{DE})$. The geometric configuration of a misplaced pair is more likely to occur when the size of the sampling radius cannot properly capture the local topology of the surface. In Section 5 we show that the well-placed property together with Lemma 1 play an important

4.2 Hierarchical LR-PDS

Let $\mathcal{H} = \{\mathcal{P}_u\}_{u=1}^{h}$ be a family of samplings over $\mathbf{S}$. We say that $\mathcal{H}$ is a hierarchy iff $\mathcal{P}_1 \subset \mathcal{P}_2 ... \subset \mathcal{P}_m$. A hierarchy $\mathcal{H}$ is said to be a *hierarchical localized restricted Poisson disk sampling* (HLR-PDS) iff each *level* $\mathcal{P}_u$ is a LR-PDS with parameters α_u and β_u. Moreover, a hierarchy $\mathcal{H}$ is said to be *graded* iff every two consecutive levels differ by exactly one point. The next proposition states that the difference between the LRDCs of the sample sets of two consecutive levels of a hierarchy is equivalent to locally changing one LRDC:

Proposition 2. *Let* $\mathcal{P}_{w-1}$ *and* $\mathcal{P}_w = \mathcal{P}_{w-1} \cup \{\mathbf{p}_w\}$ *be two consecutive levels of a graded HLR-PDS and, denote their respective LR-DCs by* $\mathcal{D}_{w-1}$ *and* $\mathcal{D}_w$*. Then, we have* $\mathcal{D}_w - D_{w-1} = St(\mathbf{p}_w, \mathcal{D}_w)$.

Proof. See supplemental material. □

A *local Delaunay simplification* of $\mathcal{D}_w$ through $St(\mathbf{p}_w, \mathcal{D}_w)$ is an operation that replaces the simplices of $St(\mathbf{p}_w, \mathcal{D}_w)$ with the ones of $(\mathcal{D}_{w-1} - \mathcal{D}_w)$. A particular case of Proposition 2 takes place when $|St(\mathbf{p}_w, \mathcal{D}_w)|$ is homeomorphic to a disk. More specifically, the local modification corresponds to a replacement of a set of triangles by another set of triangles that satisfies conditions (LM1) and (LM2) (see Section 3.2). In Figure 5 we show an example of local Delaunay simplification. This observation motivated us to devise the simple MT construction described in Section 4.3.

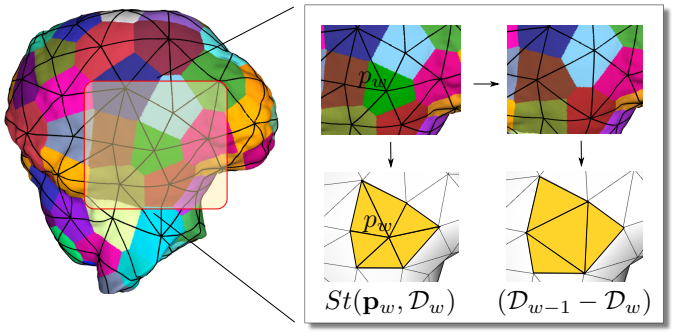

Fig. 5: Local Delaunay Modification.

function $\{\mathcal{T}_j\} \leftarrow$ HPDS-MT$(M, \mathcal{H})$

//M: Triangulation of a surface $\mathbf{S}$
//$\mathcal{H}$: HLR-PDS over $\mathbf{S}$

$\mathcal{T}_0 \leftarrow \mathcal{D}_h$;
$\mathcal{Q} \leftarrow \mathcal{P}_h$;
$w \leftarrow h$;
$j \leftarrow 1$;
while $w \geq 2$ **do**
  $\text{cond}_1 \leftarrow [St(\mathbf{p}_w, \mathcal{D}_w) = St(\mathbf{p}_w, \bigoplus_{s=0}^{j-1} \mathcal{T}_s)]$;
  if *cond*$_1$ **then**
    $\text{cond}_2 \leftarrow true$;
    for *each* $\mathbf{p}_l$ *adjacent to* $\mathbf{p}_w$ **do**
      if *not* TOPODISKPROPERTY$(\mathbf{p}_l)$ **then**
        $\text{cond}_2 \leftarrow$ **false**;
      end if;
    end for;
    if *cond*$_2$ **then**
      $\mathcal{T}_j \leftarrow \text{Cl}\,(\mathcal{D}_{w-1} - \mathcal{D}_w)$;
      $\mathcal{Q} \leftarrow \mathcal{Q} - \{\mathbf{p}_w\}$;
      $j \leftarrow j + 1$;
    end if;
  end if;
  $w \leftarrow w - 1$;
end while;
return $\{\mathcal{T}_j\}$;

Algorithm 1: Algorithm for HPDS-MT construction.

4.3 HPDS-MT Construction

Let $\mathcal{H} = \{\mathcal{P}_u\}_{u=1}^{h}$ be a graded HLR-PDS over $\mathbf{S}$ such that LRVD$(\mathcal{P}_h)$ is a cell decomposition of $\mathbf{S}$. The *hierarchical Poisson disk sampling multi-triangulation* (HPDS-MT) is the unique MT sequence of modifications $\{\mathcal{T}_s\}_{s=1}^{h'}$, $h' \leq h$, generated by Algorithm 1. The main challenge in constructing the HPDS-MT is that $\mathcal{D}_w$ *is not always a triangulation*. Nevertheless, we can still locally use $\mathcal{D}_w$ to generate the sequence of modifications. The idea is to preserve conflicting vertices and continue the simplification (in Figure 6 we show a configuration where conflicting vertices are not removed). As can be seen in Algorithm 1, for each point $\mathbf{p}_w$ to be decimated we first check whether the combination $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$ and the Delaunay complex $\mathcal{D}_w$ match at the vicinity of $\mathbf{p}_w$. If so, we ensure topology preservation by validating the *topological disk property* for every point $\mathbf{p}_l$ adjacent to $\mathbf{p}_w$ in $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$ (see Section 6.4.4).

5 Quality Measures and Guarantees

5.1 Aspect Ratio

Let the *aspect ratio* of a triangle $\mathbf{t}$ be the dimensionless quantity $Q(\mathbf{t}) = \frac{\sqrt{3}L^2}{4A}$, where L is the longest edge of $\mathbf{t}$ and A is its area [8]. Its smallest possible value is 1, and this value is attained by equilateral triangles only. Furthermore, the skinnier a triangle is, the larger its aspect ratio is. Proposition 3 gives an upper bound on the aspect ratio in terms of the parameters α and β of a LR-PDS:

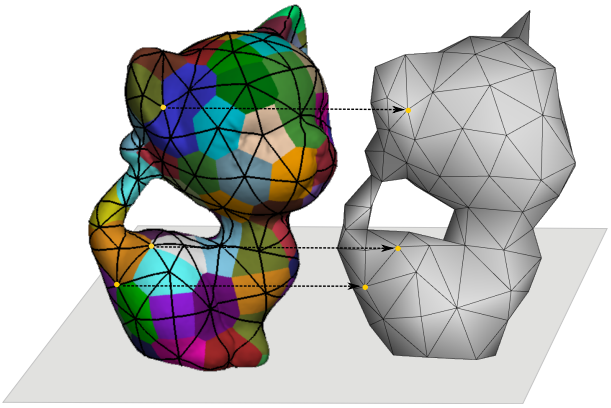

Fig. 6: The LRVD regions in colors and the combination $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$. The edges of the combination $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$ also overlap the LRVD. Notice that the complex $\mathcal{D}_w$ (which is dual graph of the LRVD) and the triangulation $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$ are not isomophic. We highlight some non-removed vertices because of matching and/or topological conflicts.

Proposition 3. *Let $\mathcal{P}_{\alpha\beta} = \{\mathbf{p}_i\}_{i=1}^{n}$ be a well-placed LR-PDS generated over $\mathbf{S}$, and let $\mathcal{D}$ be its LRDC. Then the aspect ratio $Q(\mathbf{t})$ of every triangle $\mathbf{t} \in \mathcal{D}$ is bounded above by $\frac{2\sqrt{3}\beta}{\alpha}\sqrt{1 - (\frac{\alpha}{2\beta})^2}$.*

Proof. See supplemental material. □

The main result of this paper can now be stated as follows:

Proposition 4. *Let $\mathcal{M}$ be an HPDS-MT over $\mathbf{S}$ and $\{\mathcal{P}_u\}$ be its corresponding graded HLR-PDS. For a positive constant c, assume that $\mathcal{P}_w$ is well-placed and that $\frac{\alpha_w}{\beta_w} \geq c$, for every level w. Then the aspect ratio $Q(\mathbf{t})$ of every triangle $\mathbf{t} \in \bigoplus \mathcal{T}_\mathcal{E}$ is bounded above by $\frac{2\sqrt{3}}{c}\sqrt{1 - (\frac{c}{2})^2}$.*

Proof. See supplemental material. □

Proposition 4 generalizes Proposition 3 for combinations $\bigoplus \mathcal{T}_\mathcal{E}$.

5.2 Hausdorff Distance

Given two subsets $\mathcal{A}$ and $\mathcal{B}$ of $\mathbb{R}^3$, the *Hausdorff distance* $\text{H}(\mathcal{A}, \mathcal{B})$ is defined to be the supremum infimum distance between $\mathcal{A}$ and $\mathcal{B}$:

$$\text{H}(\mathcal{A}, \mathcal{B}) = \sup_{\mathbf{a} \in \mathcal{A}} \inf_{\mathbf{b} \in \mathcal{B}} \|\mathbf{a} - \mathbf{b}\| .$$

The Hausdorff distance, $\text{H}(\mathbf{S}, |\mathcal{T}|)$, from a smooth surface $\mathbf{S}$ to the underlying space $|\mathcal{T}|$ of triangulation $\mathcal{T}$ of $\mathbf{S}$ is commonly used for surface approximation analysis. As stated in the next proposition, the approximation error defined by the Hausdorff distance between $\mathbf{S}$ and the underlying space of the LRDT is bounded by the disk radius β.

Proposition 5. *Let $\mathbf{S}$ be a triangulable surface and $\mathcal{P}_{\alpha\beta} = \{\mathbf{p}_i\}_{i=1}^{n}$ a LR-PDS generated over $\mathbf{S}$ such that its LRVD is a cell decomposition. Let $\mathcal{T}$ be the LRDT of $\mathcal{P}_{\alpha\beta}$. Then $H(\mathbf{S}, |\mathcal{T}|) \leq \beta$.*

Proof. See supplemental material. □

Proposition 5 can be generalized for combinations $\bigoplus \mathcal{T}_\mathcal{E}$:

Proposition 6. *Let $\mathcal{M}$ be an HPDS-MT over $\mathbf{S}$ and $\{\mathcal{P}_u\}$ be its corresponding graded HLR-PDS. Let z be the maximum index such that $\mathcal{P}_z \subset \bigoplus \mathcal{T}_\mathcal{E}$. Then $H(\mathbf{S}, |\bigoplus \mathcal{T}_\mathcal{E}|) \leq \beta_z$.*

Proof. See supplemental material. □

6 Implementation Details

The pipeline of the HPDS-MT construction is as follows: *generate the HLR-PDS → grade the HLR-PDS → incrementally simplify*. The graded HLR-PDS $\mathcal{H} = \{\mathcal{P}_u\}_{u=1}^h$ over an input surface $\mathbf{S}$ must be such that $\text{LRVD}(\mathcal{P}_h)$ is a cell decomposition of $\mathbf{S}$. This is a requirement for the correctness of Algorithm 1. From now on, we assume that $\mathbf{S}$ is the underlying space, $|M|$, of a given triangle mesh M.

6.1 Generating the HLR-PDS

Let M be a triangle mesh whose underlying space $|M|$ is a surface in $\mathbb{R}^3$. Let $r_1 > r_2 > ... > r_h$ be a sequence of decreasing radii. Our ultimate goal is *to build a hierarchy of samplings $\{\mathcal{P}_w\}_{w=1}^h$ such that each level $\mathcal{P}_w$ is a (α, β)-PDS of radius $\alpha_w = \beta_w = r_w$*. To that end, we build the sampling hierarchy using the algorithm from [34] combined with the maximal Poisson disk sampling algorithm from [21]. The latter is used to obtain a maximal PDS at every level of the hierarchy. More specifically, for each level w, we cover the points in

$$|M| - \bigcup_{s \in \mathcal{P}_{(w-1)}} \mathcal{N}_{r_w}(s)\,,$$

which is the complement set of the union of all localized restricted neighborhoods of radius r_w and center at the samples at level $(w-1)$.

6.2 Initial Mesh

Algorithm 1 requires that $\text{LRVD}(\mathcal{P}_h)$ be a cell decomposition of $\mathbf{S}$, i.e., that $\text{LRDT}(\mathcal{P}_h)$ be a triangulation whose underlying space is homeomorphic to $|M|$. A straightfoward way to fulfill this requirement is to establish a criteria for choosing the value of parameter r_h and some properties of the input mesh M so that $|\text{RDT}_{\big||M|}(\mathcal{P}_h)|$ is homeomorphic to $|M|$. Indeed, this condition implies that $\text{LRDT}(\mathcal{P}_h) = \text{RDT}_{\big||M|}(\mathcal{P}_h)$, as we showed in Section 3.4.1. Below, we describe an algorithm that simultaneously chooses r_h and ensures the topological equivalence of $|\text{RDT}_{\big||M|}(\mathcal{P}_h)|$ and $|M|$.

A set $\mathcal{E}$ of points on a surface $\mathbf{S}$ is an *ϵ-sampling* iff for any given point $\mathbf{x} \in \mathbf{S}$, there exists a point $\mathbf{p} \in \mathcal{E}$ such that $\|\mathbf{x} - \mathbf{p}\| < \epsilon \cdot \mathbf{lfs}(\mathbf{p})$, where ϵ is a positive number and $\mathbf{lfs}$ corresponds to the *local feature size function*. For sufficiently smooth surfaces, the ϵ-sampling theorem from [3] tells us that $|\text{RDT}_{|\mathbf{S}}(\mathcal{E})|$ is a surface homeomorphic to $\mathbf{S}$, for every choice of $\epsilon \in \mathbb{R}$ such that $\epsilon \leq 0.3$. If $|M|$ were a sufficiently smooth surface, then there would be a small value of r_h such that $|\text{RDT}_{\big||M|}(\mathcal{P}_h)|$ is homeomorphic to $|M|$. Indeed, since $\mathcal{P}_h$ is a *maximal* Poisson disk sampling over $|M|$, the ϵ-sampling theorem ensures that our claim holds for any choice of r_h, with $r_h < \inf_{\mathbf{p} \in |M|} \mathbf{lfs}(\mathbf{p})$. However, $|M|$ is piecewise-linear, and hence the hypothesis of the ϵ-sampling theorem does not hold for $|M|$.

Fortunately, if certain conditions on the proximity of the vertices of M and on its dihedral angles are satisfied (see [5], [15]), then it is still possible to find a finite set R of points from $|M|$ such that $|\text{RDT}_{\big||M|}(\mathcal{P}_h \cup R)|$ is a surface homeomorphic to $|M|$ if $|\text{RDT}_{\big||M|}(\mathcal{P}_h)|$ is not. Here, we assume that the proximity and angle conditions on M derived in [15] hold. To obtain $\mathcal{P}_h$, we first choose an initial value for r_h. We follow [1] and compute the poles [2] of vertices of the input mesh M. The distance from a point in $|M|$ to its nearest pole is then used as an approximation to $\mathbf{lfs}$ at the point. Finally, we let $r_h := \inf_{\mathbf{p} \in M^0} \mathbf{lfs}(\mathbf{p})$, where M^0 is the 0-skeleton of M (i.e., its set of vertices), and then build $\{\mathcal{P}_w\}_{w=1}^h$ as described in Section 6.1. If $|\text{RDT}_{\big||M|}(\mathcal{P}_h)|$ is homeomorphic to $|M|$, we are done. Otherwise, we compute the set R above.

To obtain R, we implemented the routine $\text{TopoRecov}(G, Q)$ from [15]. This routine receives a triangle mesh G and a point set $Q \subset \mathbb{R}^3$ as input. In our case, we let $G := M$ and $Q := \mathcal{P}_h$. Under the aforementioned proximity and angle conditions on M [15], the output is guaranteed to be a RDT of a point set from $|M|$, namely $\text{RDT}_{\big||M|}(\mathcal{P}_h \cup R)$, such that $R \subset |M|$ and $|\text{RDT}_{\big||M|}(\mathcal{P}_h \cup R)|$ and $|M|$ are homeomorphic. Finally, we let $\mathcal{P}_{h+1} := \mathcal{P}_h \cup R$ and increment $h = h + 1$. As a result, $\mathcal{P}_h$ is a (α, β)-PDS such that α_h is the smallest distance r_h between any two vertices of $\text{RDT}_{\big||M|}(\mathcal{P}_h)$ and $\beta_h = r_{h-1}$.

6.3 Grading the Hierarchy

Given a HLR-PDS $\mathcal{H} = \{\mathcal{P}_u\}_{u=0}^h$, our next goal is *to generate a graded hierarchy $\tilde{\mathcal{H}}$ that contains $\mathcal{H}$*. Graded hierarchy $\tilde{\mathcal{H}}$ enables us to iteratively decimate $\text{RDT}_{\big||M|}(\mathcal{P}_h)$ — using Algorithm 1 — in order to construct the MT. Moreover, we must be able to pick values for parameters α and β, for each level of $\tilde{\mathcal{H}}$, so that the quality bounds from Section 5 hold.

To grade $\mathcal{H} = \{\mathcal{P}_u\}_{u=0}^h$, we properly sort $\Delta_w = \mathcal{P}_w - \mathcal{P}_{w-1} = \{\mathbf{p}_{w_1}, \mathbf{p}_{w_2}, ..., \mathbf{p}_{w_{n_w}}\}$, for every level w. To this end we follow a simple idea from [34]. Given the set $\{\mathcal{N}_{r_w}(\mathbf{p})\}_{\mathbf{p} \in \mathcal{P}_w}$ of r-neighborhoods, increase the radius of each $\mathcal{N}_{r_w}(\mathbf{p})$ continuously. At a certain time, a point $\mathbf{q} \in \Delta_w$ will be reached by some r-neighborhood $\tilde{\mathcal{N}}_{\tilde{r}_w}(\mathbf{p})$, such that $\mathbf{q} \in \partial\tilde{\mathcal{N}}_{\tilde{r}_w}(\mathbf{p})$ and $\tilde{r}_w = ||\mathbf{p} - \mathbf{q}|| > r_w$. Thus, $\mathcal{P}_w - \{\mathbf{q}\}$ becomes a new LR-PDS with radius $\tilde{r}_w$. We repeat this process until every sample $\mathbf{p} \in \Delta_w$ has been reached. We obtain a sorted sequence $\mathbf{p}_{w_{\sigma(1)}}, \mathbf{p}_{w_{\sigma(2)}}, ..., \mathbf{p}_{w_{\sigma(n_i)}}$ when the process ends, where σ is a permutation of $\{1, 2..., n_w\}$ (see [34] for more details). It can be shown that the grading algorithm still preserves the blue noise characteristics of the sampling [34]. In addition, we have the following result:

Proposition 7. *Every new level $\mathcal{P}_{\tilde{w}}$ of $\tilde{\mathcal{H}}$ inserted between $\mathcal{P}_{w-1}$ and $\mathcal{P}_w$ in $\mathcal{H}$ is a LR-PDS with $\alpha_{\tilde{w}} = r_w$ and $\beta_{\tilde{w}} = r_{w-1}$.*

Proof. See supplemental material. □

Proposition 7 enables us to take advantage of the quality bounds stated in Proposition 4 by restricting the sequence

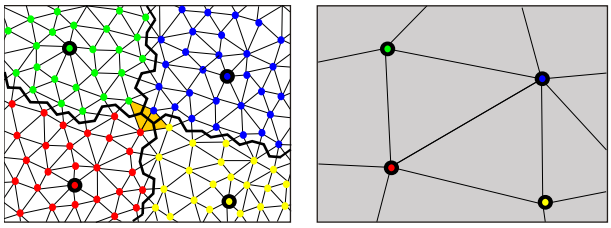

Fig. 7: Vertex Clustering. (Left) Voronoi cells with vertices labeled by colors. There are two highlighted triangles which correspond to two triangles in the Delaunay Complex (Right).

of radii r_i to a geometric progression. More specifically, we have the following result:

Corollary 1. *If r_i is a geometric progression whose common ratio is c, then $\tilde{\mathcal{H}}$ is a graded hierarchy such that $\frac{\alpha_{\tilde{w}}}{\beta_{\tilde{w}}} \geq c$, for every level $\tilde{w}$.*

6.4 Algorithm 1

This section provides some more details of Algorithm 1. Recall that the input of the algorithm is a triangle mesh M and a graded HLR-PDS $\mathcal{H} = \{\mathcal{P}_w\}_{w=1}^{h}$ over $|M|$ such that $\mathrm{LRVD}(\mathcal{P}_h)$ is a cell decomposition of $|M|$. The output is a multi-triangulation $\mathcal{M}$. Two structures guide the construction of $\mathcal{M}$: the localized restricted Delaunay complex (LRDC) $\mathcal{D}_w$ and the triangle mesh $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$.

6.4.1 Computing $\mathcal{D}_w$

To compute $\mathcal{D}_w$, we could follow [45] to obtain the exact LRVD of $|M|$ with respect to $\mathcal{P}_w$. This computation amounts to calculating the exact intersection between $|M|$ and the supporting planes of the two-dimensional faces of each Voronoi cell of the Voronoi diagram of $\mathcal{P}_w$. To speed up the computation of the LRVD and avoid numerical robustness issues, we resorted to a simpler approach that relies on a dense mesh, M_d. This mesh is initially defined by refining $\mathrm{RDT}_{||M|}(\mathcal{P}_h)$, and it can be further refined (if necessary and locally) to ensure the correctness of the computation of $\mathcal{D}_w$. Mesh M_d enables us to compute Voronoi cells using *vertex clustering* and *breadth-first traversals* across its vertices and edges.

Vertex Clustering. Let $\mathbf{X} = \{\mathbf{x}_i\}$ be a subset of the set of vertices of M_d. We follow [7] and approximate each Voronoi cell $\mathcal{V}_i$ of a vertex $\mathbf{x}_i \in \mathbf{X}$ by an equivalence class V_i of vertices in M_d. Define $\mathbf{x}_i$ as the *proxy* of the vertices in region V_i. Two elements $\mathbf{v}_j$ and $\mathbf{v}_k$ in V_i are equivalent iff they have the same proxy $\mathbf{x}_i$. The corresponding LRDC of $\mathbf{X}$ is induced by the distribution of the proxies in each triangle of M_d. Let $\mathbf{x}_i$, $\mathbf{x}_j$ and $\mathbf{x}_k$ be the proxies of the triangle $\mathbf{t} = \mathbf{v}_i\mathbf{v}_j\mathbf{v}_k$. There are three cases: 1) if $\mathbf{x}_i = \mathbf{x}_j = \mathbf{x}_k$, $\mathbf{t}$ is a vertex in $\mathrm{LRDC}(\mathbf{X})$, 2) if $\mathbf{x}_i = \mathbf{x}_j \neq \mathbf{x}_k$ then $\mathbf{t}$ is a segment of $\mathrm{LRDC}(\mathbf{X})$ and 3) if $\mathbf{x}_i \neq \mathbf{x}_j \neq \mathbf{x}_k$ then $\mathbf{t}$ is a triangle of $\mathrm{LRDC}(\mathbf{X})$ (see Figure 7).

Breadth-First Traversal. As done in [45] and [12], we run a region growing algorithm that mimics the inflating balls construction described in Section 3.3. For every proxy $\mathbf{x}_i \in \mathbf{X}$, we insert its adjacent vertices $\mathbf{v}_l$ in M_d into a priority queue. The priority associated with vertex $\mathbf{v}_l$ is set to the distance $||\mathbf{x}_i - \mathbf{v}_l||$. A tag is also associated with $\mathbf{v}_l$ in the priority queue to keep track of proxy $\mathbf{x}_i$ from which the distance was computed. The region growing process is then performed by popping off vertices with least distances until the priority queue is empty. For each vertex $\mathbf{v}$ popped off the queue, we check whether it has been already assigned to a region. If so, nothing is done and we consider the next vertex in the queue. Otherwise, vertex $\mathbf{v}$ is assigned to the region indicated by its associated tag, and the unassigned vertices adjacent to $\mathbf{v}$ are inserted into the queue. These (adjacent) vertices are labeled with the same tag as $\mathbf{v}$.

Cluster Refinement. If M_d is not sufficiently dense, complex $\mathrm{LRDC}(\mathbf{X})$ may be incorrectly computed. So, to obtain a good approximation of the $\mathrm{LRDC}(\mathbf{X})$, we can locally refine M_d. More specifically, for each vertex $\mathbf{x}_i$ of $\mathbf{X}$, let M_i be the connected submesh of M_d, which consists of triangles with at least one vertex whose proxy is $\mathbf{x}_i$. We refine M_i so that we can improve the accuracy of the computation of $\mathrm{LRDC}(\mathbf{X} - \{\mathbf{x}_i\})$. This is done by the following steps:

1) If there exists at least one new triangle $\mathbf{t}$ in $\mathrm{LRDC}(\mathbf{X} - \{\mathbf{x}_i\})$ such that $Q(\mathbf{t})$ does not satisfy the upper bound in Proposition 4, go to step (2); otherwise stop.
2) Reset the cluster proxies of every vertex in V_i to $\mathbf{x}_i$ and apply $\sqrt{3}$-subdivision [27] (without smoothing) to M_i, making M_d a denser mesh.
3) Reassign the proxies of vertices in V_i using breadth-first traversal, as described above, and go back to step (1).

Figure 8 shows the resulting triangulation after one refinement step. The maximum number of refinement levels is given by a parameter n_r. The higher n_r is, the better the precision of the LRVC. In addition, memory consumption increases. Experiments conducted with the data set used in this paper indicate that $n_r = 4$ suffices to reach the upper bound property in Proposition 4.

6.4.2 Initializing $\mathcal{T}_0$ and $\mathcal{D}_h$

The first step of Algorithm 1 is to create M_d. This is done by refining $\mathrm{RDT}_{|M}(\mathcal{P}_h)$ with two consecutive subdivision steps of the $\sqrt{3}$-subdivision scheme [27]. Next, a breadth-first traversal takes place, starting at the vertices of $\mathcal{P}_h$, to assign each new vertex of M_d to one of the proxies in $\mathcal{P}_h$. Finally, an initial vertex clustering $\{V_u\}_{u=1}^{h}$ of M_d is obtained. This clustering is a discrete approximation of $\mathrm{LRVD}(P_h)$ over M_d, and it induces a triangulation $\mathcal{D}_h = \mathrm{LRDT}(P_h)$ over M_d. Triangulation $\mathcal{D}_h$ is isomorphic to $\mathrm{RDT}_{||M|}(\mathcal{P}_h)$, and triangle mesh $\mathcal{T}_0$ is initialized with a new copy of $\mathcal{D}_h$.

6.4.3 Simplifying $\mathcal{D}_w$ to $\mathcal{D}_{w-1}$

For a fixed level w, complex $\mathcal{D}_w$ is represented by the clusters $\{V_u\}_{u=1}^{w}$ which consist of a discrete approximation to $\mathrm{LRVD}(\mathcal{P}_w)$ over M_d. As point $\mathbf{p}_w$ is removed from $\mathcal{D}_w$, the neighboring regions of $\mathcal{V}_w$ are locally updated. Accordingly, since vertex cluster V_w is eliminated from the

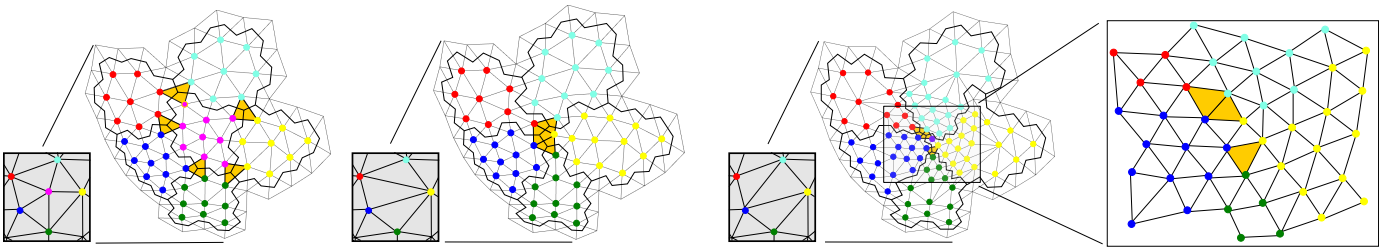

Fig. 8: Visualizing cluster refinement via LRDC and LRDV. (a) The $\mathrm{LRDC}(\mathbf{X})$ before decimation of proxy $\mathbf{x}_i$ (pink). (b) The $\mathrm{LRDC}(\mathbf{X} - \{\mathbf{x}_i\})$ after decimation. In this case a triangle contains one angle below the theoretical bound. (c) Refinement and a new configuration of $\mathrm{LRDC}(\mathbf{X} - \{\mathbf{x}_i\})$.

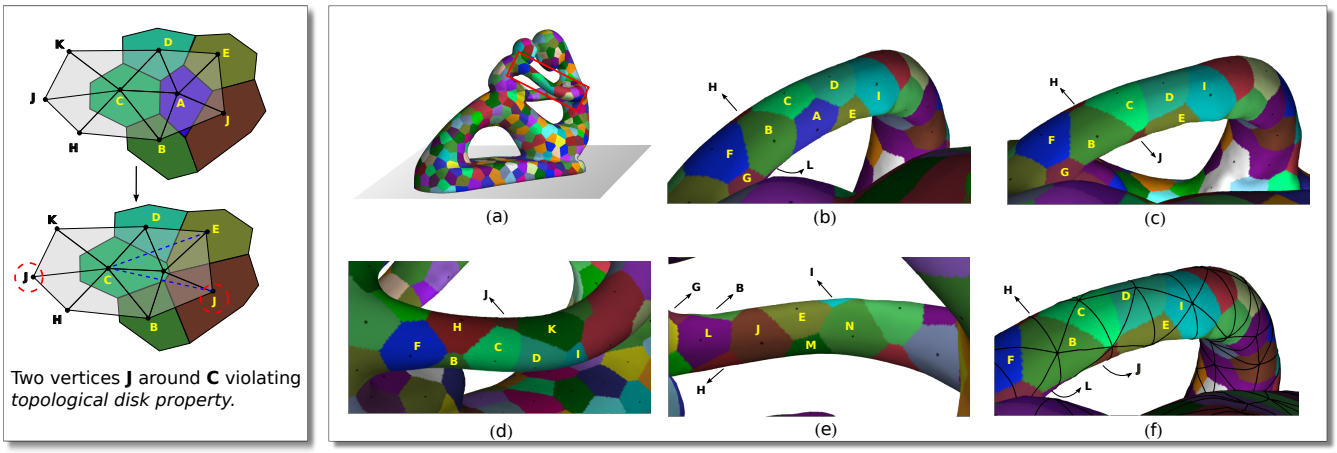

Fig. 9: Checking the topological disk property. (Left-top) Vertex $\mathbf{A} = \mathbf{p}_w$ is a candidate vertex to be removed from triangulation $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$ because it locally matches with the complex $\mathcal{D}_w$ (here represented by the Voronoi diagram). Vertex $\mathbf{C} = \mathbf{p}_l$ is the neighbor of $\mathbf{p}_w$ against which the topological disk property will be verified. (Left-bottom) After removal we have that $(F_l - F_w) \bigcup NF_l$ has four adjacent triangles to the edge $\mathbf{CJ}$. They are: $\mathbf{JCK}$, $\mathbf{CJH}$, $\mathbf{CJE}$ and $\mathbf{JCB}$. Therefore $\mathbf{p}_w$ will not be decimated. (a) A real example where this configuration hold. Here, $\mathbf{J}$ wraps around to be adjacent to $\mathbf{C}$ on both sides after the decimation of $\mathbf{A}$. (b) Delaunay complex $\mathcal{D}_w$ (c) Delaunay complex $\mathcal{D}_{w-1}$. (d)-(e) Two different views of (b). (f) $\mathcal{D}_{w-1}$ and triangulation $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$ overlaped.

LRVD of $\mathcal{P}_w$, each vertex of V_w is reassigned to another proxy in the vicinity of $\mathcal{V}_w$ using the breadth-first traversal described in Section 6.4.1. Notice that this traversal ensures that the updated regions are connected and do not overlap with each other.

6.4.4 Topological Disk Property

We must also verify whether the *topological disk property* (see Section 3.1) is preserved when triangles in $St(\mathbf{p}_w, \bigoplus_{s=0}^{j-1} \mathcal{T}_s)$ are replaced with new triangles in $(\mathcal{D}_{w-1} - \mathcal{D}_w)$. Recall from Section 3.1 that a triangulation $\mathcal{T}$ without boundary is a 2-dimensional simplicial complex in which $|St(\mathbf{v}, \mathcal{T})|$ is an open disk, for every vertex v. This means that each edge incident to v is adjacent to exactly two triangles. Thus, for each vertex $\mathbf{p}_l$ adjacent to $\mathbf{p}_w$ in $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$, let F_l and F_w be the set of triangles incident to $\mathbf{p}_l$ and $\mathbf{p}_w$ in $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$, respectively. Refer to Figure 9 and denote the set of *new* triangles incident to vertex $\mathbf{p}_l$ in $(\mathcal{D}_{w-1} - \mathcal{D}_w)$ by NF_l. Then, the topological disk property is preserved if and only if the following holds: every edge incident to $\mathbf{p}_l$ that belongs to a triangle in $(F_l - F_w) \bigcup NF_l$ is adjacent to exactly two triangles in $(F_l - F_w) \bigcup NF_l$. Line 14 of Algorithm 1 tests this condition for a given $\mathbf{p}_l$ by invoking function TopoDiskProperty(). This function verifies the topological disk property as we just described before. Notice that it is not necessary to check if p_l has only one umbrella because $\bigoplus_{s=0}^{j-1} \mathcal{T}_s$ is a manifold by construction.

7 Experimental Results

In this section, we present several results obtained from the HPDS-MT representation. We run the experiments on an Intel Core i7 Quad 2.67GHz CPU with 8 GB memory and a 64-bit Ubuntu operating system. We use CGAL [43] for 3D Delaunay triangulation. We computed the LR-PDS of the following models: *Chinese dragon*, *Fertility*, *Bunny*, *Botijo*, *Kitten* and *Egea*.

First, we set $r_i = 0.95 \cdot r_{i-1}$, which we experimentally found to yield a good balancing between runtime and blue noise characteristics [34]. By Corollary 1, we conclude that,

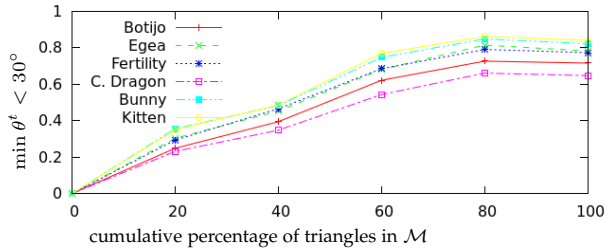

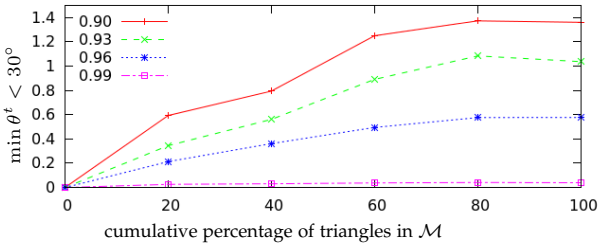

dataset	$\#T_b$	$\#T$	$\min \theta^t$	time(s)
Botijo	494	173k	28.45	53
Egea	16	238k	28.54	85
Fertility	340	357k	28.61	111
C. Dragon	138	393k	28.48	129
Bunny	148	395k	28.51	137
Kitten	170	429k	28.56	146

Botijo	$\#L$	$\min \theta^t$	θ_{min}
$r_i = 0.90 \cdot r_{i-1}$	39	27.05	26.74
$r_i = 0.93 \cdot r_{i-1}$	55	27.96	27.71
$r_i = 0.96 \cdot r_{i-1}$	98	28.93	28.69
$r_i = 0.99 \cdot r_{i-1}$	398	29.71	29.67

Fig. 10: Triangle shape quality measurements confirm the theoretical bounds from Section 5.1. The plots show the percentage of triangles whose minimum angle is below 30° (y-axis). This percentage is relative to the accumulated number of triangles of MT (x-axis), from the coarsest (0) to the finest (100) mesh. Below the plot on the left, we show the smallest angle $\min \theta^t$ among all triangles $t \in \mathcal{M}$. Column $\#T_b$ shows the number of triangles of the coarsest meshes. Column $\#T$ shows the total number of triangles in $\mathcal{M}$. The rightmost column of the left table (time(s)) shows the time Algorithm 1 took to generate the coarsest meshes. Below the plot on the right, we show the number of levels of the HPDS $\mathcal{H}$ (Column $\#L$) for the Botijo model, as well as the minimum angle ($\min \theta^t$) found in the corresponding coarse mesh and the theoretical lower bound (θ_{min}) from Section 5.1.

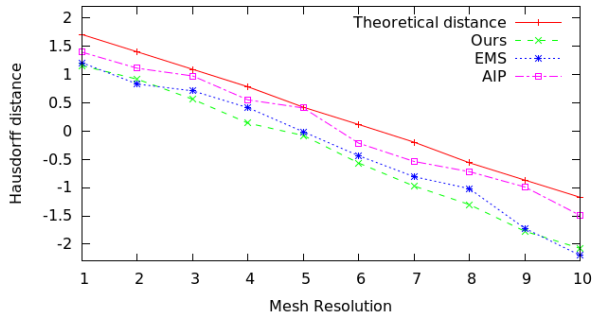

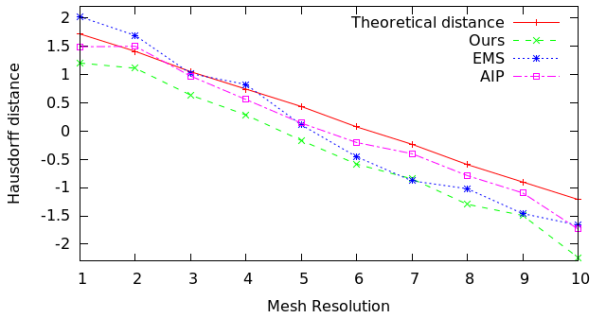

Fig. 11: Absolute Hausdorff distance measurements for the Fertility (Left) and Bunny (Right) models. The x-axis represents meshes sampled at fixed resolutions, varying from coarse (1) to fine (10). The y-axis is in $\log$ scale.

for each level $\tilde{w}$ of a graded LR-PDS, the ratio $\frac{\alpha_{\tilde{w}}}{\beta_{\tilde{w}}}$ is 0.95. This implies that the worst triangle aspect ratio is ~ 2.79, and the theoretical extrema of the angles are $\theta_{max} = 123.28^\circ$ and $\theta_{min} = 28.36^\circ$, respectively. As shown in Figure 10 (left), the shape quality has low variation among the models, which indicates that the algorithm is not strongly influenced by the geometry nor the topology of the input surface, $|M|$. The minimum angles presented in the table (column $\min \theta^t$) are all above the theoretical lower bound, which confirms our estimates. At each simplification step, the running time is dominated by the computation of the Voronoi diagrams. In fact, we need to propagate each neighbor cluster to fill the removed cell. Therefore, the decimation time may not be the same for each vertex. More specifically, the decimation time is proportional to the number of triangles to be removed from the cell. This number increases as the number of removed vertices increases.

In Figure 10 (right), we focus on the Botijo model. As the ratio r_i/r_{i-1} between the levels of the HPDS increases, the amount of triangles with minimum angles below 30° decreases. Again, all triangles have angles above the theoretical lower bound, that is, $\min \theta^t > \theta_{\min}$ which agrees with our expectations.

Figure 11 shows the Hausdorff distance measurements for the Fertility and Bunny models. As expected, the measured distances are also below the theoretical bound for both models.

We selected two algorithms to compare with our algorithm for extracting meshes from the HPDS-MT. The chosen algorithms simplify dense meshes to produce coarser ones. The first one is called Edge Midpoint Simplification (EMS), a simple implementation case of the Lindstrom-Turk strategy (EMS) [32]. This algorithm is a naïve simplification method based on iterative contractions of vertex pairs. The optimal contraction pair is the one that minimizes the the edge length cost. Then, the pair of points contracts to the midpoint of the corresponding shortest edge. The second algorithm is the AIP [38], which is representative of a class of algorithms more concerned with shape quality preservation. The simplification is also based on collapsing the

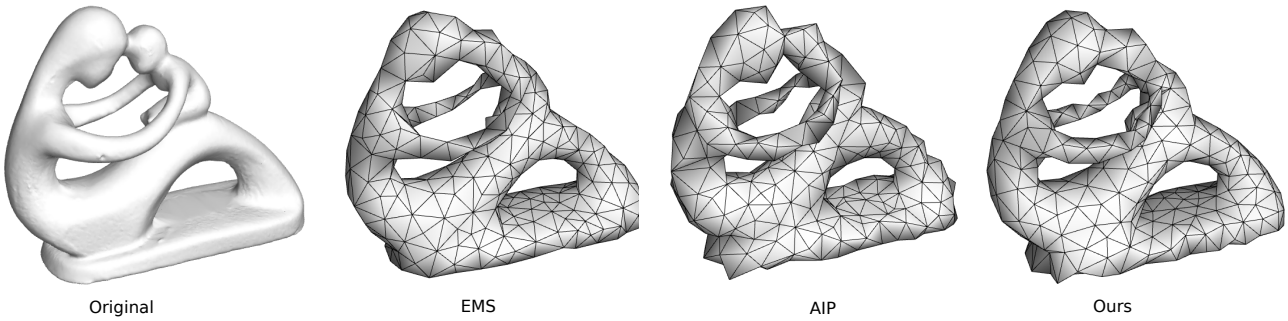

Fig. 12: A visual comparison of results corresponding to the rows (778*) in Table 1.

multi-triangulation			*angle*						*aspect-ratio*			*distance*			*valence*		
			min			max											
dataset	#T	method	min	μ	σ	max	μ	σ	max	μ	σ	max	μ	σ	max	μ	σ
Chinese Dragon	138	EMS	23.09	44.01	7.28	117.45	79.97	12.30	10.97	1.60	0.85	11.24%	1.94%	1.65%	**9**	5.83	**1.23**
		AIP	7.89	43.96	8.61	161.73	80.39	15.02	4.91	1.52	0.49	8.48%	1.47%	1.19%	**9**	5.86	1.51
		Ours	**32.27**	**46.02**	**6.09**	**115.35**	**77.75**	**9.42**	**2.73**	**1.45**	**0.26**	**6.50%**	**1.35%**	**1.08%**	**9**	**5.83**	1.34
	382	EMS	24.21	45.04	7.07	128.36	77.90	11.22	3.59	1.47	0.34	5.75%	0.97%	0.81%	**8**	5.93	1.16
		AIP	19.90	44.79	7.66	132.10	78.62	12.30	4.02	1.49	0.39	4.78%	1.02%	0.78%	9	5.93	1.33
		Ours	**32.63**	**46.36**	**6.38**	**109.63**	**76.97**	**10.12**	**2.45**	**1.42**	**0.27**	**4.05%**	**0.71%**	**0.62%**	9	**5.93**	**1.09**
	720	EMS	23.77	45.04	7.01	124.26	78.27	11.71	3.35	1.48	0.35	3.63%	0.65%	0.53%	**8**	5.96	1.08
		AIP	24.67	46.10	7.18	129.81	76.80	11.45	3.69	1.43	0.35	4.24%	0.77%	0.61%	11	5.96	1.29
		Ours	**32.13**	**47.39**	**6.07**	**106.19**	**75.78**	**9.36**	**2.30**	**1.39**	**0.25**	**3.31%**	**0.50%**	**0.44%**	9	**5.96**	**0.99**
Fertility	340	EMS	6.48	43.03	9.46	154.38	81.29	15.24	10.11	1.65	0.81	5.03%	1.08%	0.79%	11	6.21	1.16
		AIP	20.32	43.93	9.09	136.30	79.01	13.39	4.33	1.54	0.47	**4.94%**	**0.90%**	**0.75%**	11	6.21	1.51
		Ours	**31.47**	**47.20**	**6.25**	**105.59**	**76.69**	**9.67**	**2.28**	**1.41**	**0.26**	5.24%	0.90%	0.79%	**10**	**6.21**	**1.10**
	480	EMS	13.48	44.45	8.12	144.57	79.38	12.97	5.70	1.53	0.48	3.38%	0.82%	0.61%	**10**	6.15	1.13
		AIP	23.30	45.05	7.33	132.14	78.06	11.47	3.90	1.47	0.36	4.02%	0.73%	0.62%	12	6.15	1.32
		Ours	**31.47**	**47.49**	**5.97**	**115.04**	**76.64**	**9.64**	**2.72**	**1.41**	**0.26**	**3.17%**	**0.53%**	**0.47%**	**10**	**6.15**	**1.04**
	778*	EMS	15.91	44.37	7.60	143.20	79.39	12.98	5.30	1.52	0.44	2.90%	0.52%	0.43%	12	6.09	1.12
		AIP	14.97	45.12	7.30	146.21	**77.95**	11.79	5.76	1.48	0.41	3.02%	0.52%	0.44%	12	6.09	1.22
		Ours	**30.89**	**46.01**	**6.56**	**118.14**	78.00	**10.53**	**2.89**	**1.45**	**0.29**	**2.49%**	**0.36%**	**0.34%**	**10**	**6.09**	**0.97**
Bunny	148	EMS	13.99	44.48	8.30	151.27	79.85	14.32	6.76	1.56	0.60	10.80%	1.46%	1.30%	10	5.86	1.27
		AIP	23.94	46.07	7.35	125.04	**75.92**	10.96	3.34	1.41	0.34	8.11%	1.36%	1.13%	12	5.86	1.51
		Ours	**33.58**	**48.29**	**6.14**	**107.20**	76.12	**10.12**	**2.24**	**1.39**	**0.26**	**4.35%**	**1.02%**	**0.76%**	**9**	**5.86**	**1.12**
	388	EMS	15.29	44.55	7.21	146.79	78.87	12.47	5.84	1.40	0.43	7.50%	0.86%	0.82%	10	5.93	1.11
		AIP	18.01	46.07	7.67	137.69	79.03	13.04	4.58	1.51	0.43	4.42%	0.88%	0.73%	10	5.93	1.34
		Ours	**34.24**	**47.96**	**5.97**	**107.93**	**75.27**	**9.45**	**2.38**	**1.37**	**0.24**	**3.31%**	**0.56%**	**0.48%**	**9**	**5.93**	**1.03**
	736	EMS	11.49	44.71	7.24	148.96	79.71	12.58	6.69	1.50	0.43	5.40%	0.55%	0.54%	**9**	5.96	1.09
		AIP	16.71	45.68	7.25	141.94	77.86	11.66	5.10	1.46	0.38	4.46%	0.60%	0.50%	10	5.96	1.14
		Ours	**30.04**	**46.78**	**6.51**	**119.41**	**77.12**	**10.28**	**2.96**	**1.43**	**0.28**	**3.02%**	**0.36%**	**0.35%**	**9**	**5.96**	**0.94**

TABLE 1: A numerical comparison with EMS and AIP. Best results are highlighted in boldface. The Hausdorff distance is relative to the bounding box diagonal of the surface.

minimal cost edges. But, the cost is inversely proportional to the sum of the square of the length of a discrete geodesic path and the areas of two surface regions sharing the path. The path is the result of mapping the edge onto the input surface, while the two regions are the result of mapping the two triangles adjacent to the edge onto the input surface. This heuristic yields triangles with good shape quality for semi-regular surface remeshing [37]. In particular, AIP outperforms previous algorithms (e.g., see [26], [30]) that also aim at preserving the shape quality of the triangles during simplification. We summarize the comparison with the EMS and AIP algorithms in Table 1.

For each dataset, we generated meshes at three different resolutions using our method. Then, we gave EMS and AIP the number of triangles of the resulting meshes as input. We ran EMS and AIP using CGAL and *Meshlab* [11], respectively. Our comparison focuses on statistics of shape metrics (maximum angle, minimum angle and aspect-ratio), mesh approximation (Hausdorff distance and average distance) and mesh connectivity (vertex valence). In all examples, our method presents the best results with respect to shape metrics. The statistics for vertex degree are very similar, except for the variance: our method yields smaller average dispersion. The approximation distance is best attained by our method in most examples. In general, greedy approximation strategies yield good results, but they do not provide bound guarantees on the approximation error. Indeed, as shown in Figure 11 for the Bunny model, both EMS and AIP present deviations above the theoretical distance while our method is more stable and keeps itself below it. One way to enhance distance approximation is to use anisotropic triangles [24] which are elongated along prescribed directions and may produce high aspect ratios. Our method offers the best tradeoff in terms of triangle quality preservation and distance approximation.

8 Applications

An immediate application of the HPDS-MT is the extraction of adaptive meshes. Using the Multi-Tesselation library [13] we can store the history of local modifications during the simplification process. We can then generate static/dynamic extractors by defining a resolution filter condition and a focus condition. Figure 1 (Left) shows an example of an

extractor that produces a magnifying lens effect. Every triangle inside the cube (focus) is refined. Figure 13 (Left) has an extractor where the resolution filter is a scalar field produced by the mean curvature. The resulting mesh is an adaptive triangulation with smooth grading guided by the mean curvature variation.

Following a similar scheme presented in MAPS [30], we construct base domains along with a simple dense-to-coarse mesh parametrization. We recall that in [30], the vertex removal is followed by the re-triangulation of the hole left by the removed vertex and its incident triangles. More specifically, the star of the vertex is flattened with the exponential map z^θ [16], and a constrained Delaunay triangulation is built. In the construction of the HPDS-MT, the re-triangulation that follows the vertex removal is preset by the LR-VD. Therefore, we only need to recompute the barycentric coordinates of the previously removed vertices. To this end, we used the re-scaled polar map [22]:

$$k_\alpha^h(z) = \frac{z^\alpha}{|z|^{(\alpha-1)h}}, \quad \alpha = \frac{6}{n},$$

where n is the degree of the vertex to be removed. For $h = 0$, we obtain the exponential map. By letting $h = 1$ the map preserves the magnitude of the input complex number and re-scales its angular component. As shown in [22], this map alleviates the radial distortion on vertices with degrees higher/lower than 6. This approach may face two issues. First, the domains of the preset re-triangulation and the re-scaled map may not overlap. Second, some triangles may happen to be flipped after star flattening. In the first case we simply flatten the star of the vertex over a unit disk, that is, we apply the map $z(r, \theta) = re^{i\theta}$ with $r = 1$. In the second case we use the solution described in [18], [30]. Surprisingly, we did not encounter the above issues in any of our experiments. The reason has to do with the small length variation of the neighboring edges around the vertex star, a feature inherited from the LR-PDS properties. Figure 1 (Right) and Figure 13 (Right) show examples of base domains and their dense mesh counterparts. The edges of the base domain were mapped onto the underlying space of the corresponding dense mesh, and then displayed as discrete geodesics. Conversely, the triangles of the dense meshes were mapped onto the underlying space of the corresponding base meshes, and the resulting images were also shown in both figures.

9 Limitations and Future Work

Recall from Section 6.1 that we compute the hierarchy of samplings $\{\mathcal{P}_w\}_{w=1}^h$ using the algorithm from [34] combined with the maximal Poisson disk sampling algorithm from [21]. Unfortunately, the algorithm resulting from this combination does not ensure that the samplings are well-placed (see Section 4.1). Nevertheless, we were still able to generate experimental results that respect our theoretical bounds from Section 5. Indeed, the well-placed property is a sufficient but not a necessary condition to obtain good shape triangles. This means that the existence of misplaced pairs does not imply that the aspect ratio of some triangle will violate the upper bound in Proposition 3. For that, there must be a triangle with aspect ratio close to the worst-case configuration (see proof of Proposition 3). However, the nonexistence of misplaced pairs is *applied* in our theoretical results. One way around this shortcoming is to provide a weaker necessary condition, which could replace the well-placed property with minor impact on our theoretical bounds. Otherwise, we need to answer questions that remained open until the conclusion of this work: Is it possible to generate a *well-placed* LR-PDS over $\mathbf{S}$, for any given surface $\mathbf{S}$? If not, for which class of surfaces is that (im)possible? Is there an (efficient) algorithm to decide whether a surface admits a well-placed LR-PDS? Whenever possible, can we give an algorithm that also provides *provable* guarantees on triangle shape quality?

For most examples described throughout this paper, set R was empty (see Section 6.2). Moreover, for those examples in which R was non-empty, the good shape quality of the triangles was preserved even though the conditions of Corollary 1 were not necessarily met (i.e. the the ratio $\frac{\alpha_h}{\beta_h}$ could be less than the value of constant c). The reason is that the point of R are generated as locally farthest points. This selection strategy tends to produce well-spaced samples, which in turn yields triangles with good quality [10], [14], [15], [36]. We are currently working on a sampling conditition that enables us to define, *without computing the additional set* R, a LRDT homeomorphic to surface $|M|$. More specifically, we are addressing the following problem: for which conditions on $|M|$ and $\mathcal{P}_h$ can we guarantee that $\mathrm{LRDT}(\mathcal{P}_h)$ is homeomorphic to $|M|$? Since the LRDT is more resilient than the RDT when it comes to topological changes under mesh simplification (see Section 3.4.1), we would like to obtain a new criteria for ensuring homeomorphism, which is less restrictive than the one in [14], [15] for the RDT.

To maximize the expressive power of the proposed HPDS-MT, one must remove topological noise, which requires the input mesh be preprocessed and modified [23]. However, there are cases in which small topological features (e.g., small tunnels and handles) must be preserved. Yet they may still be very small compared to the volume enclosed by the mesh. In these cases the simplified drain mesh of the MT is likely to consist of a large number of triangles, which prevents further simplication. In turn, the simplification step of our approach does not depend on the sampling strategy. So, in principle, one can use any sampling algorithm for generating a sorted dense cloud of samples. Here, we provided the HPDS which gives good theoretical bounds on the shape quality of the triangles and on the Hausdorff distance to the original surface. These bounds can be tuned by the ratio of the radii between levels of the hierarchy. Moreover, because PDS is a versatile sampling method, one can take advantage of extensions of this method from the literature to improve the results presented in this paper. Finally, our approach is limited to the class of surfaces without boundary. We are also currently working on an extension of the HPDS-MT construction to handle surfaces with boundary.

10 Conclusion

We described a new class of multi-triangulations denoted by HPDS-MT. We showed that given a surface $\mathbf{S}$ with an empty boundary, the HPDS-MT preserves triangle quality in

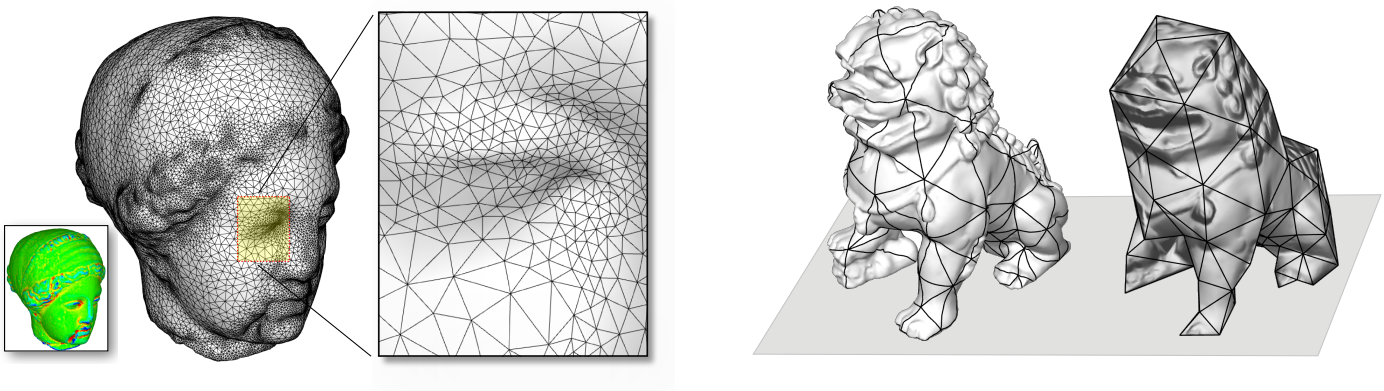

Fig. 13: (Left) Adaptive mesh obtained by a mean curvature-based resolution filter. Green regions are flat, red regions have negative mean curvature, and blue regions have positive mean curvature. (Right) Base domain and dense-to-coarse parametrization.

the varying resolution meshes that approximate the surface. Moreover, the geometric approximation is controlled by distribution of the Poisson disks radii over $\mathbf{S}$. We provided an algorithm to compute the HPDS-MT, and we described its implementation for the case in which $\mathbf{S}$ is the underlying space, $|M|$, of a triangle mesh M. We evaluated the performance of our algorithm for a few examples and compared the results with the ones provided by similar algorithms. Finally, we showed practical applications of the HPDS-MT representation; in particular, its application to level-of-detail (LOD) operations and dense-to-coarse mesh parametrization. We are currently working on a new criteria on PDS samplings to ensure topological equivalence between the underlying space of the given mesh and the one of the LRDT of the given sampling. We are also working on an extension of the HPDS-MT to represent surfaces with boundary.

Acknowledgment

The 3d input meshes used in all experiments described in this paper were freely obtained from the Aim@Shape publicly open repository. During the development of this work, Dr. Marcelo Siqueira was partially funded by CNPq Grant 305845/2012-8 and CNPq Grant 486951/2012-0.

References

[1] P. Alliez, D. Cohen-Steiner, M. Yvinec, and M. Desbrun. Variational tetrahedral meshing. *ACM Transactions on Graphics*, 24(3):617–625, July 2005. doi: 10.1145/1073204.1073238

[2] N. Amenta and M. Bern. Surface reconstruction by voronoi filtering. *Discrete & Computational Geometry*, 22(4):481–504, 1999. doi: 10.1007/PL00009475

[3] N. Amenta, M. Bern, and M. Kamvysselis. A new voronoi-based surface reconstruction algorithm. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, pp. 415–421. ACM, New York, NY, USA, 1998. doi: 10.1145/280814.280947

[4] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of surfaces. *Graphical Models*, 67(5):405–451, September 2005. doi: 10.1016/j.gmod.2005.01.004

[5] J.-D. Boissonnat and S. Oudot. Provably good sampling and meshing of lipschitz surfaces. In *Proceedings of the Twenty-second Annual Symposium on Computational Geometry*, SCG '06, pp. 337–346. ACM, New York, NY, USA, 2006.

[6] M. Botsch and L. Kobbelt. A remeshing approach to multiresolution modeling. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on Geometry Processing*, SGP '04, pp. 185–192. ACM, New York, NY, USA, 2004. doi: 10.1145/1057432.1057457

[7] T. Boubekeur and M. Alexa. Mesh simplification by stochastic sampling and topological clustering. *Computers & Graphics*, 33(3):241–249, 2009. {IEEE} International Conference on Shape Modelling and Applications 2009. doi: 10.1016/j.cag.2009.03.025

[8] S.-W. Cheng. On the sizes of delaunay meshes. *Computational Geometry: Theory and Applications*, 33(3):130–138, February 2006. doi: 10.1016/j.comgeo.2005.08.002

[9] S.-W. Cheng, T. K. Dey, E. A. Ramos, and T. Ray. Sampling and meshing a surface with guaranteed topology and geometry. *SIAM Journal on Computing*, 37(4):1199–1227, 2007. doi: 10.1137/060665889

[10] L. P. Chew. Guaranteed-quality triangular meshes. Technical Report TR-89-983, Comp. Science Dept., Cornell University, 1989.

[11] P. Cignoni, M. Corsini, and G. Ranzuglia. Meshlab: an open-source 3d mesh processing system. *ERCIM News*, 73:45–46, April 2008.

[12] D. Cohen-Steiner, P. Alliez, and M. Desbrun. Variational shape approximation. *ACM Trans. Graph.*, 23(3):905–914, August 2004. doi: 10.1145/1015706.1015817

[13] L. de Floriani, P. Magillo, and E. Puppo. The (mt) multi-tesselation (mt) package. *Available at http://www.disi.unige.it/person/MagilloP/MT/*, 1999.

[14] T. K. Dey, J. A. Levine, and A. Slatton. Localized delaunay refinement for sampling and meshing. *Computer Graphics Forum*, 29(5):1723–1732, 2010. doi: 10.1111/j.1467-8659.2010.01781.x

[15] T. K. Dey and T. Ray. Polygonal surface remeshing with delaunay refinement. *Engineering with Computers*, 26(3):289–301, 2010. doi: 10.1007/s00366-009-0162-1

[16] T. Duchamp, A. Certain, A. Derose, and W. Stuetzle. Hierarchical computation of PL harmonic embeddings. Technical report, University of Washington, 1997.

[17] M. S. Ebeida, S. A. Mitchell, A. A. Davidson, A. Patney, P. M. Knupp, and J. D. Owens. Efficient and good delaunay meshes from random points. *Computer-Aided Design*, 43(11):1506–1515, 2011. Solid and Physical Modeling 2011. doi: 10.1016/j.cad.2011.08.012

[18] M. Eck, T. DeRose, T. Duchamp, H. Hoppe, M. Lounsbery, and W. Stuetzle. Multiresolution analysis of arbitrary meshes. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '95, pp. 173–182. ACM, New York, NY, USA, 1995. doi: 10.1145/218380.218440

[19] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. In *Proceedings of the Tenth Annual Symposium on Computational Geometry*, SCG '94, pp. 285–292. ACM, New York, NY, USA, 1994. doi: 10.1145/177424.178010

[20] J. Gallier and D. Xu. *A Guide to the Classification Theorem for Compact Surfaces*. Springer, 2013.

[21] J. Guo, D.-M. Yan, X. Jia, and X. Zhang. Efficient maximal Poisson-disk sampling and remeshing on surfaces. *Computer Graphics*, 46(C):72–79, February 2015. doi: 10.1016/j.cag.2014.09.015

[22] I. Guskov. Manifold-based approach to semi-regular remeshing. *Graphical Models*, 69(1):1–18, 2007. doi: 10.1016/j.gmod.2006.05.001

[23] I. Guskov and Z. J. Wood. Topological noise removal. In *Proceedings of Graphics Interface 2001*, GI '01, pp. 19–26. Canadian Information Processing Society, Toronto, Ontario, Canada, 2001.

[24] P. S. Heckbert and M. Garland. Optimal triangulation and quadric-based surface simplification. *Comput. Geom. Theory Appl.*, 14(1-3):49–65, Nov. 1999. doi: 10.1016/S0925-7721(99)00030-9

[25] H. Kang, H. Jang, C.-S. Cho, and J. Han. Multi-resolution terrain rendering with GPU tessellation. *The Visual Computer*, 31(4):455–469, 2014. doi: 10.1007/s00371-014-0941-6

[26] A. Khodakovsky, N. Litke, and P. Schröder. Globally smooth parameterizations with low distortion. *ACM Transactions on Graphics*, 22(3):350–357, July 2003. doi: 10.1145/882262.882275

[27] L. Kobbelt. $\sqrt{3}$-subdivisionn. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '00, pp. 103–112. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 2000. doi: 10.1145/344779.344835

[28] L. Kobbelt, S. Campagna, J. Vorsatz, and H.-P. Seidel. Interactive multi-resolution modeling on arbitrary meshes. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, pp. 105–114. ACM, New York, NY, USA, 1998. doi: 10.1145/280814.280831

[29] A. Lagae and P. Dutré. A comparison of methods for generating Poisson disk distributions. *Computer Graphics Forum*, 27(1):114–129, March 2008. doi: 10.1111/j.1467-8659.2007.01100.x

[30] A. W. F. Lee, W. Sweldens, P. Schröder, L. Cowsar, and D. Dobkin. Maps: Multiresolution adaptive parameterization of surfaces. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, pp. 95–104. ACM, New York, NY, USA, 1998. doi: 10.1145/280814.280828

[31] P. Lindstrom, D. Koller, W. Ribarsky, L. F. Hodges, N. Faust, and G. A. Turner. Real-time, continuous level of detail rendering of height fields. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '96, pp. 109–118. ACM, New York, NY, USA, 1996. doi: 10.1145/237170.237217

[32] P. Lindstrom and G. Turk. Fast and memory efficient polygonal simplification. In *Proceedings of the Conference on Visualization '98*, VIS '98, pp. 279–286. IEEE Computer Society Press, Los Alamitos, CA, USA, 1998.

[33] M. McCool and E. Fiume. Hierarchical Poisson disk sampling distributions. In *Proceedings of the Conference on Graphics Interface '92*, pp. 94–105. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1992.

[34] E. Medeiros, L. Ingrid, S. Pesco, and C. Silva. Fast adaptive blue noise on polygonal surfaces. *Graphical Models*, 76(1):17–29, January 2014. doi: 10.1016/j.gmod.2013.10.004

[35] E. Medeiros, L. Velho, H. C. V. Lopes, and T. M. Lewiner. On 2d solid alpha-complex of Poisson disk sampling. *International Journal of Shape Modeling*, 15:77–91, 2009.

[36] G. L. Miller, D. Talmor, S.-H. Teng, N. Walkington, and H. Wang. Control volume meshes using sphere packing: Generation, refinement and coarsening. In *Proceedings of the 5th International Meshing Roundtable*, IMR '96, pp. 47–61, October 10-11 1996.

[37] F. Payan, C. Roudet, and B. Sauvage. Semi-regular triangle remeshing: A comprehensive study. *Computer Graphics Forum*, 34(1):86–102, 2015. doi: 10.1111/cgf.12461

[38] N. Pietroni, M. Tarini, and P. Cignoni. Almost isometric mesh parameterization through abstract domains. *Visualization and Computer Graphics, IEEE Transactions on*, 16(4):621–635, July 2010. doi: 10.1109/TVCG.2009.96

[39] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag New York, Inc., New York, NY, USA, 1985.

[40] E. Puppo. Variable resolution triangulations. *Computational Geometry: Theory and Applications*, 11(3-4):219–238, December 1998. doi: 10.1016/S0925-7721(98)00029-7

[41] E. Puppo and R. Scopigno. Simplification, LOD and multiresolution principles and applications. In *EUROGRAPHICS'97*, vol. 16, 1997.

[42] J. R. Shewchuk. What is a good linear element? interpolation, conditioning, and quality measures. In *Proceedings of the 11th International Meshing Roundtable*, IMR '02, pp. 115–126, September 15-18 2002.

[43] The CGAL Project. *CGAL User and Reference Manual*. CGAL Editorial Board, 4.2 ed., 2013.

[44] L. Velho and J. Gomes. Variable resolution 4-k meshes: Concepts and applications. *Computer Graphics Forum*, 19(4):195–212, 2000. doi: 10.1111/1467-8659.00457

[45] D. Yan, G. Bao, X. Zhang, and P. Wonka. Low-resolution remeshing using the localized restricted voronoi diagram. *IEEE Transactions on Visualization and Computer Graphics*, 20(10):1418–1427, 2014. doi: 10.1109/TVCG.2014.2330574

[46] D. Yan, J. Guo, B. Wang, X. Zhang, and P. Wonka. A survey of blue-noise sampling and its applications. *Journal of Computational Science and Technology*, 30(3):439–452, 2015. doi: 10.1007/s11390-015-1535-0

[47] D.-M. Yan and P. Wonka. Gap processing for adaptive maximal Poisson-disk sampling. *ACM Transactions on Graphics*, 32(5):148:1–148:15, October 2013. doi: 10.1145/2516971.2516973

Esdras Medeiros received the BSc (computer science), in 2000, from Universidade Federal do Ceará (UFC), Brazil. He also received the master degree in 2003 and the DSc degree in 2009, both in mathematics from Instituto de Matemática Pura e Aplicada, Brazil. In 2009 he joined the Department of Mathematics, Universidade Federal do Ceará, as Adjoint Professor. During the period of 2012-2013, he visited the Polytechinic Institute of New York University for a postdoctoral program under supervision of Prof. Claudio Silva. His work focuses on topics related to Geometric Modeling, Computational Topology, Topological Data Strucures, Geometry Processing and 3D Scanning.

Marcelo Siqueira received the BSc (computer science), in 1992, from Universidade Federal do Rio Grande do Norte. He received the master degree, in 1994, from Universidade de São Paulo, and the PhD degree, in 2006, from University of Pennsylvania. He is currently a full professor at the Department of Mathematics of Universidade Federal do Rio Grande do Norte. Dr. Siqueira conducts research in Mesh Generation, Modeling of Curves and Surfaces, and Digital Topology.