

Adaptive Maximal Poisson-disk Sampling on Surfaces

Dong-Ming Yan*
KAUST

Peter Wonka†
KAUST

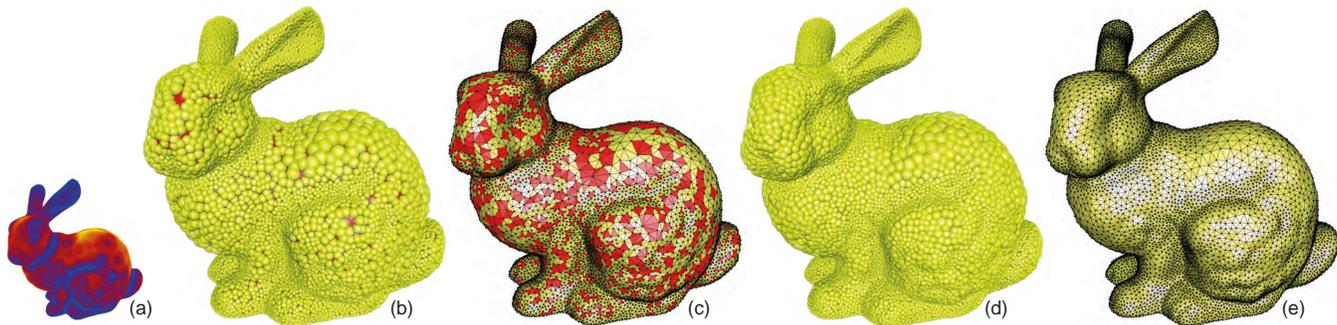


Figure 1: Adaptive maximal Poisson-disk sampling and remeshing of the Bunny model. (a) The density map, where cooler colors correspond to a smaller radius while warmer colors correspond to a larger radius; (b) non-maximal sampling results in gaps (red regions); (c) triangles of the remeshing that are affected by gaps (red triangles); (d) optimized maximal sampling; and (e) remeshing of (d).

Abstract

In this paper, we study the generation of maximal Poisson-disk sets with varying radii on surfaces. Based on the concepts of power diagram and regular triangulation, we present a geometric analysis of gaps in such disk sets on surfaces, which is the key ingredient of the adaptive maximal Poisson-disk sampling framework. Moreover, we adapt the presented sampling framework for remeshing applications. Several novel and efficient operators are developed for improving the sampling/meshing quality over the state-of-the-art.

CR Categories: I.3.6 [Computer Graphics]: Methodology and Techniques—Poisson-disk Sampling;

Keywords: Adaptive maximal Poisson-disk sampling, gaps, regular triangulation, power diagram, remeshing

1 Introduction

Sampling is one of the most fundamental research topics in computer graphics, which has many applications such as rendering, anti-aliasing, texture synthesis and geometry processing. Research shows that *Maximal Poisson-disk Sampling* (MPS) exhibits not only excellent blue-noise characteristics in the spectral domain, but also high meshing quality from a geometric point of view [Chew 1989; Ebeida et al. 2011a].

In this paper we make two major contributions: 1) On the theoretical side, we extend the existing work in 2D MPS with two new components: adaptive sampling and sampling on surfaces. 2) On

the application side, we present a high-quality surface remeshing approach based on the *Adaptive Maximal Poisson-disk Sampling* (AMPS). The key problem of our approach is how to efficiently detect and extract uncovered regions, called *gaps* from a non-maximal sampling on surfaces. Our approach is based on the concepts of the *power diagram* [Aurenhammer 1987] and the *regular triangulation*.

1.1 Related work

We briefly review techniques for random points generation and surface remeshing.

Poisson-disk sampling: Most recent algorithms for efficient maximal Poisson-disk sampling maintain a data structure of gap primitives. Dunbar and Humphreys [2006] use the scalloped sectors to track the uncovered regions. Ebeida et al. [2011b] propose a hybrid approach that first uses grid and later convex polygons bounding the intersections of a grid and multiple circles as gap primitives. The most related work to ours is [Jones 2006], which uses a Voronoi diagram to extract gaps. In contrast, we use a power diagram and the dual regular triangulation for our analysis, which results in a more general solution.

Surface sampling/remeshing: Cline et al. [2009] propose dart throwing algorithms on surfaces based on a hierarchical triangulation. Bowers et al. [2010] extend the parallel sampling to mesh surfaces and introduce a spectrum analysis for uniform surface sampling algorithms as well. Corsini et al. [2012] present an algorithm for surface blue noise sampling based on a space subdivision combined with a pre-generation of the samples. Fu and Zhou [2008] generalize the scalloped sector based sampling [Dunbar and Humphreys 2006] to 3D mesh surfaces, and present an isotropic remeshing algorithm by extracting a mesh from the samples. Lloyd iterations are used to further smooth the resulting mesh. Yan et al. [Yan et al. 2009] propose an exact algorithm for computing the *Restricted Voronoi Diagram* (RVD) on surfaces, based on which they developed a surface remeshing framework based on *Centroidal Voronoi Tessellation* (CVT). Chen et al. [2012] combine the CVT [Yan et al. 2009] and the capacity area constraint for blue noise sampling on surfaces. None of the above mentioned surface sampling methods discuss the maximal sampling properly.

*e-mail:dongming.yan@kaust.edu.sa

†e-mail:peter.wonka@kaust.edu.sa

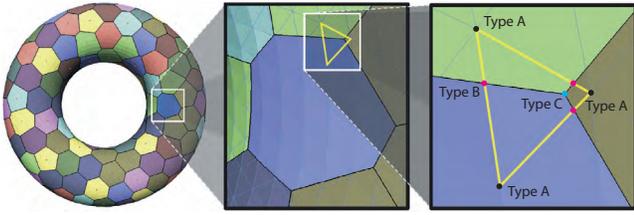


Figure 2: Illustration of the restricted power diagram on a surface (left); middle: a restricted power cell; right: a mesh triangle is split into polygons shared by the incident cells.

2 Problem Formulation

Given a compact domain Ω and a smooth density function ρ defined on the domain, each point $\mathbf{p} \in \Omega$ is equipped with a radius $r(\mathbf{p})$ adapting to the density function, e.g., local feature size of a surface. We would like to sample the domain based on the density function ρ , so that the sampling disk set $\mathcal{D} = \{(\mathbf{p}_i, r_i)\}_{i=1}^n$ satisfies the following properties: (1) Minimal distance property: none of the disk centers \mathbf{p}_i can be covered by other disks, i.e., for $\forall i, j (i \neq j), \|\mathbf{p}_i - \mathbf{p}_j\| \geq \max(r_i, r_j)$. (2) Maximal property: the domain Ω is fully covered by the disks and (3) Unbiased sampling property: the probability of sampling a point in a subregion Ω' is proportional to $\int_{\Omega'} \rho(\mathbf{x}) d\mathbf{x} / \int_{\Omega} \rho(\mathbf{x}) d\mathbf{x}$.

An efficient framework for this purpose involves two stages: First, an initial sampling stage which performs classic dart throwing on the domain. Second, a following gap filling stage which extracts and samples the tiny gaps. The gaps are decomposed and approximated by a set of convex polygons, called *gap primitives*. The second stage is repeated until there is no gap left. The bottleneck of the AMPS is the gap computation in the second stage, i.e., how to efficiently detect and extract gaps. Assuming that the current sampling set \mathcal{D} is non-maximal, which means that if we draw a disk at each center \mathbf{p}_i with radius r_i , the domain Ω will not be fully covered. We are interested in the properties of the uncovered region (gaps), which is defined as $\Omega - \mathcal{D}$, as shown in Fig. 1(b). In the following, we first give the definition of the power diagram and its dual, regular triangulation, then we analyze the geometric properties of gaps on surfaces in the next section.

Power diagram and regular triangulation: The disk set \mathcal{D} is represented by a set of weighted points $\mathcal{P}_w = \{(\mathbf{p}_i, w_i)\}_{i=1}^n$, where $w_i = r_i^2$. The power of two weighted points is defined as $\Pi(\mathbf{p}_i, w_i, \mathbf{p}_j, w_j) = \|\mathbf{p}_i - \mathbf{p}_j\|^2 - w_i - w_j$. Then the power diagram \mathcal{PD} of \mathcal{P}_w is a set of non-overlapping power cells $\{\Omega_i\}_{i=1}^n$ such that

$$\Omega_i = \{\mathbf{x} \in \Omega \mid \Pi(\mathbf{p}_i, w_i, \mathbf{x}, 0) \leq \Pi(\mathbf{p}_j, w_j, \mathbf{x}, 0), \forall j \neq i\},$$

A vertex of the power diagram is called a *power vertex* and an edge is called a *power edge*.

3 Gap Computation on Surfaces

Suppose that the input domain $\Omega = \{f_j\}_{j=1}^m$ is a triangle mesh surface and \mathcal{P}_w is a set of weighted samples lying on the surface, we define the *Restricted Power Diagram (RPD)* as the intersection of the 3D power diagram and the surface, i.e., $\mathcal{RPD}(\mathcal{P}_w) = \mathcal{PD}(\mathcal{P}_w) \cap \Omega = \{\Omega_i \cap \Omega\}_{i=1}^n$. In this case, the intersection of a 3D power cell Ω_i and the mesh surface is called a *restricted power cell*, i.e., $\Omega_i \cap \Omega = \{\cup\{\Omega_i \cap f_j\}, \forall f_j \in \Omega\}$. Fig. 2 shows an example of the \mathcal{RPD} of a torus. There are three types of vertices in the \mathcal{RPD} of a mesh surface, Type A: the original vertices of

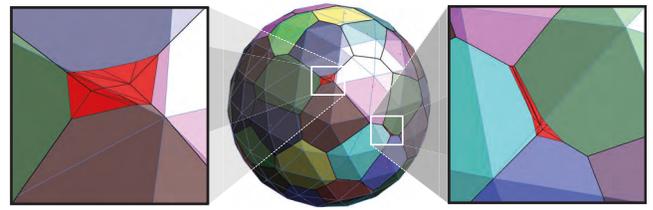


Figure 3: Gap computation on surfaces by triangle-sphere clipping. Gaps are shown in red.

the mesh; Type B: the intersection of a mesh edge and a bi-sector power plane, and Type C: the intersection of a power edge and a mesh triangle (see Fig. 2(right)). The Type C vertex is also called a *restricted power center*. Note that each restricted power center corresponds to a triangle of the so called regular triangulation (the dual triangle of the power edge which intersects the surface). The collection of these triangles is called the *Restricted Regular Triangulation (RRT)* on the surface. The triangles of \mathcal{RRT} are not actually lying on the surface, which is a linear interpolation of the input mesh. We use the algorithm presented in [Yan et al. 2009] for the \mathcal{RPD} computation.

Once we have computed the \mathcal{RPD} of the surface for a set of initial samples, it becomes straightforward to test whether there exists any gap or not. We extend Jones [2006]’s approach to 3D surfaces: The surface is fully covered if and only if each restricted power cell $\Omega_i \cap \Omega$ of a sample \mathbf{p}_i is fully covered by the disk centered at \mathbf{p}_i with radius r_i . The proof is trivial and we omit it here.

The gap primitives are computed by clipping the restricted power cell by the sphere centered at each weighted point. Each restricted power cell can be split into a set of triangles. Then the clipping problem is reduced to a triangle-sphere intersection problem. The clipped regions, i.e., gap primitives, are approximated by a set of triangles and associated with their incident gaps, as shown in Fig. 3.

4 Surface Sampling and Remeshing

In this section, we first describe a framework for adaptive maximal Poisson-disk sampling on surfaces. Then we present a high quality surface remeshing algorithm, as well as a randomized mesh optimization algorithm built on top of the sampling framework, which greatly improves the sampling/meshing quality.

4.1 Adaptive sampling on surfaces

As input we use a mesh $\Omega = \{f_i\}_{i=1}^m$, a minimal sampling radius r_{min} , a maximal sampling radius r_{max} (default value $16 r_{min}$, used to clamp the big radius), as well as a density function $\rho(\mathbf{x})$ defined on the mesh surface. A voxel grid is built for accelerating the sampling process. We first voxelize the mesh surface with voxels whose sizes are equal to $\frac{r_{min}}{\sqrt{3}}$. Each voxel records the indices of samples that fully cover it. A voxel is *valid* if it is not fully covered by any disk (sphere in 3D). We follow the two steps sampling strategy presented in [Ebeida et al. 2011b]: 1) dart throwing with a grid and 2) gap filling.

In the first step, we perform classic dart throwing on surfaces. Each triangle f is associated with a weight $w_f = \rho(\mathbf{c}_f)|f|$, where \mathbf{c}_f is the barycenter and $|f|$ is the area of the triangle. The cumulative probability density function (cpdf) of the weights is stored in a flat array. Each time a dart is generated by first selecting a triangle from the cpdf, then randomly generating a point \mathbf{p} in the triangle, as well as a radius $r = \min(r_{max}, \max(r_{min}, 1/\sqrt{\rho(\mathbf{p})}))$ associated

with the point. The new dart is tested against the grid. The dart is accepted if it is not contained by previous samples and does not contain any other existing samples. The index of the new sample is recorded by the voxels that are fully contained by this sample. The first step is terminated when k consecutive rejections are observed ($k = 300$ in our implementation).

Then we switch to the second step. i.e., iterative gap updating and gap filling. At each iteration, all gap primitives are triangulated and each triangle is associated with a weight as before. We create a cpdf for all the triangles of gap primitives. Note that in the gap filling step, the newly generated disk in the gap may cover existing samples. In this case, we set the radius of the new disk as the distance to the nearest sample, so that it will not contain any previous samples.

4.2 Surface remeshing and optimization

We extract the mesh from surface samples using the algorithm presented in [Yan et al. 2009].

Property of uniform sampling: The uniform blue noise remeshing has many nice properties. Given a constant sampling radius r , the meshes generated from blue noise sampling exhibit the following bounds: the angle bound $[30^\circ, 120^\circ]$, the edge length bound $[r, 2r]$ and the area bound $[\frac{\sqrt{3}}{4}r^2, \frac{3\sqrt{3}}{4}r^2]$ as its 2D counterpart [Chew 1989]. All these properties are desired in many applications, especially the angle bound, which is crucial for FEM applications. However, in the case of the adaptive sampling, the above theoretical bounds do not hold any more. To improve the meshing quality, we introduce a novel and simple randomized optimization algorithm, i.e., angle bound optimization and valence optimization.

Valence and angle optimization: The optimization iteratively removes the sample points with unsatisfactory properties and their neighborhoods and then re-fills the gaps. In angle bound optimization, the vertices with one triangle angle less than a minimal angle threshold or larger than a maximal threshold are removed. In valence optimization, the vertices whose degrees are less than 5 or larger than 7 are removed. These two optimization criteria can either be performed separately, or jointly. The valence/angle optimization terminates when the required criteria are met or the maximal iteration number (25 in our implementation) is reached. In a joint valence and angle optimization, a global optimization is performed interleaving between valence and angle optimization. Typically it takes 5-10 global iterations to meet both quality requirements. During the optimization, the $\mathcal{RR}\mathcal{T}$ and \mathcal{RPD} are locally updated.

5 Experimental Results

We use CGAL for the 3D regular triangulation. The experimental results are conducted on an Intel X5680 Dual Core 3.33GHz CPU with 4GB memory and a 64-bit Windows 7 operating system.

Uniform sampling/remeshing: We show several experimental results of surface remeshing and optimization. Fig. 4 shows the results of uniform sampling on a 2D square and a 3D sphere. In this example, we generate three meshes for each model using: non-maximal sampling (dart throwing without gap filling), maximal sampling, and optimized sampling. The statistics of the meshes are shown in Table 1. We can see that the theoretical bounds do not hold for a non-maximal sampling (first experiment), while these bounds are guaranteed by maximal sampling (second experiment). For optimized (uniform) sampling, the desired angle bound is

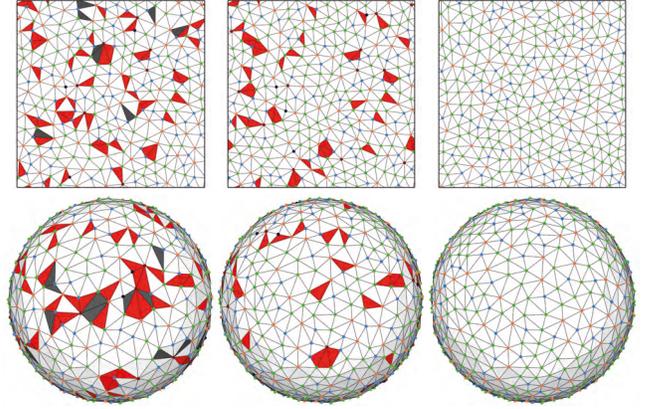


Figure 4: Uniform sampling and remeshing ($r_{min} = 0.05$ for visualization). Left column: dart throwing; middle: maximal sampling; right: optimized sampling (angle and valence). Vertex colors: V5: blue, V6: green, V7: orange. Darker points correspond to higher (> 7) or lower (< 5) valences. The triangles with $\theta_{min} < 30^\circ$ or $\theta_{max} > 120^\circ$ are shown in dark gray, triangles with $\theta_{min} \in [30^\circ, 35^\circ]$ or $\theta_{max} \in [105^\circ, 120^\circ]$ are shown in red.

set to $[35^\circ, 105^\circ]$. This third experiment shows that the geometric properties are greatly improved. The spectral analysis of the optimized sampling is given in Fig. 5, which shows that our optimization framework preserves the blue noise properties well.

Model	#v	θ_{min}	θ_{max}	$ e'_{max} $	$ t'_{min} $	$ t'_{max} $	v_{567}
Square1	26.8k	21.6	130.3	1.458	0.947	2.043	94.2
Square2	34.5k	30.1	118.6	0.999	1.001	0.998	96.4
Square3	35.1k	35.0	104.9	0.999	1.001	0.997	100
Sphere1	23.5k	22.1	129.3	1.618	0.993	2.182	94.5
Sphere2	29.7k	30.2	119.0	0.998	1.002	0.991	96.6
Sphere3	30.4k	35.0	105.0	0.998	1.002	0.996	100

Table 1: Statistics of uniform sampling/remeshing. The minimal sampling radius $r_{min} = 4.5 \times 10^{-3}$. $|e'_{min}| = \frac{|e_{min}|}{r_{min}}$ is the ratio of minimal edge length over the theoretical minimal edge length bound, same for $|e'_{max}| = \frac{|e_{max}|}{2r_{min}}$, $|t'_{min}| = \frac{|t_{min}|}{\frac{\sqrt{3}}{4}r_{min}^2}$, and $|t'_{max}| = \frac{|t_{max}|}{\frac{3\sqrt{3}}{4}r_{min}^2}$. $|e'_{min}| = 1$ for all the tests, which is omitted in the table. v_{567} is the percentage of vertices with valence 5, 6, or 7. For each model, we show the remeshing quality of 1) classic dart throwing, 2) maximal sampling and 3) optimized sampling.

We compare the uniform surface sampling with the most recent approach [Corsini et al. 2012]. As shown in Fig. 7(a), we are able to detect the gaps from their output and show that this competing result is not maximal.

Adaptive sampling/remeshing: We applied our adaptive remeshing/optimization to various models. We use *local feature size* (lfs) [Amenta et al. 1998] as the density function, i.e., $\rho(\mathbf{x}) = 1/\text{lfs}(\mathbf{x})^2$. We compare our remeshing algorithm with the state-of-the-art remeshing approaches [Valette et al. 2008] (CVD), [Yan et al. 2009] (CVT) and [Chen et al. 2012] (CAP), in terms of the min quality (Q_{min}), min/max angle ($\theta_{min}/\theta_{max}$), Hausdorff (Hdist) and RMS distances (% of the diagonal length of the bounding box), the ratio of angles smaller than 30° and the ratio of the valence 5,6,7 vertices (see Table 2). In adaptive remeshing, the desired angle bound is set to $[32^\circ, 115^\circ]$. Similar to the uniform remeshing case, we can observe that our approach exhibits better Q_{min} and θ_{min} , as well as high approximation quality to the input surface.

Efficiency: Fig. 6(a) shows the timing curve of our algorithm with

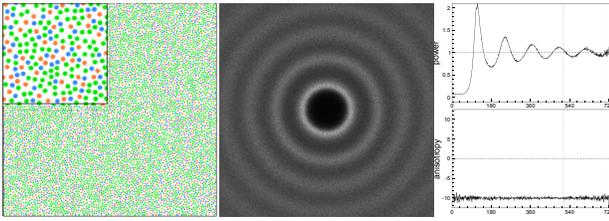


Figure 5: Spectral analysis of the optimized sampling (10k samples) in the plane with angle bounds of $[35^\circ, 105^\circ]$ and only vertices of valence 5, 6, 7.

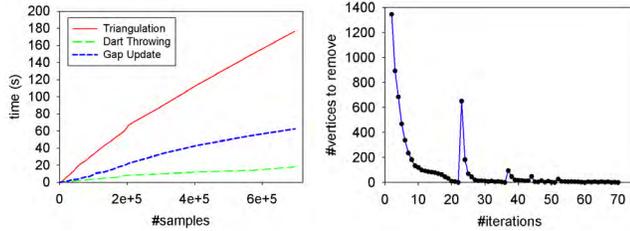


Figure 6: Left: Timing curves of our sampling algorithm on the Bunny model (Fig. 1). Right: convergence curve of the valence/angle optimization of Fig. 1. Each peak of the curve corresponds to a switch between valence/angle optimization.

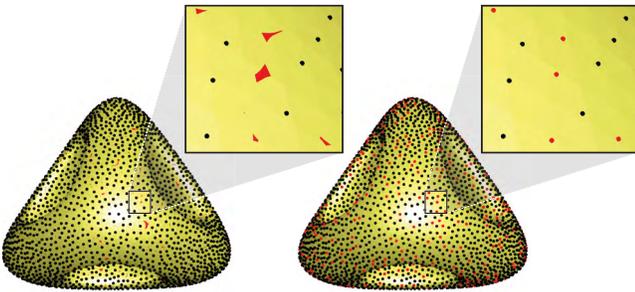


Figure 7: Comparison with [Corsini et al. 2012]. Left: uniform sampling result of [Corsini et al. 2012] with 3909 samples ($r \approx 0.012$), gaps are detected by our technique (red regions); right: maximal sampling (4560 samples) by filling the gaps using our approach, and the new sampled points are shown in red.

an increasing number of sample points, and Fig. 6(b) shows the convergence of the valence/angle optimization on the Bunny model shown in Fig. 1. For this example, our algorithm takes 12.4s for initial sampling, 6.4s for gap filling and 4.9s for optimization (23.7s in total), while CVT takes 182s and CAP takes 391s for the same number of samples, respectively.

6 Conclusion and Future Work

We have presented an efficient framework for adaptive maximal Poisson-disk sampling on surfaces, as well as a remeshing algorithm and a mesh optimization framework. In the future, we would like to generalize the current approach to more general sampling such as anisotropic sampling.

References

AMENTA, N., BERN, M., AND KAMVYSSELIS, M. 1998. A new Voronoi-based surface reconstruction algorithm. In *Proc. ACM SIGGRAPH*, 415–421.

Model	#v	Q_{min}	θ_{min}	θ_{max}	Hdist	RMS	$\theta < 30^\circ$	v567
B[CVDT]	12k	0.15	9.6	160.1	0.34	0.028	0.71	96.0
B[CVT]	12k	0.36	17.8	133.2	0.20	0.029	0.38	96.4
B[CAP]	12k	0.20	7.48	137.8	0.48	0.038	4.99	97.5
B[MPS]	9.8k	0.36	19.2	133.2	0.46	0.042	0.90	94.7
B[OUR]	12k	0.51	32.0	114.6	0.37	0.035	0	100
T[CVDT]	9k	0.01	0.46	179.1	0.59	0.050	4.37	94.1
T[CVT]	9k	0.39	15.5	127.3	0.23	0.030	0.29	99.2
T[CAP]	9k	0.41	21.2	126.4	0.30	0.031	0.43	97.2
T[MPS]	9k	0.29	13.5	138.6	0.48	0.062	1.23	92.0
T[OUR]	9k	0.51	32.0	114.8	0.46	0.062	0	100

Table 2: Statistics of the adaptive remeshing quality compared with pervious work. Model ‘B’ refers to the Bunny and ‘T’ refers to the Triceratops, respectively. Q_{min} is the minimal triangle quality [Yan et al. 2009], where $Q(t) = \frac{6}{\sqrt{3}} \frac{|t|}{p(t)h(t)}$, (where $|t|$ is the area of t , $p(t)$ is the half-perimeter of t and $h(t)$ the longest edge length of t). MPS stands for the results of our method without optimization. The best result is highlighted with bold font.

AURENHAMMER, F. 1987. Power diagrams: Properties, algorithms and applications. *SIAM J. Comput.* 16, 789–800.

BOWERS, J., WANG, R., WEI, L.-Y., AND MALETZ, D. 2010. Parallel Poisson disk sampling with spectrum analysis on surfaces. *ACM Trans. on Graphics (Proc. SIGGRAPH Asia)* 29, 6, 166:1–166:10.

CHEN, Z., YUAN, Z., CHOI, Y.-K., LIU, L., AND WANG, W. 2012. Variational blue noise sampling. *IEEE Trans. on Vis. and Comp. Graphics* 18, 10, 1784–1796.

CHEW, L. P. 1989. Guaranteed-quality triangular meshes. Department of computer science tech report 89-983, Cornell University.

CLINE, D., JESCHKE, S., RAZDAN, A., WHITE, K., AND WONKA, P. 2009. Dart throwing on surfaces. *Computer Graphics Forum (Proc. EGSR)* 28, 4, 1217–1226.

CORSINI, M., CIGNONI, P., AND SCOPIGNO, R. 2012. Efficient and flexible sampling with blue noise properties of triangular meshes. *IEEE Trans. on Vis. and Comp. Graphics* 18, 6, 914–924.

DUNBAR, D., AND HUMPHREYS, G. 2006. A spatial data structure for fast Poisson-disk sample generation. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 25, 3, 503–508.

EBEIDA, M. S., MITCHELL, S. A., DAVIDSON, A. A., PATNEY, A., KNUPP, P. M., AND OWENS, J. D. 2011. Efficient and good Delaunay meshes from random points. *Computer-Aided Design* 43, 11, 1506–1515.

EBEIDA, M. S., PATNEY, A., MITCHELL, S. A., ANDREW DAVIDSON, P. M. K., AND OWENS, J. D. 2011. Efficient maximal Poisson-disk sampling. *ACM Trans. on Graphics (Proc. SIGGRAPH)* 30, 4, 49:1–49:12.

FU, Y., AND ZHOU, B. 2008. Direct sampling on surfaces for high quality remeshing. In *ACM symposium on Solid and physical modeling*, 115–124.

JONES, T. R. 2006. Efficient generation of Poisson-disk sampling patterns. *Journal of Graphics Tools* 11, 2, 27–36.

VALETTE, S., CHASSERY, J.-M., AND PROST, R. 2008. Generic remeshing of 3D triangular meshes with metric-dependent discrete Voronoi diagrams. *IEEE Trans. on Vis. and Comp. Graphics* 14, 2, 369–381.

YAN, D.-M., LÉVY, B., LIU, Y., SUN, F., AND WANG, W. 2009. Isotropic remeshing with fast and exact computation of restricted Voronoi diagram. *Computer Graphics Forum* 28, 5, 1445–1454.