

Improving Spatial Coverage while Preserving the Blue Noise of Point Sets

Mohamed S. Ebeida^a Muhammad A. Awad^b Xiaoyin Ge^c Ahmed H. Mahmoud^b Scott A. Mitchell^a
Patrick M. Knupp^a Li-Yi Wei^d

^aComputing Research, Sandia National Laboratories, United States

^bAlexandria University, Egypt

^cComputer Science and Engineering, Ohio State University, United States

^dComputer Science, The University of Hong Kong, Hong Kong

Abstract

We explore the notion of a Well-spaced Blue-noise Distribution (WBD) of points, which combines two desirable properties. First, the point distribution is random, as measured by its spectrum having blue noise. Second, it is well-spaced in the sense that the minimum separation distance between samples is large compared to the maximum coverage distance between a domain point and a sample, i.e. its Voronoi cell aspect ratios $2\beta^i$ are small. It is well known that maximizing one of these properties destroys the other: uniform random points have no aspect ratio bound, and the vertices of an equilateral triangular tiling have no randomness. However, we show that there is a lot of room in the middle to get good values for both. Maximal Poisson-disk sampling provides $\beta = 1$ and blue noise. We show that a standard optimization technique can improve the well-spacedness while preserving randomness.

Given a random point set, our Opt- β^i algorithm iterates over the points, and for each point locally optimizes its Voronoi cell aspect ratio $2\beta^i$. It can improve β^i to a large fraction of the theoretical bound given by a structured tiling: improving from 1.0 to around 0.8, about half-way to 0.58, while preserving most of the randomness of the original set. In terms of both β and randomness, the output of Opt- β^i compares favorably to alternative point improvement techniques, such as centroidal Voronoi tessellation with a constant density function, which do not target β directly. We demonstrate the usefulness of our output through meshing and filtering applications. An open problem is constructing from scratch a WBD distribution with a guarantee of $\beta < 1$.

Key words: optimization, Voronoi aspect ratio, disk, triangulation,

1. Introduction

Many applications desire a distribution of points that are not too close to one another, yet are evenly distributed throughout the domain. Such points are *well-spaced*, or *separated yet dense*. The separation distance improves efficiency; close points often add little information, yet they consume time and memory. Well-spaced points reduce the interpolation and simulation error in scientific applications. Points spread uniformly can help reduce noise or variance. Graphics applications often desire randomness to help reduce aliasing or bias. Blue noise refers to distributions that are roughly uniform random with no preferred inter-point directions or distances. Note that well-spacedness is a local measure of nearby points. In contrast blue noise is a global measure dependent on the distances between far points. Some methods for well-spaced points often do not produce

blue noise: e.g. periodic tilings such as lattices and structured meshes; raw low discrepancy sequences; and Delaunay Refinement (DR).

In the 1990s, the fracture mechanics community studied the effects of mesh structure in finite element fracture simulations where the crack directions are limited to mesh edges [5]. They concluded that uniform random points lead to uniform random edge orientations, and more physically realistic simulations. Both Delaunay and dual Voronoi elements are of interest. The initial techniques for generating such meshes included simple jittering, perturbing the position of points. Interest in the topic recently revived, and researchers have developed more sophisticated techniques, including random sampling [3,4,13,14,34]. Computer graphics has a long-standing interest in techniques for blue-noise distributions, for applications such as rendering [10,30], animation [35], modeling [42,31], and imaging [1,17]. Graphics provides standardized methods for determining if a distribution has blue noise [36].

The well-spacedness of a distribution can be characterized by two measures: *conflict*, which is about the spacing between samples, and *coverage*, which is about how samples collectively cover the underlying domain. The conflict distance r_f is

Email address: msebeid@sandia.gov (Mohamed S. Ebeida).

URL: <http://msebeida.net/> (Mohamed S. Ebeida).

J Computer Aided-Design, vol no, pages, 2013. Definitive version available from Elsevier at <http://www.sciencedirect.com/science/article/pii/S0010448513001553>. Reposting on author website with permission from Elsevier as U.S. government funded work. <http://dx.doi.org/10.1016/j.cad.2013.08.015>.

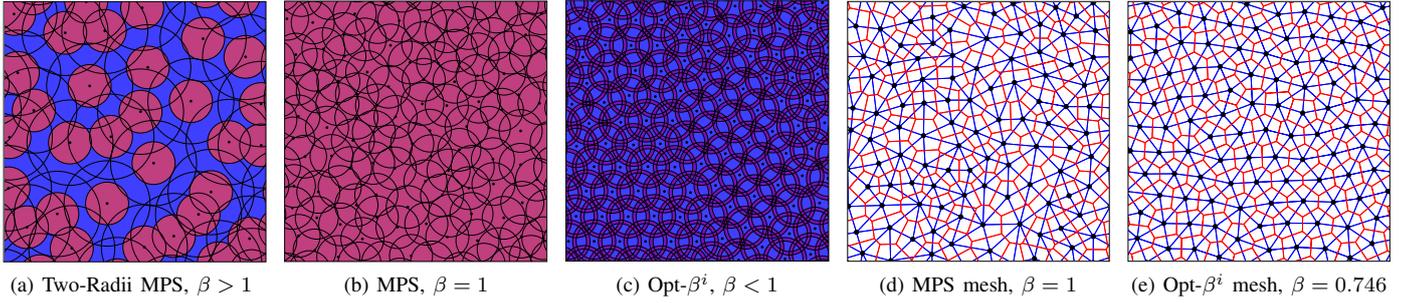


Fig. 1. Maximal Poisson-disk sampling (MPS) with different uniformity as measured by $\beta = \frac{r_c}{r_f}$. In Poisson-disk sampling, each black sample point must be outside all other points' red conflict disks with radii r_f . A distribution is maximal if the blue coverage disks with radii r_c cover the entire domain. Differentiating the two types of disks results in three sampling scenarios: (a) Two-radii MPS [27] where $r_f < r_c$. The distribution is less uniform and more random than MPS. (b) Maximal Poisson-disk Sampling (MPS) [16] where $r_f = r_c$. (c) Opt- β^i , with $r_f > r_c$; this work. Opt- β^i is more uniform and less random than MPS, improving the well-spacedness while retaining most of the blue-noise properties, as seen in the meshes in (d) and (e).

the minimum distance between two samples, the coverage distance r_c is the maximum distance between a domain point and a sample, and $\beta = r_c/r_f$ is our spacing measure. The allowed (or target) values in an algorithm and the achieved values in its output are different concepts. In typical algorithms, these distances are enforced a priori by placing disks around sample points. On output, sample points might be farther apart than the enforced minimum, and domain points might be closer to a sample than the enforced maximum. See Figs. 1–3 for illustrations and Section 4 for formulations.

Different methods achieve various degrees of conflict and coverage. Dart throwing [10] enforces r_f , and attempts to have output r_c close to it, but usually falls short. Maximal Poisson-disk Sampling (MPS) [16] enforces both and has $r_f = r_c$. Two-radii MPS [27] also enforces both, but encourages more randomness (and less uniformity) by using $r_c > r_f$. In practice, for even a few thousand points, MPS output r_f is just slightly larger than the enforced r_f , and output r_c is just slightly smaller than the enforced r_f , because a local case similar to Fig. 3b or 3c is very likely. That is, MPS guarantees output $\beta \leq 1$, and usually achieves β very close to 1. Two-radii MPS guarantees output $\beta \leq$ target β , and typically achieves something close to it.

There is some correlation between coverage/conflict and randomness, but the two are not proportional. At one extreme, uniform random points (also known as white noise or Poisson Sampling PS without disks) has unbounded coverage and conflict. Spectrum plots of MPS and Two-radii MPS distributions [27] and Opt- β^i (Fig. 9) show a strong correlation between larger achieved β and more randomness.

However, this is not the full story. In two dimensions, the vertices of the equilateral-triangle tiling achieve the minimum possible β , with $r_c = r_f/\sqrt{3}$. It is the only distribution achieving this value, and distributions coming close to this value bear some resemblance to it. The vertices of its dual, the hexagonal tiling, have $r_c = r_f$. The square tiling has $\beta = 1/\sqrt{2}$, somewhere in between. By adding or moving a single point of any of these one can generate a periodic tiling with any larger β desired, $\infty > \beta > 1/\sqrt{3}$. Delaunay Refinement (DR) [8,38] is a deterministic process for generating a triangulation. Its out-

put vertices achieve $r_c \leq r_f$, and its spectra show features of both regularity and MPS. None of the tilings or algorithms mentioned in this paragraph are random. However, a random algorithm such as MPS or even PS could in principle produce one of these, albeit with very small probability. The possible β achieved by different distribution types are shown in Fig. 2.

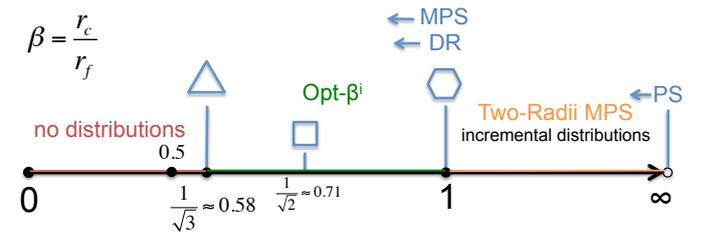


Fig. 2. Distributions by aspect ratio β , in two dimensions. Here Δ , \square , and \circ are the vertices of the equilateral triangle, square, and regular hexagonal tilings of the plane. PS is Poisson Sampling without conflict disks. DR is Delaunay Refinement. MPS is Maximal Poisson-disk Sampling. No point distribution can achieve $\beta < 1/\sqrt{3}$ in two dimensions, or $\beta < 0.5$ in any dimension.

Any random sampling method that only considers local criteria when placing points, and never moves them, cannot be guaranteed to achieve any target value of $\beta_t < 1$. The alternatives are to consider global constraints, or to adjust point locations, as Opt- β^i does in this paper. The reason that local decisions and fixed locations are insufficient is that it is easy to “paint yourself into a corner,” meaning that selecting points independently can create a subregion that is impossible to cover without placing points too close together. This can easily happen for any value of $\beta < 1$; see Fig. 3. Achieving $\beta < 1$ is actually more challenging than achieving $\beta = 1$ [16] or $\beta > 1$ [27].

To summarize,

- $\beta > 1$ is produced by Two-radii MPS [27], and the user may select β . Uniform sampling without disks produces white noise, which may have β approaching ∞ .
- $\beta = 1$ is produced by traditional maximal Poisson-disk sampling [16]. Stopping short of maximality produces $\beta > 1$.

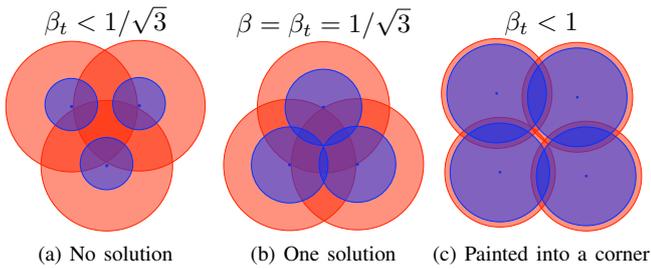


Fig. 3. Trying to achieve target ratio β_t using fixed r_c and r_f disks. In a solution the red r_f conflict disks must be free of other samples, and the blue r_c coverage disks must cover the entire domain. In (a), any new blue disk centered outside the red disks would be too small to cover the point equidistant from the three extant samples. In (b), there is only one solution: vertices forming equilateral triangles. In (c), algorithms that incrementally add points and never move them can easily get stuck even if β_t is very close to 1. The red center is too isolated to be coverable by blue disks, just as in (a).

- $\beta < 1$ is produced by Opt- β^i and it tends to make the distribution more uniform. Incremental insertion rarely achieves β significantly below 1.
- The smaller the β , the more uniform a point set tends to be, but this is a feature of both the random process generating the points and β . In two dimensions, $\beta = 1/\sqrt{3}$ indicates a periodic equilateral triangle tiling. A periodic structure can be tuned to any larger value of β . The minimum achievable β is dimension dependent, but $\beta \geq 1/\sqrt{3}$ in two-dimensions, and $\beta \geq 0.5$ by definition.

2. Spacing measures in other contexts

There are prior formulations quantifying both conflict and coverage, such as minimal spacing r_f in dart throwing [10], normalized spacing ρ [22], and coverage radius r_c [27]. Discrepancy measures spatial uniformity by comparing the fraction of points in any box to the fraction of the domain (volume) in that box [39]. In spatial statistics, in the hard-core Strauss disc process, r_f is known as the inhibition radius and r_c is known as the coverage radius. In biological system analysis, e.g. forestry image analysis [6], one usually observes a (non-maximal) point distribution that is assumed to come from this process, then seeks to estimate the hidden generating parameters.

In computational geometry, r_f^i is twice the inradius of the i th Voronoi cell. And r_c^i is known as the outradius of the Voronoi cell, the distance from a seed point to the furthest Voronoi vertex of its cell. The outradius/inradius ratio is the Voronoi cell aspect ratio, $A = 2\beta^i$. Since outradius \geq inradius, $A \geq 1$ and $\beta^i \geq 0.5$ in any dimension. A point set with maximum aspect ratio A is called an A -well-spaced set [40]. Much effort [19,26] has been devoted to generating A -well-spaced sets, for both uniform and spatially varying densities. Usually the upper bound on A is greater than 2, but we strive for $A < 2$.

In graphics a common measure of point distributions is the relative radius ρ [22]. Its definition starts with non-overlapping disks with constant $r_f/2$ radius. The packing density η is the

fraction of the domain area covered by these disks; η is itself a popular measure in physics. Given a unit-area periodic domain containing N disks, r_{\max} is the radius achieved by the densest (in terms of η) known arrangement, assuming that N divides evenly into that tiling. In two dimensions, $r_{\max} = \text{const}/\sqrt{N}$. We define $\rho = r_f/(2r_{\max})$, where $\rho \in [0, 1]$. Thus ρ is directly related to r_f , but only indirectly related to r_c through N . For example, if we create a gap by removing one vertex (disk) from an equilateral triangle tiling, then ρ decreases by a factor of $\sqrt{N-1}/\sqrt{N}$, a negligible change if N is large, but r_c and β almost double from $1/\sqrt{3}$ to 1. In our context the number of points N is fixed, so we have $\beta \propto r_c/\rho$. Area coverage and r_{\max} is a global averaging over the disks, whereas r_c^i is a maximum achieved in a local neighborhood of a point, its Voronoi cell. We think that β^i is a better measure of local spacing than ρ because β measures local gaps and ρ does not.

3. Comparison methods

We compare our method against three alternative sampling techniques. While they all strive to improve the spatial quality of a distribution, none optimize β^i directly as we do.

3.1. Centroidal Voronoi Tessellation CVT

A CVT [12] is defined as a distribution in which each point is at the center of mass of its Voronoi cell. The earliest technique to achieve a CVT is Lloyd's iteration [25], iteratively moving each point to the center of its Voronoi cell. Recently Liu et al. [24] improved the speed of converging to a CVT by recognizing the second-order smoothness of its energy function, and optimizing it using a quasi-Newton method. But the converged solution is the same. The CVT converges to an equilateral triangle tiling, absent boundary effects. For this reason, many applications stop short of convergence.

Based on the iterative method definition, one might expect CVT to improve r_c^i , and not degrade r_f^i , at least in the average sense. Our experience is that CVT tends to decrease the maximum coverage radius r_c , but does not provide any control over the minimum disk-free radius r_f .

3.2. DistMesh

DistMesh is a popular point relocation technique that aims to make the Delaunay edge length distribution match a user-specified sizing function [33]. This technique simply treats each Delaunay edge as a spring that can expand or contract; but if its length is less than the user-specified sizing function it can only expand. For the uniform case, the sizing function is constant and is usually chosen to be 20% more than the desired edge length. While this technique is good for improving the quality

of the majority of the elements, it is not as good for improving the worst-quality elements.

In contrast to CVT, DistMesh tends to increase the disk-free radius r_f of an MPS, however, it does not provide any control over the coverage radius r_c .

3.3. Farthest neighbor (Far-Point)

Farthest neighbor [37,7] is an iterative technique that moves each point to the Voronoi vertex farthest from its immediate neighbors. As such it locally optimizes r_f^i , and ignores r_c^i . The algorithm works by finding the Delaunay neighbors $\mathcal{N} = \{x^j\}$ of sample x^i , computing the Voronoi vertices of \mathcal{N} (without x^i), then reinserting x^i at the Voronoi vertex with the maximum minimum distance to \mathcal{N} . If \mathcal{N} are not in convex position, there will be some Voronoi vertices far from the convex hull of \mathcal{N} ; these vertices are ignored.

This technique tends to increase all Delaunay edge lengths, not only the short ones.

3.4. Our method, Opt- β^i

We propose Opt- β^i , local optimization of the β^i of the initial MPS sample positions, stopping when a user-specified ratio $\beta_t = r_c/r_f$ is achieved. By tuning β , our method can produce sample distributions with a level of uniformity and randomness as desired by the users. There is no guarantee that our method can achieve any value of $\beta < 1$, or even that β will not increase, but in practice we can obtain $\beta \approx 0.75$.

In the sampling methods of Section 1, at least one of r_c or r_f is a fixed enforced value. Opt- β^i is different, in that these are both measured quantities of the output. We simultaneously adjust both r_c and r_f in pursuit of optimizing their ratio β . Moving points might increase r_f or decrease r_c . Experiments show that despite this possibility, we usually simultaneously achieve both an increased r_f and a reduced r_c , while maintaining good spectra; see Table 1 and Fig. 5. We compare the methods over a periodic square in Fig. 4.

Our paper has the following main contributions.

- Introducing and formalizing the notion of Well-spaced Blue-noise Distributions (WBDs).
- Proposing an algorithm, Opt- β^i , to generate a WBD with controllable uniformity and randomness through β_t .
- Demonstrating the benefits of a WBD and Opt- β^i through applications in geometry meshing and image filtering.

4. Formal definitions

Conflict radius r_f . The minimum achieved spacing between any pair of samples in the set \mathcal{S} . For example, for MPS, at least the enforced Poisson-disk radius.

Coverage radius r_c . Any point in the domain Ω will be at most r_c away from the nearest sample in \mathcal{S} . Equivalently, the set of spheres with radius r_c will cover the entire domain.

Ratio β . The ratio between conflict and coverage is measured by $\beta = r_c/r_f$. This is the ultimate objective of our method.

Initial radii. The input (MPS) point set has $r_{\text{MPS}} = r_f = r_c$.

During the local optimization procedure, we only consider the local values of these quantities, i.e.

Adjusted point x^i . We move one point at a time, $x^i \in \mathcal{S}$.

Conflict radius r_f^i . Distance from x^i to its nearest $x^j \in \mathcal{S}$.

Coverage radius r_c^i . Distance from x^i to the farthest vertex of its Voronoi cell, i.e. the distance to the domain point farthest from x^i that is not closer to some other sample x^j .

Local ratio β^i . This is r_c^i/r_f^i . Note that

$$\beta = \frac{\max_i r_c^i}{\min_j r_f^j} \geq \max_i (\beta^i) = \max_i \left(\frac{r_c^i}{r_f^i} \right). \quad (1)$$

5. Opt- β^i algorithm

We now describe our algorithm in detail. We start with an MPS, or some other distribution with blue noise, with $\beta = 1$. (Recall that the MPS output might have $\beta < 1$, but in practice β is so close to 1 that we just say that $\beta = 1$.) It is possible to start with other distributions, such as the output of DR, but, since our method removes noise, the output will be limited by the noise in the input. In principle it is also possible to start with a distribution with $\beta > 1$; in this paper we assume that such distributions are enriched to maximality with an MPS algorithm so that $\beta = 1$.

The user specifies the target $\beta = \beta_t$, the desired trade-off between spacing and randomness. For two dimensions, $\beta_t \in [\frac{1}{\sqrt{3}}, 1]$.

Let $\beta^i = r_c^i/r_f^i$ denote the local coverage/conflict ratio. Ideally, one might wish to simultaneously adjust the location of all points to find a global minimum for β , or at least to minimize the maximum of β^i over all x^i . A more practical method is to iteratively move each x^i towards its local optimum [20].

We iteratively reposition each sample point x^i to locally minimize its β^i , until globally $\beta < \beta_t$.

To reposition x^i , we perform ten iterations of Nelder–Mead [28]. Nelder–Mead is a downhill-traveling optimization

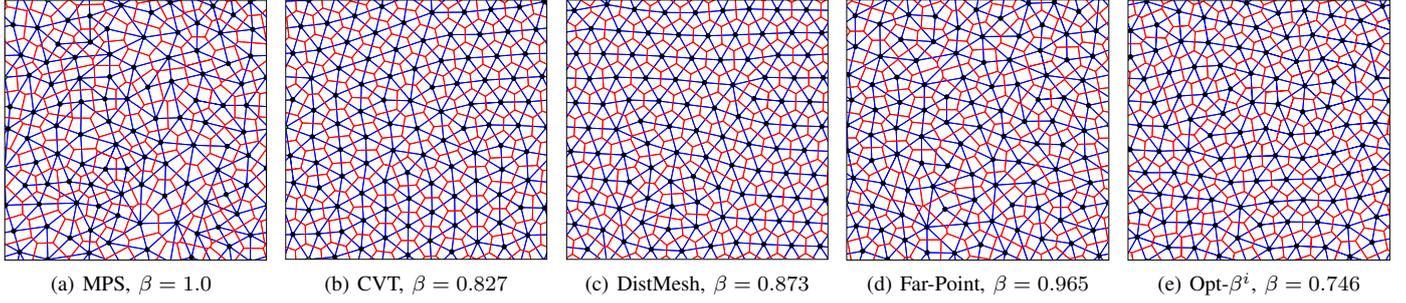


Fig. 4. Final mesh for a periodic unit box after applying various methods to the input MPS in (a). While CVT and DistMesh improve the quality of the majority of the Voronoi cells, they tend to lose randomness at larger values of β . Far-Point on the other hand tends to violate the coverage condition. None of the four methods were able to achieve $\beta < 0.7$ in general.

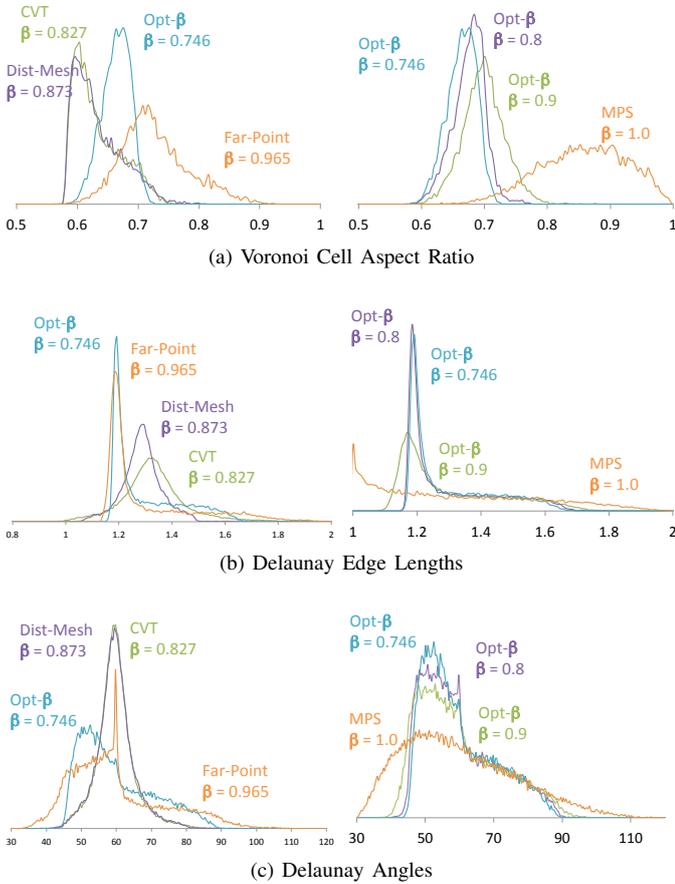


Fig. 5. Distribution of quality measures. The left column compares the different methods, and the right column compares Opt- β^i for different β .

heuristic that requires only function values, not derivatives. Its state consists of three candidate points $c_{\{1,2,3\}}$, and the β^i function values if x^i was repositioned at those locations. To get started we select the first candidate point c_1 at x^i , the second c_2 at x^i shifted in the x-axis direction by $r_{\text{MPS}}/10$, and the third c_3 at x^i shifted in the y-axis direction by $r_{\text{MPS}}/10$. These $c_{\{1,2,3\}}$ define a triangle, tilted in three dimensions by assigning the respective height β^i to its corners. Nelder–Mead replaces the high corner by some point on the other side of its opposite triangle edge, and recomputes a new triangle and its tilt. We flip-flop through ten triangles and stop.

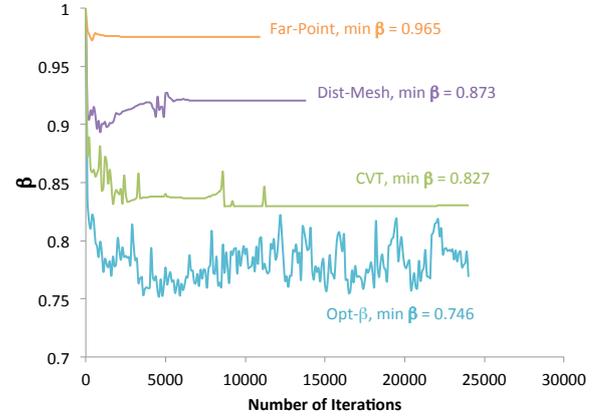


Fig. 6. Variation of β as the iteration proceeds.

Each candidate may have different Delaunay neighbors than the original point; we calculate β^i using the candidate’s Delaunay neighbors, rather than the original points’ neighbors. Sometimes, especially in initial iterations, x^i moves outside of the convex hull of its initial Delaunay neighbors: r_c^i decreases as x^i moves towards the convex hull, but it may happen that r_f^i increases more rapidly, yielding a local improvement in β^i . We experimented with preventing this from happening, but found it was faster to allow this to happen.

Often, improving β^i causes a nearby point’s β^j to get worse. This is a common phenomenon in mesh smoothing algorithms, and the common approach is to allow this to happen, and attempt to remedy it later during the optimization of x^j .

We sweep over all x^i , locally optimizing and updating their positions. Each update happens immediately. That is, for a neighbor x^j of x^i , we use the updated position of x^i rather than its initial position at the start of the sweep. Our input comes from Simple MPS [15]. The output of that algorithm provides a nice ordering to the points, by scan lines. Simple MPS divides the domain into boxes of side length r_{MPS} , and each box has at most one point. On output the points are lexicographically ordered by the box they lie in, first by row then by column. This is the order the points are visited during our Opt- β^i sweep, regardless of where they later migrate to. This is better than considering points in random order.

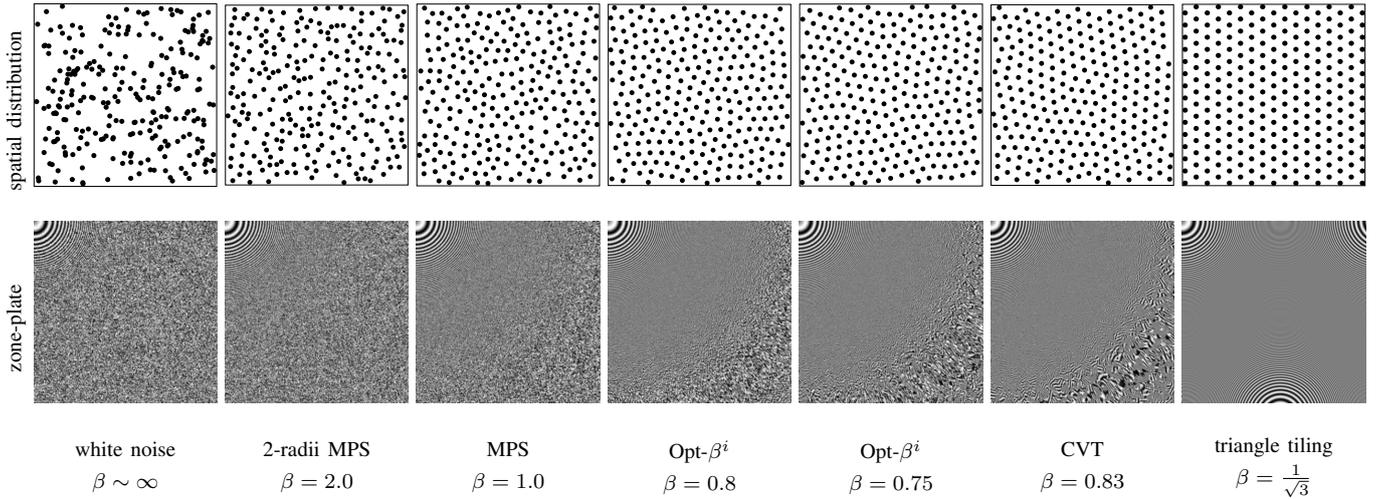


Fig. 7. Spatial analysis via zone-plate sampling. Each uses about 122,500 samples (about 1 sample per pixel) and a 3-pixel-wide Gaussian kernel.

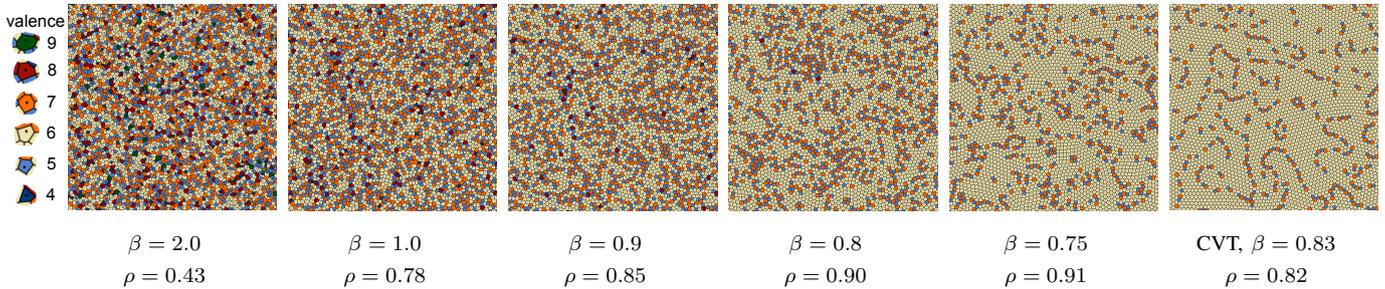


Fig. 8. Voronoi cell valence and ρ for different β achieved by $\text{Opt-}\beta^i$ and CVT. Recall that $\beta \propto r_c/\rho$. See also Fig. 8 in “Blue Noise...” [11].

In each sweep we optimize the position of all x^i , even if their local β^i is already small. We tried iterating over just the large- β^i points, but this tended to get us stuck in a local minimum.

For large targets, e.g. $\beta_t > 0.9$, we often get far below it in one or two iterations. To prevent that, we use a damping factor, moving x^i only about half-way (0.6 to 0.8) of the distance from its initial to its Nelder–Mead optimized position.

Local patch smoothing is common for unstructured meshes [18]. Unlike patch smoothing, the connectivity of the implied mesh, the set of Delaunay neighbors of each sample, changes throughout our algorithm, and in our comparison methods.

6. Analysis

6.1. Angle and edge bounds

Here, we recall the relationship between $\beta = r_c/r_f$ and the edge lengths $|e|$, empty circumcircle radius R , and angles α in a Delaunay Triangulation (DT) of the point set. The relationships are well known, and they form the basis for the Delaunay refinement [8] family of algorithms. We restate the succinct summary from Mitchell et al. [27]:

Proposition 6.1 $|e| \in [r_f, 2R]$ and $R \leq r_c$.

Proposition 6.2 $\sin \alpha \geq |e|/2R$.

Normalizing by r_f , and noting that the largest angle in a triangle is the supplement of the two smaller ones, we get the following.

Proposition 6.3 $|e|/r_f \in [1, 2\beta]$

Proposition 6.4 $\alpha \in [\arcsin 1/(2\beta), 180^\circ - 2 \arcsin 1/(2\beta)]$

For example, $\beta < 0.75$ gives $|e|/r_f \in [1, 1.5]$ and $\alpha \in [41.8^\circ, 96.4^\circ]$. A value of $\beta < 1/\sqrt{2} \approx 0.71$ provides a non-obtuse triangulation.

These hold locally; i.e., if the local value of β^i for x^i is smaller, it provides tighter bounds for the edges and triangles of x^i .

6.2. β ratio

The main results are that $\text{Opt-}\beta^i$ generates the smallest β values, and the best blue-noise (the most-random spectra). We get a smaller β simultaneously with a better spectrum. $\text{Opt-}\beta^i$ is the best, followed by CVT, DistMesh, and then Far-Point.

Unfortunately, the speed of *one iteration* roughly reverses this order: DistMesh is the fastest, followed by Lloyd’s iteration (CVT without Liu et al.’s faster optimization [24]), Far-Point, and then $\text{Opt-}\beta^i$. *One iteration* of $\text{Opt-}\beta^i$ is $15 \times$ slower than one iteration of standard CVT. For all methods, the overall runtime is roughly linear in the number of iterations. However, $\text{Opt-}\beta^i$ is competitive with respect to the overall runtime. For

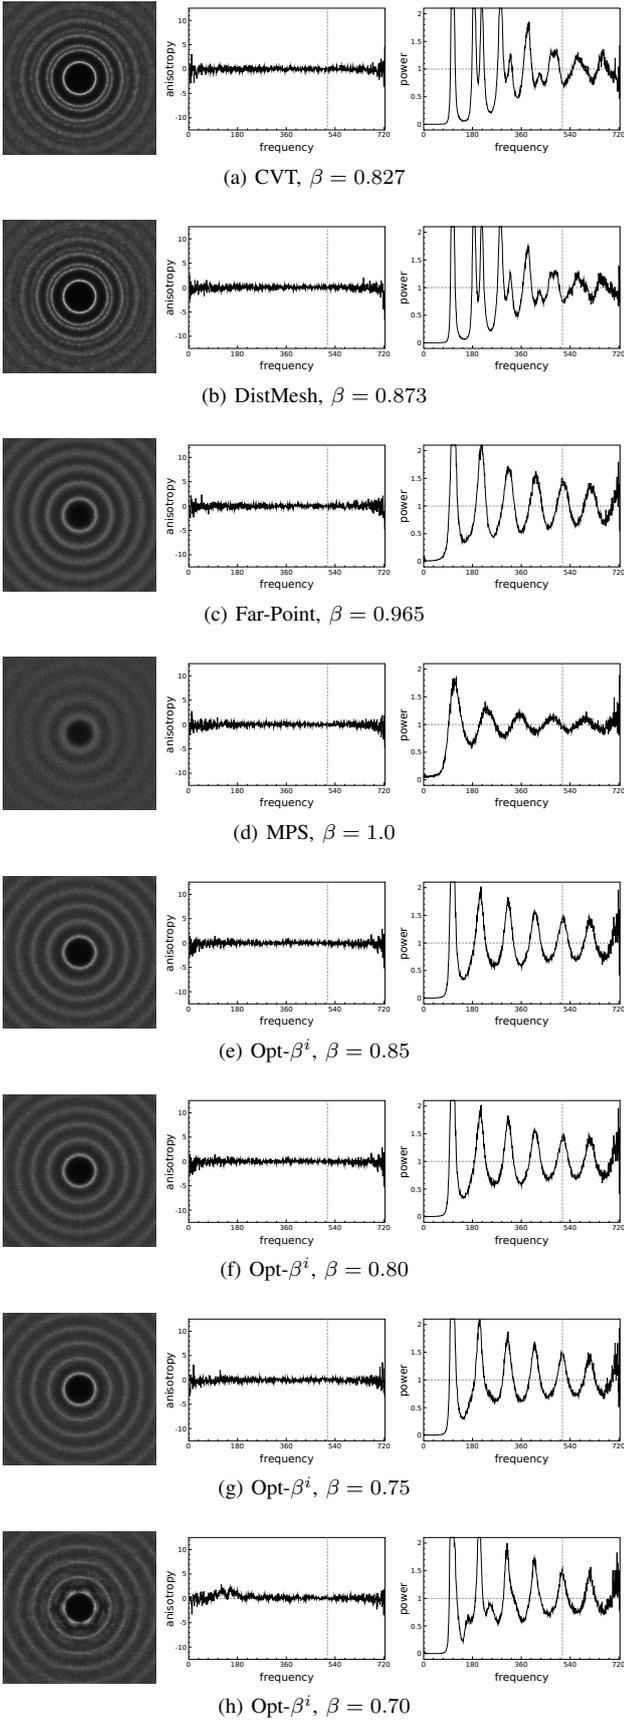


Fig. 9. Spectral analysis. Shown here are the Fourier spectrum analyses of various sample patterns. Each case contains $\sim 10K$ samples in a unit box.

10,000 points in the periodic cube, to reach $\beta = 0.85$, $\text{Opt-}\beta^i$ took 3.3 seconds and CVT took 10 seconds, because CVT required more iterations.

Since we only locally optimize the local $\beta^i = r_c^i/r_f^i$, a somewhat surprising result is that the global $r_f = \min_i r_f^i$ is actually improved above r_{MPS} . Since we improved β , this implies that we have also reduced the maximum r_c^i . These features can be seen from the third and fourth columns of Table 1.

Each β^i is smaller than β , so how small β can be and still have randomness depends on the sample size. A larger sample size is more likely to have a local case stuck at a larger value of β . One alternative would be to consider average β^i .

6.3. Spatial and spectral analysis

Figs. 7 and 9 analyze the $\text{Opt-}\beta^i$ output for different β values, over the periodic unit box domain. We use common spatial and spectral measures [29,22,43,17,44,11]. For spatial analysis in Fig. 7, the zone-plate patterns are input, and we produce Gaussian-filtered output via our sample sets. Since the pattern contains a variety of spatial frequencies in different directions, it is a good stress test to detect any sampling pattern anomalies. We also plot the spatial samples directly for visual inspection. $\text{Opt-}\beta^i$ exhibits the classic noise/alias trade-off: the smaller the β , the more uniform the distribution, and thus there is less noise but more aliasing. Fig. 9 shows the Fourier spectra, radial mean, and anisotropy. We start to lose blue noise between $\beta = 0.75$ and $\beta = 0.7$. The noise/alias trade-off does not apply across methods: for example, even though CVT has a higher β than our $\text{Opt-}\beta^i$ $\beta = 0.8$ case, it still produces more aliasing.

6.4. Convergence effect of small β_t

Our experience is that, for values of β_t above about $1/\sqrt{2} \approx 0.71$, the repositioning problem is mostly a geometry problem. For most local configurations changing point positions to achieve this β_t is eventually possible. Below about 0.71, we see oscillatory behavior in Fig. 6, and the problem becomes a discrete configuration problem. To achieve $\beta_t = 1/\sqrt{3}$, each point must have exactly six neighbors, evenly spaced on a circle around it. The six-neighbor pattern appears at about $\beta = 0.7$ in the hexagonal shape of the rings in the spectrum, bottom left of Figure 9. Even with five or seven neighbors, it is difficult to obtain a β close to $1/\sqrt{3}$. In the rightmost $\text{Opt-}\beta^i$ example in Fig. 8, with $\beta = 0.75$, all vertices have valence five, six, or seven.

Besides this local discrete configuration constraint, small β_t introduces global geometric constraints. The equilateral-triangle tiling has points aligned along three sets of parallel lines. These lines constrain points that are far from one another. This is another reason for slow (or no) convergence for small β_t . Fig. 11

shows that little progress is made beyond a certain number of iterations.

We perform ten Nelder-Mead steps for each local optimization. Using seven or fewer steps causes the overall algorithm to get stuck at a larger value of β . Using more steps requires more run-time; ten steps ensures both convergence and efficiency.

7. Applications

7.0.1. Meshing planar non-convex domains

We generated Well-spaced Blue-noise Distributions (WBDs) for the bounded non-convex domain in Fig. 10. This is more challenging than for a periodic square, and all four methods suffered from boundary constraints. All methods tended to improve the quality of most of the Voronoi cells. However, CVT and DistMesh violated the disk-free condition. On the other hand, Far-Point violated the coverage condition. Only Opt- β^i was able to reduce the value of β while preserving both properties. The other methods increased β over unity within their first few iterations. The problem is harder for the coarse mesh, because the boundary has an increased effect. Statistics after convergence (or more precisely, after stagnation) are given in Table 1. For Opt- β^i we solved a one-dimensional version of Nelder-Mead optimization to keep points on the boundary.

Algorithm	β	$\frac{r_c}{r_{\text{MPS}}}$	$\frac{r_f}{r_{\text{MPS}}}$	min α	max α
MPS	1.0	1.0	1.0	31	115
CVT	1.226	0.931	0.759	24	103
DistMesh	1.089	1.012	0.929	31	114
Far-Point	1.048	1.043	0.996	32	106
Opt- β^i	0.988	0.995	1.007	32	110

(a) Coarse mesh, $r_{\text{MPS}} = 0.0314$, Fig. 10 left column.

Algorithm	β	$\frac{r_c}{r_{\text{MPS}}}$	$\frac{r_f}{r_{\text{MPS}}}$	min α	max α
MPS	1.0	1.0	1.0	30	117
CVT	1.02	0.989	0.852	33	96
DistMesh	1.07	0.869	0.925	34	107
Far-Point	1.06	1.106	1.047	31	113
Opt- β^i	0.932	0.939	1.008	34	99

(b) Fine mesh, $r_{\text{MPS}} = 0.0157$, Fig. 10 right column.

Table 1
Statistics of meshes in Fig. 10 mesh; α denotes triangle angles.

7.0.2. Meshing curved surfaces

We applied Opt- β^i to a sphere and the Fertility sculpture. The disk-free and coverage distances are geodesic. For the uniform sphere, Fig. 14, Opt- β^i reduced β significantly, from 1.0 to 0.73. This is reflected in more regularity in the Voronoi and Delaunay meshes. Using CVT, the minimum β achieved was 0.81. Fig. 13 shows the quality measure distributions for valence, Delaunay angle, and Delaunay edge length.

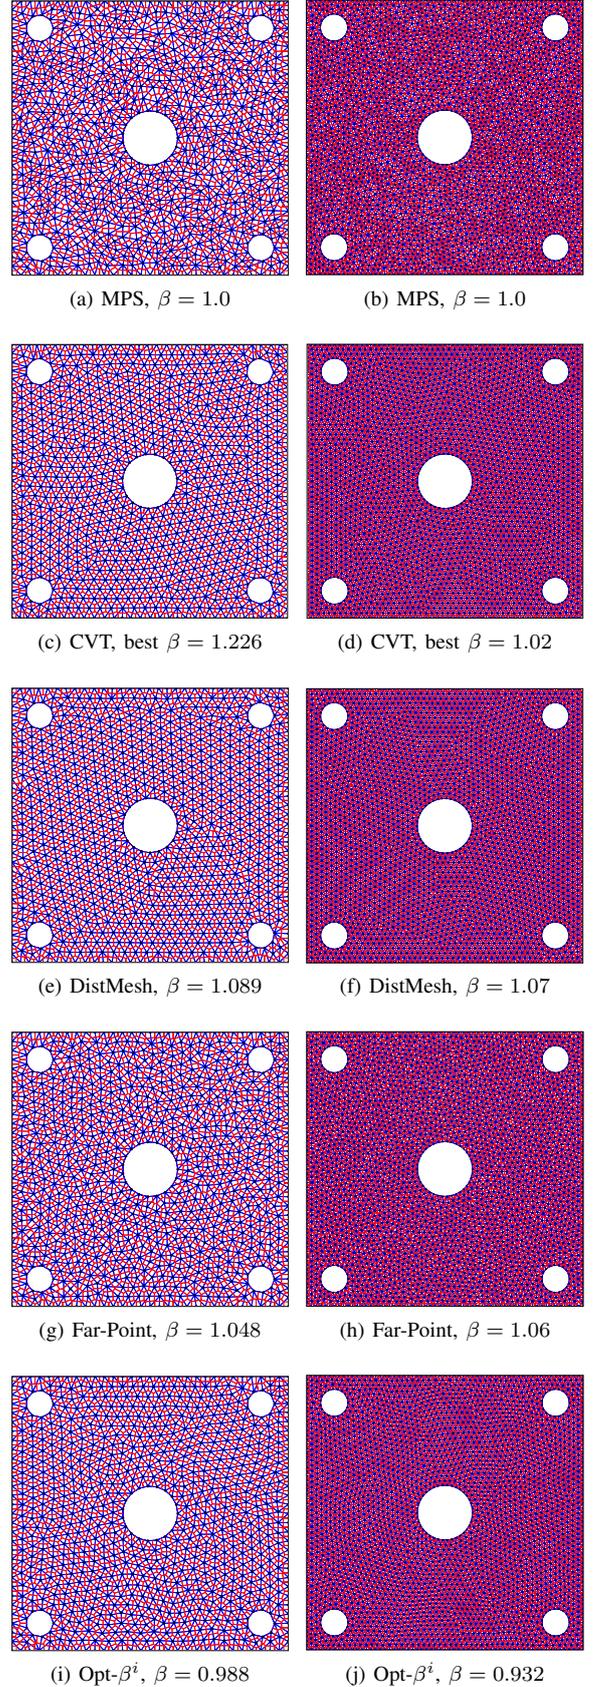


Fig. 10. Applying various methods to a coarse (left) and fine (right) mesh of a non-convex two-dimensional domain. For the finer mesh the boundary effects are less constraining, and a smaller β was achieved.

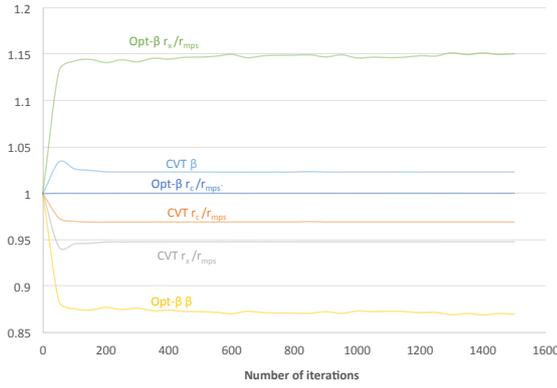
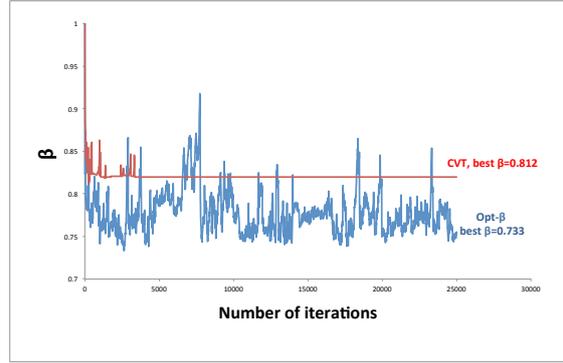
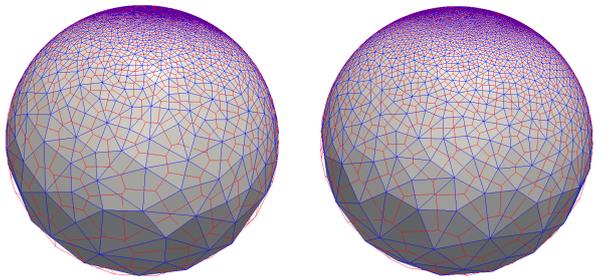


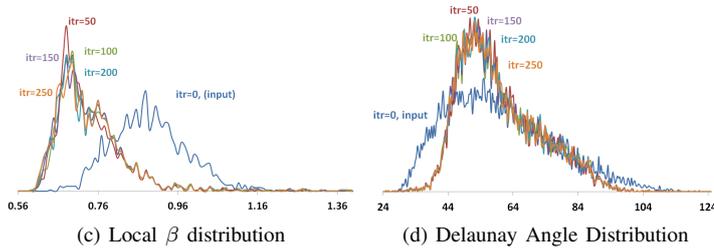
Fig. 11. Typical progress in β , r_c/r_{MPS} , and r_f/r_{MPS} for CVT and $\text{Opt-}\beta^i$.



(a) Variation of β as the iteration proceeds



(a) Input mesh (b) $\text{Opt-}\beta^i$ mesh after convergence



(c) Local β distribution (d) Delaunay Angle Distribution

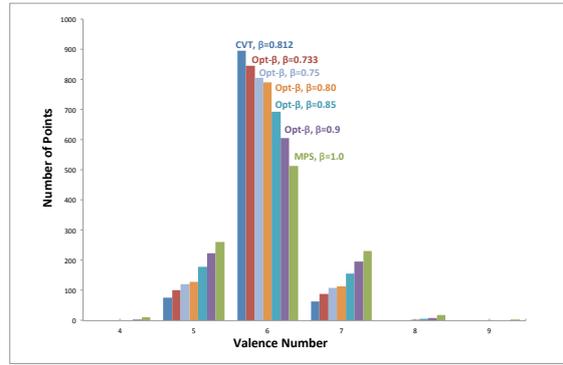
Fig. 12. Applying $\text{Opt-}\beta^i$ on non-uniform surface mesh with linear sizing function in the z -direction. The angles distribution and local β distribution are shown for number of iterations until convergence (or stagnation) occurs.

For the non-uniform sphere, Fig. 12, we used a linear sizing function in the z -direction, varying from 0.21 to 0.01. We define the conflicts based on the **smaller-disk** criteria [27]. Fig. 12 shows that neither the Delaunay angle bounds nor the maximum global β have changed much; instead their distributions have changed within about the same intervals.

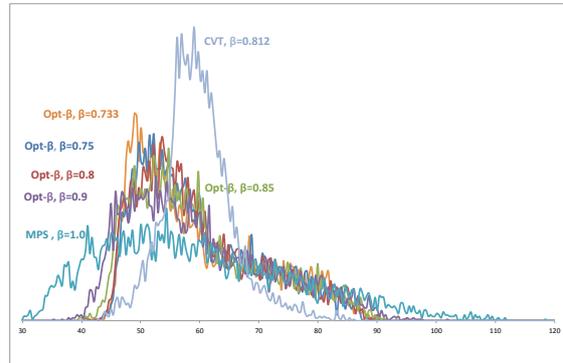
We took the same approach with the Fertility model, a more complex shape. Fig. 15 shows that we reach convergence after a few iterations, and achieve a better angle and local β distribution.

7.1. Image filtering

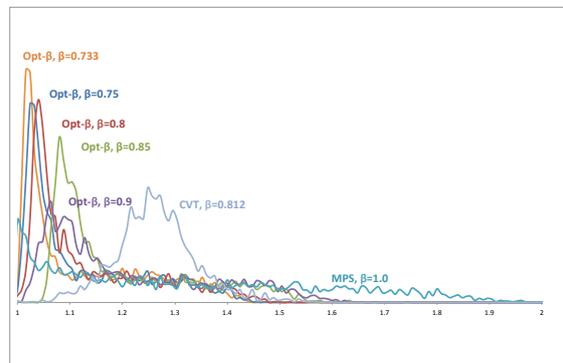
Bilateral filtering is a core filtering method with a variety of applications in graphics, vision, and image processing [41]. Similar to traditional linear filtering, e.g. Gaussian blur, bilateral filtering strives to reduce noise through blending nearby pix-



(b) Valence for different β for $\text{Opt-}\beta^i$ and CVT



(c) Delaunay Angles



(d) Delaunay Edge Length

Fig. 13. β by iteration, and qualities by β , for $\text{Opt-}\beta^i$ and CVT. The input was 1,000 points on a sphere.

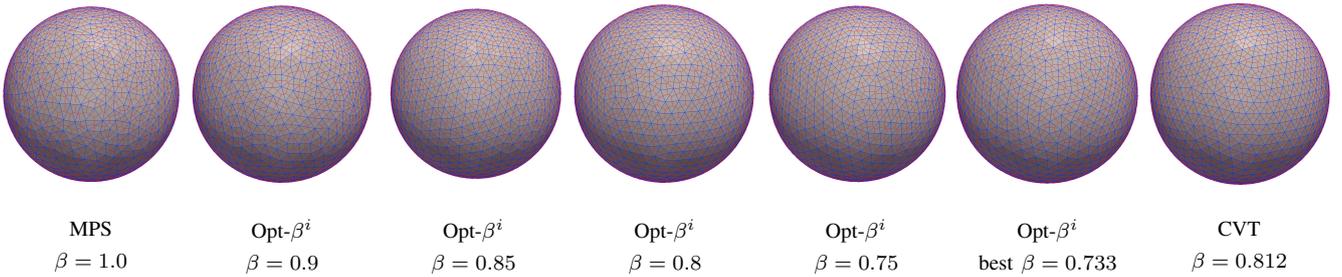


Fig. 14. Meshing for various β values over a sphere surface using 1,000 samples.

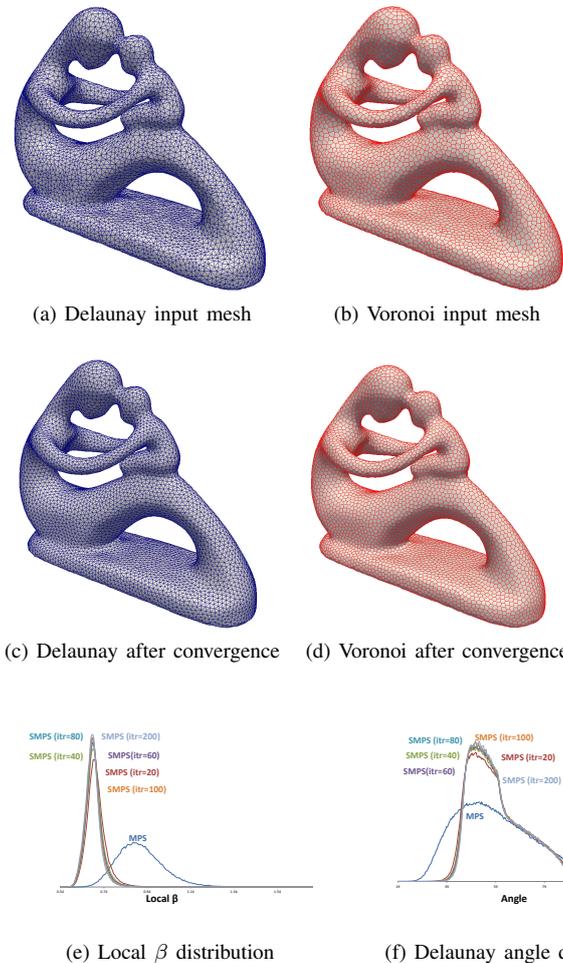


Fig. 15. Applying $\text{Opt-}\beta^i$ on the Fertility model with a non-uniform sizing function.

els. However, unlike linear filtering which assigns filter weights based on domain information (spatial pixel locations) only, bilateral filtering also considers range information, such as pixel colors. This allows better preservations of image features such as region boundaries.

The downside is that bilateral filtering is slower and more difficult to accelerate than linear filtering. Many methods have been proposed to accelerate it [32]. Among these, subsampling [2] is simple and effective. Instead of using all pixels, Banterle et al. [2] select only a subset of the kernel pixels. A properly chosen subset yields faster computation without a noticeable loss

in output quality. The authors of [2] experimented with various subsampling methods, and concluded that Poisson-disk sampling can be an excellent choice, due to its uniformity (reducing filtering noise) and randomness (low aliasing). Their paper suggests that *maximal* Poisson-disk sampling could be even better.

Here, we explore maximal Poisson-disk sampling with different β for subsampling-accelerated bilateral filtering. Intuitively, since a smaller β will produce more uniform and more regular distributions, it has the potential to reduce noise but increase aliasing in the filtering output. This is confirmed in Figure 16, in which we experiment through a range of β , using $\text{Opt-}\beta^i$ ($\beta < 1$) and other methods ($\beta \geq 1$) [27]. In our experiments it seems $\beta \in [0.85, 0.75]$ strikes the right balance between uniformity (reducing noise) and randomness (avoiding bias).

Measure	White noise	$\text{Opt-}\beta^i$, by β			Triangle tiling	
		2	1	0.85	0.75	
Cattle skull						
mean error	0.013	0.010	0.0093	0.0089	0.0087	0.012
RMSE	0.020	0.014	0.013	0.013	0.012	0.017
Reflection						
mean error	0.020	0.013	0.012	0.011	0.010	0.016
RMSE	0.029	0.020	0.019	0.017	0.017	0.020

Table 2

Mean error and Root Mean Squared Error (RMSE) for various schemes over models from Fig. 16.

8. Conclusions

This paper introduced a Well-spaced Blue-noise Distribution (WBD), with $\beta = r_c/r_f$ measuring coverage uniformity or well-spacedness. We proposed the $\text{Opt-}\beta^i$ algorithm to change a random point set to a WBD; blue noise is preserved up to $\beta \approx 0.75$. We demonstrated $\text{Opt-}\beta^i$'s efficacy in geometry meshing and image filtering applications. We believe that the main contribution of this paper is on the introduction and measurement side, and we envision fruitful future directions for both algorithm and application development. For example, generating a WBD with $\beta < 1$ from scratch is an open problem.

Adaptive and anisotropic sampling. We kept r_c and r_f constant throughout the domain, except for the sphere. In principle, we could extend both to spatially varying functions of sample positions, for adaptive or anisotropic sampling [23]. The spatial

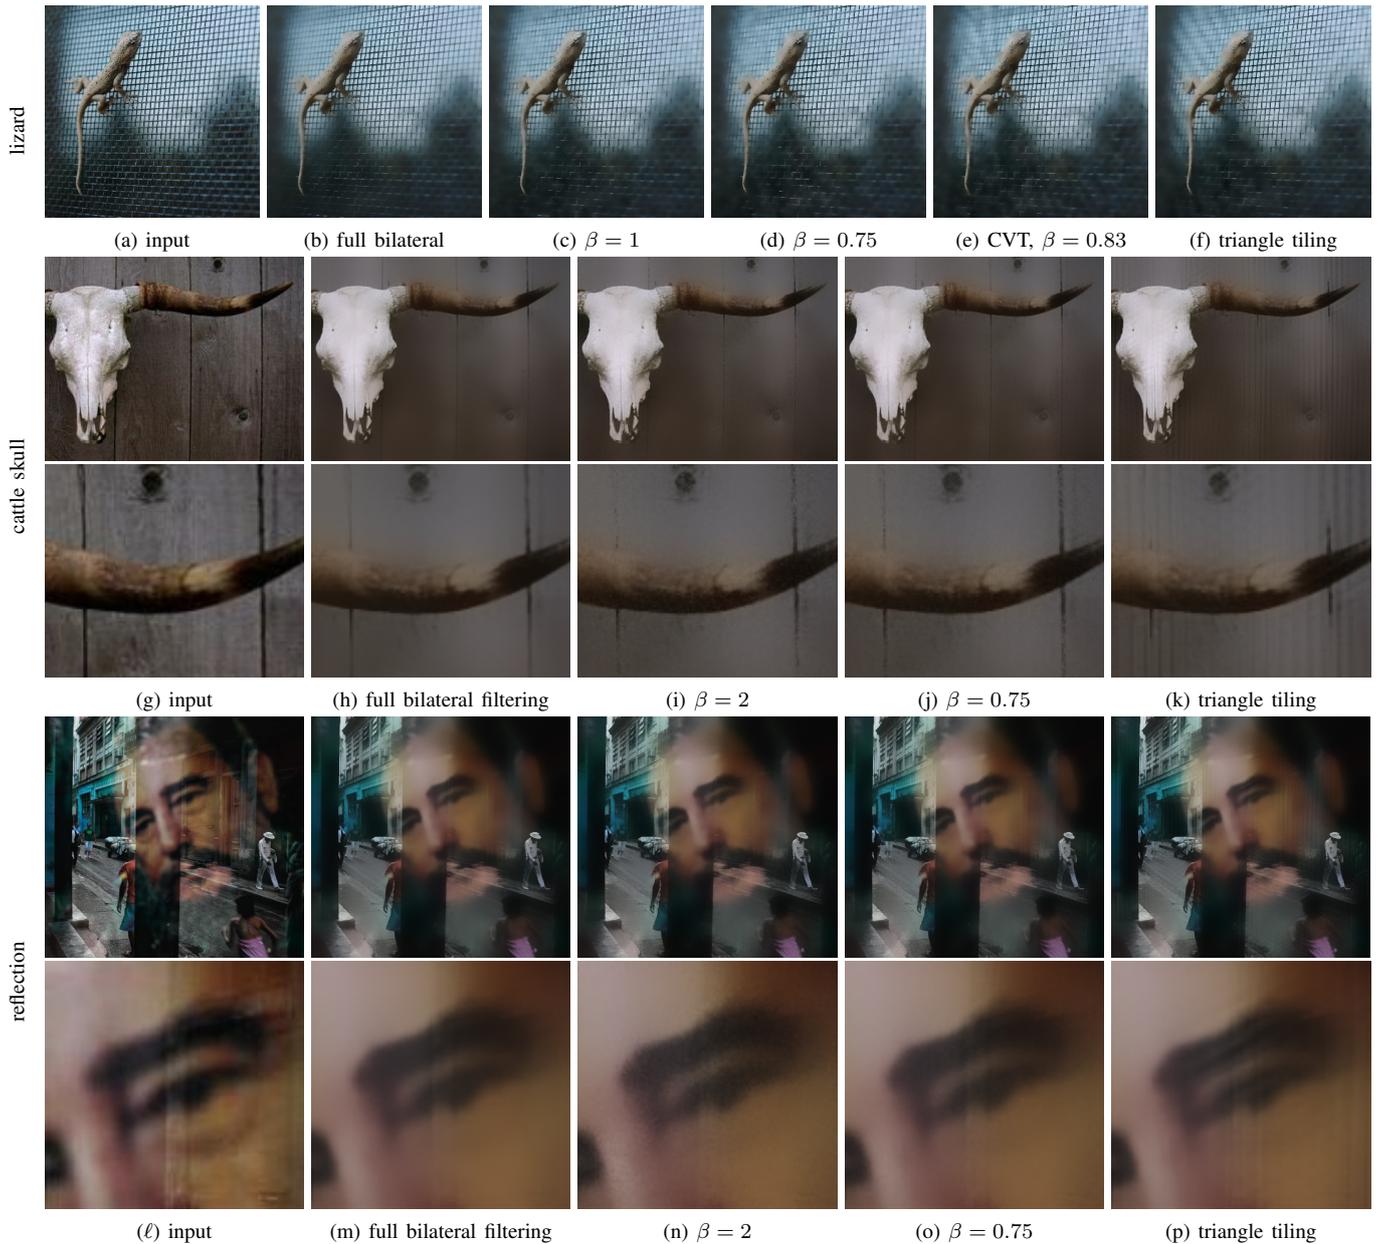


Fig. 16. Subsampling-accelerated bilateral filtering. $\beta = 0.75$ achieves the right balance between uniformity (reducing noise) and randomness (avoiding aliasing). Notice the noisier results with less uniform sampling ($\beta = 2.0$) and more aliasing with more regular sampling (CVT and triangle tiling). For the skull and reflection cases, we show both the whole images and partial zoom-ins. Table 2 lists the accuracy by β . All results are produced with a kernel width of 40 pixels and 30 samples per kernel, the recommended setting in [2].

and spectral properties of such adaptive/anisotropic distributions can be analyzed by the differential domain method [44] through local warping. It remains to formulate an optimization objective function.

Tiling. We currently compute entire sample sets. A potential extension is generating sample tiles [9,21,29] for acceleration.

Higher dimensions and meshing. We would like to address $d > 2$. Recall that MPS [15] can produce $\beta = 1$ in any dimension, but the lower limit achieved by the densest packing is dimension dependent. In this paper we focused on planar meshing. Volumetric and general curved surface meshing would be interesting and important extensions. Lower β provides better angles

in $d = 2$, and better radius-edge condition in $d \geq 2$, which are helpful for a variety of scientific and engineering applications.

Lower β . Our current algorithm can reach $\beta \approx 0.75$. We would like to investigate the theoretical potential of reaching even lower values, even though these distributions may have limited utility due to excessive regularity, as is evident from the hexagonal spectrum bias of the $\beta = 0.70$ case in Fig. 9. Also, the local optimization approach may require too much computation for small β .

Acknowledgements

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

- [1] Michael Balzer, Thomas Schlomer, and Oliver Deussen. Capacity-constrained point distributions: A variant of Lloyd's method. In *SIGGRAPH '09*, pages 86:1–8, 2009.
- [2] Francesco Banterle, Massimiliano Corsini, Paolo Cignoni, and R. Scopigno. A low-memory, straightforward and fast bilateral filter through subsampling in spatial domain. *Computer Graphics Forum*, 31(1):19–32, 2012.
- [3] Joseph E. Bishop. Simulating the pervasive fracture of materials and structures using randomly close packed Voronoi tessellations. *Comput. Mech.*, 44:455–471, 2009.
- [4] Joseph E. Bishop, Mario J. Martinez, and Pania Newell. A finite-element method for modeling fluid-pressure induced discrete-fracture propagation using random meshes. In *Proceedings of the 46th U.S. Rock Mechanics/Geomechanics Symposium*, June 2012. American Rock Mechanics Association, Document ID 2012-190.
- [5] John E. Bolander and Shigehiko Saito. Fracture analyses using spring networks with random geometry. *Eng. Fracture Mech.*, 61(5-6):569–591, 1998.
- [6] Lennart Bondesson and Jessica Fahn. Mean and variance of vacancy for hard-core disc processes and applications. *Scandinavian Journal of Statistics*, 30(4):pp. 797–816, 2003.
- [7] Renjie Chen and Craig Gotsman. Parallel blue-noise sampling by constrained farthest point optimization. *Computer Graphics Forum*, 31(5):1775–1785, 2012.
- [8] L. Paul Chew. Guaranteed-quality triangular meshes. Technical Report 89-983, Department of Computer Science, Cornell University, 1989.
- [9] Michael F. Cohen, Jonathan Shade, Stefan Hiller, and Oliver Deussen. Wang tiles for image and texture generation. In *SIGGRAPH '03*, pages 287–294, 2003.
- [10] Robert L. Cook. Stochastic sampling in computer graphics. *ACM Trans. Graph.*, 5(1):51–72, 1986.
- [11] Fernando de Goes, Katherine Breeden, Victor Ostromoukhov, and Mathieu Desbrun. Blue noise through optimal transport. In *SIGGRAPH Asia '12*, pages 171:1–171:11, 2012.
- [12] Qiang Du, Vance Faber, and Max Gunzburger. Centroidal Voronoi tessellations: Applications and algorithms. *SIAM Review*, 41(4):637–676, 1999.
- [13] Mohamed S. Ebeida and Scott A. Mitchell. Uniform random Voronoi meshes. In *20th Int. Meshing Roundtable*, pages 258–275, 2011.
- [14] Mohamed S. Ebeida, Scott A. Mitchell, Andrew A. Davidson, Anjul Patney, Patrick M. Knupp, and John D. Owens. Efficient and good Delaunay meshes from random points. *Computer-Aided Design*, 43(11):1506 – 1515, 2011. Solid and Physical Modeling 2011.
- [15] Mohamed S. Ebeida, Scott A. Mitchell, Anjul Patney, Andrew A. Davidson, and John D. Owens. A simple algorithm for maximal Poisson-disk sampling in high dimensions. *Comp. Graph. Forum*, 31(2pt4):785–794, 2012.
- [16] Mohamed S. Ebeida, Anjul Patney, Scott A. Mitchell, Andrew Davidson, Patrick M. Knupp, and John D. Owens. Efficient maximal Poisson-disk sampling. In *SIGGRAPH '11*, pages 49:1–12, 2011.
- [17] Raanan Fattal. Blue-noise point sampling using kernel density model. In *SIGGRAPH '11*, pages 48:1–12, 2011.
- [18] Peter Hansbo. Generalized Laplacian smoothing of unstructured grids. *Comm. in Numerical Methods in Engineering*, 11(5):455–464, 1995.
- [19] Ravi Jampani and Alper Üngör. Construction of sparse well-spaced point sets for quality tetrahedralizations. In *Int. Meshing Roundtable*, pages 63–80, 2008.
- [20] Patrick M. Knupp. Achieving finite element mesh quality via optimization of the Jacobian matrix norm and associated quantities. Part I: a framework for surface mesh optimization. *IJNME*, 48:401–420, 2000.
- [21] Johannes Kopf, Daniel Cohen-Or, Oliver Deussen, and Dani Lischinski. Recursive Wang tiles for real-time blue noise. In *SIGGRAPH '06*, pages 509–518, 2006.
- [22] Ares Lagae and Philip Dutré. A comparison of methods for generating Poisson disk distributions. *Comp. Graphics Forum*, 21(1):114–129, 2008.
- [23] Hongwei Li, Li-Yi Wei, Pedro Sander, and Chi-Wing Fu. Anisotropic blue noise sampling. In *SIGGRAPH Asia '10*, pages 167:1–12, 2010.
- [24] Yang Liu, Wenping Wang, Bruno Lévy, Feng Sun, Dong-Ming Yan, Lin Lu, and Chenglei Yang. On centroidal Voronoi tessellation — energy smoothness and fast computation. *ACM Transactions on Graphics*, 2009.
- [25] Stuart P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [26] Gary L. Miller, Donald R. Sheehy, and Ameya Velingker. A fast algorithm for well-spaced points and approximate Delaunay graphs. In *Symp. on Computational Geometry*, SoCG '13, pages 289–298, 2013.
- [27] Scott A. Mitchell, Alexander Rand, Mohamed S. Ebeida, and Chadrajit Bajaj. Variable radii Poisson-disk sampling, extended version. In *Canadian Conference on Computational Geometry*, pages 1–9, 2012.
- [28] John A. Nelder and Roger Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, 1965.
- [29] Victor Ostromoukhov. Sampling with polyominoes. In *SIGGRAPH '07*, page 78, 2007.
- [30] Victor Ostromoukhov, Charles Donohue, and Pierre-Marc Jodoin. Fast hierarchical importance sampling with blue noise properties. In *SIGGRAPH '04*, pages 488–495, 2004.
- [31] A. Cengiz Öztireli, Marc Alexa, and Markus Gross. Spectral sampling of manifolds. In *SIGGRAPH ASIA '10*, pages 168:1–8, 2010.
- [32] Sylvain Paris, Pierre Kornprobst, Jack Tumblin, and Frédo Durand. Bilateral filtering: Theory and applications. In *Foundations and Trends in Computer Graphics and Vision*, 2009.
- [33] Per-Olof Persson and Gilbert Strang. A simple mesh generator in MATLAB. *SIAM review*, 46(2):329–345, 2004.
- [34] Julian J. Rimoli and Juan J. Rojas. Meshing strategies for the alleviation of mesh-induced effects in cohesive element models. *arXiv:1302.1161 [physics.comp-ph]*, pages 1–19, Feb 2013.
- [35] Hagit Schechter and Robert Bridson. Ghost SPH for animating water. *ACM Trans. Graph.*, 31(4):61:1–61:8, 2012.
- [36] Thomas Schlömer. PSA point set analysis. <http://code.google.com/p/psa/>, 2011.
- [37] Thomas Schlömer, Daniel Heck, and Oliver Deussen. Farthest-point optimized point sets with maximized minimum distance. In *HPG '11*, pages 135–142, 2011.
- [38] Jonathan R. Shewchuk. Triangle: Engineering a 2d quality mesh generator and Delaunay triangulator. *Applied Computational Geometry Towards Geometric Engineering*, pages 203–222, 1996.
- [39] Peter Shirley. Discrepancy as a quality measure for sample distributions. In *Eurographics '91*, pages 183–194, 1991.
- [40] Dafna Talmor. *Well-Spaced Points for Numerical Methods*. PhD thesis, Carnegie Mellon University, 1997. Tech Report CMU-CS-97-164.
- [41] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Int. Conf. on Computer Vision*, pages 839–846, 1998.
- [42] Greg Turk. Re-tiling polygonal surfaces. In *SIGGRAPH '92*, pages 55–64, 1992.
- [43] Li-Yi Wei. Parallel Poisson disk sampling. In *SIGGRAPH '08*, pages 20:1–9, 2008.
- [44] Li-Yi Wei and Rui Wang. Differential domain analysis for non-uniform sampling. In *SIGGRAPH '11*, pages 50:1–10, 2011.