SAND2012-1153C

Variable Radii Poisson-Disk Sampling

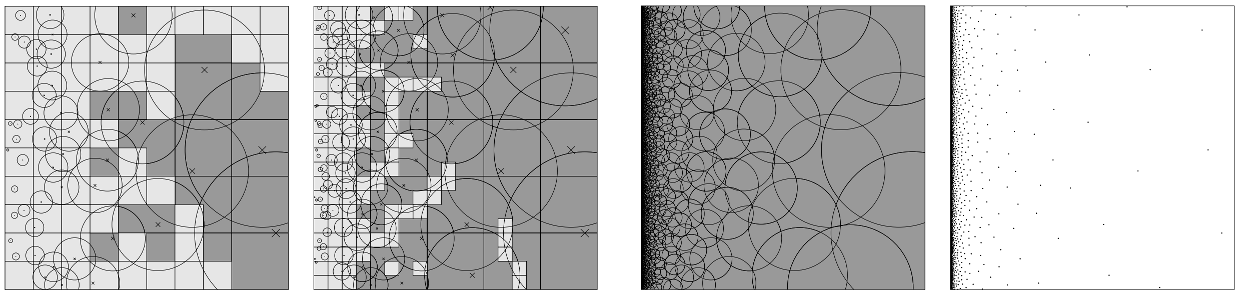

Figure 1: *Sampling using prior-disks. Right, the first two iterations, showing flat quadtree refinement and active squares (light). Left, the output disks and samples. The sizing function spans three orders of magnitude over the unit box:* $r(x, y) = 0.001 + 0.3x$.

Abstract

We introduce two natural and well-defined generalizations to the definition of the Poisson-disk sampling problem. The first is to decouple the disk-free (inhibition) radius from the maximality (coverage) radius. By scaling these radii by an abstract parameter (e.g. time), we may generate hierarchical samples with more randomness than if a single radius is used. The radial power of the FFT is more uniform than for classical MPS: the oscillating ring pattern is attenuated.

The second generalization is to allow the radii to vary spatially, according to a sizing function. Our main contributions there are a formal characterization of sizing functions (radii) for non-uniform Poisson-disk samples and a generic algorithm for sample generation for a wide variety of applications. These results hold in all dimensions. We contrast the results to Delaunay refinement. Our definitions and algorithms do not depend on a maximum and minimum radii, but rather the rate at which the radii can change. This rate fundamentally determines the quality of the resulting point cloud. It provides bounds on the distances to neighboring points: specifically bounds on the ratio of lengths of edges sharing a sample in a Delaunay triangulation of the sampling. We provide experimental results.

CR Categories: I.3.5 [Computing Methodologies]: Computer Graphics—Computational Geometry and Object Modeling

Keywords: Poisson-disk, sampling, variable radius, spatial variation, inhibition radius, hierarchical sampling

1 Classic MPS Definition

A *sampling* is a set of ordered points taken from a domain at random. In Poisson-disk sampling, each point has an associated disk. No other point may be inside this disk. Points are chosen uniformly outside the prior points' disks. The sampling is maximal if the entire domain is covered by disks. Together these define maximal Poisson-disk sampling (MPS), a.k.a. the Matérn second process [1960].

More formally, a sampling $X = (\mathbf{x}_i)_{i=1}^n$, $\mathbf{x}_i \in \Omega$ satisfies the *inhibition* or *empty disk* property if

$$\forall i < j \leq n, |\mathbf{x}_i - \mathbf{x}_j| \geq r. \tag{1}$$

The set of *uncovered points* is defined to be

$$S(X) = \{\mathbf{y} \in \Omega : |\mathbf{y} - \mathbf{x}_i| \geq r,\ i = 1..n\}. \tag{2}$$

A sampling X is *maximal* if $S(X)$ is empty:

$$S(X) = \emptyset. \tag{3}$$

Given a non-maximal sampling, the next sample is *bias-free* if the probability of selecting it from any uncovered subregion is proportional to the subregion's area, i.e.,

$$\forall A \subset S(X) : P(\mathbf{x}_{n+1} \in A \,|\, X) = \frac{|A|}{|S(X)|}, \tag{4}$$

where $|\cdot|$ denotes area. The sampling *process* is bias-free if all of its sample points were selected according to a bias free random process. This criteria is in contrast to measuring the *output* of one run of the process, e.g. the FFT spectrum of the pairwise distances between the points [Schlömer 2011]. Some algorithms are highly parallel and fast [Wei 2008], but follow an inherently biased process. Whether this makes a difference for applications is unclear, as definitive requirements and measures of whether point sets meet these requirements are not generally available. Because of this, and to have a clear frame of reference, we focus on bias-free processes.

In the following we will generalize these equations and elucidate the consequences. For simplicity our figures and language focus on two dimensional domains, but the definitions and algorithms are general dimensional. We consider three specific generalizations of the MPS problem: (1a) decoupling the radii in the disk-free and maximality conditions, (1b) a hierarchical constructions for progressively denser samples, and (2) sampling with spatially varying point cloud density based on a variable sizing function. These generalizations are necessary to utilize the benefits of uniform MPS for problems in computer graphics and mesh generation.

We define sizing functions over the domain that determine the radius of the Poisson-disks. A sizing function f is a simple function (e.g. a local max) of an underlying radius function, r, or $r(p)$ to emphasize its dependence on the position p. We require that r is L-Lipschitz, i.e., for all $\mathbf{x}, \mathbf{y} \in \Omega$, $|r(\mathbf{x}) - r(\mathbf{y})| \leq L\,|\mathbf{x} - \mathbf{y}|$ for some constant L. The field of mesh generation, notably Delaunay refinement, has a long history of considering sizing functions with $L < 1$. In particular, for polygonal domains, the *local feature size* is the radius of the smallest ball that contains two disjoint faces of the boundary, e.g. two edges that do not share a common vertex. The local feature size has $L < 1$. Delaunay refinement can be made to follow any sizing function less than the local feature size that also has $L < 1$.

For clarity and ease of language, we will call the disk-free radius the *inhibition radius* R_f; and the maximality radius the *coverage radius* R_c throughout. As a function of another parameter, we use r_f and r_c.

2 Motivation and Previous Work

Maximal Poisson-disk sampling (MPS) is a popular topic in computer graphics [Lagae and Dutré 2008]. MPS is used for texture generation. The random nature of the point cloud avoids visual artifacts that would arise if the distances between points had repeating patterns. Given a fixed budget of points, inhibition disks and maximality help to use it efficiently: they prevent points from being too close together while ensuring that sample points lie throughout the entire domain. Uniform sampling with fixed radius disks are traditional. However, variable-density samplings have several uses.

In adaptive level-of-detail renderings, we wish to use a finer sample locally as the camera zooms closer to an object. Often one switches between a discrete series of prescribed samples as the camera-to-object distance crosses some thresholds. One must avoiding visible artifacts at these steps. We suggest that our hierarchical sampling is well suited to this application, especially the continuous variation which adds one random point at a time. In real-time applications such as games or data exploration [Ljung 2006], this approach avoids scene jumps (frame coherence) due to lengthy computations or memory fetches [Vanderhaeghe et al. 2007].

Spatially varying samplings are useful for objects with both sharp curvature and large flat regions, or other non-constant visualization gradients [Kopf et al. 2006; Bowers et al. 2010].

Because of these potential uses, computer graphics publications have occasionally extended their MPS methods to spatially varying inhibition/maximality radii. Often these have been presented as an extension to a result with a different focus, and there is little description of the underlying definitions and algorithm requirements: e.g. how quickly the point density is allowed to vary. Sometimes only a picture of an example output is given.

Wei [2008] provides adaptive sampling. A quadtree is refined proportional to the local sizing function. A set of distant squares are sampled from concurrently. Sets are chosen hierarchically and randomly to balance speed with the measurable bias in the output. The distribution is non-maximal but with some probability approximates the local density. This is extended [Bowers et al. 2010] to sampling triangulated surfaces. The sampling can be non-uniform, but the spatial sizing function must stay close to its maximum value. (These works consider symmetric conflict conditions akin to our smaller-radius and bigger-radius conditions.)

In relaxation dart throwing [McCool and Fiume 1992], an initial MPS has the sampling disks radii reduced by a scaling parameter. The parts of the domain that are no longer maximal are filled in with samples. (In the original description, the sampling is not actually maximal, and the scaling occurs when classical dart throwing has a high miss rate.) A variation [Vanderhaeghe et al. 2007] for deforming point clouds coarsens to remove points that are too close together, and refines to re-achieve maximality. For coarsening the disk-free and maximal criteria hold approximately, subject to a tolerance band. Blending and sliding heuristics try to minimize the visual artifacts that arise from the discrete stages of the deformation.

A time-evolving sizing function defined at sample points, extended to a 1-Lipschitz function over the entire domain, can be used to both refine and coarsen a mesh of a sphere packing with overlap tolerances [Li et al. 1998]. The refinement/coarsening can be done deterministically, for example by Delaunay refinement. It is very fast to generate importance samplings using deterministic methods. A hierarchy of refinements based on aperiodic tiling gives point clouds with reasonable blue-noise properties [Ostromoukhov et al. 2004; Kopf et al. 2006; Ostromoukhov 2007]. Given a non-uniform sampling, Wei [2011] provides a way to measure it.

In many MPS algorithms, generating a sample in small uncovered areas to achieve maximality is difficult. Quadtree [Gamito and Maddock 2009; White et al. 2007; Ebeida et al. 2012] based methods might have to refine their squares down to numerical precision to represent them. The coverage radius is often relaxed by an epsilon value. This is different than decoupling the radii as we do here.

Variable-radii samplings also appear in other fields. In physics, in random sequential absorption [Dickman et al. 1991], the Poisson-disk radius is the Van der Waals radius, which is different for different types of atoms. Physicists have various recipes for generating point clouds, depending on the desired density. For example, some recipes produce point clouds that resemble gas, liquid, glass, and solid states. However, Van der Waals radii vary little compared to the orders of magnitudes we address here. In environmental sciences, individual organisms can be modeled as sample points, with disks corresponding to their territory [Renshaw 2010].

In mesh generation, maximal samplings satisfying the empty disk property are commonly desired since they yield provably good quality Delaunay triangulations [Chew 1989; Miller et al. 1996; Ebeida et al. 2011a]. Delaunay refinement algorithms construct a maximal sample by incrementally adding circumcenters of Delaunay triangles [Chew 1989; Ruppert 1995] and without considering the entire set of uncovered points. When Delauany refinement is applied to generate quality graded meshes [Ruppert 1995; Shewchuk 2002], a user defined mesh sizing function must satisfy a Lipschitz property and be dominated by the local feature size. Beyond considering only circumcenters, the set of potential new vertices can be expanded to produce meshes with larger angles or fewer vertices [Chew 1997; Li 2003; Chernikov and Chrisochoides 2009; Erten and Üngör 2009]. However, randomly sampling the entire uncovered set (i.e., producing unbiased samples) has received less attention due to the difficulty and cost of generating suitable point sets, and analysis. Random meshes with time and spatially varying densities are preferred for certain fracture mechanics, because cracks propagate in more physically realistic directions [Bolander and Saito 1998; Ebeida and Mitchell 2011].

A hierarchy of meshes has many uses [Miller et al. 1996; Li et al. 1998; Devillers 2002]: point location, multigrid, coarsening and refinement in adaptivity, and convergence studies.

3 Different inhibition and coverage radii

Here we relax the condition that the coverage and inhibition radii are equal. We focus on a particular relaxation that proves useful when generating hierarchical point sets and smoothing the FFT radial power. Contrast to Equations 1– 4. Let $R_f \leq R_c$.

The *empty disk* property is

$$\forall i < j \leq n, |\mathbf{x}_i - \mathbf{x}_j| \geq R_f. \quad (5)$$

The set of *free* points is defined to be

$$S(X) = \{\mathbf{y} \in \Omega : |\mathbf{y} - \mathbf{x}_i| \geq R_f,\, i = 1..n\}. \quad (6)$$

The set of *uncovered* points is defined to be

$$U(X) = \{\mathbf{y} \in \Omega : |\mathbf{y} - \mathbf{x}_i| \geq R_c,\ i = 1..n\}. \quad (7)$$

In order for this variation to be useful and different than the single radius case, the sampling is *maximal* if $U(X)$ is empty,

$$U(X) = \emptyset. \quad (8)$$

We take samples from S, but restrict to the points that are close enough to U in order to reduce it.

$$T(X) = S(X) \cap \{U(X) + R_c\}. \quad (9)$$

We call the sampling *bias-free* if we sample from $T(X)$ uniformly.

This process is useful to add randomness to initial and parameterized (hierarchical) samples. This process provides samplings that are less uniform, with greater variation in inter-sample distances, than classical MPS. In particular, this process avoids the visible rings in the FFT spectrum of the output. The radial power does not have the low frequency oscillations that are so characteristic of classical MPS. See Figure 7 for an example.

Samplings will likely have points that could be removed and still meet the coverage condition Equation 8. There are more extra points the smaller R_f is compared to R_c.

3.1 Algorithm for Two Radii Sampling, and Other Variations

Algorithm 1 MPS using generic conflict and coverage tests.

initialize kd-tree $\mathcal{T} = \emptyset$, $i = 0$, $\mathcal{C}^i = \mathcal{C}^o$
while $|\mathcal{C}^i| > 0$ **do**
  {throw darts}
  for all $A|\mathcal{C}^i|$ (constant) dart throws **do**
    select an active cell $\mathcal{C}^i_c$ from $\mathcal{C}^i$ uniformly at random
    throw candidate dart $\mathbf{y}$ into $\mathcal{C}^i_c$, uniformly at random
    if $\mathbf{y}$ does not conflict **then**
      {promote dart to sample}
      add $\mathbf{y}$ to $\mathcal{T}$ as an accepted sample $\mathbf{x}$
    end if
  end for
  {coverage}
  for all active cells $\mathcal{C}^i$ **do**
    if $i <$bits_of_precision subdivide $\mathcal{C}^i_c$ into 2^d subcells
    retain uncovered (sub)cells as $\mathcal{C}^{i+1}$
  end for
  increment i
end while

The basic algorithm is simple and a variation of the MPS algorithm in [Ebeida et al. 2012]; see Algorithm 1. A batch of darts are thrown, and each one checks whether it *conflicts* with any nearby sample. We store samples in a kd-tree $\mathcal{T}$, so that we can retrieve those that are nearby. A flat quadtree $\mathcal{C}$ tracks the uncovered area where the next sample may arrive. At a given iteration i, all the cells $\mathcal{C}^i$ have been refined to the same size. After a batch of throws, squares are refined and discarded if they are *covered*, i.e., too close to a prior sample to contain a new one. (Prior works store the samples in the quadtree, but this is not as efficient for us because of our variable radii.)

The batch size is the number of quadtree cells times A, an empirically derived constant depending on the dimension of the domain and algorithm variation. The exact value of A is unimportant as long as it is above a threshold. $A \approx 0.5$ for $d < 5$. Uncovered cells are represented by their indices, which are stored in an array; covered cells have no data whatsoever. This is the key to the memory-efficiency of the method [Ebeida et al. 2012]. The squares $\mathcal{C}^i_c$ are checked for coverage by prior samples. The algorithm terminates when no squares are uncovered.

We use this algorithm outline for all variations in this paper. The differences are the conditions for the conflict and coverage checks.

Specialization for Two Radii The pool of active quadtree squares $\mathcal{C}^i$ at iteration i are an outer approximation to our sampling set T. For determining if a square $\mathcal{C}^i_c$ is covered, we check whether its corners are inside some $D(\mathbf{x}_i, R_f)$, the d-dimensional disk (ball) centered at $\mathbf{x}_i$ of radius R_f. Such squares have all their points outside S and are discarded.

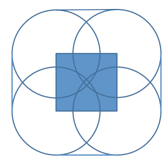

Figure 2: *The outer perimeter encompasses the points of T that a sample in a square might cover.*

In addition, we wish to discard squares $\mathcal{C}^i_c$ outside $U + R_c$. The set of points that a sample in a square $\mathcal{C}^i_c$ might cover is a larger, rounded square $V = \mathcal{C}^i_c + D(R_c)$ as in Figure 2. If V is already covered by samples' disks, the cell may be discarded. Since V is large it cannot be covered by a single sample's disk. To check if it is covered by a collection of disks we subdivide it and check the sub-pieces. We take its bounding box, and subdivide it as in a flat quadtree. All boxes at a given level are checked as in the following paragraph. The checks may decide that $\mathcal{C}^i_c$ (likely) contains a point of T. The checks may discard some boxes. If no boxes remain, then V does not contain any point of T and $\mathcal{C}^i_c$ may be discarded. If the situation is ambiguous, we subdivide all the boxes and recheck.

If a box has all four corners outside V, we discard it. If a box has some corners inside V and outside V, it is ambiguous. If a box has all four corners inside V, then there are three subcases. (1) If any corner is outside all nearby sample's disks, then we decide that the box contains a point of T; such points might actually be inside the R_f disks, but a false positive is acceptable. (2) If all corners are inside a single sample's disk, then discard the box. (3) Else it is ambiguous and the box has to be subdivided to decide.

(Checking V is similar to the BiX box-in-disk check in Section 6. An alternative that may be faster, but use more memory, is to store a flat quadtree that represents T. Keep this quadtree at the same level i as $\mathcal{C}^i$, and check its cells as described above.)

We select darts by selecting a square in the pool then selecting a point in the square.

For a dart $\mathbf{y}$ to be conflict free, we require

1. $\mathbf{y} \notin D(\mathbf{x}_i, R_f)$ $\forall$ prior sample $\mathbf{x}_i$.
2. The dart's large disk $D(\mathbf{y}, R_c)$ must cover an uncovered point.

Checking Condition 1 is standard using the kd-tree of the samples. To check Condition 2 we have two options; their relative efficiency depends on the ratio of R_f to R_c, see Figure 3.

First option. Form a local bounding box of $D(\mathbf{y}, R_c)$. Refine it as a quadtree: discard squares covered by $D(\mathbf{x}_i, R_c)$ or outside $D(\mathbf{y}, R_c)$, otherwise refine them. If any refined square corner is discovered that is outside all $D(\mathbf{x}, R_c)$ and inside $D(\mathbf{y}, R_c)$, accept the dart. Reject the dart if all squares are discarded. This test is accurate up to the roundoff error desired.

Second option. Throw a constant number of random points into $D(\mathbf{y}, R_c)$ and accept the dart if any are outside $D(\mathbf{x}_i, R_c)\ \forall i$. This works well if $R_f \ll R_c$, but introduces some bias.

4 Hierarchical Sampling

4.1 Parameterized radii

Consider a maximal sampling, from either a single disk radius or decoupled inhibition and coverage radii. We parameterize these radii by a scaling parameter t; e.g., t could be time. For simplicity we consider only linear scaling, so that for any particular value of t we have $r_f(t) = tR_f$ and $r_c(t) = tR_c$. We now consider constructing a family of samplings over this parameter.

4.2 Continuous Decrease Refinement

Consider decreasing t continuously from 1 to 0. The sampling becomes non-maximal when $U(X) \neq \emptyset$; recall Equation 8. A new sample is needed. Assume for simplicity that the t when this occurs are distinct, so that $U(X)$ grows by a single point u. This position u is the circumcenter of a Delaunay sphere through some nearby samples.

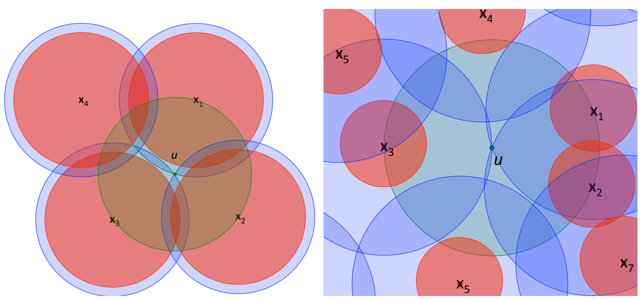

(a) $r_c \approx r_f$ (b) $r_c \ll r_f$.

Figure 3: *Possible void shapes for two radii. u is the circumcenter of $\triangle \boldsymbol{x}_1\boldsymbol{x}_2\boldsymbol{x}_3$. The circumcircle is green, r_f disks are red, r_c disks are blue.*

If $r_f = r_c$ then there is only one place to put the sample, at u, so the process is deterministic.

Otherwise, we insert a random point. We have two simple and efficient solutions, depending on the relative value of r_f and r_c. if r_f is much less than r_c, we sample uniformly from the sphere of radius $r_c = tR_c$ centered at u. If the sample point is closer than r_f to a nearby point, we resample. The sphere of radius $r_c - r_f$ centered at u is free and all samples from it will be a hit. Therefore since $r_f \ll r_c$ the probability of a hit is high and this scheme is efficient enough.

If r_f is close to r_c, then to achieve efficiency the free region may be (approximately) constructed using polygons [Ebeida et al. 2011b] or a quadtree [Ebeida et al. 2012; Gamito and Maddock 2009]. See Figure 3 for examples void shapes.

In 2d, we observe that u will always be either the circumcenter of a non-obtuse Delaunay triangle or inside a boundary Delaunay triangle. If u were interior and obtuse, then there is a Delaunay triangle on the other side of the longest edge, and it will have a larger circumsphere, and so its center would be uncovered for a smaller value of t. (The circumcenter of a non-obtuse triangle lies inside the triangle.)

In any case, our refinement process can be implemented as a discrete event simulation. The circumcenters of Delaunay triangles, and their associated t values, are the events. When an event occurs, a sample is generated, which creates new triangles and destroys some old ones, so the queue must be updated. This is essentially the generic Delaunay refinement algorithm with a largest-first queue priority for inserting circumcenters. The main difference is that when an event occurs, we insert a nearby random point, but DR inserts the point itself. The usual analysis of Delaunay refinement makes no restrictions on the circumcenter insertion order, and the Triangle code [Shewchuk 2002] takes the opposite approach: processing the smallest triangles first.

4.3 Discrete Decrease Refinement

Consider decreasing t in discrete jumps. For a new value of t, the sample will be non-maximal, and the same algorithm that generated the initial sampling can be continued to achieve maximality.

Comparison The prior approaches [McCool and Fiume 1992; Li et al. 1998] achieve maximality and disk-free up to some wide tolerance band. In contrast, if our inhibition and coverage radii are the same, then we achieve the maximality and disk-free conditions exactly. If two radii are used, then these take the place of the tolerance band. In a sense the effective tolerance band can be tuned by their ratio.

4.4 Edge Length and DT Angle Bounds

We consider a Delaunay triangulation (DT) of our point cloud. Two samples are Delaunay neighbors if they share an edge e in a DT. The inhibition radius bounds the shortest edge length. The coverage radius bounds the largest empty Delaunay circumcircle. The longest edge length is at most the diameter of that circle. To summarize:
Proposition 4.1. $|e| \in [R_f, 2R]$ *and* $R \leq R_c$, *where* R *is the radius of a Delaunay circumcircle.*

The Central Angle Theorem provides a relation between the smallest angle α in a triangle to a lower bound on the shortest edge length $|e|$ and an upper bound on the circumradius R of the triangle. This relation is fundamental to the inception of Delaunay refinement [Chew 1989].
Proposition 4.2. $\sin\alpha \geq |e|/2R$.

For example, in a DT of a point set with $R_c = R_f$, we have $\alpha > 30°$. If $R_c = 2R_f$, then $\alpha > 14.4°$.

5 Spatially Varying Radii

We aim to produce spatially varying point density according to a sizing function $r(\mathbf{x}) : \Omega \to (0, \infty)$. A sample satisfies the *empty disk* property (vs. Equation 1) if

$$\forall i < j \leq n, |\mathbf{x}_i - \mathbf{x}_j| \geq f(\mathbf{x}_i, \mathbf{x}_j), \tag{10}$$

and the set of *uncovered points* is (vs. Equation 2)

$$S(X) = \{\mathbf{y} \in \Omega : |\mathbf{y} - \mathbf{x}_i| > f(\mathbf{x}_i, \mathbf{y}),\ i = 1..n\}. \tag{11}$$

Here $f(\mathbf{x}_i, \mathbf{y})$ is a function of $r(\cdot)$ evaluated at a previously accepted sample and a later sample or candidate. We have four criteria vari-

Method	Distance Function	Order Independent	Full Coverage	Conflict Free	Edge Min	Edge Max	Sin Angle Min	Max L
Prior	$r(\mathbf{x})$	no	no	no	$1/(1+L)$	$2/(1-2L)$	$(1-2L)/2$	$1/2$
Current	$r(\mathbf{y})$	no	no	no	$1/(1+L)$	$2/(1-L)$	$(1-L)/2$	1
Bigger	$\max(r(\mathbf{x}), r(\mathbf{y}))$	yes	no	yes	1	$2/(1-2L)$	$(1-2L)/2$	$1/2$
Smaller	$\min(r(\mathbf{x}), r(\mathbf{y}))$	yes	yes	no	$1/(1+L)$	$2/(1-L)$	$(1-L)/2$	1

Table 1: *Summary of results for spatially varying radii. The distance function f determines conflicts. Order independence means that if a sampling X satisfies the empty disk property, then so do permutations of X. Full coverage means that every point of the domain is inside a sample's disk. Conflict free means that no sample is inside another sample's disk. Edge max and min bound the possible lengths of an edge containing sample $\boldsymbol{x}$ in a Delaunay triangulation of X, as a factor of $r(\boldsymbol{x})$. Max L is the largest Lipschitz constant for which the algorithm is guaranteed to be robust and produce correct output.*

ations:

$$f(\mathbf{x}, \mathbf{y}) := r(\mathbf{x}) \qquad \textbf{(Prior-disks)},$$
$$f(\mathbf{x}, \mathbf{y}) := r(\mathbf{y}) \qquad \textbf{(Current-disks)},$$
$$f(\mathbf{x}, \mathbf{y}) := \max(r(\mathbf{x}), r(\mathbf{y})) \qquad \textbf{(Bigger-disks)},$$
$$f(\mathbf{x}, \mathbf{y}) := \min(r(\mathbf{x}), r(\mathbf{y})) \qquad \textbf{(Smaller-disks)}.$$

The f are equivalent for a fixed sampling radius r, but are all distinct for spatially-varying radii. Each approach has certain advantages in terms of edge length ratios, order independence, how quickly the sizing function may vary, and simplicity of implementation; see Table 1 for a summary.

There is a limit to how quickly $r(\cdot)$ is allowed to vary. We require that r is L-Lipschitz, i.e., for all $\mathbf{x}, \mathbf{y} \in \Omega$, $|r(\mathbf{x}) - r(\mathbf{y})| \leq L\,|\mathbf{x} - \mathbf{y}|$ for some constant L. Some approaches require $L < 1$; for others $L < 1/2$. In all cases, as L approaches zero, the quality guarantees smoothly approach those in the uniform case. Appendix A provides the proofs for the different cases.

For constant radii, edges in a Delaunay triangulation (DT) of the sampling are bounded between r and $2r$. (If edges are less than r, the inhibition distance is violated; if greater than $2r$, then the sampling is not maximal.) For spatially varying radii, the length of edge $\overline{\mathbf{x}_i\mathbf{x}_j}$ could be a smaller or larger fraction of $r(\mathbf{x}_i)$, because the strategies also depend on $r(\mathbf{x}_j)$. How much $r(\mathbf{x}_j)$ may differ depends on L.

By symmetry the bigger-disk and smaller-disk constructions are order independent, i.e., any valid sampling with the order of samples permuted still satisfies the empty disk property.

Bias-free The standard bias-free definition can be used for non-uniform sampling and is the basis of all the algorithms we have implemented. However, an alternative is to locally weight the uncovered set based on the sizing function, i.e., the desired output density. Specifically, the weight of a region can be defined by,

$$w(S) = \int_S \frac{1}{r(\mathbf{x})^d}\,\mathrm{d}\mathbf{x},$$

where d is the spatial dimension. Discrete approximations are possible. Then the "weighted-bias-free" property is:

$$\forall A \subset S(X) : P(\mathbf{x}_{n+1} \in A \,|\, X) = \frac{w(A)}{w(S(X))}. \qquad (12)$$

6 Spatially Varying Radii Algorithms

Again we use Algorithm 1, with the following nuances.

The diagonal-length of the base quadtree level is proportional to the maximum sampling radius over the domain, or the bounding box of the entire domain if there is no known maximum. Since the radii vary spatially, we can not efficiently use a uniform base grid for finding nearby samples for checking for conflicts as is common [Gamito and Maddock 2009; Ebeida et al. 2011b; Ebeida et al. 2012]. Instead we use a kd-tree $\mathcal{T}$ to store and find nearby samples [Mount and Arya 2010]. When searching the tree, the triangle inequality for Euclidean distance is used to determine if both branches could contain conflicting samples or only the near branch needs to be plumbed. This same tree and the triangle quality is also used when checking if cells are covered.

A variation would be to use a full quadtree as in [Gamito and Maddock 2009]. Samples would be stored in the quadtree at a depth related to their radii. Finding potentially conflicting samples would involve walking the quadtree. We have not implemented this variation, as we expect its memory requirements to be larger.

6.1 Primitives

The conflict and coverage checks rely on proximity primitives. Let $D(p, r)$ be the disk centered at p with radius r. Here $D(\mathbf{x})$ is the disk (ball) of radius $r(\mathbf{x})$ centered $\mathbf{x}$.

- **PiX point-in-disk, conflict.** Is a point in a sample's disk? Given p, $?\exists \mathbf{x} : p \in D(\mathbf{x})$?
- **XiP sample-in-disk, conflict.** Does a point's disk contain a sample? Given p, $?\exists \mathbf{x} : \mathbf{x} \in D(p)$?
- **BiX box-in-disk, coverage.** Is a square inside a sample's disk? Given cell $\mathcal{C}_c^i$, $?\exists \mathbf{x} : p \in D(\mathbf{x})\ \forall p \in \mathcal{C}_c^i$?
- **XiB sample-in-box's-disks, coverage.** Do all the disks of a square contain a sample? Given cell $\mathcal{C}_c^i$, $?\exists \mathbf{x} : \mathbf{x} \in D(p)\ \forall p \in \mathcal{C}_c^i$?

6.2 Conflict and Coverage Tests

Which proximity primitives are relevant to the conflict and coverage checks depends on the variation.

Method	Conflict	Coverage
Prior-disks	PiX	BiX
Current-disk	XiP	XiB
Bigger-disk	PiX or XiP	BiX or XiB
Smaller-disk	PiX and XiP	BiX and XiB

Primitive	"Yes" if $\exists \mathbf{x}$:	Prune Branch if
PiX point-in-disk	$\lvert p-\mathbf{x}\rvert \le r(\mathbf{x})$	$\lvert p-p_\perp\rvert^2(1-L^2) > r(p_\perp)^2$
XiP sample-in-disk	$\lvert p-\mathbf{x}\rvert \le r(p)$	$\lvert p-p_\perp\rvert > r(p)$
BiX box-in-disk	$\lvert c-\mathbf{x}\rvert + h/2 \le r(\mathbf{x})$	$(\lvert c-c_\perp\rvert + s/2)^2(1-L^2) > r(c_\perp)^2$
XiB sample-in-square's-disk	$\lvert c-\mathbf{x}\rvert \le r(c)-(1+L)h/2$	$\lvert c-c_\perp\rvert > r(c)-(1+L)h/2$

Table 2: *Conflict and coverage primitive criteria. Answer "yes" if such a sample is found. Prune (do not search) branches of the kd-tree where the prune branch condition holds. For a square, c is its center, s its side length and h its diagonal length. Here $p_\perp$ is the projection of p onto the hyperplane subdividing the two branches of the kd-tree.*

6.3 Primitive Implementation

Since the radii vary, it is insufficient to find the nearest sample(s) to the query point. Instead, we search the kd-tree, pursuing branches depending on whether the triangle inequality and the Lipschitz condition indicate that the branch may contain a close enough sample. We project the point p onto the dividing hyperplane of $\mathbf{x}_i$ in $\mathcal{T}$, obtaining $p_\perp$. If L and either $r(p_\perp)$ or $r(p)$ are sufficiently small compared to $|p-p_\perp|$, then the far branch of the kd-tree is pruned. For the conflict primitives, we must search any branch that might contain a conflicting sample. For the coverage primitives we can pursue branches less aggressively, because the algorithm will still be correct, albeit less efficient, if we do not detect that a square is covered.

Table 2 summarizes the branching conditions. Appendix B provides the proofs.

7 Experimental Results

Our first example is for different inhibition and coverage radii, extended to a hierarchy, as developed in Sections 3 and 4. Figure 4 shows a very coarse sampling with two radii. Figure 5 extends it to a discrete hierarchy of samplings. Observe that the new samples are sometimes inside the covered region, but nonetheless reduce the uncovered region.

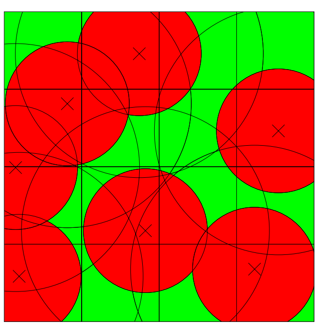

Figure 4: *Two-radii MPS, $\sqrt{2}/4 = r_f = r_c/2$.*

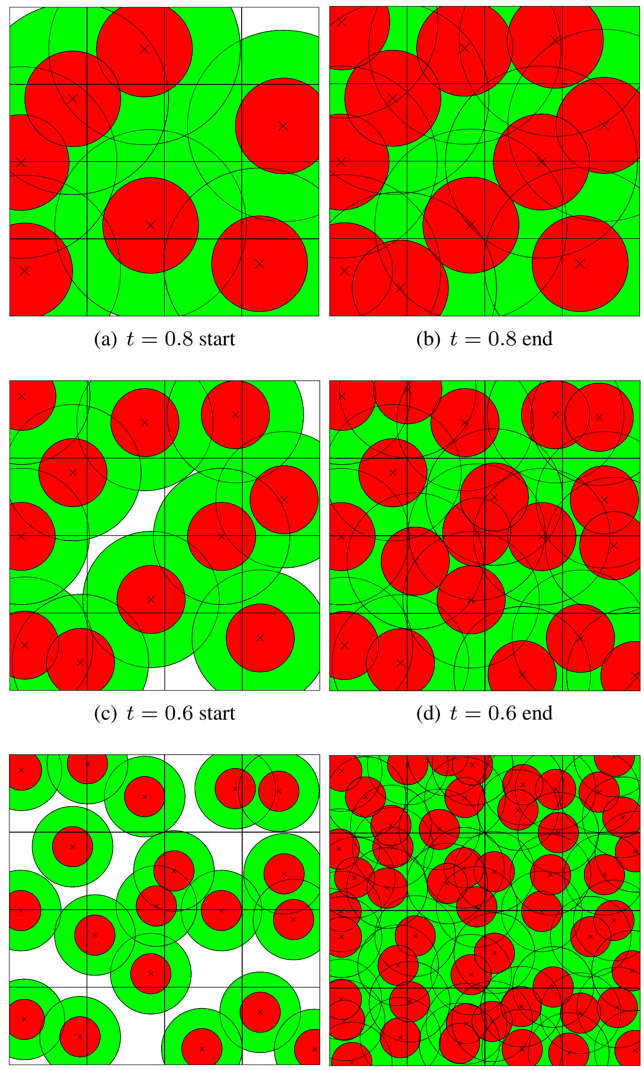

(a) $t = 0.8$ start (b) $t = 0.8$ end

(c) $t = 0.6$ start (d) $t = 0.6$ end

(e) $t = 0.4$ start (f) $t = 0.4$ end

Figure 5: *A discrete hierarchy of samplings with $t = 0.8$, $t = 0.6$, $t = 0.4$.*

Here we confirm our theoretical expectations of the output quality with experimental results. We generated point clouds in a 2d unit box. We analyzed them using the Point Set Analysis [Schlömer 2011] tool. PSA generates standardized spectral diagrams for 2d point distributions, aiding direct comparison. The first panel is the point set. The second panel is the FFT spectrum of the point set, with the DC component removed. MPS typically generates spectra with a dark central disk surrounded by alternating light and dark rings rippling out from the center, decreasing in magnitude. The third panel is the radial mean power, which measures the average variation of the second panel's rings' magnitudes. The fourth panel is the anisotropy, which measures the variance along the rings' circumferences.

Figure 6 contains this analysis for a uniform MPS point cloud with $r = 0.01$. Figure 7 shows the PSA results for a point cloud generated using different inhibition and coverage radii after a hierarchical construction. The ringing artifacts in the FFT spectrum are dramat-

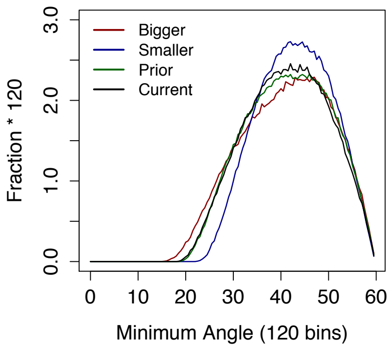

Figure 11: *Angle histograms for samplings using the four spatially varying radii strategies, over the radial sizing function (as in Figure 10).*

ically reduced in this setting. The smallest angle in the Delaunay triangulation of this point cloud is $15.1°$ consistent with the theoretical guarantee of $14.4°$.

For comparison, Figure 8 contains PSA for two points clouds generated using Delaunay refinement (a deterministic method) in the software Triangle [Shewchuk 1996-2005]. Refinement with boundary edge protection was performed on a square of side length three. Only points in the middle, in a centered square of side length one, are analyzed, to avoid FFT artifacts from the boundary bias. Steiner vertices were inserted until no edges longer than 0.02 remained. The first example results from traditional circumcenter insertion [Chew 1989] and produces an unbiased spectrum. The second example results from off-center vertex insertion using a $30°$ target angle [Rand 2011] and leads to a biased point cloud. Biased results from Delaunay refinement generally cannot be attributed to a single cause, but several combinations of input symmetry, the types of Steiner points used, queue ordering under which the triangles are processed and size/quality requirements can yield biased point clouds.

To demonstrate the spatially varying radii method, Section 5 and 6, Figure 9 shows sampling a simple linear function using all four strategies.

Figure 10 contains PSA for a point cloud generated for a nonuniform sizing function $r(\mathbf{x}) = r_m + (r_M - r_m)\,|\sin(8\pi d)|$ where $r_m = 0.015, r_M = 0.00015$, and $d = ||\mathbf{x} - (.5, .5)||$. Point clouds generated for this sizing function using bigger, smaller, prior, and current disks can be found in Appendix D Figure 18. Spectrum plots for anisotropic point sets provide limited insight. Figure 11 compares histograms of angles in the Delaunay triangulations of the point clouds for the four different variable radius constructions. The experimental results match the theory: the smaller-disk construction yields a larger minimum angle.

Figure 12 shows our resampling of a spatially varying image [Kopf et al. 2006; Wei 2008]. See Appendix D for a comparison of the four strategies.

8 Conclusions

We have provided definitions, requirements, and algorithms to perform maximal Poisson-disk sampling with spatially varying radii. The key requirement is a limit to the rate at which the radius function changes. We provided four variations. We suggest that the smaller-disks approach has the weakest requirements and provides the best output. The prior-disks method is the easiest to implement, especially as it is a minor change to existing algorithms. However, it has the most restrictions on the input and provides the weakest output guarantees.

We have provided a definition and algorithm for decoupling the disk-free radius from the coverage radius. The algorithm may be used to create a hierarchy of refined meshes, either adding one point at a time or a batch of points based on a scaling parameter. Two radii provides additional randomness over classical MPS; the FFT spectrum of the output does not have the alternating ring pattern, and the radial power is almost uniform beyond the minimum radius threshold. The continuous approach may be viewed as a way to randomize deterministic Delaunay refinement to avoid artifacts and bias.

Figure 12: *Spatially varying radii output. The input was a point cloud from Wei (originally Kopf et al.) that we scanned, grayscaled, smoothed for L, then resampled.*

The requirements of the Lipschitz constant, $L < 1$, for the algorithm to be correct are quite mild in the sense that without it, any algorithm using the same conflict and coverage conditions might produce output with unbounded jumps in the spacing of points.

In the future, spatially varying radii may be combined with two-radii and hierarchical sampling.

Predictive tools for the output of Poisson-disk sampling over spatially-varying functions would be very useful, because they could be used to verify the output of an algorithm. (Wei [2011] can measure output, but currently it is difficult to say whether the measure is what one would predict from the input, and whether other inputs would produce the same output.) Verification would be especially useful, because it would allow the community to accept faster and less memory intensive algorithms that deviate from the pure MPS process.

References

BOLANDER, J. E., AND SAITO, S. 1998. Fracture analyses using spring networks with random geometry. *Engineering Fracture Mechanics 61*, 5-6, 569 – 591.

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

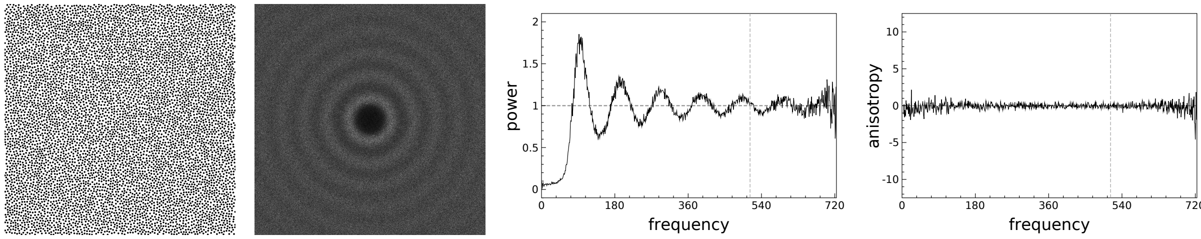

(a) Uniform Point Set (b) FFT Spectrum (c) Radial Power (d) Anisotropy

Figure 6: *Unbiased uniform MPS output using the PSA tool.* $r = 0.01$.

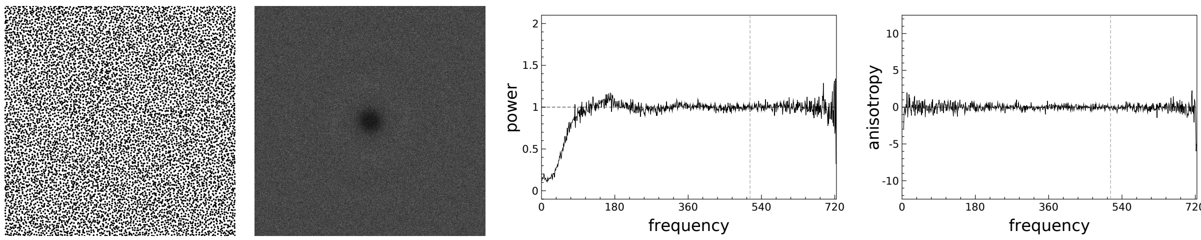

(a) Hierarchy Point Set (b) FFT Spectrum (c) Radial Power (d) Anisotropy

Figure 7: *Final sampling in a hierarchy of different inhibition and coverage radii, with* $r_c = r$ *from Figure 6. The initial sampling used* $R_f = 0.025$ *and* $R_c = 0.05$*. Discrete decrease refinement (Section 4.3) was performed using* $t = 0.8, 0.6, 0.4, 0.2$*. Analysis of the final point set (*$t = 0.2$*,* $r_f = 0.005$*,* $r_c = 0.01$*) is shown.*

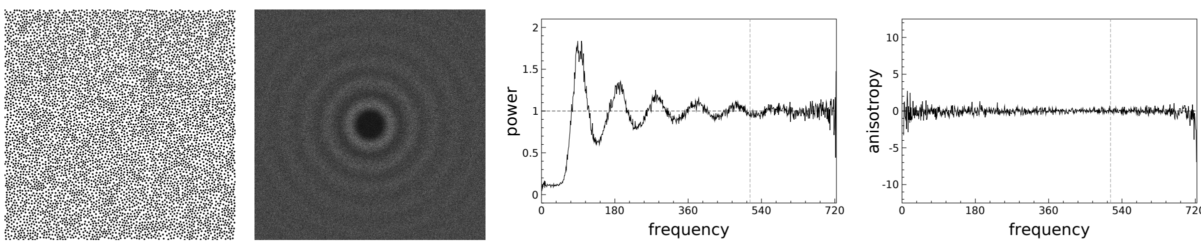

(a) Good DR Point Set (b) FFT Spectrum (c) Radial Power (d) Anisotropy

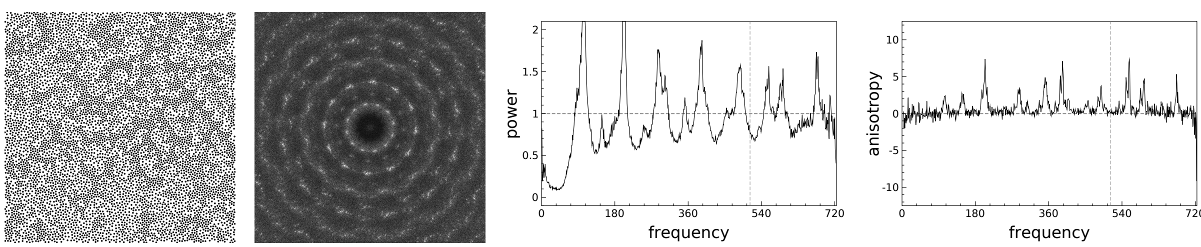

(e) Bad DR Point Set (f) FFT Spectrum (g) Radial Power (h) Anisotropy

Figure 8: *Two example of Delaunay refinement (DR) output, one with little bias and one with a lot of bias, as measured by PSA. We do not know how to ensure or predict the output bias.*

CHENG, S.-W., DEY, T. K., EDELSBRUNNER, H., FACELLO, M. A., AND TENG, S.-H. 1999. Sliver exudation. In *Proceedings of the Fifteenth Annual Symposium on Computational Geometry*, SCG '99, 1–13.

CHERNIKOV, A., AND CHRISOCHOIDES, N. 2009. Generalized two-dimensional Delaunay mesh refinement. *SIAM Journal on Scientific Computing 31*, 3387–3403.

CHEW, L. P. 1989. Guaranteed-quality triangular meshes. Tech. Rep. 89-983, Department of Computer Science, Cornell University.

CHEW, L. P. 1997. Guaranteed-quality delaunay meshing in 3d (short version). In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry*, ACM, New York, NY, USA, SCG '97, 391–393.

DEVILLERS, O. 2002. The delaunay hierarchy. *International Journal of Foundations of Computer Science 13*, 163–180.

DICKMAN, R., WANG, J.-S., AND JENSEN, I. 1991. Random sequential adsorption: Series and virial expansions. *Journal of Chemical Physics 94*, 8252–8257.

EBEIDA, M. S., AND MITCHELL, S. A. 2011. Uniform random Voronoi meshes. In *Proceedings of the 20th International Meshing Roundtable*, 258–275.

EBEIDA, M. S., MITCHELL, S. A., DAVIDSON, A. A., PATNEY,

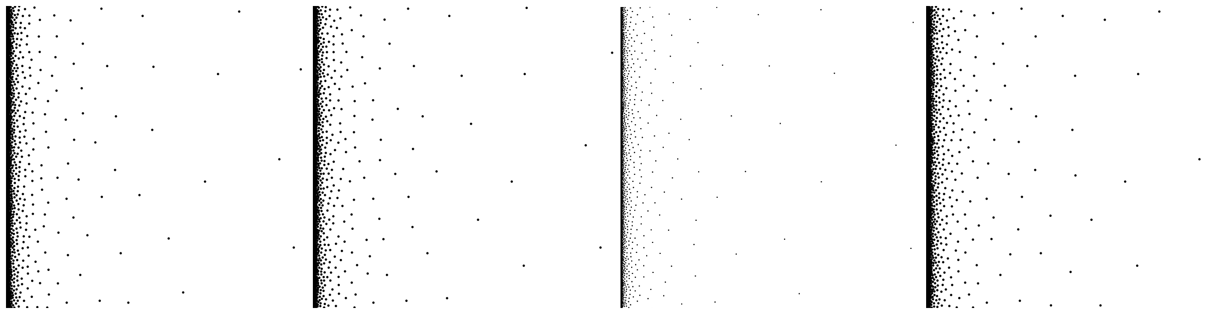

(a) prior-disks (b) current-disks (c) bigger-disks (d) smaller-disks

Figure 9: *Variable radii samplings with linear-ramp function,* $r(x, y) = 0.001 + 0.3 * x$*, using the same random number seed.*

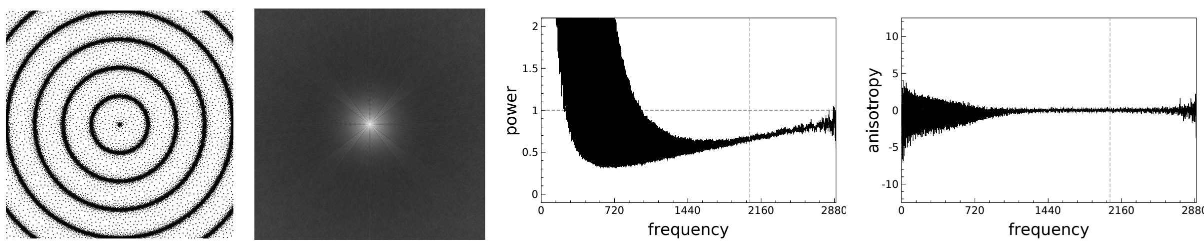

(a) Bigger-Disk Point Set (b) FFT Spectrum (c) Radial Power (d) Anisotropy

Figure 10: *Bigger-disk sampling for a non-uniform sizing function,* $r(\boldsymbol{x}) = r_m + (r_M - r_m)\,|\sin(8\pi d)|$ *where* $r_m = 0.015$, $r_M = 0.00015$*, and* $d = ||\boldsymbol{x} - (.5, .5)||$.

A., KNUPP, P. M., AND OWENS, J. D. 2011. Efficient and good Delaunay meshes from random points. *Computer-Aided Design 43*, 11, 1506 – 1515. Solid and Physical Modeling.

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

EBEIDA, M. S., MITCHELL, S. A., PATNEY, A., DAVIDSON, A. A., AND OWENS, J. D. 2012. A simple algorithm for maximal Poisson-disk sampling in high dimensions. *Computer Graphics Forum, Proc. Eurographics 31*, 2, tbd.

EDELSBRUNNER, H., LI, X.-Y., MILLER, G., STATHOPOULOS, A., TALMOR, D., TENG, S.-H., ÜNGÖR, A., AND WALKINGTON, N. 2000. Smoothing and cleaning up slivers. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing*, STOC '00, 273–277.

ERTEN, H., AND ÜNGÖR, A. 2009. Quality triangulations with locally optimal Steiner points. *SIAM Journal on Scientific Computing 31*, 2103.

GAMITO, M. N., AND MADDOCK, S. C. 2009. Accurate multidimensional Poisson-disk sampling. *ACM Transactions on Graphics 29*, 1 (Dec.), 8:1–8:19.

JONES, T. R., AND KARGER, D. R. 2011. Linear-time Poisson-disk patterns. *Journal of Graphics, GPU, & Game Tools* (to appear). arXiv:1107.3013v1.

KOPF, J., COHEN-OR, D., DEUSSEN, O., AND LISCHINSKI, D. 2006. Recursive Wang tiles for real-time blue noise. *ACM Transactions on Graphics 25*, 3, 509–518.

LAGAE, A., AND DUTRÉ, P. 2008. A comparison of methods for generating Poisson disk distributions. *Computer Graphics Forum 27*, 1 (Mar.), 114–129.

LI, X., AND TENG, S. 2001. Generating well-shaped Delaunay meshed in 3d. In *Proceedings of the Twelth Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, 28–37.

LI, X.-Y., TENG, S.-H., AND ÜNGÖR, A. 1998. Simultaneous refinement and coarsening: Adaptive meshing with moving boundaries. In *Proceedings of the 7th International Meshing Roundtable*, 201–210.

LI, X. 2003. Generating well-shaped d-dimensional delaunay meshes. *Theoretical Computer Science 296*, 1, 145–165.

LJUNG, P. 2006. Adaptive sampling in single pass, GPU-based raycasting of multiresolution volumes. Eurographics Association, Boston, Massachusetts, USA, R. Machiraju and T. Möller, Eds., 39–46.

MATÉRN, B. 1960. Spatial variation. *Meddelanden från Statens Skogsforskningsinstitut 49*, 1–140.

MCCOOL, M., AND FIUME, E. 1992. Hierarchical Poisson disk sampling distributions. In *Graphics Interface '92*, 94–105.

MILLER, G. L., TALMOR, D., TENG, S.-H., WALKINGTON, N. J., AND WANG, H. 1996. Control volume meshes using sphere packing: Generation, refinement and coarsening. In *Proceedings of the Fifth International Meshing Roundtable*, 47–61.

MOUNT, D. M., AND ARYA, S., 2010. ANN: A library for approximate nearest neighbor searching. Version 1.1.2, `http://www.cs.umd.edu/~mount/ANN/`.

OSTROMOUKHOV, V., DONOHUE, C., AND JODOIN, P.-M. 2004. Fast hierarchical importance sampling with blue noise properties. *ACM Transactions on Graphics 23*, 3 (Aug.), 488–495.

OSTROMOUKHOV, V. 2007. Sampling with polyominoes. *ACM Transactions on Graphics 26*, 3 (July), 78:1–78:6.

RAND, A. 2011. Where and how Chew's second Delaunay refinement algorithm works. In *Proceedings of the 23rd Canadian Conference on Computational Geometry*, 157–162.

RENSHAW, E. 2010. Spatial-temporal marked point processes: a spectrum of stochastic models. *Environmetrics 21*, 3-4, 253–269.

RUPPERT, J. 1995. A Delaunay refinement algorithm for quality 2-dimensional mesh generation. *Journal of Algorithms 18*, 3, 548–585.

SCHLÖMER, T., 2011. PSA point set analysis. Version 0.2.2, `http://code.google.com/p/psa/`.

SHEWCHUK, J. R., 1996-2005. Triangle: a two-dimensional quality mesh generator and Delaunay triangulator. `http://www.cs.cmu.edu/~quake/triangle.html`.

SHEWCHUK, J. R. 2002. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry 22*, 1–3, 86–95.

TALMOR, D. 1997. *Well-Spaced Points for Numerical Methods*. PhD thesis, Carnegie Mellon University, Pittsburgh. CMU CS Tech Report CMU-CS-97-164.

ÜNGÖR, A. 2009. Off-centers: A new type of Steiner points for computing size-optimal quality-guaranteed Delaunay triangulations. *Compututational Geometry: Theory and Applications 42* (February), 109–118.

VANDERHAEGHE, D., BARLA, P., THOLLOT, J., AND SILLION, F. 2007. Dynamic point distribution for stroke-based rendering. In *Rendering Techniques 2007 (Proceedings of the Eurographics Symposium on Rendering)*, 139–146.

WEI, L.-Y., AND WANG, R. 2011. Differential domain analysis for non-uniform sampling. *ACM Transactions on Graphics 30* (Aug.), 50:1–50:10.

WEI, L.-Y. 2008. Parallel Poisson disk sampling. *ACM Transactions on Graphics 27*, 3 (Aug.), 20:1–20:9.

WHITE, K. B., CLINE, D., AND EGBERT, P. K. 2007. Poisson disk point sets by hierarchical dart throwing. In *RT '07: Proceedings of the 2007 IEEE Symposium on Interactive Ray Tracing*, 129–132.

A Theoretical Guarantees for Spatially Varying MPS

A.1 Prior-disk

Proposition A.1. *Suppose that sample X satisfies the empty disk property. Then for all i,j, $|\boldsymbol{x}_i - \boldsymbol{x}_j| \geq \frac{r(\boldsymbol{x}_i)}{1+L}$.*

Proof. If $i < j$, the empty-disk definition implies $|\mathbf{x}_i - \mathbf{x}_j| \geq r(\mathbf{x}_i)$. Otherwise we apply the Lipschitz property and the fact that $\mathbf{x}_i$ satisfies the empty-disk property when it is inserted:

$$r(\mathbf{x}_i) \leq r(\mathbf{x}_j) + L\,|\mathbf{x}_i - \mathbf{x}_j| \leq |\mathbf{x}_i - \mathbf{x}_j| + L\,|\mathbf{x}_i - \mathbf{x}_j|\,. \qquad \square$$

Proposition A.2. *Suppose that sample X is maximal and T is a resulting Delaunay triangle. Then $R_T \leq \min\left(\frac{r(\boldsymbol{y})}{1-L}, \frac{r(\boldsymbol{x})}{1-2L}\right)$ where R_T is the circumradius, $\boldsymbol{y}$ is the circumcenter and $\boldsymbol{x}$ is any triangle vertex.*

Proof. Since X is maximal, $|\mathbf{z} - \mathbf{y}| \leq r(\mathbf{z})$ for some vertex $\mathbf{z} \in X$ which is not necessarily a vertex of T; see Figure 13(a). Now applying the Lipschitz property gives

$$|\mathbf{z} - \mathbf{y}| \leq r(\mathbf{z}) \leq r(\mathbf{y}) + L\,|\mathbf{z} - \mathbf{y}|\,.$$

Rearranging gives $R_T \leq |\mathbf{z} - \mathbf{y}| \leq \frac{r(\mathbf{y})}{1-L}$. Now we apply the Lipschitz property again:

$$R_T = |\mathbf{x} - \mathbf{y}| \leq |\mathbf{z} - \mathbf{y}| \leq \frac{r(\mathbf{y})}{1-L} \leq \frac{1}{1-L}\left(r(\mathbf{x}) + L\,|\mathbf{x} - \mathbf{y}|\right).$$

Again rearranging completes the proof. $\square$

Corollary A.3. *Suppose that sample X is maximal. Then $|\boldsymbol{x}_i - \boldsymbol{x}_j| \leq \frac{2r(\boldsymbol{x}_i)}{1-2L}$.*

Lemma A.4. *Suppose X is a maximal sample satisfying the empty dist property. Then all the angles in the Delaunay triangulation are at least* $\arcsin\left(\frac{1-2L}{2}\right)$.

Proof. Let α be an angle in the Delaunay triangulation of X and let $\mathbf{x}$ be the vertex on the edge opposite of α which was inserted *first*. Then this opposite edge has length at least $r(\mathbf{x})$. Then applying Propositions 4.2 and A.2 for a vertex opposite angle angle α in the triangulation

$$\sin\alpha \geq \frac{r(\mathbf{x})}{2r(\mathbf{x})/(1-2L)} = \frac{1-2L}{2}. \qquad \square$$

Proposition A.5. *Suppose that sample X is maximal. Then for all $\boldsymbol{y} \in \Omega$, $\min_i |\boldsymbol{x}_i - \boldsymbol{y}| \leq \frac{r(\boldsymbol{y})}{1-L}$.*

Proof. The maximal definition requires the existence of a vertex $\mathbf{x}_k$ such that $|\mathbf{x}_k - \mathbf{y}| \leq r(\mathbf{x}_k)$. Then using the Lipschitz property:

$$|\mathbf{x}_k - \mathbf{y}| \leq r(\mathbf{x}_k) \leq r(\mathbf{y}) + L\,|\mathbf{x}_k - \mathbf{y}|\,.$$

Rearranging terms completes the proof. $\square$

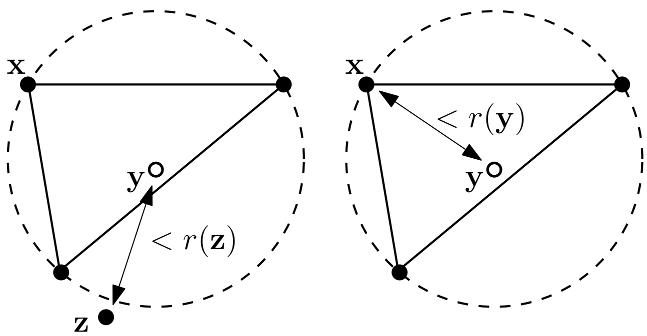

(a) Propositions A.2 and A.12 (b) Propositions A.7 and A.17

Figure 13: *Notation for proofs of estimates of circumradii in the Delaunay triangulation of maximal samples.*

A.2 Current-disk

Proposition A.6. *Suppose that sample X satisfies the empty disk property. The for all i,j, $|\boldsymbol{x}_i - \boldsymbol{x}_j| \geq \frac{r(\boldsymbol{x}_i)}{1+L}$.*

Proof. If $i > j$, the empty-disk definition implies $|\mathbf{x}_i - \mathbf{x}_j| \geq r(\mathbf{x}_i)$. Otherwise we apply the Lipschitz property and the fact that $\mathbf{x}_j$ satisfies the empty-disk property when it is inserted:

$$r(\mathbf{x}_i) \leq r(\mathbf{x}_j) + L\,|\mathbf{x}_i - \mathbf{x}_j| \leq |\mathbf{x}_i - \mathbf{x}_j| + L\,|\mathbf{x}_i - \mathbf{x}_j|\,. \quad \square$$

Proposition A.7. *Suppose that sample X is maximal and T is a resulting Delaunay triangle. Then $R_T \leq \min\left(r(\boldsymbol{y}), \frac{r(\boldsymbol{x})}{1-L}\right)$ where R_T is the circumradius, $\boldsymbol{y}$ is the circumcenter and $\boldsymbol{x}$ is any triangle vertex.*

Proof. Since X is maximal, $R_T = |\mathbf{x} - \mathbf{y}| \leq r(\mathbf{y})$ for any vertex $\mathbf{x}$ of T; see Figure 13(b). Now applying the Lipschitz property gives $R_T \leq r(\mathbf{y}) \leq r(\mathbf{x}) + L\,|\mathbf{x} - \mathbf{y}|$. □

Corollary A.8. *Suppose that sample X is maximal and $\boldsymbol{x}_i, \boldsymbol{x}_j \in X$ are Delaunay neighbors. Then $|\boldsymbol{x}_i - \boldsymbol{x}_j| \leq \frac{2r(\boldsymbol{x}_i)}{1-L}$.*

Lemma A.9. *Suppose X is a maximal sample satisfying the empty dist property. Then all the angles in the Delaunay triangulation are at least* $\arcsin\left(\frac{1-L}{2}\right)$.

Proof. Let α be an angle in the Delaunay triangulation of X and let $\mathbf{x}$ be the vertex on the edge opposite of α which was inserted *last*. Then this opposite edge has length at least $r(\mathbf{x})$. Then applying Propositions 4.2 and A.7 for a vertex opposite angle angle α in the triangulation

$$\sin\alpha \geq \frac{r(\mathbf{x})}{2r(\mathbf{x})/(1-L)} = \frac{1-L}{2}. \quad \square$$

Proposition A.10. *Suppose that sample X is maximal. Then for all $\boldsymbol{y} \in \Omega$, $\min_i |\boldsymbol{x}_i - \boldsymbol{y}| \leq r(\boldsymbol{y})$.*

Proof. This is exactly the definition of maximal sample that we are using. □

A.3 Bigger-disk

Proposition A.11. *Suppose that sample X satisfies the empty disk property. Then for all i, j, $|\boldsymbol{x}_i - \boldsymbol{x}_j| \geq r(\boldsymbol{x}_i)$.*

Proof. Immediate from the empty disk definition. □

Proposition A.12. *Suppose that sample X is maximal and T is a resulting Delaunay triangle. Then $R_T \leq \min\left(\frac{r(\boldsymbol{y})}{1-L}, \frac{r(\boldsymbol{x})}{1-2L}\right)$ where R_T is the circumradius, $\boldsymbol{y}$ is the circumcenter and $\boldsymbol{x}$ is any triangle vertex.*

Proof. Since X is maximal, $|\mathbf{z} - \mathbf{y}| \leq \max(r(\mathbf{z}), r(\mathbf{y}))$ for some vertex $\mathbf{z} \in X$ which is not necessarily a vertex of T; see Figure 13(a). So if $|\mathbf{z} - \mathbf{y}| > r(\mathbf{y})$ then $|\mathbf{z} - \mathbf{y}| \leq r(\mathbf{z})$. Now applying the Lipschitz property gives

$$|\mathbf{z} - \mathbf{y}| \leq r(\mathbf{z}) \leq r(\mathbf{y}) + L\,|\mathbf{z} - \mathbf{y}|\,.$$

Rearranging gives $R_T \leq |\mathbf{z} - \mathbf{y}| \leq \frac{r(\mathbf{y})}{1-L}$. Now we apply the Lipschitz property again:

$$R_T = |\mathbf{x} - \mathbf{y}| \leq |\mathbf{z} - \mathbf{y}| \leq \frac{r(\mathbf{y})}{1-L} \leq \frac{1}{1-L}\left(r(\mathbf{x}) + L\,|\mathbf{x} - \mathbf{y}|\right).$$

Again rearranging completes the proof. □

Corollary A.13. *Suppose that sample X is maximal and $\boldsymbol{x}_i, \boldsymbol{x}_j \in X$ are Delaunay neighbors. Then $|\boldsymbol{x}_i - \boldsymbol{x}_j| \leq \frac{2r(\boldsymbol{x}_i)}{1-2L}$.*

Lemma A.14. *Suppose X is a maximal sample satisfying the empty dist property. Then all the angles in the Delaunay triangulation are at least* $\arcsin\left(\frac{1-2L}{2}\right)$.

Proof is nearly identical to Lemma A.4.

Proposition A.15. *Suppose that sample X is maximal. Then for all $\boldsymbol{y} \in \Omega$, $\min_i |\boldsymbol{x}_i - \boldsymbol{y}| \leq \frac{r(\boldsymbol{y})}{1-L}$.*

Proof. The maximal definition requires the existence of a vertex $\mathbf{x}_k$ such that $|\mathbf{x}_k - \mathbf{y}| \leq \max(r(\mathbf{x}_k), r(\mathbf{y}))$. Thus either $\min_i |\mathbf{x}_i - \mathbf{y}| \leq r(\mathbf{y})$ or $\min_i |\mathbf{x}_i - \mathbf{y}| \leq r(\mathbf{x}_k)$. In the latter case the Lipschitz property gives

$$|\mathbf{x}_k - \mathbf{y}| \leq r(\mathbf{x}_k) \leq r(\mathbf{y}) + L\,|\mathbf{x}_k - \mathbf{y}|\,.$$

Rearranging terms completes the proof. □

A.4 Smaller-disk

Proposition A.16. *Suppose that sample X satisfies the empty disk property. Then forall i, j, $|\boldsymbol{x}_i - \boldsymbol{x}_j| \geq \frac{r(\boldsymbol{x}_i)}{1+L}$.*

Proof. The empty disk requirement immediately implies that $|\mathbf{x}_i - \mathbf{x}_j| \geq \min(r(\mathbf{x}_i), r(\mathbf{x}_j))$. If $r(\mathbf{x}_i) > r(\mathbf{x}_j)$, then we can apply the Lipschitz property:

$$r(\mathbf{x}_i) \leq r(\mathbf{x}_j) + L\,|\mathbf{x}_i - \mathbf{x}_j| \leq |\mathbf{x}_i - \mathbf{x}_j| + L\,|\mathbf{x}_i - \mathbf{x}_j|\,. \quad \square$$

Proposition A.17. *Suppose that sample X is maximal and T is a resulting Delaunay triangle. Then $R_T \leq \min\left(r(\boldsymbol{y}), \frac{r(\boldsymbol{x})}{1-L}\right)$ where R_T is the circumradius, $\boldsymbol{y}$ is the circumcenter and $\boldsymbol{x}$ is any triangle vertex.*

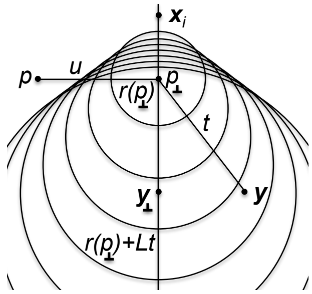

Figure 14: *Possible disks for samples* $\mathbf{y}$ *on the far side of the kd-tree branch for* $\mathbf{x}_i$*. Based solely on* L*, the samples that might overlap* $\overline{pp_\perp}$ *the most lie on the branching hyperplane, i.e.* $\mathbf{y} = \mathbf{y}_\perp$*. If* p *is far enough away, we know it can not lie in any of these disks. In this figure* $L = 0.8$*. The supporting line to the family of hypothetical disks is more vertical for smaller* L*, leading to less overlap.*

Proof. Since X is maximal, $|\mathbf{z} - \mathbf{y}| \le \min(r(\mathbf{z}), r(\mathbf{y})) \le r(\mathbf{y})$ for some vertex $\mathbf{z} \in X$ which is not necessarily a vertex of T; see Figure 13(b). Then applying the Lipschitz property completes the proof:

$$R_T = |\mathbf{x} - \mathbf{y}| \le |\mathbf{z} - \mathbf{y}| \le r(\mathbf{y}) \le r(\mathbf{x}) + L\,|\mathbf{x} - \mathbf{y}|\,. \qquad \square$$

Corollary A.18. *Suppose that sample* X *is maximal and* $\boldsymbol{x}_i, \boldsymbol{x}_j \in X$ *are Delaunay neighbors. Then* $|\boldsymbol{x}_i - \boldsymbol{x}_j| \le \frac{2r(\boldsymbol{x}_i)}{1-L}$.

Lemma A.19. *Suppose* X *is a maximal sample satisfying the empty dist property. Then all the angles in the Delaunay triangulation are at least* $\arcsin\left(\frac{1-L}{2}\right)$.

Proof is identical to Lemma A.9.

Proposition A.20. *Suppose that sample* X *is maximal. Then for all* $\boldsymbol{y} \in \Omega$, $\min_i |\boldsymbol{x}_i - \boldsymbol{y}| \le r(\boldsymbol{y})$.

Proof. The maximal definition requires the existence of a vertex $\mathbf{x}_k$ such that $|\mathbf{x}_k - \mathbf{y}| \le \min(r(\mathbf{x}_k), r(\mathbf{y}))$. Thus

$$\min_i |\mathbf{x}_i - \mathbf{y}| \le |\mathbf{x}_k - \mathbf{y}| \le \min(r(\mathbf{x}_k), r(\mathbf{y})) \le r(\mathbf{y}). \qquad \square$$

B Primitive Implementation Proofs

PiX point-in-disk. Search the kd-tree for $\mathbf{x}$ with $p \in D(\mathbf{x}) \Leftrightarrow |p - \mathbf{x}| < r(\mathbf{x})$. If such a sample is found, then answer "yes." Prune a branch of the kd-tree if the following holds:

- $|p - p_\perp|^2(1 - L^2) > r(p_\perp)^2$.

If $r_{\max}$ is available and either L is close to 1 or there are not good estimates for L, the following conditions may be more convenient.

- $r(p_\perp) > r_{\max}(1 - L^2)$ and $|p - p_\perp|^2 \ge r_{\max}^2 - (r_{\max} - r(p_\perp))^2/L^2$; the latter holds if $|p - p_\perp|^2 \ge r(p_\perp)(2r_{\max} - r(p_\perp))$.
- $r(p_\perp) \le r_{\max}(1 - L^2)$ and $|p - p_\perp|^2 \ge r(p_\perp) r_{\max}$.

We provide the proof that the branching condition is sufficient. Consider Figure 14. To show that $p \notin D(\mathbf{y})$ for any sample on the far side of the branch, it is sufficient to show that $|p - \mathbf{y}| > r(\mathbf{y})$. Let $u = |p - p_\perp|$ and $t = |p_\perp - \mathbf{y}|$. Since $\mathbf{y}$ is on the far side of the kd-tree branch, by the law of cosines we have $|p - \mathbf{y}|^2 \ge u^2 + t^2$. By the Lipschitz condition $r(\mathbf{y}) \le r(p_\perp) + Lt$. Thus a sufficient condition for pruning the branch is $r(\mathbf{y})^2 < |p - \mathbf{y}|^2 \Leftarrow (r(p_\perp) + Lt)^2 < u^2 + t^2$. Or

$$u^2 - r(p_\perp)^2 \ge t(2Lr(p_\perp) - (1 - L^2)t) \forall t \ge 0 \qquad (13)$$

Using basic calculus, the maximum of the right hand side is $L^2 r(p_\perp)^2/(1 - L^2)$, achieved for $t_* = Lr(p_\perp)/(1 - L^2)$. At that value Equation 13 is equivalent to $u^2 \ge r(p_\perp)^2/(1 - L^2)$.

If the maximum value of $r(\cdot) = r_{\max}$ is known, we need only consider t up to $r(p_\perp) + Lt \le r_{\max}$ or $t_o = (r_{\max} - r(p_\perp))/L$. Now $t_o < t_*$ is equivalent to $(r_{\max} - r(p_\perp))(1 - L^2) < L^2 r(p_\perp) \Leftrightarrow r(p_\perp) \ge r_{\max}(1 - L^2)$. If this holds then substituting $t = t_o$ into Equation 13 yields the sufficient condition $u^2 > r_{\max}^2 - (r_{\max} - r(p_\perp))^2/L^2$. Using $L \le 1$ a weaker sufficient condition is $u^2 > r(p_\perp)(2r_{\max} - r(p_\perp))$.

Now if $t_o \ge t_*$ we get the third condition. A sufficient condition is $u^2 - r(p_\perp)^2 \ge t_o(2Lr(p_\perp) - (1 - L^2)t_*)$ which reduces to $u^2 \ge r(p_\perp) r_{\max}$.

The proofs for the other tests use the same principles, and are less involved.

XiP sample-in-disk. Given a query point p, we search the kd-tree for a sample $\mathbf{x}$ with $|p - \mathbf{x}| < r(p)$, pursuing branches with $|p - p_\perp| < r(p)$.

BiX square-in-disk. We seek a sample whose disk contains the square. One solution is to apply the point-in-disk primitive for all the corners of a square. This is easy to describe but not particularly efficient. A faster solution is to apply the point-in-disk primitive to the cell center, and prune if the disk radius is not large enough to possibly encompass the entire square. For a cell, c is its center, h its diagonal. Apply the point-in-disk pruning criteria, with $p = c$ and $|p - p_\perp|$ replaced by $|p - p_\perp| + s/2$ where s is the side length of a cell.

The following shows these criteria are correct. A sample's disk can only cover the cell if $r(\mathbf{x}) \ge |c - \mathbf{x}| + s/2$ because squares and disks are convex, and the point exterior to the square nearest to its center is at distance $s/2$ of the center. We prune kd-tree branches if we can show any sample on the branch has $r(\mathbf{x}) < |c - \mathbf{x}| + s/2$. Thus we prune if $(r(p_\perp) + Lt)^2 < (u + s/2)^2 + t^2$. This is equivalent to point-in-disk with u replaced by $u + s/2$.

For any unpruned sample, we may either check if the disk covers all 2^d corners, or a sufficient condition for coverage is $|c - c_\perp| + h/2 < r(\mathbf{x})$.

XiB sample-in-square's-disk. Search the kd-tree for a sample with $|c - \mathbf{x}| \le r(c) - (1 + L)h/2$, pursuing branches with with half-spaces closer than that to c. If any such sample is found, answer "yes."

Proposition B.1. *If* $r(c) - |c - \boldsymbol{x}_i| \ge (1 + L)h/2$ *then the disk of every candidate dart in* c*'s cell contains* $\boldsymbol{x}_i$*.*

Proof. Let p be any other point in the cell, then we wish to show $r(p) \ge |p - \mathbf{x}_i|$. By the Lipschitz condition $r(p) \ge r(c) - L|p - c| \ge r(p) - Lh/2$. By the triangle inequality, $|p - \mathbf{x}_i| \le |c - \mathbf{x}_i| + |p - c| \le |c - \mathbf{x}_i| + h/2$. Therefore a sufficient condition for $r(p) \ge |p - \mathbf{x}_i|$ is $r(c) - Lh/2 \ge |c - \mathbf{x}_i| + h/2$. $\square$

C DR and MPS contrasted

Delaunay refinement (DR) [Chew 1989; Ruppert 1995; Shewchuk 2002] can be viewed as a deterministic variant of disk packing. MPS places the next point at random in any disk-free region. DR places points exactly at the centers of large empty circumspheres. These are the Voronoi vertices of the sample-so-far, and are thus a subset of the points that disk-free. Variants of DR use random selection regions, small spheres around circumcenters, which targets more of the domain. Offcenters [Üngör 2009] is a technique for selects points nearer to short edges than the circumcenter. This improves mesh grading and uses fewer points than circumradii. A *sliver* is a flat tetrahedra that has a good edge lengths and circumradii, but small dihedral angles. Given three points, in order to form a sliver the fourth point must lie in a particular region near the circumcircle of the first three points. The restrictive nature of this configuration means that slivers can be removed by perturbing vertices [Edelsbrunner et al. 2000] or Delaunay weights [Cheng et al. 1999] or randomly sampling during Delaunay refinement [Chew 1997; Li and Teng 2001]. Despite practical success, the guarantees of these results are quite limited: the theoretical lower bound on dihedral angles in the triangulation is much smaller than 1° in each case.

DR and MPS produce point clouds that have similar properties [Talmor 1997; Miller et al. 1996; Ebeida and Mitchell 2011]. Both DR and MPS point clouds satisfy the empty-disk Equation 1 and maximal Equation 3 properties; only MPS points are unbiased Equation 4. (Some methods relax these conditions for points near the boundary, or to remove slivers.)

DR tracks uncovered regions using Voronoi vertices, whereas MPS tracks voids using uniform grids plus either quadtrees or polygonal approximations [Ebeida et al. 2011b]. (A third category of MPS methods models the arrival time of points within a uniform grid [Jones and Karger 2011].) In the limit that a void is single point, MPS and DR will add the same point. If DR selects a point at random from a small sphere around an empty circumcenter, the inhibition condition may have to be relaxed. This variant is similar to what we present in Section 3 and Section 4 [Chernikov and Chrisochoides 2009].

While DR produces point clouds that share many features with MPS point clouds, there are a few key statistical differences. Since DR is a family of algorithms and methods, variants can often be devised that duplicate certain properties of the MPS point clouds, but no known variant has been shown to produce unbiased output. To study uniform point clouds, we consider two DR variants: Chew's first DR algorithm which inserts the circumcenter of any triangle with circumradius larger than the target size and a second variant that inserts the circumcenter of any triangle with maximum edge longer than twice the target size. In both cases, point clouds are taken from a subset of the domain meshes so that the artificial boundaries used do not directly impact the placement.

While DR point clouds often appear to satisfy the blue noise criteria, other statistics do not match the MPS output. Specifically, we consider histograms of each triangle's smallest angle and edge lengths for the Delaunay triangulation in the point clouds using $r = 0.001$ in the unit square as shown in Figure 15. In each case, the resulting distributions are different. While the edge-length histogram for Chew's first DR algorithm is quite similar to that of MPS, a large enough sample has been used that the differences, especially near length ratio 1.0, are meaningful. Specifically, the MPS point clouds contains 697,529 vertices while those for Chew's first DR algorithm and the second DR variant we consider have $700,093$ and $642,173$ vertices, respectively.

D More Experimental Results

Figure 16 shows our resampling of a spatially varying image [Kopf et al. 2006; Wei 2008] using all four strategies. The input image [Wei 2008] (originally [Kopf et al. 2006]) was scanned, grayscaled, smoothed for L, then resampled.

Figure 17 shows PSA analysis of the output of all four variations, over a linear sizing function. The limitations of FFT analysis for non-uniform samples is apparent.

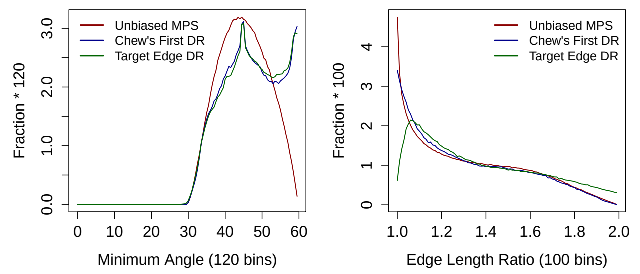

Figure 15: *Comparison of minimum angle and edge-length histograms resulting from uniform triangulation using maximal Poisson disk sampling and Delaunay refinement.*

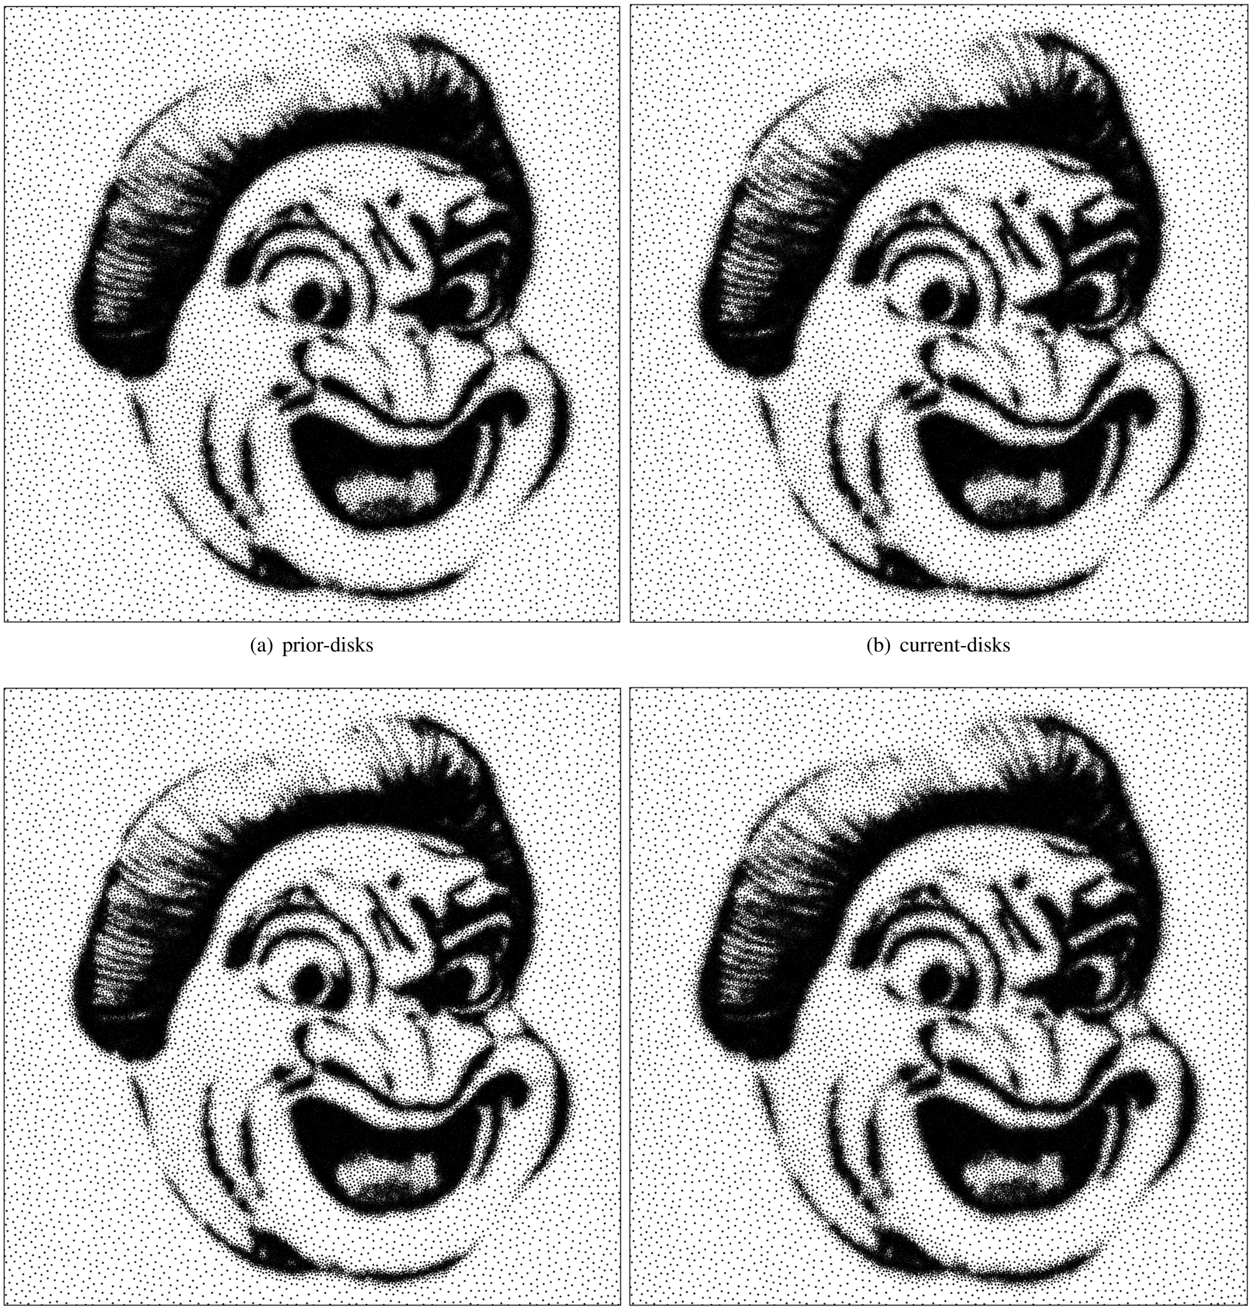

(a) prior-disks

(b) current-disks

(c) bigger-disks

(d) smaller-disks

Figure 16: *Resampling a cartoon face with varying radii.*

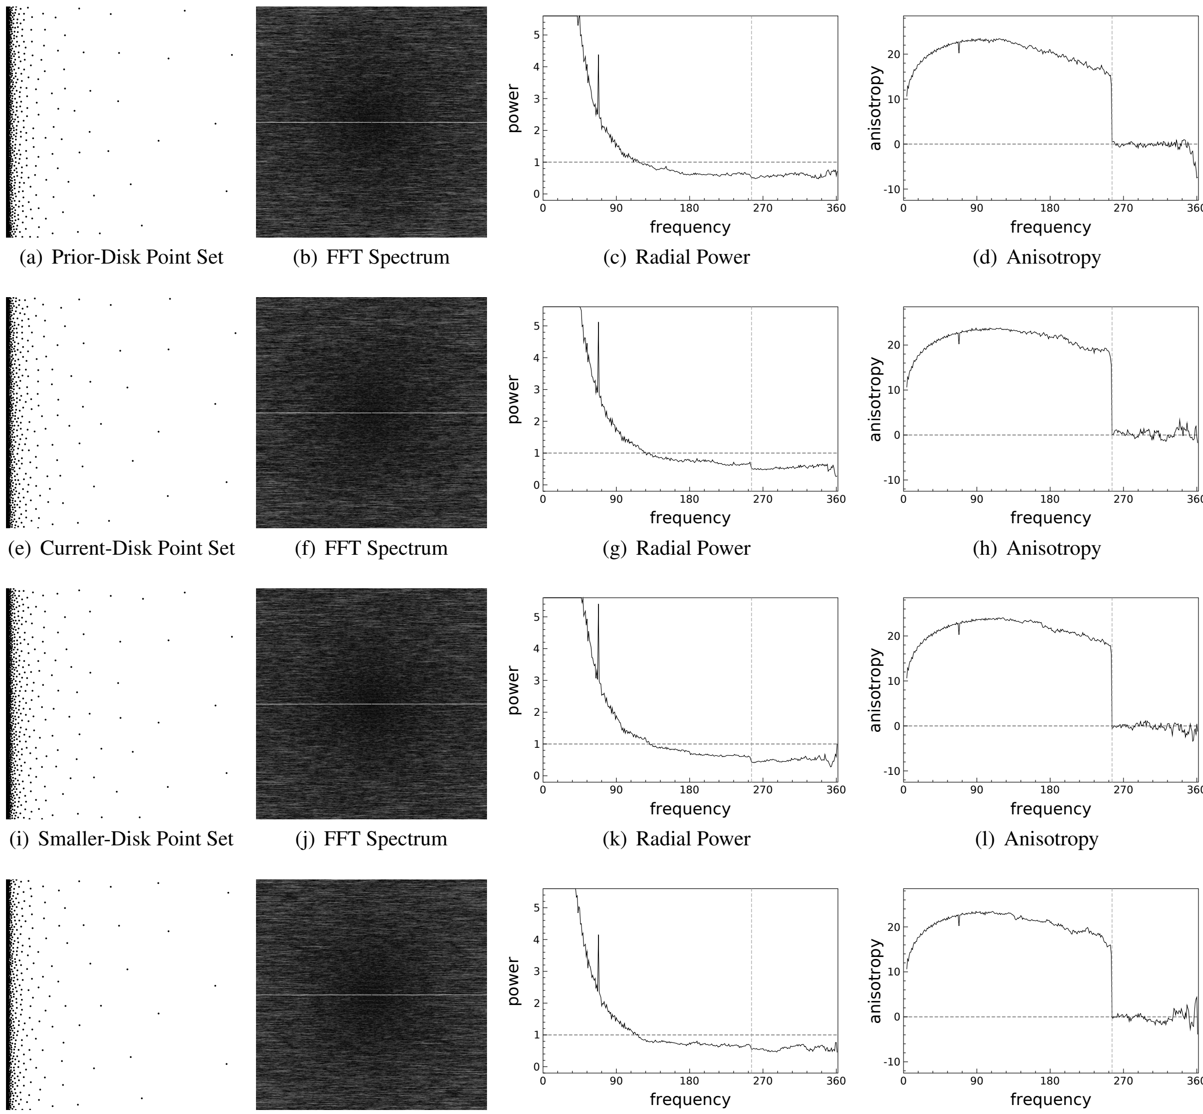

(a) Prior-Disk Point Set (b) FFT Spectrum (c) Radial Power (d) Anisotropy

(e) Current-Disk Point Set (f) FFT Spectrum (g) Radial Power (h) Anisotropy

(i) Smaller-Disk Point Set (j) FFT Spectrum (k) Radial Power (l) Anisotropy

(m) Bigger-Disk Point Set (n) FFT Spectrum (o) Radial Power (p) Anisotropy

Figure 17: *Linear Ramp Example. PSA for a non-uniform sizing function* $r(x, y) = 0.001 + .3x$ *over the unit square using the four approaches with the same random number seed.*

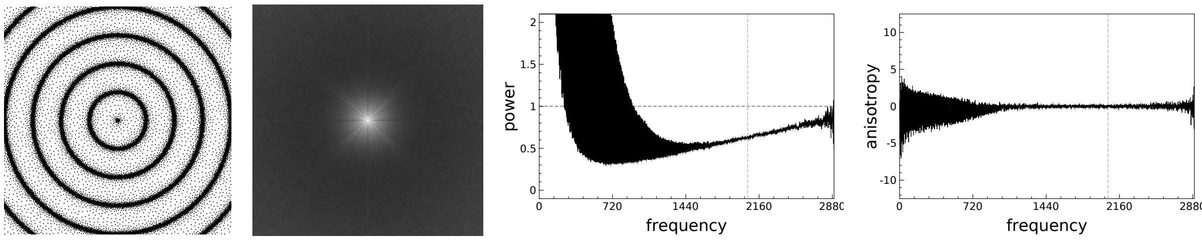

(a) Prior-Disk Point Set (b) FFT Spectrum (c) Radial Power (d) Anisotropy

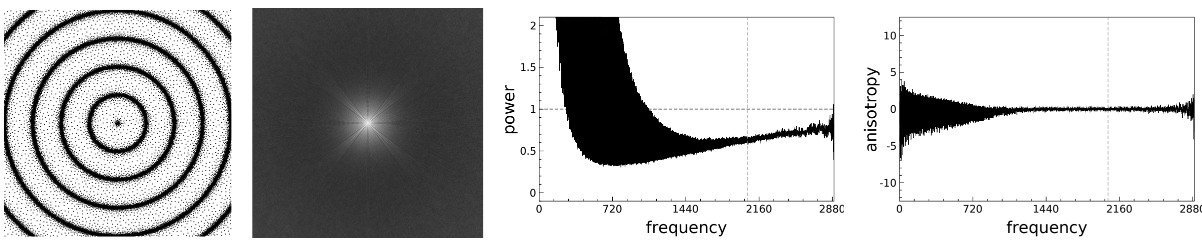

(e) Current-Disk Point Set (f) FFT Spectrum (g) Radial Power (h) Anisotropy

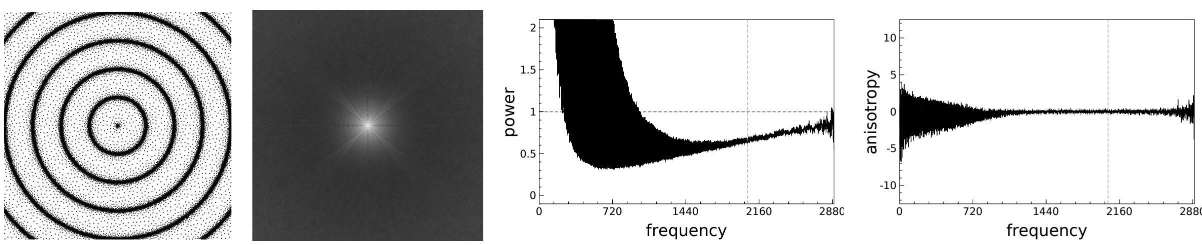

(i) Bigger-Disk Point Set (j) FFT Spectrum (k) Radial Power (l) Anisotropy

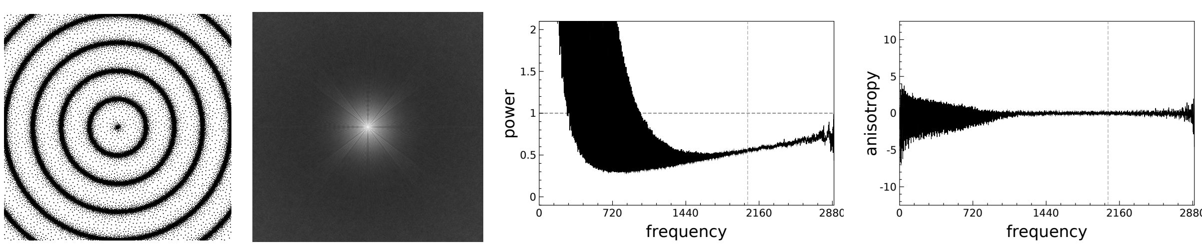

(m) Smaller-Disk Point Set (n) FFT Spectrum (o) Radial Power (p) Anisotropy

Figure 18: *Samplings for a non-uniform sizing function,* $r(\boldsymbol{x}) = r_m + (r_M - r_m)\,|\sin(8\pi d)|$ *where* $r_m = 0.015$*,* $r_M = 0.00015$*, and* $d = ||\boldsymbol{x} - (.5, .5)||$.