# FAST HIGH-DIMENSIONAL NODE GENERATION WITH VARIABLE DENSITY 

O. VLASIUK, T. MICHAELS, N. FLYER, AND B. FORNBERG


#### Abstract

We present an algorithm for producing discrete distributions with a prescribed nearestneighbor distance function. Our approach is a combination of quasi-Monte Carlo (Q-MC) methods and weighted Riesz energy minimization: the initial distribution is a stratified Q-MC sequence with some modifications; a suitable energy functional on the configuration space is then minimized to ensure local regularity. The resulting node sets are good candidates for building meshless solvers and interpolants, as well as for other purposes where a point cloud with a controlled separation-covering ratio is required. Applications of a three-dimensional implementation of the algorithm, in particular to atmospheric modeling, are also given.


## 1. Introduction

1.1. RBFs and meshless methods. In a number of important applications, usefulness of meshless methods in general, and of radial basis functions (RBFs) in particular, is well-known. They have found their way into high-dimensional interpolation, machine learning, spectral methods, vectorvalued approximation and interpolation, just to name a few [53, 16, 26, 12, 49, RBFs have multiple advantages, most importantly extreme flexibility in forming stencils (in the case of RBF-FD) and high local adaptivity; allowing spectral accuracy on irregular domains; the fact that the corresponding interpolation matrix (denoted by $\boldsymbol{A}$ below) is positive definite for several types of radial functions and does not suffer from instability phenomena characteristic of some of the alternative interpolation methods.

Applying RBF-FD stencils to building solvers requires an efficient way of distributing the nodes of basis elements in the domain, which can be either a solid or a surface. The tasks of modeling and simulation often call for massive numbers of nodes, so it is important to ensure that the distribution process is easily scalable. One further has to be able to place the RBFs according to a certain density, as a method of local refinement, for example, at the boundary, or in regions of special interest. Yet another challenge arises when it is necessary to deal with complex or non-smooth domains and/or surfaces.

Recall [28] that an RBF is a linear combination of the form

$$
\begin{equation*}
\mathcal{S}(\boldsymbol{x})=\sum_{k=1}^{K} c_{k} \phi\left(\left\|\boldsymbol{x}-\boldsymbol{x}_{k}\right\|\right) \tag{1}
\end{equation*}
$$

where $\phi(\cdot)$ is a radial function, and $\boldsymbol{x}_{k}, k=1, \ldots, K$, is a collection of pairwise distinct points in $\mathbb{R}^{d}$. A common choice of $\phi$ is the Gaussian $\phi(r)=e^{-(\epsilon r)^{2}}$, although one may also use $1 /\left(1+(\epsilon r)^{2}\right)$, $r^{2 p} \log (r), p \in \mathbb{N}$, etc. In this discussion, we are not concerned with the distinctions between the

[^0]different radial kernels, so the reader can assume that $\phi(r)=e^{-(\epsilon r)^{2}}$. In contrast to pseudospectral methods [26], RBF-FD approach means that to obtain a useful approximation of a function, or a differential operator, the nodes in expressions like (1) must be in the vicinity of the point $\boldsymbol{x}$, and therefore a large number of stencils are constructed throughout the underlying set. It is well-known that the matrix
\[

\boldsymbol{A}=\left[$$
\begin{array}{cccc}
\phi\left(\left\|\boldsymbol{x}_{1}-\boldsymbol{x}_{1}\right\|\right) & \phi\left(\left\|\boldsymbol{x}_{1}-\boldsymbol{x}_{2}\right\|\right) & \ldots & \phi\left(\left\|\boldsymbol{x}_{1}-\boldsymbol{x}_{K}\right\|\right)  \tag{2}\\
\phi\left(\left\|\boldsymbol{x}_{2}-\boldsymbol{x}_{1}\right\|\right) & \phi\left(\left\|\boldsymbol{x}_{2}-\boldsymbol{x}_{2}\right\|\right) & \ldots & \phi\left(\left\|\boldsymbol{x}_{2}-\boldsymbol{x}_{K}\right\|\right) \\
\vdots & \vdots & & \vdots \\
\phi\left(\left\|\boldsymbol{x}_{K}-\boldsymbol{x}_{1}\right\|\right) & \phi\left(\left\|\boldsymbol{x}_{K}-\boldsymbol{x}_{2}\right\|\right) & \ldots & \phi\left(\left\|\boldsymbol{x}_{K}-\boldsymbol{x}_{K}\right\|\right)
\end{array}
$$\right]
\]

is positive definite if the nodes $\boldsymbol{x}_{1} \ldots \boldsymbol{x}_{K}$ are all distinct 48, and so under this assumption there exist $K$-point RBF interpolants for any function data. A different question, however, is whether the matrix $\boldsymbol{A}$ will be well-conditioned: it is not the case, for example, when the nodes are placed on a lattice and $\epsilon \rightarrow 0,30$. The other extreme, having low regularity, also does not provide a reliable source of nodes, as can be seen on the example of the Halton sequence [30]. Furthermore, node clumping can lead to instability of PDE solvers, [26]. To avoid this, one must guarantee that the nodes are separated. In effect, generally the quasi-uniform node sets generated by the present algorithm, or, for example, the one constructed by the third and fourth authors [29], perform better than either lattice or the Halton sequence.

In many applications, one has to ensure that the distance from a node $\boldsymbol{x}$ to its nearest neighbor behaves approximately as a function of the position of the node [29. Prescribing this function, $\rho(\boldsymbol{x})$, which we call the radial density, is a natural way to treat the cases when a local refinement is required in order to capture special features of the domain. In the present paper we will describe a method of node placement for which the actual distance to the nearest neighbor, denoted by $\Delta(\boldsymbol{x})=\min _{\boldsymbol{x}^{\prime} \neq \boldsymbol{x}}\left\|\boldsymbol{x}^{\prime}-\boldsymbol{x}\right\|$, satisfies the above description. To summarize, we are interested in a procedure for obtaining discrete configurations inside a compact set that will:

- guarantee that $\Delta(\boldsymbol{x}) \asymp \rho(\boldsymbol{x})$ (that is, differ only up to a constant factor) for a given function $\rho(\boldsymbol{x})$ with a reasonably wide choice of $\rho$;
- be suitable for mesh-free PDE discretizations using RBFs, i.e., produce well-separated configurations without significant node alignment;
- result in quasi-uniform node distributions also on the surface boundaries of the domain;
- be computationally efficient, easily scalable, and suitable for parallelization.
1.2. Notation and layout. The bold typeface is reserved for vectors in $\mathbb{R}^{d}$. With few exceptions, letters of the Greek alphabet denote functions, calligraphic letters and $\Omega$ denote sets, and the regular Roman typeface is used for scalar variables. The symbolic notation employed throughout the paper is summarized in Table 1.

The paper is structured as follows: Section 2.1 outlines the RBF-FD method using Gaussian and Polyharmonic Spline kernels; Sections 2.2 and 2.3 introduce the two essential components of our approach, Riesz energy functionals and quasi-Monte Carlo methods. The main algorithm and its discussion are the subjects of Sections 3.1 and 3.2 , respectively. Sections 4.1 , 4.3 offer applications of the algorithm; the corresponding run times are summarized in Section 4.4 Section 5.1 contains comparisons of the condition numbers of RBF-FD matrices with stencils on periodic Riesz minimizers, Halton nodes, and the Cartesian grid; Section 5.2 discusses the range of dimensions where the present method is applicable. The Appendix is dedicated to numerical experiments with the mean and minimal separation distance of Riesz minimizers and irrational lattices.

| Symbol | Description |
| :---: | :---: |
| $\bar{x}$ | mean value of $x$ |
| $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{d-1}$ | fixed linearly independent irrational numbers unit cube $[0,1]^{d}$ |
| $C_{1}, C_{2}$ | positive constants in (8) and 10 |
| $\boldsymbol{c}_{m}$ | the closest to the origin corner of $\mathcal{V}_{m}$ |
| $\chi(\cdot ; \Omega)$ | characteristic function of the set $\Omega$ |
| D | $m \in \mathcal{D}$ if $\mathcal{V}_{m}$ has nonempty neighbors on Step 1 |
| $\Delta(\boldsymbol{x})$ | distance from the node $\boldsymbol{x}$ to the nearest neighbor |
| $\Delta\left(\left\{\boldsymbol{x}_{1}, \ldots \boldsymbol{x}_{N}\right\}\right)$ | separation of the configuration |
| $\begin{gathered} \Delta^{k}(\boldsymbol{x}), k=1, \ldots, K \\ \Delta_{n} \\ \bar{\Delta}_{n} \end{gathered}$ | distance from $\boldsymbol{x}$ to the $k$-th nearest neighbor minimal separation of periodized $\mathcal{L}_{n}, \mathcal{M}_{n}$, p. 9 mean separation of the periodized $\mathcal{L}_{n}, \mathcal{M}_{n}$, p. 9 |
| $\mathcal{E}$ | $m \in \mathcal{E}$ if $\mathcal{V}_{m}$ is empty after Step 2 |
| $\begin{aligned} & E_{s}, E_{s}^{\kappa} \\ & \operatorname{frac}(x) \end{aligned}$ | Riesz $s$-energy and weighted $s$-energy, (5)- 7 fractional part of the nonnegative number $x$ |
| $\mathcal{H}_{d}$ | $d$-dimensional Hausdorff measure |
| $\kappa$ | kernel of the weighted Riesz energy, (7) |
| $\lambda$ | interpolated inverse of $\bar{\Delta}_{n}$, p. 9 |
| $\mathcal{L}_{n}, n \geq 1$ | $n$-point irrational lattice, (8) |
| $\mathcal{M}_{n}, n \geq 1$ | $n$-point periodic Riesz minimizer, p. 8 |
| $\begin{gathered} \mathcal{L}_{n}^{\prime}, \mathcal{M}_{n}^{\prime} \\ n_{m} \end{gathered}$ | translated and rescaled $\mathcal{L}_{n}$ and $\mathcal{M}_{n}, 11-(12)$ number of nodes in $\mathcal{V}_{m}$ |
| $\Omega$ | target distribution support |
| $\Omega_{\text {etopo }}, \Omega_{\text {shell }}$ | the underlying sets in Sections 4.1 and 4.3 |
| $\begin{gathered} (r, a, p) \\ \rho(\boldsymbol{x}) \end{gathered}$ | spherical coordinates, p. 13 <br> objective radial density, pp. 2, 8 |
| $\mathcal{V}_{m}, m=1, \ldots, M^{d}$ | cube-shaped voxel in $\mathbb{R}^{d}$, p. 7 |
| $\boldsymbol{x} ; \boldsymbol{x}_{i}, i=1, \ldots, N$ | points in $\mathbb{R}^{d}$; nodes of the configuration |
| $\boldsymbol{x}_{j(i, k)}, k=1, \ldots, K$ | the $k$-th nearest neighbor of $\boldsymbol{x}_{i}, \sqrt{7}$ center of $\mathcal{V}_{m}$ |

Table 1. Symbolic notation employed throughout the paper.

## 2. Choice of method

2.1. RBF-FD approximations. In this section we shall outline the common practices involving RBFs, in order to motivate the requirements that have to be imposed on the node distribution used in the respective computations. For a more in-depth discussion see one of [28, 53, 13, 21]. A significant portion of the RBF approach hinges on the theory of positive definite functions.

Suppose we need to approximate a linear operator $\mathfrak{L}$ acting on sufficiently smooth functions supported on $\Omega$, given locally by their values at the nodes $\boldsymbol{x}_{k}, k=1, \ldots, K$. More specifically, we need to compute the value $\mathfrak{L} \psi\left(\boldsymbol{x}_{0}\right)$ for some fixed $\boldsymbol{x}_{0} \in \Omega$ and a variable function $\psi$. A generalization of the standard [27] finite-difference (FD) approach consists in constructing weights $w_{k}, k=1, \ldots, K$,
that recover the value of $\mathfrak{L}$ at $\boldsymbol{x}_{0}$ in the form

$$
\begin{equation*}
\mathfrak{L S}\left(\boldsymbol{x}_{0}\right)=\sum_{k=1}^{K} w_{k} \mathcal{S}\left(\boldsymbol{x}_{k}\right) \tag{3}
\end{equation*}
$$

for every $\mathcal{S}$ from a convenient functional space; the $\mathcal{S}$ is then chosen to interpolate $\psi$ at the given nodes $\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{K}$. In our case, $\mathcal{S}$ is spanned by RBFs, so by analogy to the 1-dimensional case this method is called RBF-FD; there exists extensive literature covering different types of kernels and different applications [28, 25, 31, 7, 24, 23, 3, 22. Note that the node localization is required due to that (i) local stencils lead to sparse matrices, and are thus much more suitable for computations, (ii) in most applications, $\mathfrak{L}$ is either an interpolation or a differential operator; both act locally, so it is natural to use local stencils.

For example, using the space of shifts of the Gaussian kernel $\phi(r)=e^{-(\epsilon r)^{2}}$, one arrives at an RBF interpolant

$$
\mathcal{S}(\boldsymbol{x})=\sum_{k=1}^{K} c_{k} \phi\left(\left\|\boldsymbol{x}-\boldsymbol{x}_{k}\right\|\right)
$$

In order to express $\mathfrak{L S}\left(\boldsymbol{x}_{0}\right)$ as a functional of $\mathcal{S}\left(\boldsymbol{x}_{k}\right), k=1, \ldots, K$, as in (3), it suffices to do so for the functions $\phi_{k}(\boldsymbol{x})=\phi\left(\left\|\boldsymbol{x}-\boldsymbol{x}_{k}\right\|\right), k=1, \ldots, K$. The weights $\left\{w_{k}\right\}$ are then obtained as the solution to

$$
\left[\begin{array}{cccc}
\phi\left(\left\|\boldsymbol{x}_{1}-\boldsymbol{x}_{1}\right\|\right) & \phi\left(\left\|\boldsymbol{x}_{1}-\boldsymbol{x}_{2}\right\|\right) & \ldots & \phi\left(\left\|\boldsymbol{x}_{1}-\boldsymbol{x}_{K}\right\|\right) \\
\phi\left(\left\|\boldsymbol{x}_{2}-\boldsymbol{x}_{1}\right\|\right) & \phi\left(\left\|\boldsymbol{x}_{2}-\boldsymbol{x}_{2}\right\|\right) & \ldots & \phi\left(\left\|\boldsymbol{x}_{2}-\boldsymbol{x}_{K}\right\|\right) \\
\vdots & \vdots & & \vdots \\
\phi\left(\left\|\boldsymbol{x}_{K}-\boldsymbol{x}_{1}\right\|\right) & \phi\left(\left\|\boldsymbol{x}_{K}-\boldsymbol{x}_{2}\right\|\right) & \ldots & \phi\left(\left\|\boldsymbol{x}_{K}-\boldsymbol{x}_{K}\right\|\right)
\end{array}\right]\left[\begin{array}{c}
w_{1} \\
w_{2} \\
\vdots \\
w_{K}
\end{array}\right]=\left[\begin{array}{c}
\mathfrak{L} \phi\left(\left\|\boldsymbol{x}_{0}-\boldsymbol{x}_{1}\right\|\right) \\
\mathfrak{L} \phi\left(\left\|\boldsymbol{x}_{0}-\boldsymbol{x}_{2}\right\|\right) \\
\vdots \\
\mathfrak{L} \phi\left(\left\|\boldsymbol{x}_{0}-\boldsymbol{x}_{K}\right\|\right)
\end{array}\right] .
$$

Observe that in order to find the interpolant $\mathcal{S}$ for $\psi,\left\{c_{k}\right\}$ are determined from the same system, with $\mathfrak{L}$ taken to be the identity map. The matrix on the left is denoted by $\boldsymbol{A}$ in 2 ; it is degenerate whenever any two of $\left\{\boldsymbol{x}_{k}\right\}$ coincide, and is ill-conditioned whenever any two are very close, due to the continuity of $\phi$. This brings us to further considerations of how the stencil $\left\{\boldsymbol{x}_{k}\right\}$ can be chosen. By the above, it is necessary that the nodes be (i) distinct and well-separated, and (ii) localized inside the domain $\Omega$. For a quasi-uniform node set, $K$ nearest neighbors of a fixed node satisfy both conditions.

Observe that for any strictly positive definite kernel $\phi$, provided $\left\{\boldsymbol{x}_{k}\right\}$ are all distinct, the interpolation matrix $\boldsymbol{A}$ is always invertible (the Gaussian is an example of such kernel). To summarize, the above expression for weights $\left\{w_{k}\right\}$ is the defining property of the RBF-FD methods with the Gaussian kernel.

Taking the limit of the shape parameter $\epsilon \rightarrow 0$ can cause the interpolant $s$ to diverge for other RBF kernels [33, 14] that contain $\epsilon$; this phenomenon however does not occur for the Gaussian $\phi(r)=e^{-(\epsilon r)^{2}}$. The motivation for considering the "increasingly flat" limit $\epsilon \rightarrow 0$ is that the resulting RBFs can be used to obtain highly accurate solutions of elliptic problems and approximants of smooth data [33, 40. We now conclude the discussion of the Gaussian kernel and turn to its novel alternative.

In the recent years, there have been noteworthy advances in RBF-FD using Polyharmonic Spline (PHS) kernels, $\phi(r)=r^{2 p-1}$ or $\phi(r)=r^{2 p} \log r, p \in \mathbb{N}$; it has has been shown [23, 3, 22] that using PHS-based RBF-FD leads to improved accuracy, stability, and eliminates the Runge phenomenon at the boundary of the domain [3], which is not the case in general 34. Another benefit from using the PHS kernel is that it does not contain the shape parameter $\epsilon$. The analytical property underlying
existence of the weights $\left\{w_{k}\right\}$ for the PHS kernels is that the PHS are conditionally strictly positive definite [21, 43] and thus need a slightly different treatment, which we shall now outline.

To ensure unisolvency (uniqueness of the weights and interpolants), we need to augment the $\mathcal{S}$ with polynomial terms: it is selected from the space defined by

$$
\mathcal{S}(\boldsymbol{x})=\sum_{k=1}^{K} c_{k} \phi\left(\left\|\boldsymbol{x}-\boldsymbol{x}_{k}\right\|\right)+\sum_{i=1}^{\binom{l+d}{l}} b_{i} \pi_{i}(\boldsymbol{x})
$$

with $\left\{c_{k}\right\}$ satisfying the constraint

$$
\sum_{k=1}^{K} c_{k} \pi_{i}\left(\boldsymbol{x}_{k}\right)=0, \quad i=1,2, \ldots,\binom{l+d}{l}
$$

where $\phi$ is now the PHS kernel, and $\binom{l+d}{l}$ is the dimension of the space of multivariate polynomials of degree up to $l$ in $\mathbb{R}^{d}$; accordingly, $\pi_{i}$ varies over the monomial basis for such polynomials.

The degree $l$ has to satisfy $l=p-1$ and $l=p$ for $\phi(r)=r^{2 p-1}$ and $\phi(r)=r^{2 p} \log r$, respectively 21, Chapter 8]. For example, when $\phi(r)=r^{2} \log r$ and $d=2$, then $l=1$ and the weights corresponding to an operator $\mathfrak{L}$ in $\mathbb{R}^{2}$ are determined by

$$
\left[\begin{array}{ccccccc} 
& & & \mid c c c c  \tag{4}\\
& \boldsymbol{A} & & \mid & x_{11} & x_{12} \\
& & & \vdots & \vdots \\
- & - & - & + & - & - & - \\
1 & 1 & 1 & \mid & & & \\
x_{11} & \ldots & x_{k 1} & \mid & & 0 & \\
x_{12} & \ldots & x_{k 2} & \mid & &
\end{array}\right]\left[\begin{array}{c}
w_{1} \\
\vdots \\
w_{k} \\
- \\
w_{k+1} \\
w_{k+2} \\
w_{k+3}
\end{array}\right]=\left[\begin{array}{c}
\mathfrak{L} \phi\left(\left\|\boldsymbol{x}_{0}-\boldsymbol{x}_{1}\right\|\right) \\
\vdots \\
\mathfrak{L} \phi\left(\left\|\boldsymbol{x}_{0}-\boldsymbol{x}_{K}\right\|\right) \\
- \\
\mathfrak{L} 1 \\
\mathfrak{L} x_{1}\left(x_{01}\right) \\
\mathfrak{L} x_{2}\left(x_{02}\right)
\end{array}\right]
$$

where the matrix $\boldsymbol{A}$ is the same as in (2) with the PHS kernel; $x_{k j}$ is the $j$-th coordinate of $\boldsymbol{x}_{k}, k=1, \ldots, K ; j=1,2$, and similarly for $\boldsymbol{x}_{0}=\left(x_{01}, x_{02}\right)^{\operatorname{tr}} ; \mathfrak{L} 1, \mathfrak{L} x_{1}, \mathfrak{L} x_{2}$ denote images of the constant and coordinate functions under $\mathfrak{L}$, respectively. Here, as before, the interpolation case is obtained by taking $\mathfrak{L}$ equal to the identity operator; compare the constraints on $\left\{c_{k}\right\}$ above. The generalization to larger values of $l$ and higher dimensions follows along the same lines, with higher-degree monomials instead of linear terms [53, Chapters 8, 11.1].

For this and the other commonly used kernels, non-degeneracy follows from a strengthened form of Micchelli's theorem 45], see also [21, 53, 17, 43; namely, the matrix in the LHS of the previous equation is non-degenerate for any unisolvent $\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{K}$. The remaining part of the discussion for the Gaussian kernel above is further applicable without any modifications. It should be noted that the optimal choice of the degree of PHS that needs to be used, does in general depend on the particular problem under consideration.
2.2. Riesz energy. To generate nodes both devoid of lattice alignment and having near-optimal local separation, we shall minimize a functional on the space of discrete subsets of $\Omega$. Equivalently, one can think of the corresponding gradient flow moving the starting configuration to a suitable position. The desired properties of the minimizing configuration will then follow from the strong repulsion imposed by the functional.

First, for a fixed $s>d$ we introduce Riesz s-energy, a functional $E_{s}: \Omega^{N} \mapsto(0, \infty)$ such that for a collection of vectors $\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{N}$ in $\Omega$,

$$
\begin{equation*}
E_{s}\left(\boldsymbol{x}_{1}, \ldots \boldsymbol{x}_{N}\right)=\sum_{i \neq j} \frac{1}{\left\|\boldsymbol{x}_{i}-\boldsymbol{x}_{j}\right\|^{s}} \tag{5}
\end{equation*}
$$

There exists extensive literature dedicated to the collections minimizing this and derived functionals for $s \geq d$, their asymptotics and limiting measures, see for example [37, 8, 10. It turns out that in the case $s \geq d$ the distribution of minimizers of $E_{s}$ coincides with the normalized Hausdorff measure on $\Omega$; practically this means that the minimizers are uniform in the volumetric sense on $\Omega$, that is, the number of nodes per unit volume is close to constant. In order to produce non-uniform nodes, we shall further add multiplicative weight to (5); this modification of (5) was first studied in [8]. The weighted Riesz s-energy with kernel $\kappa: \Omega \times \Omega \rightarrow(0, \infty)$ is the functional $E_{s}^{\kappa}: \Omega^{N} \mapsto(0, \infty)$ defined by

$$
E_{s}^{\kappa}\left(\boldsymbol{x}_{1}, \ldots \boldsymbol{x}_{N}\right)=\sum_{i \neq j} \frac{\kappa\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{j}\right)}{\left\|\boldsymbol{x}_{i}-\boldsymbol{x}_{j}\right\|^{s}}
$$

It has been shown [8, that the counting measures of the minimizers of the weighted energy converge to the probability measure with volumetric density proportional to $\kappa(\boldsymbol{x}, \boldsymbol{x})^{-d / s} d \mathcal{H}_{d}(\boldsymbol{x})$, with $\mathcal{H}_{d}$ denoting the $d$-dimensional Hausdorff measure. More precisely, for any $\mathcal{B} \subset \Omega$ with boundary of zero measure there holds

$$
\frac{1}{N} \sum_{i=1}^{N} \chi\left(\boldsymbol{x}_{i} ; \mathcal{B}\right) \longrightarrow \frac{1}{Z(\kappa, \Omega)} \int_{\mathcal{B}} \kappa^{-d / s}(\boldsymbol{x}, \boldsymbol{x}) d \mathcal{H}_{d}(\boldsymbol{x}), \quad N \rightarrow \infty
$$

where $\chi(\cdot ; \mathcal{B})$ is the indicator function of $\mathcal{B}$ and $Z(\kappa, \Omega)$ is the normalization constant. Of course, for most applications the set $\mathcal{B}$ will have zero-measure or even differentiable boundary.

It is worth noting that the previous equation shows that the distribution of nodes produced by minimizing $E_{s}^{\kappa}$ depends only on the diagonal values of $\kappa$ for large enough $N$. Indeed, this has been explored in [9], where it is shown that omitting interactions of points at least $r_{N}$ apart in $E_{s}^{\kappa}$ and minimizing the resulting expression leads to the same distribution when $N \rightarrow \infty$; the sequence $r_{N}, N \geq 2$ here satisfies $r_{N} N^{1 / d} \rightarrow \infty, N \rightarrow \infty$. Following [9], a weight $\kappa(\boldsymbol{x}, \boldsymbol{y})$ vanishing whenever $\|\boldsymbol{x}-\boldsymbol{y}\|>r_{N}$ is said to be truncated.

Configurations that minimize $E_{s}^{\kappa}$ over $\Omega^{N}$ for a compact $\Omega$ are well-separated, that is, the quantity $\Delta\left(\left\{\boldsymbol{x}_{1}, \ldots \boldsymbol{x}_{N}\right\}\right)=\min _{i} \Delta\left(\boldsymbol{x}_{i}\right)$ satisfies

$$
\begin{equation*}
\Delta\left(\left\{\boldsymbol{x}_{1}, \ldots \boldsymbol{x}_{N}\right\}\right) \geq C N^{-1 / d}, \quad N \geq 2 \tag{6}
\end{equation*}
$$

We shall assume that the terms in $E_{s}^{\kappa}$ for which $\boldsymbol{x}_{j}$ is not among the $K$ nearest nodes to $\boldsymbol{x}_{i}$ are zero, a condition equivalent to truncating $\kappa$, provided the nodes are well-separated. Under this assumption, the expression for $E_{s}^{\kappa}$ can be rewritten as

$$
\begin{equation*}
E_{s}^{\kappa}\left(\boldsymbol{x}_{1}, \ldots \boldsymbol{x}_{N}\right)=\sum_{i=1}^{N} \sum_{k=1}^{K} \frac{\kappa\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{j(i, k)}\right)}{\left\|\boldsymbol{x}_{i}-\boldsymbol{x}_{j(i, k)}\right\|^{s}} \tag{7}
\end{equation*}
$$

where nodes $\boldsymbol{x}_{j(i, k)}, k=1, \ldots, K$, are the $K$ nearest neighbors of $\boldsymbol{x}$. That minimizers of (7) are well-separated can be shown by the standard argument from one of [37, 8, 9]. This further implies that they are quasi-uniform, which is the key property for us in view of the discussion in Section 2.1. As the form (7) makes clear, for the truncated kernel $\kappa$ the $E_{s}^{\kappa}$ can be computed in $O(N K)$ operations, unlike the $O\left(N^{2}\right)$ operations required to evaluate the functional for a non-vanishing $\kappa$. This, and
the fact that (7) requires constant size memory for storage makes this form of $E_{s}^{\kappa}$ the most useful for our purposes.

The value of the exponent $s$ is chosen so that $s \geq d$ to ensure that the energy functional is sufficiently repulsive; it is known from the classical potential theory that for $s<d$ the minimal energy configurations are not necessarily uniform, and their local structure depends on the shape of the domain [39]. Property (6) holds for any $s>d$, when minimizing the energy over any fixed compact set $\Omega \subset \mathbb{R}^{d}$. While setting $\kappa(\boldsymbol{x}, \boldsymbol{y})=f(\|\boldsymbol{x}-\boldsymbol{y}\|)$ for an arbitrary positive definite radial function $f(r)$ that grows fast enough when $r \rightarrow 0$ would produce similar results, we chose the Riesz energy because the properties of its minimizers are well understood.

Note that simply looking for minimizers of $E_{s}^{\kappa}$ does not provide node sets satisfying $\Delta(\boldsymbol{x}) \asymp \rho(\boldsymbol{x})$ for every $\boldsymbol{x}$; in fact, boundary nodes of such minimizers will often have smaller separations than desired. Since in such cases the boundary has a lower Hausdorff dimension, it does not influence the volumetric density, which agrees with the results above. With this motivation in mind, we are ready to introduce the second component of our method.
2.3. Quasi-Monte Carlo methods. To facilitate convergence of whichever optimization algorithm is used to find minimizers of (7), we can initialize it with a configuration that approximates the limiting measure. One has to rule out Monte Carlo methods due to the separation requirement: random points exhibit clustering [11], which makes deterministic post-processing, in particular by energy minimization, costly. Similarly, mitigating the clustering by purely probabilistic approaches, as for example thinning discussed in 42, or by sampling from a random process with repulsive properties [1, 4], does not generally yield satisfying results, since the separation can only be guaranteed on average. Instead, we turn to the quasi-Monte Carlo (Q-MC) approach. As has been pointed out at the end of Section 2.1, in order to ensure convergence of RBF-FD interpolants, the underlying node set must be (locally) unisolvent; for our purposes this just means that the nodes are in a generic position with respect to each other. The latter is clearly not the case for lattice-like Q-MC configurations, which explains why we resort to energy minimization. On the other hand, we choose not to use other popular Q-MC sequences, such as Halton nodes, as they do not necessarily lead to the best conditioned systems, see [32, Figure 5.1] and Figure 10 below, and are harder to handle when the distribution support $\Omega$ has complex geometry.

The key element of our construction lies in distributing the node set in a deterministic way so as to guarantee low discrepancy between the desired and the obtained radial densities. This is achieved by a Q-MC analog of the stratification of the Monte Carlo method [15]: nodes are distributed with piecewise constant (radial) density that approximates the desired one. We consider two different Q-MC sequences to draw from with (near-)constant radial density: irrational lattices and periodic Riesz minimizers. After dividing the set $\Omega$ into cube-shaped voxels, each voxel is filled with nodes obtained in one of the two ways, appropriately scaled, then the weighted s-energy 7 of the whole node set is minimized. Although we discuss the radial density in the present paper, an argument for the volumetric density can be produced along the same lines.

Yet another reason to make use of a Q-MC sequence is to avoid recursive data structures, which in some cases can be detrimental to the overall performance. Even though such structures have seen significant developments over the years, both dynamic update and parallelization for them remain challenging, 47, 54. The approach of the present paper should therefore be understood as almost opposite to the well-known "quadtree" algorithm [36], that indeed has been used for meshless node generation [51]. Namely, as outlined above, our algorithm places nodes en masse inside the voxels to produce a rough approximation of the target distribution, and subsequently adjusts them by a gradient flow, which is straightforward to parallelize. Although this does involve the computation of
the nearest neighbors in (7), which in practice will be done by constructing a k-d tree, by initializing the node configuration with a stratified $\mathrm{Q}-\mathrm{MC}$ sequence we ensure the indices $j(i, k)$ in $(7)$ will not undergo significant changes during the energy minimization stage, so the k-d tree will not require intensive updates.

An irrational lattice (IL) is defined as a discrete subset of the $d$-dimensional unit cube $[0,1]^{d}$

$$
\begin{equation*}
\mathcal{L}_{n}=\left\{\left(\operatorname{frac}\left(C_{1}+i / n\right), \operatorname{frac}\left(i \alpha_{1}\right),\left\{i \alpha_{2}\right\}, \ldots, \operatorname{frac}\left(i \alpha_{d-1}\right)\right): i=1, \ldots, n\right\} \tag{8}
\end{equation*}
$$

where $\operatorname{frac}(x)=\bmod (x, 1)=x-\lfloor x\rfloor$ denotes the fractional part of $x, C_{1}>0$ is fixed, and $\alpha_{1}, \alpha_{2}, \ldots, \alpha_{d-1}$ are irrational numbers, linearly independent over the rationals. This terminology seems to be accepted in the low-discrepancy community [5], while closely related objects, when used for Q-MC purposes, are known as Korobov/lattice point sets 41.

The motivation for using an IL in this context is due to the existing results on the discrepancy of ILs. It is known for example, that the two-dimensional ILs have the optimal order of $L^{2}$ discrepancy, [6, 5]. Furthermore, in all dimensions ILs are uniformly distributed [38, Chapter 1.6], that is, the fraction of lattice points inside any rectangular box with faces parallel to the coordinate planes converges to its volume. The simple linear structure of ILs makes them especially attractive for SIMD-parallelization.

Another Q-MC sequence that has proven to suit our purposes consists of periodic Riesz minimizers on the unit flat torus, that is, $n$-point collections $\mathcal{M}_{n}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{n}\right\}$ that minimize (5) on $\left([0,1]^{d}\right)^{n}$ with the Euclidean distance $\|\cdot\|$ replaced by the periodic metric

$$
\begin{equation*}
\|\boldsymbol{x}-\boldsymbol{y}\|_{\sim}^{2}=\Pi\left(x_{1}-y_{1}\right)+\Pi\left(x_{2}-y_{2}\right)+\Pi\left(x_{3}-y_{3}\right) \tag{9}
\end{equation*}
$$

where $\Pi(x)=\min \left(x^{2},(1-x)^{2}\right), 0 \leq x \leq 1$. It follows from 37] that such configurations have optimal separation and asymptotically uniform volumetric density. It follows from the numerics also, that in this case the nearest neighbor distances vary very little from node to node; this and that minimizing configurations do not suffer from the lattice-like alignment, makes their rescaled copies good candidates for the stratification.

The number of nodes in individual voxels is defined by the function $\rho$, so the resulting collection has piecewise constant density; refining the voxel partition leads to an improved piecewise approximation of the desired (e.g., smooth) density. In practice, the dependence between the number of nodes contained in the unit cube, and average/minimal nearest neighbor distance is tabulated in advance, and then inverted during the construction of the node set.

## 3. The algorithm

The interested reader will find a Matlab codebase implementing the algorithm described here, as well as the sources for all the figures contained in the present paper, at 52 .
3.1. Formulation. If the nodes must be restricted to a certain compact set $\Omega$, for example, support of a given indicator, we will refer to the set as density support, and to the indicator function as point inclusion function. We may further assume that $\Omega$ is contained in the $d$-dimensional unit cube $\mathcal{C}=[0,1]^{d}$ (see Figure 1 for some of the notation involved) ; the case of an arbitrary compact set then follows by choosing a suitable enclosing cube and applying scaling and translation. Suppose the radial density is prescribed by a Lipschitz-1 function; i.e., $|\rho(\boldsymbol{x})-\rho(\boldsymbol{y})| \leq\|\boldsymbol{x}-\boldsymbol{y}\|$. The reason for this assumption is the respective property of $\Delta(\boldsymbol{x})$, and is explained in further detail in the following


Figure 1. An illustration of some of the symbolic notation used in the algorithm below, in the case $d=2$.
section. Recall that we use an exponent $s>d$. We summarize the discussion in Section 2 into the following algorithm for generating nodes with radial density $\rho$ :

## Initial node layout.

Step 0 Choose one of the two Q-MC sequences described in Section $2,\left\{\mathcal{L}_{n}: n \geq 1\right\}$ or $\left\{\mathcal{M}_{n}\right.$ : $n \geq 1\}$, draw configurations with up to $n_{\text {max }}$ nodes from it, and determine the average nearest neighbor distance for its periodization by the integer lattice, denoted by $\bar{\Delta}_{n}$ for $n$ nodes. Let $\lambda:(0, \infty) \rightarrow\left\{0,1,2, \ldots, n_{\max }\right\}$ be the interpolated inverse to $\bar{\Delta}_{n}:\left\{1,2, \ldots, n_{\max }\right\} \rightarrow(0,1] 1_{1}^{1}$
Step 1 Partition $\mathcal{C}$ into $M^{d}$ equal cube-shaped voxels $\mathcal{V}_{m}, m=1, \ldots, M^{d}$ of side length $1 / M$, with faces parallel to the coordinate planes. Let $\left\{\mathcal{V}_{m}: m \in \mathcal{D}\right\}$ be the subset for which at least one of the adjacent (i.e., sharing a face) voxels has a vertex inside $\Omega$.
Step 2 Let $\bar{\rho}_{m}$ be the average value of $\rho$ at the $2^{d}$ vertices of a voxel $\mathcal{V}_{m}, m \in \mathcal{D}$. Place inside $\mathcal{V}_{m}$ a scaled and translated version of the $n_{m}$-point IL (8), or of the $n_{m}$-point periodic Riesz minimizer, using $n_{m}$ defined by

$$
n_{m}=\lambda\left(\bar{\rho}_{m} M\right)
$$

Repeat for each $m \in \mathcal{D}$.

## Saturation and cleanup.

Step 3 Let $\mathcal{E} \subset \mathcal{D}$ be the set of $m$ for which voxels $\left\{\mathcal{V}_{m}: m \in \mathcal{E}\right\}$ contain no nodes and the centers $\left\{\boldsymbol{z}_{m}: m \in \mathcal{E}\right\}$ satisfy $\Delta\left(\boldsymbol{z}_{m}\right)>\rho\left(\boldsymbol{z}_{m}\right)$. Sort $\mathcal{E}$ by the increasing values of $\rho\left(\boldsymbol{z}_{m}\right)$. Repeat until $\mathcal{E}$ is empty: for every $m \in \mathcal{E}$ place a node in $\boldsymbol{z}_{m}$; recompute $\mathcal{E}$.
Step 4 For all nonempty voxels, remove nodes outside $\Omega$.

## Repel-type iterations, boundary detection.

Step 5 Perform $T$ iterations of the partial gradient descent on the weighted $s$-energy functional (7) with $\kappa(\boldsymbol{x}, \boldsymbol{x})=c \rho(\boldsymbol{x})^{s}$, using the $K$ nearest neighbors of each node: Let the initial configuration be the 0 -th iteration, $\boldsymbol{x}_{i}^{(0)}=\boldsymbol{x}_{i}, i=1, \ldots, N$, with $N$ denoting the total number of nodes distributed. On the $t^{t h}$ iteration, $1 \leq t \leq T$, given a node $\boldsymbol{x}_{i}^{(t)}$ with $K$

[^1]nearest neighbors $\boldsymbol{x}_{j(i, k)}^{(t)}, k=1, \ldots, K$, form the weighted vector sum
$$
\boldsymbol{g}_{i}^{(t)}=s \rho\left(\boldsymbol{x}_{i}^{(t)}\right)^{s} \sum_{k=1}^{K} \frac{\boldsymbol{x}_{i}^{(t)}-\boldsymbol{x}_{j(i, k)}^{(t)}}{\left\|\boldsymbol{x}_{i}^{(t)}-\boldsymbol{x}_{j(i, k)}^{(t)}\right\|^{s+2}}, \quad 1 \leq i \leq N
$$
and the new node position can now be expressed as
\[

\boldsymbol{x}_{i}^{(t+1)}= $$
\begin{cases}\boldsymbol{x}_{i}^{(t)}+\frac{\Delta\left(\boldsymbol{x}_{i}^{(t)}\right)}{t+C_{2}} \frac{\boldsymbol{g}_{i}^{(t)}}{\left\|\boldsymbol{g}_{i}^{(t)}\right\|} & \text { if this sum is inside } \Omega ; \quad 1 \leq i \leq N  \tag{10}\\ \boldsymbol{x}_{i}^{(t)}, & \text { otherwise }\end{cases}
$$
\]

where $C_{2}$ is a fixed offset chosen to control the step size between $\boldsymbol{x}_{i}^{(t)}$ and $\boldsymbol{x}_{i}^{(t+1)}$. If a "pullback" function is provided from a neighborhood of $\Omega$ to its boundary, the condition of $\boldsymbol{x}_{i}^{(t+1)}$ being inside $\Omega$ is replaced with applying the pullback; furthermore, if the radial density has an easily computable gradient, or is changing rapidly, an additional term must be included in (10) (see discussion below).
Update the neighbor indices $j(i, k)$ after every few iterations.
Step 6 If no boundary node set/pullback function is prescribed, define the boundary nodes as follows. Evaluate the point inclusion function for $\boldsymbol{x}_{i} \pm \Delta\left(\boldsymbol{x}_{i}\right) \boldsymbol{e}_{l}, l=1, \ldots, d, i=1, \ldots, N$, where $\boldsymbol{e}_{l}$ is the $l$-th basis vector. If at least one such point lies outside $\Omega$, the $\boldsymbol{x}_{i}$ is considered to be a boundary node.
3.2. Discussion. Our assumption of $\rho$ being Lipschitz- 1 is natural, since $\Delta(\cdot)$ is always Lipschitz-1, if viewed as a function of position. To see this, consider any two nodes $\boldsymbol{x}, \boldsymbol{y}$, and let $\boldsymbol{x}^{\prime}, \boldsymbol{y}^{\prime}$ be their nearest neighbors, respectively, so that $\left\|\boldsymbol{x}-\boldsymbol{x}^{\prime}\right\|=\Delta(\boldsymbol{x})$ and $\left\|\boldsymbol{y}-\boldsymbol{y}^{\prime}\right\|=\Delta(\boldsymbol{y})$. It follows, $\Delta(\boldsymbol{x}) \leq\left\|\boldsymbol{x}-\boldsymbol{y}^{\prime}\right\| \leq\|\boldsymbol{x}-\boldsymbol{y}\|+\Delta(\boldsymbol{y})$, which by symmetry implies $\Delta$ is Lipschitz-1.

Initial node layout. In the parts of the density support with constant $\rho$, the nodes will locally look like a periodization of the initial Q-MC sequence; hence the average neighbor distance in Step 0 is tabulated for the periodized version. Observe that there is some freedom in the notion of voxel adjacency used to define $\left\{\mathcal{V}_{m}: m \in \mathcal{D}\right\}$ in Step 1. for example, in the case of a highly non-convex domain $\Omega$, it might be reasonable to denote the $3^{d}-1$ voxels sharing a vertex with a given $\mathcal{V}_{m}$ as adjacent to it, rather than only the $2 d$ voxels that it has a common face with. This would then ensure that no part of $\Omega$ will be omitted in the node allocation; imagine a long and thin peninsula in $\Omega$ containing no corners of $\mathcal{V}_{m}, m=1, \ldots, M^{d}$. We have found however, that the subsequent repel iterations will guarantee that such a peninsula is adequately filled with nodes even when using only the face-adjacent voxels.

If the IL sequence is chosen in Step 0, the $n_{m}$-node set placed in voxel $\mathcal{V}_{m}$ at Step 2 is an adjusted version of $(8)$ as follows. Let for every $\mathcal{V}_{m}$ the corner with the smallest absolute value be $\boldsymbol{c}_{m}$; the points $\boldsymbol{c}_{m}$ are then vertices of a lattice. Before scaling and translating $\mathcal{L}_{n_{m}}$, apply a random permutation to the coordinates of each node in it, so that to remove long-range lattice structure from the distribution; we will denote such an operation by $\sigma$. Then the IL in voxel $\mathcal{V}_{m}$ becomes

$$
\begin{equation*}
\mathcal{L}_{n_{m}}^{\prime}=\boldsymbol{c}_{m}+\frac{f}{M} \sigma\left(\mathcal{L}_{n_{m}}\right)+\frac{\boldsymbol{h}}{M} \tag{11}
\end{equation*}
$$

where

$$
f=1-c_{d}\left(n_{\max }\right)^{-1 / d}, \quad \boldsymbol{h}=\frac{1-f}{2} \cdot(1,1, \ldots, 1)^{\operatorname{tr}}
$$

with $c_{d}$ depending only on the dimension. The quantities $f$ and $\boldsymbol{h}$ ensure that the lattice points in $\mathcal{L}_{m}$ are inset into the voxel by about half the separation distance, avoiding poorly separated points along the voxel interfaces.

When the periodic minimizer sequence is selected in Step 0, the inset is defined in a similar way, but the permutation is just an identity, $\sigma \equiv \mathrm{id}$, as the minimizers don't have the lattice structure. Likewise, the scaling factor and translation are

$$
f_{m}=1-c_{d}\left(n_{m}\right)^{-1 / d}, \quad \boldsymbol{h}_{m}=\frac{1-f_{m}}{2} \cdot(1,1, \ldots, 1)^{\mathrm{tr}}
$$

The analog of (11) thus takes the form

$$
\begin{equation*}
\mathcal{M}_{n_{m}}^{\prime}=\boldsymbol{c}_{m}+\frac{f_{m}}{M} \mathcal{M}_{n_{m}}+\frac{\boldsymbol{h}_{m}}{M} \tag{12}
\end{equation*}
$$

As one would expect, the average separation for the sequence $\left\{\mathcal{M}_{n}\right\}$ is larger than that of $\left\{\mathcal{L}_{n}\right\}$ for the respective values of $n$. While the inset for the latter is necessary to account for the node proximity after periodization, for the former it serves to mitigate the effects of interfacing voxels containing different number of nodes. This is further discussed in the Appendix.

Saturation and cleanup. Observe that after Step 2, voxels in $\left\{\mathcal{V}_{m}: m \in \mathcal{D}\right\}$ satisfying $\bar{\rho}_{m} M>1$ do not contain any nodes. The goal of Step 3 is therefore to remove any redundant sparsity that may be present whenever the radial density $\rho$ is larger than $1 / M$, as in this case the function $\lambda$ in Step 2 is set to zero. More careful geometric considerations would lead one to set $\lambda(x)>0$ when $0<x<\sqrt{d} / M$, the length of a voxel diagonal, and thus make $\lambda$ dependent on the dimension; on the other hand, using the interval $0<x<1 / M$ as we did appears to suffice due to correction of density in Step 5

Note that in practice, when recomputing $\mathcal{E}$ in Step 3, to verify $\Delta\left(\boldsymbol{z}_{m_{0}}\right)>\rho\left(\boldsymbol{z}_{m_{0}}\right)$ for a fixed $m_{0} \in \mathcal{E}$ it is enough to check $\left\|\boldsymbol{z}_{m_{0}}-\boldsymbol{z}_{m}\right\|>\rho\left(\boldsymbol{z}_{m_{0}}\right)$ for the previously selected $\boldsymbol{z}_{m}$ with $\rho\left(\boldsymbol{z}_{m}\right)<\rho\left(\boldsymbol{z}_{m_{0}}\right)$. Indeed, let $\boldsymbol{z}_{m_{0}}$ be the center of $\mathcal{V}_{m_{0}}$. Then, by the definition of $\lambda$ in Step 2 , the radial density $\rho\left(\boldsymbol{z}_{m_{0}}\right)=(1+D) / M$ for some $D>0$, so the Lipschitz- 1 property implies, for any $\boldsymbol{x}$ such that $\rho(\boldsymbol{x}) \leq 1 / M$ there holds $\left\|\boldsymbol{z}_{m_{0}}-\boldsymbol{x}\right\| \geq\left|\rho\left(\boldsymbol{z}_{m_{0}}\right)-\rho(\boldsymbol{x})\right| \geq D / M$. This ensures that distances from $\boldsymbol{z}_{m_{0}}$ to the nodes produced on Step 2 satisfy

$$
\left\|\boldsymbol{z}_{m_{0}}-\boldsymbol{x}\right\| \geq \frac{D}{1+D} \rho\left(\boldsymbol{z}_{m_{0}}\right)
$$

This shows, when $D \gg 1$, not checking the inequality $\left\|\boldsymbol{z}_{m_{0}}-\boldsymbol{x}\right\|>\rho\left(\boldsymbol{z}_{m_{0}}\right)$ leads to at most a bounded factor error. On the other hand, for $D \ll 1$ distances from $\boldsymbol{z}_{m_{0}}$ to the nodes from Step 2 are also controlled: it follows from (11)-(12) that for $n_{m}$ small, nodes in the voxel $\mathcal{V}_{m}$ have larger inset (depending on $c_{d}$ ). This analysis is certainly not rigorous; however, applying the partial gradient descent in Step 5, we are able to ensure that the ratio $\rho / \Delta$ is close to 1 , as desired.

Observe that in Step 2 the nodes are only placed in $\mathcal{V}_{m}$ 's for which either of the adjacent voxels has corners inside the density support $\Omega$, so removing nodes outside $\Omega$ in Step 4 does not lead to much overhead. Furthermore, since the density is evaluated at the corners only and not at individual nodes, the total number of evaluations may be significantly reduced, which is especially useful when $\rho$ is computationally expensive. It is essential here that due to the Lipschitz- 1 property, $\rho$ is well estimated by its values at the corners $\boldsymbol{c}_{m}$; specifically, $\left|\rho(\boldsymbol{x})-\rho\left(\boldsymbol{c}_{m}\right)\right| \leq \sqrt{d} / 2 M$ with $\boldsymbol{c}_{m}$ the nearest voxel corner to $\boldsymbol{x}$.

Repel-type iterations, boundary detection. The equality $\kappa(\boldsymbol{x}, \boldsymbol{x})=c \rho(\boldsymbol{x})^{s}$ can be justified by observing that each node $\boldsymbol{x}$ of the target distribution must be contained in a ball of radius $\rho(\boldsymbol{x})$, not containing any of the other nodes, hence, the volumetric density must be inverse proportional to $\rho(\boldsymbol{x})^{d}$. On the other hand, minimizers of 77 converge to the distribution with volumetric density $\kappa(\boldsymbol{x}, \boldsymbol{x})^{-d / s}$; hence $\kappa(\boldsymbol{x}, \boldsymbol{x})^{d / s}=c \rho(\boldsymbol{x})^{d}$.

The vector sum in Step 5 is the partial $\boldsymbol{x}$-gradient of the weighted Riesz $s$-energy (7) in the sense that a single summand of $(7)$ is $e(\boldsymbol{x}, \boldsymbol{y})=\kappa(\boldsymbol{x}, \boldsymbol{y})\|\boldsymbol{x}-\boldsymbol{y}\|^{-s}$, and thus its complete $\boldsymbol{x}$-gradient is equal to

$$
\nabla_{\boldsymbol{x}} e(\boldsymbol{x}, \boldsymbol{y})=-s \kappa(\boldsymbol{x}, \boldsymbol{y})(\boldsymbol{x}-\boldsymbol{y})\|\boldsymbol{x}-\boldsymbol{y}\|^{-s-2}+\nabla_{\boldsymbol{x}} \kappa(\boldsymbol{x}, \boldsymbol{y})\|\boldsymbol{x}-\boldsymbol{y}\|^{-s} .
$$

For our purposes, the $\boldsymbol{y}$ here is one of the $K$ nearest nodes to $\boldsymbol{x}$, and, since due to the Q-MC initialization there will be few isolated nodes, and since off-diagonal values of $\kappa(\boldsymbol{x}, \boldsymbol{y})$ do not influence the limiting distribution (for details see [9]), we assume $\kappa(\boldsymbol{x}, \boldsymbol{y}) \approx \kappa(\boldsymbol{x}, \boldsymbol{x})$ to rewrite the previous equation as

$$
\begin{equation*}
\nabla_{\boldsymbol{x}} e(\boldsymbol{x}, \boldsymbol{y})=-s \rho(\boldsymbol{x})^{s}(\boldsymbol{x}-\boldsymbol{y})\|\boldsymbol{x}-\boldsymbol{y}\|^{-s-2}+s \nabla_{\boldsymbol{x}} \rho(\boldsymbol{x}) \rho(\boldsymbol{x})^{s-1}\|\boldsymbol{x}-\boldsymbol{y}\|^{-s} \tag{13}
\end{equation*}
$$

As has been pointed out at the beginning of Section 3.2, in order to be meaningful as a radial density, the function $\rho$ must be Lipschitz- 1 . Then by the Rademacher's theorem, $\nabla_{\boldsymbol{x}} \rho$ exists almost everywhere; this validates the use of it in 13 as well as the approximation $\kappa(\boldsymbol{x}, \boldsymbol{y}) \approx \kappa(\boldsymbol{x}, \boldsymbol{x})$. The ratio of the second term to the first one in 13 ) is bounded by $\nabla_{\boldsymbol{x}} \rho(\boldsymbol{x})\|\boldsymbol{x}-\boldsymbol{y}\| / \rho(\boldsymbol{x})$ and, provided that the distances from $\boldsymbol{x}$ to its nearest neighbors are close to the value of $\rho(\boldsymbol{x})$, is at most $c \nabla_{\boldsymbol{x}} \rho(\boldsymbol{x})$ for a constant $c$. This condition is satisfied because the chosen Q-MC sequences have very regular local structure. In practice, the node distance is small on the scale of the support and varies slowly, so the second term will have negligible impact on the direction of the gradient after normalization; besides, precise gradient computation may prove costly. For these reasons we omit the second term in equation 10 . If it is necessary to deal with a fast-changing radial density, a trade-off between the computational costs and the resulting distribution properties must be sought.

It doesn't matter which minimization method is applied to the weighted $s$-energy, rather the gradient descent is chosen due to its simplicity. Note, the second case in (10), leading to shrinking of the line stepping distance, can be thought of as a simplistic backtracking line search; it turns out to be sufficient for our purposes. Furthermore, applying a more involved line search may significantly degrade performance for complicated or nonsmooth domains.

The number of nearest neighbors $K$ in (7), and the number of iterations $T$ in Step 5 can be adjusted to achieve a trade-off between execution speed/memory consumption/local separation. In our experiments ${ }^{2}$ even relatively small values of $K$ and $T$ produce good results: we used $K \approx T \approx 30$ for 1.36 million nodes with constant density in Section 4.1, and $K=30, T=200$ for 0.58 million and 0.36 million nodes with variable densities in Sections 4.2 and 4.3, respectively.

## 4. Sample applications

4.1. Atmospheric node distribution using surface data. We use the geodata [2] from the collection of global relief datasets produced by NOAA (National Oceanic and Atmospheric Administration), which contains a 1 arc-minute resolution model. We generate a sample configuration consisting of $1,356,566$ nodes distributed uniformly inside an atmospheric-type shell $\Omega_{\text {etopo }}$ : the outer boundary of $\Omega_{\text {etopo }}$ is spherical, the inner one is an interpolation of the relief from ETOPO1 data, exaggerated by a factor of 100 . The scale is chosen so that the average Earth radius, assumed to be

[^2]$6,371,220$ meters, has unit length; the radius of the outer boundary is set to 1.1 , which corresponds to the height of 6,371 meters above the average radius, given the exaggeration factor.

The ETOPO1 dataset stores relief as a $21600-$ by-10 800 array of elevations above the sea level; equivalently, of radial coordinates that correspond to the spherical angles defined by the array's indices. The data points are equispaced on lines of constant azimuth/inclination with angular distance $B=\pi / 10800$ between them. To determine whether a given node $\boldsymbol{x}=\left(r_{\boldsymbol{x}}, a_{\boldsymbol{x}}, p_{\boldsymbol{x}}\right)$ belongs to $\Omega_{\text {etopo }}$, its radial coordinate $r_{\boldsymbol{x}}$ was compared with a linear interpolation of the values of radii of three ETOPO1 points with the nearest spherical coordinates. For example, assume that such three points have the spherical coordinates $\left(r_{j}, a_{j}, p_{j}\right), j=1,2,3$, where $0 \leq a \leq 2 \pi$ and $0 \leq p \leq \pi$ are the azimuth and polar angle, respectively, and

$$
\begin{array}{cccc}
a_{1}=l B, & a_{2}=(l+1) B, & a_{3}=l B, & 0 \leq l \leq 21599 \\
p_{1}=m B, & p_{2}=m B, & p_{3}=(m+1) B, & 0 \leq m \leq 10799
\end{array}
$$

Without loss of generality, the inequalities

$$
l B \leq a_{\boldsymbol{x}}<(l+1) B ; \quad m B \leq p_{\boldsymbol{x}}<(m+1) B
$$

hold true. The point inclusion function is defined in this case as

$$
\chi\left(\boldsymbol{x} ; \Omega_{\text {etopo }}\right)= \begin{cases}1, & r_{1}+\frac{a_{\boldsymbol{x}}-a_{1}}{B}\left(r_{2}-r_{1}\right)+\frac{p_{\boldsymbol{x}}-p_{1}}{B}\left(r_{3}-r_{1}\right)<r_{\boldsymbol{x}}<1.1 \\ 0, & \text { otherwise }\end{cases}
$$

with 1.1 being the radius of the outer sphere in the chosen scale. In effect, the algorithm for evaluating the $\chi\left(\cdot ; \Omega_{\text {etopo }}\right)$ described here coincides with the star-shaped point location algorithm from 46, Section 2.2], applied to the interpolated Earth surface and the outer spherical boundary.

Our node set consists of $1,356,566$ nodes with the nearest-neighbor separation close to the constant $\rho(\boldsymbol{x})=0.01124$, and our top priority was to ensure the low variance of the radial separation across the configuration, especially on the surface; the general view of the set is given in Figure 2, We used the piecewise IL with golden-ratio derived parameters $\alpha_{1}=\sqrt{2}, \alpha_{2}=(\sqrt{5}-1) / \sqrt{2}$; regarding these $\alpha_{1}, \alpha_{2}$ see also the discussion in Appendix. Several statistics of the resulting set are presented in the following table; here again we used the common notation $\bar{x}$ for the averaged value of a quantity $x$. Notation $\Delta^{k}$ stands for the distance to the $k$-th nearest neighbor.

|  | Whole node set | Surface nodes |
| :--- | :---: | :---: |
| $\overline{\Delta^{12}(\boldsymbol{x}) / \Delta^{2}(\boldsymbol{x})}$ | 1.3674 | 2.0353 |
| $\overline{\Delta^{4}(\boldsymbol{x}) / \Delta^{1}(\boldsymbol{x})}$ | 1.0859 | 1.34019 |
| 99 th percentile of $\left\{\Delta\left(\boldsymbol{x}_{i}\right)\right\}$ | 0.012143 | 0.014444 |
| $\overline{\Delta(\boldsymbol{x})}$ | 0.011243 | 0.010879 |
| 1 st percentile of $\left\{\Delta\left(\boldsymbol{x}_{i}\right)\right\}$ | 0.009652 | 0.009340 |

Figure 3 illustrates the distribution of nodes close to the surface of $\Omega_{\text {etopo }}$. No pullback function has been used, just the inclusion check performed as in 10 . Observe that the near-surface nodes display no artifacts, and the spacing does not significantly depend on the local surface shape. The left subplot in the Figure 4 illustrates the effect of Step 5 on the distribution of distances to the nearest neighbor. In the right subplot, we have collected distances to the 12 nearest neighbors for the whole configuration, and separately for the surface subset. The histogram also contains the distribution of hole radii, that is, distances from the Voronoi centers of the entire node configuration to their respective nearest nodes. It is a well-known fact that the Voronoi centers are local maxima


Figure 2. Left: a general view of a uniform node distribution in an atmospheric-like shell. Right: a separate view of the Western hemisphere.


Figure 3. Surface subset: a fragment of the Western coast of South America. The nodes on the right are color-coded using heights.


Figure 4. The effects of the repel procedure and hole radii. Left: probability distribution of the nearest-neighbor distances in the atmospheric node set, before (blue) and after (red) executing the repel subroutine. Right: distribution of distances to the 12 nearest neighbors for the whole configuration (color only), for the surface subset (contours), the hole radii (black dashed contour on the left).


Figure 5. Distribution of distances to the 12 nearest neighbors for the atmospheric node configuration; medians and the 25 th and 75 th percentiles are shown. Left: the surface subset. Right: the whole set. Scales are the same in both subplots.
of the distance from the node set [19], considered as a function on the whole space $\mathbb{R}^{3}$. Note that all the histograms on the right are normalized by probability, not by the node count.

The pair of plots in Figure 5 shows in detail the distribution of distances to the nearest neighbors in the sample node set. It has been produced using the standard Matlab routine boxplot. For each of the blue boxes corresponding to a specific nearest neighbor, the central mark is the median, the edges of the box denote the 25 th and 75 th percentiles. The red crosses mark outliers.
4.2. Point cloud. To demonstrate a nonuniform node distribution using our algorithm, we fix a collection of 100 points, $\mathcal{P}_{100}$, inside the cube $[-1,1]^{3}$, and consider the following radial density function:

$$
\rho(\boldsymbol{x})=\left(\Delta\left(\boldsymbol{x} ; \mathcal{P}_{100}\right)+\Delta^{2}\left(\boldsymbol{x} ; \mathcal{P}_{100}\right)\right) / 20
$$

where, as above, $\Delta^{k}$ is the distance to the $k$-th nearest neighbor. A possible interpretation of this density is a distribution that concentrates about a set of points $\mathcal{P}_{100}$, which are of particular interest for a certain model.

We proceed as in the Step 5 of the algorithm, not using the full gradient expression described in 13 . In fact, it is instructive to note that computing the second term in 13 would be quite cumbersome here in view of $\rho$ being a piecewise-defined function. One could thus consider the density recovery for this distribution, Figure 7, as a validation of the gradient truncation approach in Step 55. cf. Section 3.2. The Q-MC voxels were drawn from the sequence $\left\{\mathcal{L}_{n}\right\}$ with the same lattice parameters as in Section 4.1, $\alpha_{1}=\sqrt{2}, \alpha_{2}=(\sqrt{5}-1) / \sqrt{2}$.

Figure 7 contains the distribution of the ratio $\rho(\boldsymbol{x}) / \Delta(\boldsymbol{x})$. The minimal and maximal values of the ratio are about 0.8099 and 1.8231 respectively; its mean value is 0.9797 , and the variance is 0.0019 . The 5 - and 95 -percentiles are 0.9208 and 1.0441, respectively; the right plot in Figure 6 highlights the outliers in the ratio distribution.


Figure 6. Left: the node set from Section 4.2. Right: node locations that contribute to the distribution of the ratio $\rho(\boldsymbol{x}) / \Delta(\boldsymbol{x})$ beyond the 5- and 95-percentiles.


Figure 7. Distribution of the ratios $\rho(\boldsymbol{x}) / \Delta(\boldsymbol{x})$ for the node set in Section 4.2
4.3. Spherical shell. The motivation for this example comes from atmospheric modeling. Representing the Earth surface by a sphere, we consider first a thin 3-dimensional shell $\Omega_{\text {shell }}$ of inner radius $R_{\text {inner }}$ and outer radius $R_{\text {inner }}+H_{\text {atm }}$ with constant target separation $h$ between points in the radial (vertical) direction, and the tangential (horizontal) separation to be $\tau(r)=C \cdot r$ at radius $r$, for some constant $C$. With typical choices of parameters, $\tau$ will be much larger than $h$, reflecting the much higher resolution needed in the vertical direction due to $H_{\text {atm }} \ll R_{\text {inner }}$. We make a radial change of variables, which can written in spherical coordinates as $(r, a, p) \mapsto(\hat{r}(r), a, p)$, so that any configuration in $\Omega_{\text {shell }}$ having the 2-directional resolutions $\tau(r)$ and $h$ will have isotropic resolution after the transformation. It is much easier to construct RBF bases in the isotropic case, hence our deliberation.


Figure 8. Left: the node set from Section 4.3 Right: node locations that contribute to the distribution of the ratio $\rho(\boldsymbol{x}) / \Delta(\boldsymbol{x})$ beyond the 5- and 95-percentiles.

Following this change of variables, the radial/tangential node separations become, respectively,

$$
\begin{align*}
& \hat{\nu}(r)=h \cdot \hat{r}^{\prime}(r) \\
& \hat{\tau}(r)=C \cdot \hat{r}(r) \tag{14}
\end{align*}
$$

Setting these two quantities to be equal, we obtain the ODE

$$
\hat{r}^{\prime}(r)=\frac{C}{h} \cdot \hat{r}(r)
$$

with initial condition $\hat{r}\left(R_{\text {inner }}\right)=1$, and its solution becomes

$$
\hat{r}(r)=\exp \left(C \cdot \frac{r-R_{\text {inner }}}{h}\right) .
$$

From the second equation in (14) follows that our goal is to generate a node set in the ( $\hat{r}, a, p$ )-space, whose separation is proportional to $\hat{r}$ and is equal in all directions: $\rho(\boldsymbol{x})=C \cdot\|\boldsymbol{x}\|$. The outer radius of the image of $\Omega_{\text {shell }}$ in the $(\hat{r}, a, p)$-space is a function of $R_{\text {inner }}$ and $H_{\text {atm }}$; our model implies $R_{\text {inner }}=6,371,220$, the mean radius of the Earth in meters, and $H_{\text {atm }}=12,000$, the thickness of the atmospheric layer we are interested in. The constant $C$ is determined by the desired tangential separation at the $r=R_{\text {inner }}$ level.

Say, we intend to generate nodes corresponding to the 2 degree resolution on the spherical "Earth surface" and $h=400$ meter vertical resolution. Due to the peculiarities of atmospheric modeling, we would like to fix two much denser sets of nodes on the inner and outer boundary of $\Omega_{\text {shell }}$; specifically, we are using 12,100 approximate Riesz energy minimizers on a sphere, appropriately rescaled. The interior nodes are generated using our algorithm, and then Step 5 is modified so as to leave the boundary subset intact. This, however, causes a difficulty: the separation distances between the interior and the surface nodes must remain large; on the other hand, our generic formulation of Step 5 does not account for the much higher concentration of nodes on the surface, which causes excessive repelling force, seen in the oscillations of the radial distribution, Figure 9. Mitigating this effect requires artificially weakening the repulsive force caused by the boundary nodes, a straightforward task using our codebase. Instead, we show in Figure 9 the performance of the generic algorithm, to illustrate complications that may arise when applying it to specialized problems.


Figure 9. Left: Distribution of the ratios $\rho(\boldsymbol{x}) / \Delta(\boldsymbol{x})$ for the node set in Section 4.3 Right: Radial node distribution, actual (blue) and the theoretic (red) continuous component; i.e., without the $\delta$-function spikes at the endpoints.

The set $\Omega_{\text {shell }}$ can be challenging for the basic form of our algorithm, as described in Section 3.1 . obtaining satisfying convergence requires using the full version of gradient descent 13 . The reasons for it being more difficult to tackle than, say, $\Omega_{\text {etopo }}$ in Section 4.1, are that due to convexity of the outer boundary, the weighted s-energy minimizers on it are denser than on the sphere with radius $\hat{r}\left(R_{\mathrm{inner}}+H_{\mathrm{atm}}\right)-10^{-3}$, for example; see also discussion at the end of Section 2.2. Getting rid of the artifacts at the endpoints of the radial distribution is done by using the full gradient, weakening the repulsion of the fixed boundary nodes, and not striving for the full convergence of a minimization method applied to the Riesz energy.

In this example, we used the $\left\{\mathcal{M}_{n}\right\}$ sequence to fill individual voxels. The left subplot in Figure 9 contains the distribution of the ratio $\rho(\boldsymbol{x}) / \Delta(\boldsymbol{x})$. The minimal and maximal values of the ratio are about 0.9165 and 1.8989 respectively; its mean value is 1.0226 , and the variance is 0.0024 . The $5-$ and 95 -percentiles are 0.9782 and 1.0717 , respectively.

| Example | $K$ | $T$ | $N$ | Q-MC distribution times, s | Repel times, s |
| :--- | :---: | :---: | :---: | :---: | :---: |
| Atmospheric nodes | 33 | 29 | $1,356,566$ | 5 | 89 |
| Point cloud | 30 | 200 | 577,321 | 4 | 840 |
| Spherical shell | 30 | 200 | 358,915 | 1 | 144 |

Table 2. Timings of the examples in Sections 4.1 4.3
4.4. Run times. The execution times (in seconds) for the above examples are summarized in Table 2, where, as before, $K$ and $T$ stand for the number of nearest neighbors and the number of iterations used in the repel procedure in Step 5, respectively. The fifth column contains times required to fill the voxels selected at Step 1 with configurations from either $\left\{\mathcal{L}_{n}\right\}$ or $\left\{\mathcal{M}_{n}\right\}$ and to remove any redundant nodes as in Step 4. All the computations were performed on a dedicated machine with 40 GB RAM and an 8-core Intel Xeon CPU. Note that the basic Q-MC node sets for


Figure 10. Average condition numbers of the joint RBF-FD PHS-based matrices for order 1- and 2-differential operators. Left: Between 20 and 100 nodes in the stencil. Right: Between 100 and 200 nodes in the stencil.
both sequences were precomputed, and the pre-computation times are not included in the table. Computation of configurations in $\left\{\mathcal{L}_{n}\right\}$ for $1 \leq n \leq 200$ took less than 1 second. An implementation of the $\left\{\mathcal{M}_{n}\right\}$ sequence for $1 \leq n \leq 200$ took 4311 seconds to generate; coordinates of the resulting minimizers as well as the corresponding average separation distances are distributed with the associated codebase 52.

## 5. Final observations and comparisons

5.1. Comparisons. Of the two Q-MC sequences we considered, the periodic Riesz minimizers appear more promising, being devoid of the lattice structure and having high space utilization. On the other hand, we have successfully applied ILs as an elementary uniform configuration. One could use different sets of irrational parameters $\alpha_{1}, \ldots, \alpha_{d-1}$ for different numbers of nodes in a voxel. Although this might be useful in mitigating the non-isotropic behavior of ILs, it makes hard to control node separation at voxel interfaces.

Another quasi-uniform node set commonly used in Q-MC methods can be constructed from the Halton sequence [41, Chapter 5.4], an example of a low-discrepancy sequence. To see how the Halton nodes compare to Riesz minimizers for RBF-FD methods, we have computed condition numbers of PHS-based RBF-FD matrix in the LHS of (4), using operators $\partial / \partial x_{i}, i=1,2,3$, and $\partial^{2} / \partial x_{i} \partial x_{j}, i \leq j=1,2,3$ as $\mathfrak{L}$; the resulting LHS then constitutes the joint system for the weights corresponding to these nine differential operators. We used the RBF kernel $\phi(r)=r^{5}$, and the polynomials in the interpolation space were of degree at most 2, see Section 2.1. The computations were performed for the Riesz and Halton nodes, and the uniform Cartesian grid. The stencils consisted of $K$ nearest neighbors of a random vector with Gaussian distribution, centered around $(0.5,0.5,0.5)^{\operatorname{tr}}$; the evaluation point $\boldsymbol{x}_{0}$ was taken equal to the random vector itself. The nearest neighbors were drawn from 1000 nodes of the respective sequence, uniformly distributed over the unit cube. The Riesz nodes were produced by minimizing periodic energy (5) with the distance (9).

Figure 10 contains a comparison of the condition numbers of RBF-FD matrices for the three sequences. The values shown are averages of the condition numbers for 500 random stencil centers
$x_{0}$; the averaging was introduced to eliminate the rather unpredictable dependence on $\boldsymbol{x}_{0}$, and to display the underlying trend. We omitted values of $K$ below 20 from the plots, as all the three node sets resulted in relatively ill-conditioned systems; this was to be expected, as the recommended stencil size is roughly twice the number of linearly independent polynomials in the interpolation space (there are 10 monomials of degree at most 2 in $\mathbb{R}^{3}$ ) [23].
5.2. Range of applications. Our method has proven very efficient for slowly varying radial densities that are small (recall that small radial density means a large number of nodes per unit volume) compared to the entire node set scale, and is capable of handling very complex underlying sets. The range of dimensions where the algorithm can be used efficiently is determined by the applicability of Q-MC initialization and the nearest neighbor searches: the repelling iterations for Riesz energy in Step 5 are largely (with a proper value of $s$ ) dimension-agnostic. A shortcoming that is common to all quasi-Monte Carlo methods (but of little practical relevance) is a much worse performance (measured by $L^{2}$ discrepancy), compared to Monte-Carlo distribution, in dimensions starting at about 15 15. Furthermore, using the uniform grid to detect the support $\Omega$, as is done in Step 1 Step 2, becomes unfeasible already for $d=10$; instead, one needs an efficient way to determine which corners of the grid are in some sense close to $\Omega$. This is certainly not a feature of our approach, but a manifestation of the curse of dimensionality: treatment of a complicated high-dimensional set is a computationally intensive task. Regarding finding nearest neighbors it should be noted that common implementations of k -d trees are efficient only up to about $d=20$; additionally, the k -d tree approach is faster than the full brute force search only if $N \gg 2^{d}$ [35]. On the other hand, as has already been noted, our repelling procedure does not require frequent updates of the search tree, as the local adjacency largely remains intact.

The suggested algorithm is very local, and it therefore must be straightforward to add multiresolution and adaptive refinement, as is widely done for grids [18, 20, yet as of this writing, our proof-of-concept implementation does not include these features. Still, we would like to observe that refining the voxel structure is indeed easier than refining a mesh, since no geometry is taken into account. This partially addresses the previous remark on detection of $\Omega$ in high dimensions.

The closest set of goals to what we have presented here, that we're aware of, is posed in the pioneering paper [50]; our method is crafted for full-dimensional domains, and apparently performs faster in this case. The bubble packing algorithm in 50 is conceptually similar to the greedy filling of centers in Step 3, while physical relaxation is an alternative to the energy minimization we employ; of course, the idea of relaxation can also be found in a number of related references, and is a well-known approach in this context, see for example [44. Our method requires computing the gradient of the desired radial density in the cases when the outer boundary of the underlying set is uniformly convex, and/or when the radial density changes quickly. Alternatively, fine partition of the set is necessary. Either solution, however, may be computationally expensive.

## 6. Appendix: separation properties of sequences $\left\{\mathcal{L}_{n}\right\}$ and $\left\{\mathcal{M}_{n}\right\}$

This Appendix deals with the results of our numerical experiments, set in the 3-dimensional space. The function $\lambda(r)$ used in Step 2 is the number of nodes in the unit cube $[0,1]^{3}$, placed according to (8), or obtained by minimizing the Riesz $s$-energy (5) with periodic metric, such that the mean separation distance of these nodes is the closest to $r$. To compute $\lambda(r)$ for the periodization of $\left\{\mathcal{L}_{n}\right\}$, we tabulate mean separations $\bar{\Delta}_{n}$ in a sample configuration comprising $\mathcal{L}_{n}$ and its $26=3^{3}-1$ copies, obtained translating $\mathcal{L}_{n}$ by the vectors $\left\{(i, j, k)^{\text {tr }}: i, j, k \in\{0, \pm 1\}\right.$ and $\left.|i|+|j|+|k|>0\right\}$. The tabulated dependence of separation on $n$ is then inverted and interpolated using a piecewise cubic Hermite interpolating polynomial. The reason to consider separation distance between configurations


Figure 11. Left: dependence of the mean separation distances on the number of nodes in the unit cube for different values of parameters $\alpha_{1}, \alpha_{2}$; the $n^{-1 / d}$ decay rate shown as a dash-dot line. Right: ratios of the mean separation distances to the minimal ones for the same configurations.



Figure 12. Left: A cross-section of the IL $\mathcal{L}_{100}$ with parameters 0.179373654819913 and 0.531793804909494 . Right: A (different) cross-section of $\mathcal{M}_{100}$.
in $3^{d}$ cubes in dimension $d$ (and not a single cube with a single instance of $\mathcal{L}_{n}$ ) is to account for the boundary effects. Likewise, to compute $\lambda(r)$ for the Riesz minimizers, the mean separation of $\mathcal{M}_{n}$ is tabulated for $1 \leq n \leq n_{\text {max }}$, then the inverse dependence is interpolated. No copies of $\mathcal{M}_{n}$ are considered alongside the original configuration, since periodicity condition is already included in the metric (9).

In general, putting too many nodes in individual voxels is justified only if the radial density function $\rho$ varies slowly. For our applications, $n_{\max } \leq 100$ was sufficient. The left plot in Figure 11 illustrates the delicate dependence of the separation distances of ILs on the lattice parameters. While any set of irrational quantities $\alpha_{1}, \ldots, \alpha_{d-1}$ in (8) that are linearly independent over rationals will give a uniformly distributed IL as $n$ grows, certain values may perform better than the others. In particular, adjustments can be made to improve the distribution for small values of $n$. For
example, it is known from [5] that a 2-dimensional IL generated by the golden ratio has optimal $L^{2}$ discrepancy. Numerical experiments have shown that its 3-dimensional analog with parameters $\alpha_{1}=\sqrt{2}, \alpha_{2}=(\sqrt{5}-1) / \sqrt{2}$ does perform well for large numbers of nodes; yet by carrying out a Monte Carlo search for the parameters maximizing separation distance in (8), we found several (necessarily rational) pairs that performed at least just as well for up to $n=200$, see Figure 11 .

Curiously enough, a pair of random numbers drawn uniformly from $[0,1]$ (shown in the legend as rand(1)), consistently performed better than the pair $\sqrt{3}$ and $\sqrt{5}$, starting at $n \approx 40$. We were able to reproduce this behavior in a number of runs; in fact, we haven't seen a random pair that wouldn't always outperform $\sqrt{3}$ and $\sqrt{5}$ after a fairly small $n$.

The second graph in Figure 11 shows the ratios of the mean to minimal separation distances $\bar{\Delta}_{n} / \Delta_{n}$ for the same range of $n$. In both subfigures, Riesz periodic minimizers clearly stand out, by having the largest mean separation (left), and by smallest ratios (right). This means, the nearest neighbor distances $\Delta(\boldsymbol{x})$ vary little from node to node in the $\left\{\mathcal{M}_{n}\right\}$ sequence. We conclude this section by presenting in Figure 12 a pair of cross-sections of the IL $\mathcal{L}_{100}$ and the configuration $\mathcal{M}_{100}$ that look remarkably similar. In fact, we found the vague resemblance between the low-energy periodic configurations and lattice structures, similar to ILs, quite interesting, given the connection between packing and Riesz energy minimization [37, and that the highest packing density in the 3 -dimensional space is achieved, in particular, by the hcp lattice [19].

## Acknowledgements

We thank Edward Saff and Douglas Hardin for their interest and useful comments. O.V. and T.M. were supported, in part, by the NSF grant DMS-1516400; they would also like to express their gratitude to the University of Colorado, Boulder, where a part of this work was completed. The National Center for Atmospheric Research is sponsored by the National Science Foundation.

## References

[1] Alishahi, K., and Zamani, M. The spherical ensemble and uniform distribution of points on the sphere. Electron. J. Probab. 20 (2015).
[2] Amante, C., and Eakins, B. W. ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis. https://doi.org/10.7289/V5C8276M 2009.
[3] Bayona, V., Flyer, N., Fornberg, B., and Barnett, G. A. On the role of polynomials in RBF-FD approximations: II. Numerical solution of elliptic PDEs. J. Comput. Phys. 332 (2017), 257-273.
[4] Beltrán, C., Marzo, J., and Ortega-CerdÀ, J. Energy and discrepancy of rotationally invariant determinantal point processes in high dimensional spheres. J. Complex. 37 (2016), 76-109.
[5] Bilyk, D. The $L^{2}$ discrepancy of irrational lattices. In Monte Carlo and Quasi-Monte Carlo Methods 2012. Springer, 2013, pp. 289-296.
[6] Bilyk, D., Temlyakov, V. N., and Yu, R. The $L^{2}$ Discrepancy of Two-Dimensional Lattices. In Springer Proc. Math. Stat., vol. 25. 2012, pp. 63-77.
[7] Bollig, E. F., Flyer, N., and Erlebacher, G. Solution to PDEs using radial basis function finite-differences (RBF-FD) on multiple GPUs. J. Comput. Phys. 231, 21 (2012), 7133-7151.
[8] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Asymptotics for discrete weighted minimal Riesz energy problems on rectifiable sets. Trans. Amer. Math. Soc. 360, 3 (2008), 1559-1580.
[9] Borodachov, S. V., Hardin, D. P., and Saff, E. B. Low Complexity Methods For Discretizing Manifolds Via Riesz Energy Minimization. Found. Comput. Math. 14, 6 (2014), 1173-1208.
[10] Brauchart, J. S., Hardin, D. P., and Saff, E. B. The next-order term for optimal Riesz and logarithmic energy asymptotics on the sphere. Contemp. Math 578 (2012), 1-31.
[11] Brauchart, J. S., Reznikov, A. B., Saff, E. B., Sloan, I. H., Wang, Y. G., and Womersley, R. S. Random point sets on the sphere-hole radii, covering, and separation. Experimental Mathematics (2016), 1-20.
[12] Broomhead, D. S. and Lowe, D. Multivariable Functional Interpolation and Adaptive Networks. Complex Syst. 2 (1988), 321-355.
[13] Buhmann, M. D. Radial Basis Functions. Cambridge University Press, Cambridge, 2003.
[14] Buhmann, M. D., Dinew, S., and Larsson, E. A note on radial basis function interpolant limits. IMA J. Numer. Anal. 30, 2 (2010), 543-554.
[15] Caflisch, R. E. Monte Carlo and quasi-Monte Carlo methods. Acta Numerica 7 (1998), 1.
[16] Chang, C., and Lin, C. LIBSVM. ACM Trans. Intell. Syst. Technol. 2, 3 (2011), 1-27.
[17] Cheney, W., and Light, W. A course in approximation theory, vol. 101 of Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2009. Reprint of the 2000 original.
[18] Clawpack Development Team. Clawpack software, 2017. Version 5.4.0.
[19] Conway, J. H., and Sloane, N. J. A. Sphere packings, lattices and groups, third ed., vol. 290 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, New York, 1999. With additional contributions by E. Bannai, R. E. Borcherds, J. Leech, S. P. Norton, A. M. Odlyzko, R. A. Parker, L. Queen and B. B. Venkov.
[20] Debreu, L., Vouland, C., and Blayo, E. AGRIF: Adaptive grid refinement in fortran. Computers $\mathcal{G}$ Geosciences 34, 1 (2008), 8-13.
[21] Fasshauer, G. E. Meshfree approximation methods with MATLAB, vol. 6 of Interdisciplinary Mathematical Sciences. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2007. With 1 CD-ROM (Windows, Macintosh and UNIX).
[22] Flyer, N., Barnett, G. A., and Wicker, L. J. Enhancing finite differences with radial basis functions: Experiments on the Navier-Stokes equations. J. Comput. Phys. 316 (jul 2016), 39-62.
[23] Flyer, N., Fornberg, B., Bayona, V., and Barnett, G. A. On the role of polynomials in RB-FD approximations: I. Interpolation and accuracy. J. Comput. Phys. 321 (2016), 21-38.
[24] Flyer, N., Lehto, E., Blaise, S., Wright, G. B., and St-Cyr, A. A guide to RBF-generated finite differences for nonlinear transport: Shallow water simulations on a sphere. J. Comput. Phys. 231, 11 (2012), 4078-4095.
[25] Flyer, N., Wright, G. B., and Fornberg, B. Radial Basis Function-Generated Finite Differences: A Mesh-Free Method for Computational Geosciences. Handb. Geomathematics (2013), 1-30.
[26] Fornberg, B. A practical guide to pseudospectral methods. Cambridge University Press, 1996.
[27] Fornberg, B. Calculation of Weights in Finite Difference Formulas. SIAM Rev. 40, 3 (1998), 685-691.
[28] Fornberg, B., and Flyer, N. A Primer on Radial Basis Functions with Applications to the Geosciences. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2015.
[29] Fornberg, B., and Flyer, N. Fast generation of 2-D node distributions for mesh-free PDE discretizations. Comput. Math. with Appl. 69, 7 (2015), 531-544.
[30] Fornberg, B., and Flyer, N. Solving PDEs with radial basis functions. Acta Numerica 24 (2015), 215-258.
[31] Fornberg, B., Larsson, E., and Flyer, N. Stable Computations with Gaussian Radial Basis Functions. SIAM J. Sci. Comput. 33, 2 (jan 2011), 869-892.
[32] Fornberg, B., Lehto, E., and Powell, C. Stable calculation of Gaussian-based RBF-FD stencils. Comput. Math. with Appl. 65, 4 (feb 2013), 627-637.
[33] Fornberg, B., Wright, G., and Larsson, E. Some Observations Regarding Interpolants in the Limit of Flat Radial Basis Functions. Comput. Math. with Appl. 47, 1 (2004), 37-55.
[34] Fornberg, B., and Zuev, J. The Runge phenomenon and spatially variable shape parameters in RBF interpolation. Comput. Math. with Appl. 54, 3 (aug 2007), 379-398.
[35] Freidman, J. H., Bentley, J. L., and Finkel, R. A. An Algorithm for Finding Best Matches in Logarithmic Expected Time. ACM Trans. Math. Softw. 3, 3 (1977), 209-226.
[36] Frey, P. J., and George, P.-L. Quadtree-octree Based Methods. ISTE, 2010, pp. 163-199.
[37] Hardin, D. P., and Saff, E. B. Minimal Riesz Energy Point Configurations for Rectifiable d-Dimensional Manifolds. Adv. Math. 193, 1 (2003), 174-204.
[38] Kuipers, L., and Niederreiter, H. Uniform distribution of sequences. Dover Publications, Mineola, N.Y, 2006.
[39] Landkof, N. S. Foundations of modern potential theory, vol. 180. Springer, 1972.
[40] Larsson, E., Lehto, E., Heryudono, A., and Fornberg, B. Stable Computation of Differentiation Matrices and Scattered Node Stencils Based on Gaussian Radial Basis Functions. SIAM J. Sci. Comput. 35, 4 (jan 2013), A2096-A2119.
[41] Lemieux, C. Monte Carlo and Quasi-Monte Carlo Sampling. Springer Series in Statistics. Springer New York, New York, NY, 2009.
[42] Link, W. A., and Eaton, M. J. On thinning of chains in MCMC. Methods Ecol. Evol. 3, 1 (feb 2012), $112-115$.
[43] Micchelli, C. A. Interpolation of scattered data: Distance matrices and conditionally positive definite functions. Constr. Approx. 2, 1 (1986), 11-22.
[44] Persson, P.-O., and Strang, G. A Simple Mesh Generator in MATLAB. SIAM Rev. 46, 2 (jan 2004), 329-345.
[45] Powell, M. J. D. Five Lectures on Radial Basis Functions. IMM Lect., December 2004 (2004), 27.
[46] Preparata, F. P., and Shamos, M. I. Computational Geometry. Springer New York, 1985.
[47] Procopiuc, O., Agarwal, P. K., Arge, L., and Vitter, J. S. Bkd-tree: A dynamic scalable kd-tree. In Advances in Spatial and Temporal Databases. Springer Berlin Heidelberg, 2003, pp. 46-65.
[48] Schoenberg, I. J. Metric Spaces and Completely Monotone Functions. Ann. Math. 39, 4 (1938), 811.
[49] Shankar, V., Wright, G. B., Kirby, R. M., and Fogelson, A. L. A radial basis function (RBF)-finite difference (FD) method for diffusion and reaction-diffusion equations on surfaces. Journal of Scientific Computing 63, 3 (2014), 745-768.
[50] Shimada, K., and Gossard, D. C. Bubble mesh. In Proc. third ACM Symp. Solid Model. Appl. - SMA '95 (New York, New York, USA, 1995), ACM Press, pp. 409-419.
[51] Varma, U. M., Rao, S. V. R., and Deshpande, S. M. Point distribution generation using hierarchical data structures. In Proc. ECCOMAS 2004 (Jyväskylä, 2004).
[52] Vlasiuk, O., and Michaels, T. Boxed lattices and Riesz minimizers for RBF computations. https://github. com/OVlasiuk/3dRBFnodes.git, 2017. [Online; accessed 12-October-2017].
[53] Wendland, H. Scattered Data Approximation. Cambridge University Press, 2004.
[54] Zhou, K., Hou, Q., Wang, R., and Guo, B. Real-time KD-tree construction on graphics hardware. ACM Trans. Graph. 27, 5 (dec 2008), 1.

Center for Constructive Approximation, Department of Mathematics, Vanderbilt University, Nashville, TN 37240, USA

E-mail address: oleksandr.vlasiuk@vanderbilt.edu
E-mail address: timothy.j.michaels@vanderbilt.edu
Analytics and Integrated Machine Learning, National Center for Atmospheric Research, Boulder, CO 80305, USA

E-mail address: flyer@ucar.edu
Department of Applied Mathematics, University of Colorado, Boulder, CO 80309, USA
E-mail address: bengt.fornberg@colorado.edu


[^0]:    Date: April 19, 2022.
    Key words and phrases. Quasi-uniformity, irrational lattice, quasi-Monte Carlo methods, Riesz energy, gradient descent, separation distance, nearest neighbors.

[^1]:    ${ }^{1}$ Note that both ILs and the minimizers can have the nearest neighbor distance of at most 1 , due to periodicity. We therefore take $\lambda(x)=0$ whenever $x>1$.

[^2]:    ${ }^{2}$ The Matlab code we provide performs naive autotuning of $K$ and $T$, using the total number of nodes to be placed. Although sufficient for demonstration purposes, there is room for improvement.

