

SANDIA REPORT

SAND2015-7976

Unlimited Release

Printed September 2015

Efficient Probability of Failure Calculations for QMU using Computational Geometry LDRD 13–0144 Final Report

Scott A. Mitchell, Mohamed S. Ebeida, Vicente J. Romero, Laura P. Swiler, Ahmad A. Rushdi, Ahmed Abdelkader

Prepared by

Sandia National Laboratories

Albuquerque, New Mexico 87185 and Livermore, California 94550

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

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



Efficient Probability of Failure Calculations for QMU using Computational Geometry LDRD 13–0144 Final Report

Scott A. Mitchell
Sandia National Laboratories
Albuquerque, NM 87185
samitch@sandia.gov

Mohamed S. Ebeida
Sandia National Laboratories
Albuquerque, NM 87185

Vicente J. Romero
Sandia National Laboratories
Albuquerque, NM 87185

Laura P. Swiler
Sandia National Laboratories
Albuquerque, NM 87185

Ahmad A. Rushdi
University of Texas at Austin
Austin, TX 78712

Ahmed Abdelkader
University of Maryland
College Park, MD 20742

Abstract

This SAND report summarizes our work on the Sandia National Laboratory LDRD project titled “Efficient Probability of Failure Calculations for QMU using Computational Geometry” which was project #165617 and proposal #13–0144. This report merely summarizes our work. Those interested in the technical details are encouraged to read the full published results, and contact the report authors for the status of the software and follow-on projects.

Acknowledgment

We thank Sandia's LDRD program, and especially the Computational and Information Sciences Investment Area Team (CIS IAT) for supporting this work. We thank our management, especially James Stewart, Michael Parks, Scott Collis and Bruce Hendrickson for championing this work and helping ensure that the work continues after this funding project ends.

The report authors wish to thank all who participated in the project or discussed related ideas with us. At Sandia these included Scott A. Mitchell, Mohamed S. Ebeida, Vicente J. Romero, Keith Dalbey, Laura P. Swiler, Michael S. Eldred, Brian M. Adams (and the entire DAKOTA team), Shawn English, Christopher Valicka, Peter Bosler, Eric Phipps, Marta D'Elia, George Orient, Joseph Bishop, and Tim Wildey.

We thank our university partners. At UC Davis these included John D. Owens, Anjul Patney, Andrew A. Davidson, Stanley Tzeng, and Saman Ashkiani. At UNC these included Dinesh Manocha, Chonyhon Park, At UT Austin these include Chandrajit Bajaj, Alexander Rand, and Ahmad Rushdi. We thank Ahmed Abdelkader Abdelrazek (Univ. Maryland student) for spending the summer of 2015 with us. We especially thank our student partners at the University of Alexandria: Mohammed A. Mohammed, Muhammad A. Awad, Ahmed Hassan Mahmoud. We thank others at various universities and companies: Don Sheehy (Univ Connecticut); Gary Miller (CMU); David Eppstein (UC-Irvine), Li-Yi Wei (HKU); Xiaoyin Ge (OSU); Dong-Ming Yan and Peter Wonka (KAUST).

Contents

Summary	9
1 Technical Results, Publications and Talks	11
1.1 Optimization and UQ Specific Results	11
1.1.1 POF-Darts	11
1.1.2 Voronoi Piecewise Surrogates	12
1.1.3 Methods Comparison	14
1.2 General Sampling, Across Applications	16
1.2.1 Hyperplane Sampling	16
1.2.2 Recursive k -d Darts	16
1.2.3 Spoke Darts	18
1.2.4 Randomness and Optimization	19
Noise and Coverage, Optimizing Voronoi Aspect Ratio	19
Noise and Cardinality, Resampling	20
Resampling for Modeling and Mesh Properties	20
1.2.5 High Dimensional Data Analysis	21
1.3 Meshing	22
1.3.1 Planar Meshing	22
Robust All-Quad Meshing of Domains with Connected Regions	22
Delaunay Quadrangulation	22
Steiner Point Reduction	24
1.3.2 Sizing Function Discovery	25

1.3.3	VoroCrust, Voronoi Meshing And Surface Reconstruction	25
1.4	Presentations without Papers	29
1.4.1	Minisymposia	30
1.4.2	Invited Talks	30
1.4.3	Regular Talks	30
1.4.4	Internal Talks	31
2	Partnerships	33
2.1	Innovative Interactions for Innovative Ideas—Meshing and Computer Graphics Relevance	33
2.2	University Partnerships	33
2.2.1	University of California, Davis, Prof. John Owens	33
2.2.2	University of North Carolina, Prof. Dinesh Manocha	34
2.2.3	University of Texas, Austin, Prof. Chandrajit Bajaj	34
2.3	University Interactions	34
3	Software	35
3.1	Implementations in DAKOTA	35
3.1.1	Probability Of Failure (POF)-Darts	35
3.1.2	Voronoi Piecewise Surrogates (VPS)	35
3.1.3	Recursive k -d Darts (RKD)	36
4	Conclusion	37
	References	38

List of Figures

1.1	Applying POF-Darts to a test problem with four failure regions. After exhausting the sampling budget a GP surrogate is built and utilized to estimate the POF. The failure isocontours from the surrogate are colored in red while the exact one are colored in blue. For this test problem in both cases the isocontours were on top of each other indicating an accurate estimate.	12
1.2	Comparing the sampling phase of LHS (top row) to POF-Darts (bottom row) on the Herbie problem. From left to right in both rows, the number of samples are 100, 500, 1000, and 5000, respectively. LHS does a great job of space-filling the sample points, but POF-Darts focuses the samples near the failure boundaries.	13
1.3	A two-dimensional example of VPS. Top row has sample points and cells. Bottom row shows neighbors and samples of the surrogate falling into cells.	14
1.4	VPS surrogates of various orders, compared to the true function and its isocontours.	15
1.5	Hyperplane (line, plane, etc.) samples are more expensive to compute, but provide more information. They are also much more likely to intersect a subregion of interest, especially if that subregion is long and thin.	17
1.6	A hyperplane sample is composed of one-dimension-lower hyperplane samples; projected onto the line of the fixed coordinates of the sub-samples, the subsamples lie on a line, and the hyperplane value is some average or surrogate through those points. This recasts general high dimensional problems as a series of one dimensional problems.	18
1.7	Spoke dart sampling and applying it to build a global optimization algorithm.	19
1.8	Final mesh for a periodic unit box after applying various methods to the input MPS in (a). While CVT and DistMesh improve the quality of the majority of the Voronoi cells, they tend to lose randomness at larger values of β , which measures the Voronoi cell aspect ratio bounds. Far-Point on the other hand tends to violate the coverage condition. None of the four methods were able to achieve more regular cells than $\beta < 0.7$ in general.	20
1.9	Sifted disks is based on local constrained resampling. We attempt to remove two adjacent points, and replace it with one. However, the replacement must not leave any part of the domain uncovered.	21

1.10	Removing obtuse angles in planar meshes using “Disk Density Tuning of a Maximal Random Packing.”	22
1.11	Removing obtuse angles in curved surface meshes using “Disk Density Tuning of a Maximal Random Packing.”	23
1.12	Fiber material modeling using “Disk Density Tuning of a Maximal Random Packing.” Tuned disks predict the failure point better than traditional hexagonal or square packings.	23
1.13	A conceptual space for sampling methods.	24
1.14	“Robust All-Quad Meshing of Domains with Connected Regions” algorithm steps.	25
1.15	“Robust All-Quad Meshing of Domains with Connected Regions” examples.	26
1.16	Delaunay quadrangulation (top) from two-color disks (bottom light and dark circles). Left: random. Right: advancing front.	27
1.17	The standard Delaunay refinement (DR) technique produces good quality elements, but many of them. By sifting the mesh, we were able to reduce the number of elements by 78%, while preserving the minimum and maximum angles.	27
1.18	Disk-packing sampling meshes discovered the local feature size, without being guided by a prespecified sizing function. Disk radii (inter-point distances) are about the constant L times the local feature size. There is a natural tradeoff between L and the minimum angle α in the mesh.	28
1.19	An example hex dominant 3D Voronoi mesh reconstructing a surface.	29
1.20	VoroCrust algorithm steps, illustrated for 2D and 3D domains.	29
3.1	DAKOTA’s software design for surrogates	36

Summary

The intent of the project was to explore using computational geometry and computer graphics techniques to discover new algorithms for Quantified Margins of Uncertainty (QMU), a form of uncertainty quantification. Our core techniques were sampling, notably hyper-line sampling and sphere packing sampling, together with Voronoi diagrams and surrogate functions. Many times the applications involved numerical integration and surrogate functions of some form. A number of mesh generation applications also emerged, using the same core sampling techniques and Voronoi diagrams.

The project was successful, in terms of spawning non-traditional and practical approaches, publishing, and gaining external visibility. The project remained relevant to Sandia's internal missions as well, broadly those involving simulation based science and engineering. The project had a strong cross-disciplinary approach. The project nominally consisted of two Center for Computing Research staff researching and developing software about half-time, together with part-time efforts from Engineering Sciences staff for method applications. We also had a university partnership in each year, which allowed us to gain insights into the techniques of other fields (e.g. computer graphics and geometric modeling), and enabled the application of our core algorithms to other fields.

Many of the techniques were deployed in stand-alone demonstration software. Some of the more practical capabilities, notably some uncertainty quantification and optimization algorithms, together with a library for analytic surrogate functions for interpolating data, were deployed in Sandia's optimization toolkit DAKOTA. A number of natural follow-on projects emerged.

Chapter 1

Technical Results, Publications and Talks

Here we list our papers, in various states of publication, including appeared, accepted, in revision, and in preparation.

1.1 Optimization and UQ Specific Results

1.1.1 POF-Darts

A major contribution of this LDRD has been the development of a method to calculate failure probabilities called “POF-Darts” (Probability of Failure Darts). POF-Darts is based on random disk-packing in the uncertain parameter space. POF-Darts uses hyperplane sampling to explore the unexplored part of the uncertain space. We use the function evaluation at a sample point to determine whether it belongs to failure or non-failure regions, and surround it with a protection sphere region to avoid clustering. Decomposing the domain into Voronoi cells around the function evaluations as seeds, our strategy to choose the radius of the protection sphere depends on the local Lipschitz continuity. For each cell, we estimate the local Lipschitz continuity using the function values of the significant Voronoi neighbors, and create a sphere to protect its neighborhood from future sampling. As sampling proceeds, regions uncovered with spheres will shrink, improving the estimation accuracy. After exhausting the function evaluation budget, we build a surrogate model using the function evaluations associated with the sample points and estimate the probability of failure by exhaustive sampling of that surrogate. In comparison to popular methods used to estimate failure probabilities, our algorithm has the following advantages:

1. Decoupling the sampling step from the surrogate construction step
2. The ability to reach target POF values with fewer samples
3. The capability of estimating the number and locations of disconnected failure regions in addition to the probability of failure estimate

We submitted a revised version of the paper titled “POF-Darts: Geometric Adaptive Sampling for Probability of Failure” [7] to the SIAM Journal of Uncertainty Quantification on June 26, 2015.

This version of the paper had a number of improvements based on the reviewers’ comments on our first version. Specifically, the revised version had a more detailed description of the algorithm choices we used. We also included a mathematical analysis which provides analytical bounds on the expected accuracy of the probability of failure estimates. Finally, the revised paper had a more comprehensive set of test problems and comparison of methods.

Figure 1.1 shows the concept of POF Darts applied to a problem with four failure regions. Figure 1.2 shows the results of POF Darts compared to Latin Hypercube Sampling (LHS) for this same problem, based on an increasing number of samples. Note how POF-Darts focuses the samples near the failure region much more efficiently than LHS.




Figure 1.1. Applying POF-Darts to a test problem with four failure regions. After exhausting the sampling budget a GP surrogate is built and utilized to estimate the POF. The failure isocontours from the surrogate are colored in red while the exact one are colored in blue. For this test problem in both cases the isocontours were on top of each other indicating an accurate estimate.

1.1.2 Voronoi Piecewise Surrogates

“VPS: Voronoi Piecewise Surrogates for Fitting Noisy High-Dimensional Data” describes a geometric technique for subdividing a domain, in order to describe the local neighborhood of a sample within a set. This enables us to build surrogate functions that match the data locally. This subproject was born out of a need for POF-Darts: once a set of disks has been placed in the domain, we need to estimate the probability of failure using them. Our initial idea was to estimate the volume of the union of failure disks for a lower bound on the probability of failure, and the complement of the volume of the union of not-failure disks for an upper bound. Since computing this exactly in high dimensions is prohibitive, we devised a computationally efficient way to estimate the volume




Figure 1.2. Comparing the sampling phase of LHS (top row) to POF-Darts (bottom row) on the Herbie problem. From left to right in both rows, the number of samples are 100, 500, 1000, and 5000, respectively. LHS does a great job of space-filling the sample points, but POF-Darts focuses the samples near the failure boundaries.

of a union of spheres. We gave a talk at the SIAM GD/SPM [12]. However, it became clear that we could do better. Instead of estimating volumes, we built a surrogate function meant to indicate whether one was closer to a failure disk (1) or not-failure disk (0), and numerically estimated the volume of the failure region, equivalent to estimating the average value of the surrogate function over the domain. We then realized that generalizing this to other types of surrogates could be a big contribution to the field, since most extant methods build a *global* surrogate that ignores locality. We assume this historical shortcoming was because it is hard to compute local neighborhoods (Voronoi cells) exactly in high dimensions.

VPS [18] introduces a new method to construct accurate global surrogate models with local credibility in high-dimensional spaces: Voronoi Piecewise Surrogate (VPS) models. The key idea is to implicitly decompose the high-dimensional parameter space by the Voronoi tessellation of the sample points. Once a sample point knows its (approximate) Voronoi neighbors, we can build a surrogate for its cell using any standard (global) surrogate technique. See Figure 1.3 and Figure 1.4.

Exact cells are intractable in high dimensions, but fortunately this is not needed. The actual queries we must support are to determine which cell an arbitrary point belongs to, and the surrogate function value at that point. VPS assigns points to cells unambiguously using a simple nearest

neighbor search. VPS finds the approximate cell boundaries and approximate neighbors via local hyperplane sampling, without constructing an explicit cell mesh, and this is sufficient to build a surrogate function. Approximate cells for building the surrogate are especially reasonable considering that the surrogate itself is only an approximation of the true function.

Advantages

- VPS is flexible, allowing any surrogate to be built on a cell, even a surrogate traditionally based on global data.
- VPS breaks down the high-order polynomial approximation problem into a set of piecewise low-order polynomial approximation problems in the neighborhood of each function evaluation, independently.
- The one-to-one mapping between the number of function evaluations and the number of Voronoi cells, regardless of the number of dimensions, eliminates the curse of dimensionality associated with standard domain decompositions.
- The Voronoi tessellation is naturally updated with the addition of new function evaluations.
- VPS accurately models functions with discontinuities by allowing discontinuities at Voronoi cell boundaries. Voronoi neighbors may ignore each other when building their surrogates if it appears that there is a large functional jump or discontinuity between them.

Status In Dakota, this has been implemented as a `domain_decomposition` option for global surrogates. The paper is in preparation for submission to SIAM Journal on Uncertainty Quantification.




Figure 1.3. A two-dimensional example of VPS. Top row has sample points and cells. Bottom row shows neighbors and samples of the surrogate falling into cells.

1.1.3 Methods Comparison

“A Set of Test Problems and Results in Assessing Method Performance for Calculating Low Probabilities of Failure” [16] was accepted to a conference. It compares several global reliability meth-




Figure 1.4. VPS surrogates of various orders, compared to the true function and its isocontours.

ods, to evaluate the cost and accuracy performance for calculating failure probabilities of magnitudes 10^{-2} to 10^{-6} in various low to moderate dimensional (2D to 10D) test problems, on some test problems and engineering applications. We also compared our new POF-darts algorithm over these problems, and found that it was competitive.

1.2 General Sampling, Across Applications

1.2.1 Hyperplane Sampling

“*k*-d Darts: Sampling by *k*-Dimensional Flat Searches” [8] was presented at SIGGRAPH to an audience of thousands, and the paper appeared in the ACM Transaction on Graphics, the ACM journal with the highest impact factor. We also presented it at UT Austin in an invited talk, and at the SIAM UQ14, in a minisymposium we organized.

It considers the general problem of sampling hierarchically by dimension, which finds application across fields. It formalizes this approach using *k*-d darts, a set of independent, mutually orthogonal, *k*-dimensional hyperplanes called *k*-d flats. A dart has *d* choose *k* flats, aligned with the coordinate axes for efficiency. (In later work, we discovered that deterministically distributing the flats amongst the combinations of axes was not required, and it was sufficient to select the orientation of each dart randomly.) We show *k*-d darts are useful for exploring a function’s properties, such as estimating its integral, or finding an exemplar above a threshold. We describe a recipe for converting some algorithms from point sampling to *k*-d dart sampling, if the function can be evaluated along a *k*-d flat.

We demonstrate that *k*-d darts are more efficient than point-wise samples in high dimensions, depending on the characteristics of the domain: for example, the subregion of interest has small volume and evaluating the function along a flat is not too expensive. We present three concrete applications using line darts (1-d darts): relaxed maximal Poisson-disk sampling, high-quality rasterization of depth-of-field blur, and estimation of the probability of failure from a response surface for uncertainty quantification. Line darts achieve the same output fidelity as point sampling in less time. For Poisson-disk sampling, we use less memory, enabling the generation of larger point distributions in higher dimensions. Higher-dimensional darts provide greater accuracy for a particular volume estimation problem.

1.2.2 Recursive *k*-d Darts

“Recursive *k*-d Darts for High-Dimensional Numerical Integration” [2] builds on *k*-d Darts, generalizing the notion of what it means to sample with a hyperplane. In particular, we define this recursively, so that a (*k*)-dimensional hyperplane “sample” (value) may be composed of “samples” from (*k* − 1)-dimensional sub-hyperplanes. At the bottom level of recursion there is always a point sample, unless the underlying function can be evaluated analytically along a line, etc. See Fig-




Figure 1.5. Hyperplane (line, plane, etc.) samples are more expensive to compute, but provide more information. They are also much more likely to intersect a subregion of interest, especially if that subregion is long and thin.

ure 1.6. It tackles one dimension at a time, transforming the problem into a set of 1-dimensional problems, which can each be solved efficiently. Recursive k -d darts is a novel method for numerical integration of high-dimensional functions, across applications. Our first demonstration application is a graphics ray tracer with depth of field blur, in part because this is a problem of interest to our university partner and the community at large.

Numerical integration methods are deployed in many problems, e.g. ray tracing in graphics, probability of failure in uncertainty quantification. For high-dimensional parameter spaces, classical integration methods are biased, suffer from the curse of dimensionality, or have slow error convergence rate. We introduce a novel method of sampling and evaluating integrations in high-dimensional domains: recursive k -d darts. Fundamentally, our method decomposes a d -dimensional integration problem into a series of one dimensional integrations, avoiding high-dimensional complexities. In specific, we build a recursion of 1-dimensional subspaces, sample each 1-dimensional subspace using a constrained random sampling approach, construct local Lagrange approximation surrogates to evaluate the integral up to that subspace level, and report the interpolation and evaluation errors at each step to the parent subspace. The locality of the surrogates enables our method to detect local features, such as functional discontinuities. Our method can have a much faster convergence rate when compared to random Monte Carlo methods; the actual performance depends on many factors, including the dimensionality of the problem, the orientation of the hyperplanes with respect to the function gradient, and how amenable the function is to an accurate surrogate model over the evaluation points. For an application with a limited sampling budget, our algorithm can suggest where to place new samples so they will provide the most information for the integration accuracy.




Figure 1.6. A hyperplane sample is composed of one-dimension-lower hyperplane samples; projected onto the line of the fixed coordinates of the sub-samples, the subsamples lie on a line, and the hyperplane value is some average or surrogate through those points. This recasts general high dimensional problems as a series of one dimensional problems.

1.2.3 Spoke Darts

Spoke Darts [13] considers a different kind of line sample, namely radial line samples. In this paper we use this to develop a blue noise sampling method that can achieve high quality and performance across different dimensions, although we consider it a general method for exploiting spherical coordinate systems and symmetry, and speculate that we could also create a recursive version analogous to recursive k -d darts.

In the blue noise application, performance is achieved by combining the locality of advancing front with the dimension-reduction of hyperplanes. Spoke Darts generates line segments through prior samples, to explore the local uncovered space suitable for the next sample; see Figure 1.7(a). Uniform point sampling from these lines produces traditional blue noise; non-uniform sampling produces step blue noise. We create step blue noise directly. In contrast, the current alternatives use compute-intensive position optimization as a post-process, or are fundamentally low dimensional. Our method is efficient enough to generate blue noise distributions in 23 dimensions for motion planning. Beyond generating blue noise point sets, we extend spoke darts to build approximate Delaunay graphs and perform global optimization more efficiently.

Of particular interest to this LDRD, the approximate Delaunay graph (dual Voronoi cells) al-

lowed us to build an enhanced version of the traditional DIRECT (division of rectangles) global optimization algorithm. The advantage is that approximate Voronoi cells more accurately finds the nearest sample and estimates functional trends than does hyper-rectangles. Our spoke darts variant often increases the convergence rate of DIRECT, and avoids many bad cases of false-convergence; see Figure 1.7(b).




Figure 1.7. Spoke dart sampling and applying it to build a global optimization algorithm.

1.2.4 Randomness and Optimization

Noise and Coverage, Optimizing Voronoi Aspect Ratio

In “Improving Spatial Coverage while Preserving the Blue Noise of Point Sets” [3] we explore the question of whether it is possible to change a spatial distribution of samples to improve some of its properties, without destroying its randomness properties. In particular, we wish to preserve the blue-noise Fourier spectrum of the points, while improving the aspect ratio of its Voronoi cells. It is well known that maximizing one of these properties destroys the other: uniform random points have no aspect ratio bound, and the vertices of an equilateral triangular tiling have no randomness. However, we show that there is a lot of room in the middle to get good values for both. We also show that optimizing the aspect ratio directly can be more effective than popular methods such as centroidal Voronoi tessellation, which optimize other criteria and change the aspect ratio as a side effect; see Figure 1.8.




Figure 1.8. Final mesh for a periodic unit box after applying various methods to the input MPS in (a). While CVT and DistMesh improve the quality of the majority of the Voronoi cells, they tend to lose randomness at larger values of β , which measures the Voronoi cell aspect ratio bounds. Far-Point on the other hand tends to violate the coverage condition. None of the four methods were able to achieve more regular cells than $\beta < 0.7$ in general.

Noise and Cardinality, Resampling

“Sifted Disks” [4] asks a similar question about some other properties. Namely, we seek to preserve randomness and Voronoi aspect ratios, while reducing the number of points needed. The Sifted Disks technique is based on local resampling. Two neighboring points are removed and we attempt to find a single random point that is sufficient to replace them both; see Figure 1.9. The resampling respects the original sizing function; In that sense it is not a coarsening. The angle and edge length guarantees of a Delaunay triangulation of the points are preserved. The Fourier spectrum is largely unchanged. We provide an efficient algorithm, and demonstrate that sifting uniform Maximal Poisson-disk Sampling (MPS) and Delaunay Refinement (DR) points reduces the number of points by about 25%, and achieves a density about 1/3 more than the theoretical minimum. We show two-dimensional stippling and meshing applications to demonstrate the significance of the concept. Sifted disks has obvious applications to efficient sample designs, as each sample point represents an expensive physical or computational experiment. Samples without randomness may introduce bias in the results, and poor Voronoi cells may also produce error or inefficiency.

Resampling for Modeling and Mesh Properties

We presented “Disk Density Tuning of a Maximal Random Packing” [9] at the USNCCM 13 in July 2015, and we submitted a full paper to the ACM TOG journal. Disk tuning is an algorithmic framework for adjusting the spatial density of disks in a maximal random packing, without changing the sizing function or radius function of the disks. As such, it is a general method that finds application across sample design for opt/UQ, meshing, and direct modeling. Starting from any maximal random packing such as the Maximal Poisson-disk Sampling (MPS), we iteratively relo-




Figure 1.9. Sifted disks is based on local constrained resampling. We attempt to remove two adjacent points, and replace it with one. However, the replacement must not leave any part of the domain uncovered.

cate, inject (add), or eject (remove) disks, using a set of three successively more aggressive local operations. We may achieve a user-defined density, either more dense or more sparse, almost up to the theoretical structured limits. The tuned samples are conflict-free, retain coverage maximality, and except in the extremes, retain blue noise randomness properties of the input. We change the density of the packing one disk at a time, maintaining the minimum disk separation distance and the maximum domain coverage distance required of any maximal packing. These properties are local, and we can handle spatially-varying sizing functions. Using fewer points to satisfy a sizing function improves the efficiency of some applications. We apply the framework to improve the quality of meshes, removing non-obtuse angles and increasing edge-valence; see Figure 1.10 and Figure 1.11. We also present an unusual application, that of actual physical modeling. The cross section of a fiber reinforced polymers material appears as a set of nonoverlapping circles (fibers) in the plane (cross section). We tuned the density of disks to match the fiber density in physical reality, and achieved better elastic and failure simulations; see Figure 1.12.

1.2.5 High Dimensional Data Analysis

We were invited to participate in a workshop on high dimensional data analysis. The workshop decided to write a book, and we contributed a chapter [6] summarizing our high dimensional sampling techniques, and speculating on how they might be used for data analysis. Through this exercise, we conceived of a unified way to consider all the degrees of freedom across sampling methods, and how many existing methods related to each other; see Figure 1.13.




Figure 1.10. Removing obtuse angles in planar meshes using “Disk Density Tuning of a Maximal Random Packing.”

1.3 Meshing

1.3.1 Planar Meshing

Robust All-Quad Meshing of Domains with Connected Regions

“Robust All-Quad Meshing of Domains with Connected Regions,” [17] was accepted to the 24th International Meshing Roundtable (IMR), Austin TX, October 2015. In this paper, we present a new algorithm for all-quad meshing of non-convex domains, with connected regions. Our method starts with a strongly balanced quadtree. In contrast to snapping the grid points onto the geometric boundaries, we move points a slight distance away from the common boundaries. Then we intersect the moved grid with the geometry. This allows us to avoid creating any flat quads, and we are able to handle two-sided regions and more complex topologies than prior methods. The algorithm is provably correct, robust, and cleanup-free; meaning we have angle and edge length bounds, without the use of any pillowing, swapping, or smoothing. Thus, our simple algorithm is also more predictable than prior art. See Figure 1.14 and Figure 1.15

Delaunay Quadrangulation

We discovered a method to produce quad meshes, based on sampling points and using a subset of Delaunay edges. Sample points can be generated randomly, or in advancing front; see Figure 1.16. Points are assigned one of two colors, and we only include the Delaunay edges between points of




Figure 1.11. Removing obtuse angles in curved surface meshes using “Disk Density Tuning of a Maximal Random Packing.”




Figure 1.12. Fiber material modeling using “Disk Density Tuning of a Maximal Random Packing.” Tuned disks predict the failure point better than traditional hexagonal or square packings.

opposite color. This provides even-sided cells, most of which are quads. Larger cells are either removed by resampling, or we can retain the original sample points and add a few more to sub-

Process randomness is a hidden axis,
merely a means to obtain spatial randomness.




Figure 1.13. A conceptual space for sampling methods.

divided then combine the large cells into several quadrilaterals. The conference proceedings [14] describes the algorithm, and the proofs are available as an appendix in a longer tech report [15].

Steiner Point Reduction

“Steiner Point Reduction in Planar Delaunay Meshes” [1] is a meshing application of a particular form of disk tuning. In particular, we developed a mesh simplification strategy that preserves angle bounds, called sifting. See Figure 1.17 for an example mesh reduction. We introduce a set of constraints for the location of the new point based on the desired minimum angle and compute an explicit representation of the solution region, which we then sample from to find the replacement point. Thanks to the angle bounds, the number of constraints is bounded by a constant and all updates are local. This strategy generalizes edge collapse, as it possibly combines edge swaps to update the mesh after replacement. Preliminary results for a sample of planar Delaunay meshes are then presented. We demonstrate significant improvements of Triangle output, in terms of the number of Steiner points needed for a required angle bound, specially for large bounds where Triangle is known to possibly perform poorly.




Figure 1.14. “Robust All-Quad Meshing of Domains with Connected Regions” algorithm steps.

1.3.2 Sizing Function Discovery

Spatial domain sampling is a core component for a variety of applications. Prior sampling methods need explicit sizing functions as input. However, specifying the ideal sizing function can be a difficult and ill-defined task for several important geometry problems. In mesh generation, Delaunay refinement offers a way to adaptively discover the local feature size function. However, this technique is inherently serial.

We present a disk-packing method to sample spatial domains and discover the local feature size. Our key idea is to place samples according to local features of the domain boundary, and adjust their radii based on nearby disks, to satisfy the required smoothness conditions, and other prescriptions such as element shapes. Applications of our method include mesh generation and Reeb graph construction in 2–3D Euclidean space, and embedded curved surfaces. Figure 1.18 illustrates meshing a domain with a convoluted boundary.

1.3.3 VoroCrust, Voronoi Meshing And Surface Reconstruction

We introduce VoroCrust [10], the first algorithm for simultaneous surface reconstruction and 3D Voronoi meshing; see Figure 1.19. Like the power crust, it creates Voronoi cells inside and outside a smooth manifold, and the reconstructed surface is their interface. Unlike the power crust, the Voronoi cells are unweighted and have good aspect ratio. Moreover, there is complete freedom of how to mesh the volume far from the surface. Most of the reconstructed surface is composed of




Figure 1.15. “Robust All-Quad Meshing of Domains with Connected Regions” examples.

Delaunay triangles with small circumcircle radius, and all samples are vertices. In the presence of slivers, the reconstruction lies inside the sliver, interpolating between its upper and lower pair of bounding triangles, and introducing Steiner vertices.

VoroCrust meshes are distinguished from the usual approach of clipping Voronoi cells by the manifold, which results in many extra surface vertices beyond the original samples, and may result in non-planar, non-convex, or even non-star-shaped cells.

As with many other reconstruction methods, in theory the initial sampling must be an ε -lfs dense sampling, and in practice a coarser sampling suffices. Here ε -lfs means the sample density must be a small constant times the “local feature size”, which measures both the local curvature and thickness of the model. However, VoroCrust samples must be weighted. We provide an sphere-packing algorithm to generate a suitable sampling.

We also consider the theoretical question of whether it is possible to use 3D Voronoi cells to reconstruct a surface geometrically matching a prescribed triangulation, given the freedom for the cells to shatter the triangles into smaller ones. We show this is possible, given mild conditions on the dihedral angles between the triangles; e.g. an ε -lfs sampling suffices.




Figure 1.16. Delaunay quadrangulation (top) from two-color disks (bottom light and dark circles). Left: random. Right: advancing front.




Figure 1.17. The standard Delaunay refinement (DR) technique produces good quality elements, but many of them. By sifting the mesh, we were able to reduce the number of elements by 78%, while preserving the minimum and maximum angles.

To see how the algorithm works, see Figure 1.20. We generate or are given a set of well-spaced manifold sample points that captures small and sharp features. We create a sphere around each sample. Some triples of overlapping spheres define a pair of intersection points, mirrored on each side of the manifold. Pairs outside all other spheres are represent weighted triangulation circumcenters and taken as Voronoi seeds. We create additional well-spaced seeds far interior




Figure 1.18. Disk-packing sampling meshes discovered the local feature size, without being guided by a prespecified sizing function. Disk radii (inter-point distances) are about the constant L times the local feature size. There is a natural tradeoff between L and the minimum angle α in the mesh.




Figure 1.19. An example hex dominant 3D Voronoi mesh reconstructing a surface.




Figure 1.20. VoroCrust algorithm steps, illustrated for 2D and 3D domains.

to the manifold; these may define a hex-dominant mesh. The manifold is reconstructed by the Voronoi facets separating the inside and outside cells. These unweighted Voronoi cells provide a simple and robust way of testing whether a query point is inside or outside the volume: simply find the nearest-neighbor seed and check whether it is an interior or exterior seed.

1.4 Presentations without Papers

Most of our significant results are described above and in papers, so for most of the additional talks we list them without elaboration.

1.4.1 Minisymposia

1. Characterizing Sample Distribution Properties and their Impact on Experimental Design. Scott A. Mitchell and Mohamed S. Ebeida organized and gave two talks. UQ14, SIAM Conference on Uncertainty Quantification, 2014
2. Voronoi Dual Meshing and Simulation. Mohamed Ebeida and Scott Mitchell organized. LDRD participants gave three talks: Ebeida, Abdelkader, Rushdi. USNCCM13, 2015

1.4.2 Invited Talks

1. Exploring High-dimensional Spaces using Well-spaced Random Points and Hyperplane Sampling with Application to Graphics, Meshes, Global optimization, Uncertainty and Robotics. ICES, UT Austin. Host: Chandrajit Bajaj. 2014
2. Exploring High Dimensional Spaces with Hyperplane Sampling, (k -d Darts: Sampling by k -Dimensional Flat Searches), UT Austin. Host: Chandrajit Bajaj 2014
3. Algorithms for well-spaced random points with application to uncertainty, meshes, optimization, graphics and robotics, VCCC Summit, KAUST. Host: Peter Wonka. 2014
4. Well-spaced Random Points for Graphics, Meshes, Optimization, Uncertainty and Robotics. Workshop on challenges in integrated computational structure-material modeling of high strain-rate deformation and failure in heterogeneous materials, Johns Hopkins University, Host: Somnath Ghosh. 2013
5. Well-spaced Random Points for Graphics, Meshes, Optimization, Uncertainty and Robotics. Computer Science dept., UNC Chapel Hill. Host: Dinesh Manocha, 2013
6. Improved Poisson-disk Sampling for Meshing, Uncertainty Quantification and Graphics Applications. UC Irvine, Host: David Eppstein, 2013
7. VoroCrust, at the Polytopal Element Methods workshop, 2015

1.4.3 Regular Talks

1. Improved Poisson-disk Sampling for Meshing Applications, at the World Congress on Computational Mechanics (WCCM XI) 2014. [5]
2. VoroCrust Algorithm: 3D Polyhedral Meshing with True Voronoi Cells Conforming to Surface Samples, key note speaker, at the 13th USNCCM, 2015.
3. VoroCrust Geometry: 3D Polyhedral Meshing with True Voronoi Cells Conforming to Prescribed Surface Points, also at the USNCCM 2015.

4. Feature-Preserving Spatial Density Tuning of a Maximal Random Disk Packing, also at the 13th USNCCM, 2015.
5. POF darts: Geometric Adaptive Sampling for Probability of Failure Estimation, at the SIAM conference on Uncertainty Quantification, 2014.
6. Balloon Darts: Estimating the Volume of the Union of d -Balls with Spoke Samples, at SIAM GD/SPM, 2013. [12]

1. “Improved Poisson-disk Sampling for Meshing Applications” [5] describes how randomness can also be to avoid mesh-induced non-physical phenomena in simulations. For example, in some fracture simulations, fractures only propagate along mesh edges, and having random edge orientations produces more realistic cracks. Even when randomness is not desired, we believe disk-packing based meshes promise several advantages over the alternatives. We describe several new approaches to Poisson-disk sampling and resampling to generate and improve simplicial meshes in d -dimensional spaces. We produce provably good tessellations, with quality bounds similar to (or better) than deterministic Delaunay refinement methods. It is inherently easier for our methods to follow a sizing function because of the close connection between Poisson-disks and the local mesh size. We show results for the uniform and the non-uniform case. We show several applications to examples to demonstrate the efficiency of our methods.

6. “Balloon Darts: Estimating the Volume of the Union of d -Balls with Spoke Samples” [12] describes randomized approximate algorithms for computing the volume of the union of d -dimensional balls. The deterministic, exact approaches do not scale well to high dimensions. However, we adapt several of these to a local sampling approach, sampling within each ball, using a polar variation of k - d darts. We sample balls with randomly-oriented lines through their centers. The sampling process is more accurate per unit work when we use line samples rather than point samples. This efficiency gain is because a line sample provides more information, and the analytic equation for a sphere makes computing a line sample almost as fast as a point sample. For the power cell decomposition, we extend this to sampling with planes (“wheels”). We compare the efficiency of these and other methods, and provide guidance on which method to choose for a given type of ball distribution.

1.4.4 Internal Talks

1. LS Polynomial Optimization for Constructing Voronoi Piecewise Surrogates, Sandia National Labs internal presentation, 2014
2. “Well-Spaced Random Point Sets for Sampling and Meshing” [11] to the visiting Sandia summer students. The goal was to introduce the topics of point distributions and sampling, and their various applications, to the students.
3. We gave numerous review talks to the program managers of the CIS LDRD IAT as part of the normal LDRD process, and to management.

Chapter 2

Partnerships

2.1 Innovative Interactions for Innovative Ideas—Meshing and Computer Graphics Relevance

Under this LDRD we researched the geometry and statistics of sample distributions, as well as algorithms to generate samples with particular geometric and statistical properties, for uncertainty quantification and optimization, and secondarily for mesh generation. The underlying sampling and geometry methods we studied for uncertainty quantification suggested many meshing improvements. In particular, the properties of well-spacedness and adaptive density of sample points are common for both domains, and also for sampling for graphics rendering.

Sampling, geometric computing, and mesh generation are central to Sandia’s modeling and simulation activities. While these topics originated in mathematics and engineering, they have now become cross disciplinary. It may be surprising to some in the computational science and engineering fields to note the following connection with computer graphics. Finding samples with particular properties is a key topic in graphics, for example spanning the space, without redundancy, or without aliasing or bias in particular frequencies, or other statistical properties. Generating samples with similar properties is a significant research topic for this LDRD, as it is used for uncertainty quantification, optimization, and mesh generation. These are all topics central to Sandia’s simulation-based science and stewardship missions. The common application in computer graphics is sampling light rays, for rendering a variety of phenomena such as blur and ambient light color. Rendering applications are not central goals of the LDRD, but the algorithms used for those applications are, because fundamentally both graphics and optimization & uncertainty applications use sparse samples to estimate a numerical integral.

2.2 University Partnerships

2.2.1 University of California, Davis, Prof. John Owens

We had a university contract with Prof. John Owens at University of California, Davis. The goal was to research, implement and demonstrate new random sampling techniques that enable estimat-

ing the probability of failure of high dimensional computer models. This is based on estimating the volume of a response (function value) below a given threshold. The broader research problem is sampling to discover the global properties of functions that are too complex and high dimensional to analyze analytically. UC Davis's focus was to research scalable sampling techniques. UC Davis brings a cross fertilization of ideas, from their expertise in sampling techniques and applications from computer graphics. UC Davis also provides expertise in GPU programming, and evaluation of sample designs.

2.2.2 University of North Carolina, Prof. Dinesh Manocha

We had a university contract with Prof. Dinesh Manocha of UNC. His student Chonhyun Park was active on the project. We researched algorithms and software to generate point distributions in high dimensions, for the efficient exploration of that space. Motivating applications include Uncertainty Quantification UQ and Robotic motion planning, with synergy between the two. Exploring the high-dimensional heterogeneous space of geometric configurations is UNC's expertise.

2.2.3 University of Texas, Austin, Prof. Chandrajit Bajaj

We had a university contract with Prof. Chandrajit Bajaj at UT Austin. Prof. Bajaj brought a depth of geometric knowledge that was particularly helpful for the VoroCrust project. His research associate Ahmed Rushdi worked on numerous projects, including DAKOTA surrogates software and publications.

2.3 University Interactions

We had a variety of research interactions with the following people. These included visits, research discussions and co-authorships. Don Sheehy (Univ Connecticut); Gary Miller (CMU); David Eppstein (UC-Irvine), Ahmad Rushdi (UTAustin); Ahmed Mahmoud (Alexandria Univ.); Alexander Rand (Cd-adapco); Andrew Davidson (UCDavis); Anjul Patney (nVidia); Chandrajit Bajaj (UTAustin); Chonhyon Park (UNC); Dinesh Manocha (UNC); John Owens (UCDavis); Li-Yi Wei (HKU); Mohammed Mohammed (Alexandria Univ.); Muhammad Awad (Alexandria Univ.); Tzeng Stanley (UCDavis); Xiaoyin Ge (OSU). Dong-Ming Yan, Peter Wonka (KAUST); Ahmed Abdelkader (Univ. Maryland, extended visit).

Chapter 3

Software

All of our publications, except the one comparing methods, involved writing software. Most of this software was written to demonstrate and verify the algorithms. Some of the opt/UQ capabilities were ported to DAKOTA and matured, under a variety of projects. We hope to build a publicly releasable meshing and spatial decomposition library as follow-on projects.

3.1 Implementations in DAKOTA

3.1.1 Probability Of Failure (POF)-Darts

POF-Darts has been implemented into Dakota as a NonD iterator method to estimate the tail probability (Probability Of Failure) based on random disk-packing in the uncertain parameter space. The POF-Darts method is identified using the keyword “pof_darts” in Dakota. POF-Darts uses hyperplane sampling to explore the unexplored part of the uncertain space. We use the function evaluation at a sample point to determine whether it belongs to failure or non-failure regions, and surround it with a protection sphere region to avoid clustering. Decomposing the domain into Voronoi cells around the function evaluations as seeds, our strategy to choose the radius of the protection sphere depends on the local Lipschitz continuity. For each cell, we estimate the local Lipschitz continuity using the function values of the significant Voronoi neighbors, and create a sphere to protect its neighborhood from future sampling. As sampling proceeds, regions uncovered with spheres would shrink, improving the estimation accuracy. After exhausting the function evaluation budget, we build a surrogate model using the function evaluations associated with the sample points and estimate the probability of failure by exhaustive sampling of that surrogate.

3.1.2 Voronoi Piecewise Surrogates (VPS)

VPS has been implemented in Dakota as an option for global surrogate models; see Figure 3.1. It currently supports polynomial Least-Squares regression, Gaussian Processes (GP), and Radial Basis Functions (RBF). From a Dakota input spec perspective, a global surrogate runs in a VPS mode by listing the keyword: `domain_decomp`. The keywords “`support_layers`”, “`cell_type`”, and “`discontinuity_detection`” are optional capabilities. VPS can also make use of derivative informa-

tion (gradients and Hessians). VPS breaks down the high-order polynomial approximation problem into a set of piecewise low-order polynomial approximation problems in the neighborhood of each function evaluation, independently. The one-to-one mapping between the number of function evaluations and the number of Voronoi cells, regardless of the number of dimensions, eliminates the curse of dimensionality associated with standard domain decompositions. The Voronoi tessellation is naturally updated with the addition of new function evaluations. Due to its piecewise construction, VPS handles functions with discontinuities accurately by partitioning the Voronoi tessellation into several sub-tessellations, trapping discontinuities along Voronoi facets. Allowing local approximations to use neighbors that belong to the same sub-tessellation only improves the accuracy of the final surrogate model.




Figure 3.1. DAKOTA's software design for surrogates

3.1.3 Recursive k -d Darts (RKD)

A research version of RKD has been implemented in Dakota. It is currently a “NonD” (non-deterministic) iterator standalone method to evaluate high-dimensional numerical integrations. The RKD-Darts method is identified using the keyword “rkd_darts” in Dakota. RKD uses the initial sample budget and recursively guides where to add available samples. After exhausting the sample budget, it builds a global surrogate to approximate the underlying function everywhere. It uses the surrogate to return the approximate numerical integration of the underlying function. RKD was made possible by the ideas generated by the LDRD, but the RKD software was funded solely by ASC.

Chapter 4

Conclusion

In conclusion, we consider this to be one of the more enjoyable and successful LDRD's we have participated in. We were encouraged by the breadth of ideas we were able to take from one field and apply to another, notably applying geometric ideas to uncertainty quantification. We were encouraged by the number of papers we were able to publish. We look forward to seeing some of the work started under the LDRD through to completion, as new publications and software. The LDRD inspired numerous promising ideas for follow-on projects.

References

- [1] Ahmed Abdelkader, Scott A. Mitchell, and Mohamed S. Ebeida. Steiner point reduction in planar Delaunay meshes. In *Symposium on Computational Geometry, Young Researchers Forum*, 2014. available from <http://www.computational-geometry.org/YRF/cgyrf2014.pdf> and http://www.cs.sandia.gov/~samitch/bibliography_2007.html#steiner-reduction.
- [2] Saman Ashkiani, Ahmad A. Rushdi, Scott A. Mitchell, John D. Owens, and Mohamed S. Ebeida. Recursive k-d darts for high-dimensional numerical integration. *in revision for resubmission*, 2015.
- [3] Mohamed S. Ebeida, Muhammad A. Awad, Xiaoyin Ge, Ahmed H. Mahmoud, Scott A. Mitchell, Patrick M. Knupp, and Li-Yi Wei. Improving spatial coverage while preserving the blue noise of point sets. *Computer-Aided Design*, 46(0):25 – 36, 2014. 2013 SIAM Conference on Geometric and Physical Modeling, SIAM GD/SPM13. Reposting on author website with permission from Elsevier as U.S. government funded work: http://www.cs.sandia.gov/~samitch/bibliography_2007.html#opt-beta.
- [4] Mohamed S. Ebeida, Ahmed H. Mahmoud, Muhammad A. Awad, Mohammed A. Mohammed, Scott A. Mitchell, Alexander Rand, and John D. Owens. Sifted disks. *Computer Graphics Forum, Proc. Eurographics*, 32(2pt4):509–518, 2013.
- [5] Mohamed S. Ebeida and Scott A. Mitchell. Improved Poisson-disk sampling for meshing applications. In *11th World Congress on Computational Mechanics (WCCM XI)*, July 2014. <http://www.wccm-eccm-ecfd2014.org/frontal/Ebook.asp>.
- [6] Mohamed S. Ebeida, Scott A. Mitchell, Anjul Patney, Andrew A. Davidson, Stanley Tzeng, Muhammad A. Awad, Ahmed H. Mahmoud, and John D. Owens. Exercises in high-dimensional sampling: Maximal Poisson-disk sampling and k-d darts. In Janine Bennett, Fabien Vivodtzev, and Valerio Pascucci, editors, *Topological and Statistical Methods for Complex Data*, Mathematics and Visualization, pages 221–238. Springer Berlin Heidelberg, 2015.
- [7] Mohamed S. Ebeida, Scott A. Mitchell, and Laura P. Swiler. POF-darts: Geometric adaptive sampling for probability of failure. *SIAM Journal of Uncertainty Quantification*, page submission under review, 2015.
- [8] Mohamed S. Ebeida, Anjul Patney, Scott A. Mitchell, Keith R. Dalbey, Andrew A. Davidson, and John D. Owens. K-d darts: Sampling by k-dimensional flat searches. *ACM Trans. Graph.*, 33(1):3:1–3:16, February 2014.

- [9] Mohamed S. Ebeida, Ahmad A. Rushdi, Muhammad A. Awad, Ahmed H. Mahmoud, Dong-Ming Yan, Shawn English, John D. Owens, Chandrajit L. Bajaj, and Scott A. Mitchell. Disk density tuning of a maximal random packing. *ACM Transactions on Graphics, TOG*, page under review, 2015.
- [10] Ahmed H. Mahmoud, Ahmad A. Rushdi, Scott A. Mitchell, Ahmed Abdelkader, Chandrajit L. Bajaj, John D. Owens, and Mohamed S. Ebeida. VoroCrust: unclipped conforming Voronoi meshing. *in revision for resubmission*, 2015.
- [11] Scott A. Mitchell. Well-spaced random point sets for sampling and meshing, 2013. 2013 CERI Summer Seminar Series, overview talk to Sandia summer students interns.
- [12] Scott A. Mitchell, Mohamed S. Ebeida, and Muhammad A. Awad. Balloon darts: Fast approximate union volume in high dimensions with line samples. SIAM GD/SPM conference talk, 2013.
- [13] Scott A. Mitchell, Mohamed S. Ebeida, Muhammad A. Awad, Chonhyon Park, Laura P. Swiler, Dinesh Manocha, and Li-Yi Wei. Spoke darts for efficient step blue noise in any dimension. *in revision for resubmission*, 2014. 5 Aug 2014 draft appears as arXiv:1408.1118v1.
- [14] Scott A. Mitchell, Mohammed A. Mohammed, Ahmed H. Mahmoud, and Mohamed S. Ebeida. Delaunay quadrangulation by two-coloring vertices. *Procedia Engineering*, 8:364–376, 2014. Proceedings of the 23rd International Meshing Roundtable. Freely available online at Science Direct. For the quad-quality proofs, see the Sandia tech report version at http://www.cs.sandia.gov/~samitch/bibliography_2007.html.
- [15] Scott A. Mitchell, Mohammed A. Mohammed, Ahmed H. Mahmoud, and Mohamed S. Ebeida. Delaunay quadrangulation by two-coloring vertices — extended version with quad-quality proofs appendix. Technical Report SAND2014-16625C, Sandia National Laboratories, 2014. Available at http://www.cs.sandia.gov/~samitch/bibliography_2007.html.
- [16] Vicente Romero, Laura P Swiler, Mohamed S. Ebeida, and Scott A. Mitchell. A set of test problems and results in assessing method performance for calculating low probabilities of failure. In *18th AIAA Non-Deterministic Approaches Conference, AIAA Science and Technology Forum and Exposition*, page to appear, jan 2016.
- [17] Ahmad A. Rushdi, Scott A. Mitchell, Chandrajit L. Bajaj, and Mohamed S. Ebeida. Robust all-quad meshing of domains with connected regions. In *International Meshing Roundtable, IMR*, page to appear, 2015.
- [18] Ahmad A. Rushdi, Laura P. Swiler, Scott A. Mitchell, and Mohamed S. Ebeida. VPS: Voronoi piecewise surrogates for high-dimensional data fitting, 2015. in preparation, to be submitted to SIAM/ASA Journal on Uncertainty Quantification.

DISTRIBUTION:

- 1 MS 0359 D. Chavez, LDRD Office, 1911
- 1 MS 0899 Technical Library, 9536 (electronic copy)

