Characterizing Sample Distribution Properties and their Impact on Experimental Design

MS17

SIAM UQ14

Monday 31 March 2014

4:30 – 6:30 pm

Ballroom A – 2nd Floor

4:30-4:55	Fourier Analysis of Stochastic Samnpling Strategies for Assessing Bias and Variance in Integration. Kartic Subr, Disney Research UK
5:00-5:25	POF-Darts: Geometric Adaptive Sampling for Probability of Failure Mohamed S. Ebeida, Sandia National Laboratories
5:30-5:55	Exploring High Dimensional Spaces with Hyperplane Sampling Scott A. Mitchell, Sandia National Laboratories
6:00-6:30	Building Surrogate Modesl with Quantifiable Accuracy Hany S. Abdel-Khalic and Congjian Wang, NC State

Organizers Scott A. Mitchell and Mohamed S. Ebeida

Sandia National Laboratories

k-d Darts: Sampling by k-Dimensional Flat Searches

MOHAMED S. EBEIDA
Sandia National Laboratories
ANJUL PATNEY
University of California, Davis
SCOTT A. MITCHELL and KEITH R. DALBEY
Sandia National Laboratories
ANDREW A. DAVIDSON and JOHN D. OWENS
University of California, Davis

Appeared in journal
ACM Transactions on Graphics
January 2014

ACM Reference Format:
Mohamed S. Ebeida, Anjul Patney, Scott A. Mitchell, Keith R. Dalbey, Andrew A. Davidson, and John D. Owens. 2014. *k*-d darts: sampling by *k*-dimensional flat searches. ACM Trans. Graph. 33, 1, Article 3 (January 2014), 16 pages.
DOI: http://dx.doi.org/10.1145/2522528

Exploring High Dimensional Spaces with Hyperplane Sampling

SIAM UQ14, MS17, talk 3
30 March 2014, 5:30-5:55pm

Speaker: Scott A. Mitchell

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company,
for the United States Department of Energy's National Nuclear Security Administration
under contract DE-AC04-94AL85000.

Major Points of Presentation

- Connection to techniques, and applications, from fields besides UQ
- Concepts
  - Hyperplane sampling
    - motivation, capture small thin regions
  - Formula for changing point sampling to flat sampling
  - Unbiased – provable
  - Variance – experiments, efficiency
    - hyperplane intersection with the object needs to be computable, efficient
    - volume estimation experiments
      - efficiency
      - dart type
  - Framework
    - function averaging, integration
    - finding a point with a function value (e.g. outside disks)
  - Three applications
    - Volume estimation
      - function integration
    - Generate a well-spaced point sampling a.k.a. Relaxed MPS
      - find domain points with function values
    - Depth of field with antialiasing
      - function integration

Motivation

Recall Problem Motivation
POF-Darts prior talk

- **Reliability calculations**
  - **Identify and measure tiny failure subspaces in a vast parametric space**
    - **10+ dimensions (parameters)**
    - **$<10^{-6}$ PoF (small volume region)**
    - **Expensive simulations – faster surrogate**
  - **POF-Darts was adaptive sampling (to find small regions with particular properties)**
  - **This talk is mainly about uniform sampling of regions (to measure them)**
- **Approach**
  - **Other sampling methods based on statistics and analysis**
  - **We borrow Computational Geometry, Graphics concepts:**
    - **line searches**
    - **sample-neighborhoods, geometric balls**
    - **functional integration**

Intuition
Who's going to hit the target (orange 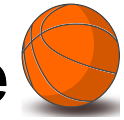)?

point sampling

line sampling

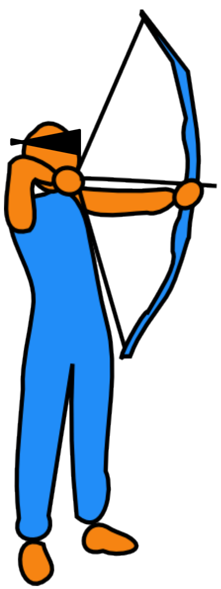

line sampling

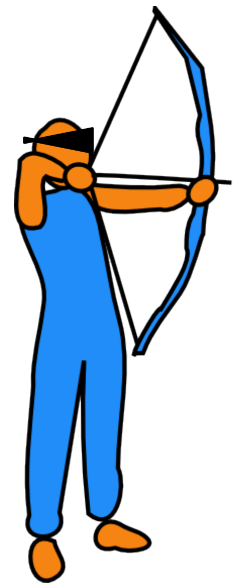

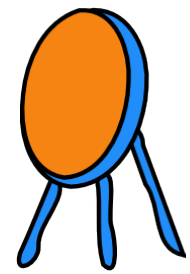

coin-shaped target
same volume, more surface

public domain clip art
Clker.com

More precisely

d=20, PoF = 10^{-6} (uniform distributions throughout talk)

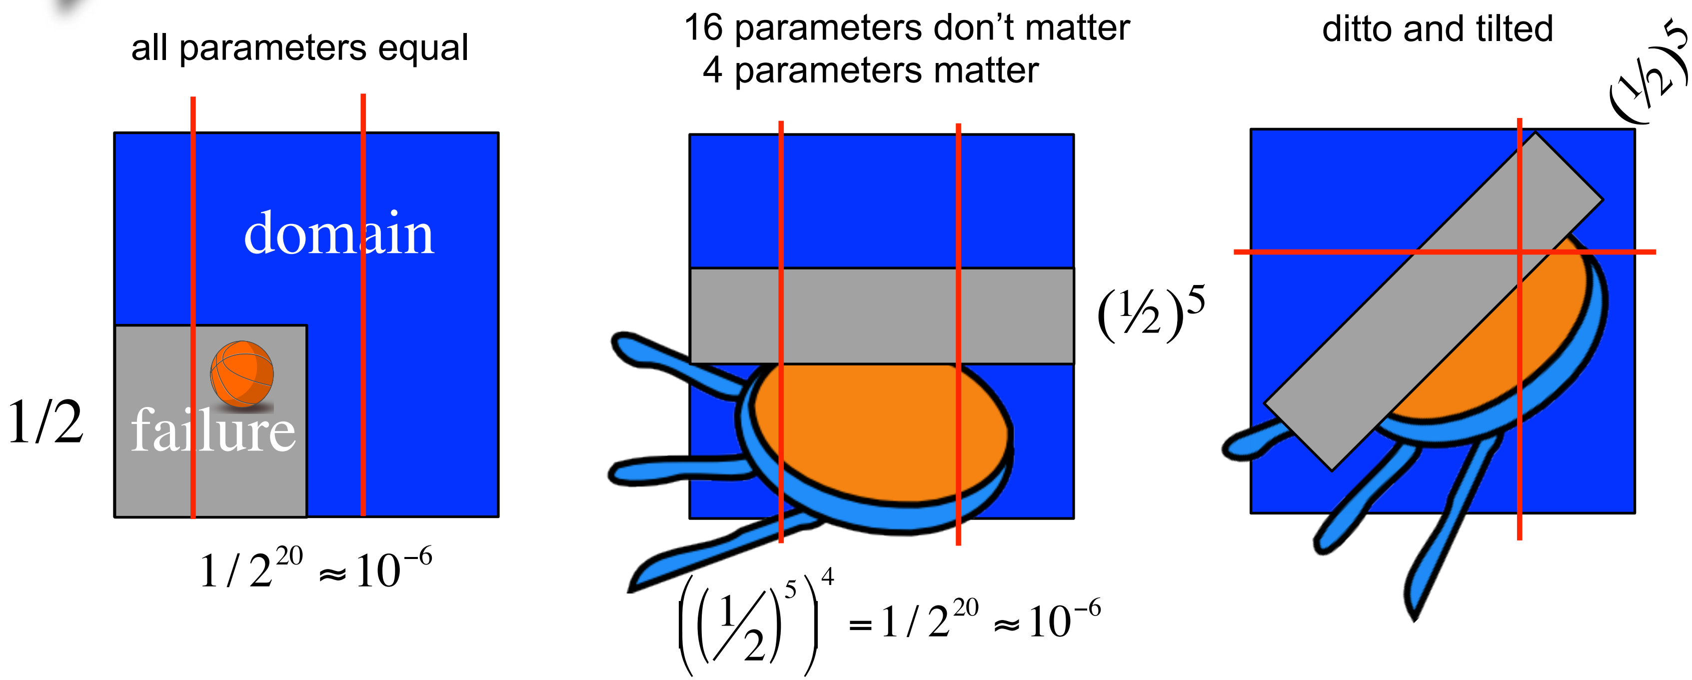

1 in 10^6 points 2 in 10^6 axis lines hit failure region	1 in 10^6 points 7 in 10^6 axis lines hit failure region	1 in 10^6 points 30 in 10^6 axis lines hit failure region

- **Lines are more likely to hit than points**
  - **better if coin-like** (bigger surface area)
  - **better if tilted** (surface area subtended by each line)
- **Intersection length is more information than binary point inside/outside**
  - **Planes are even better, hyperplanes...**

Sandia National Laboratories

Hyperplane sampling to hit regions

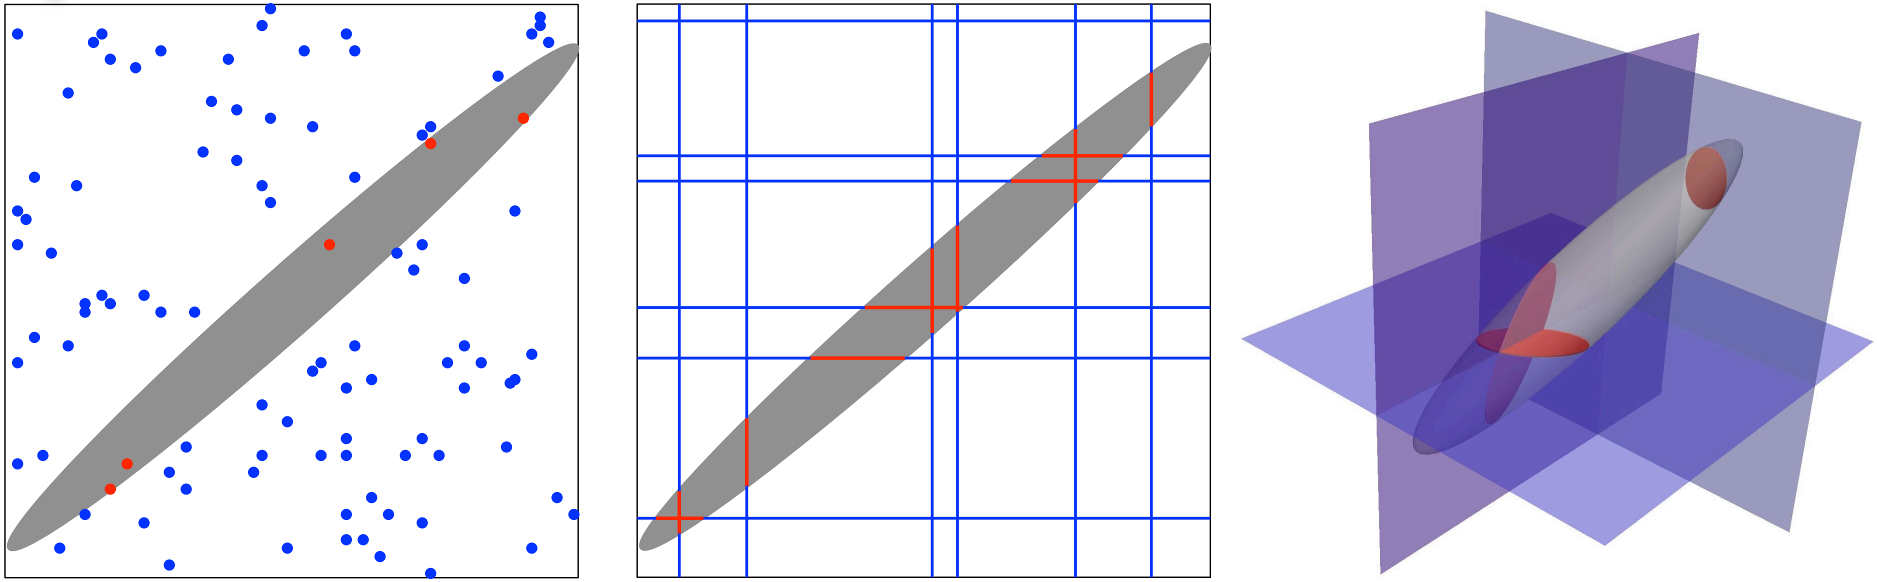

99 points | 6 line darts | 1 plane dart

Lines, hyperplanes, are more likely to intersect these regions, and they give more information

But they are more expensive.

Is it worth it?

The point of our paper is to answer “yes.”

Sandia National Laboratories

Approach, definitions

k-d Dart

- **k-dimensional hyperplanes (flats)**
  k free coordinates
  d-k fixed coordinates
- **dart = (d choose k) flats, one for every possible axis-aligned orientation**
  - **free coordinates (orientations) deterministically uniform**
  - **fixed coordinates (positions) uniform random, identically and independently distributed**

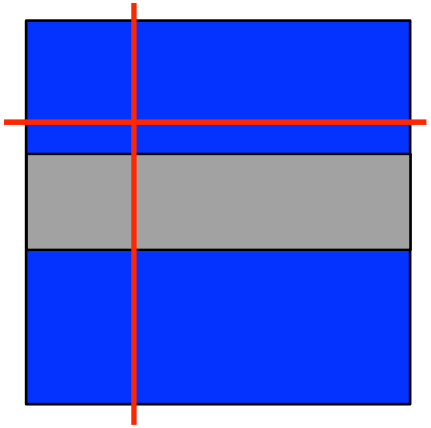

one 1-d dart in a 2-d domain
= two 1-d lines: x-aligned, y-aligned

k-d darts are unbiased

... probably obvious to UQ14 audience

- **I.e. the mean estimate is the true mean**
- **Because each flat is unbiased**
  - **because uniform point sampling of a height function is unbiased**

$$\bar{u}\int_{\mathcal{D}} 1 = \int_{\mathcal{D}} f = \int_{\mathcal{D}\text{fixed}}\int_{\mathcal{D}\text{free}} f = \int_{\mathcal{D}\text{fixed}} f_{\mathcal{D}\text{fixed}}.$$

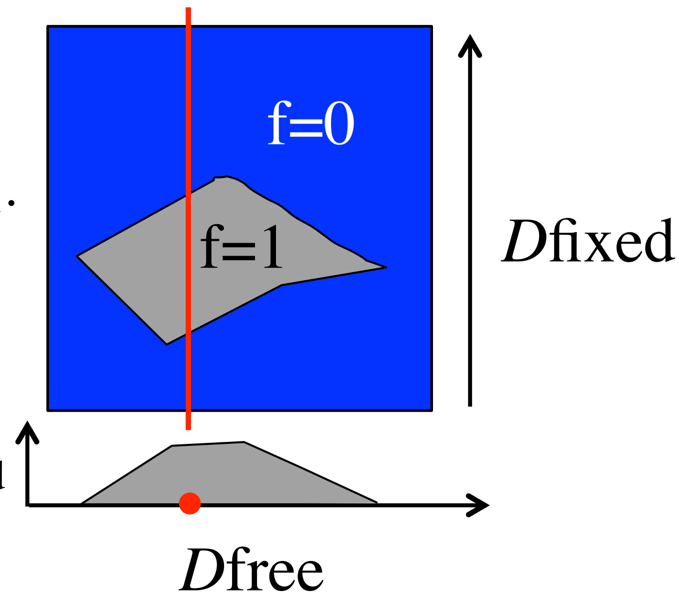

$$u_{estimate} = \frac{\sum \text{WeightedVolume}(\textit{flat})}{\sum \text{Volume}(\textit{flat})} = \frac{\sum \text{Length}(\textit{Line inside grey})}{\sum \text{Length}(\textit{Line})}$$

Variance? Efficiency?

- We have no formal proof for the variance
- Test problem:
  - estimating the volume of an ellipsoid
  - known analytic volume.
- Results: variance is well behaved
  - dropping as 1/number_of_samples^2
  - dropping by k

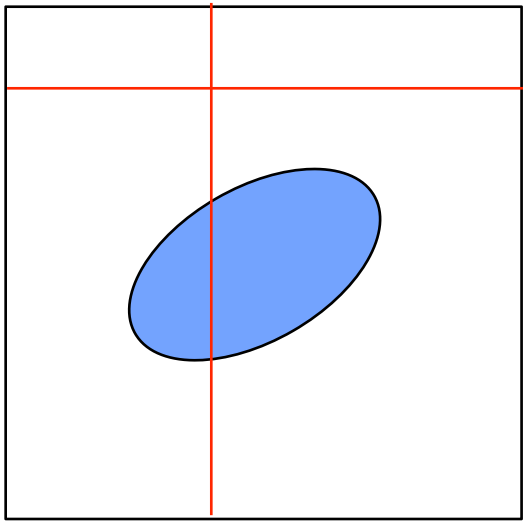

Mean error reduction by 1 / # samples^2

Vary 10-d domain aspect ratio, orientation; k-dart-dimension.

lower error for higher-k darts

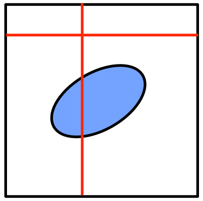

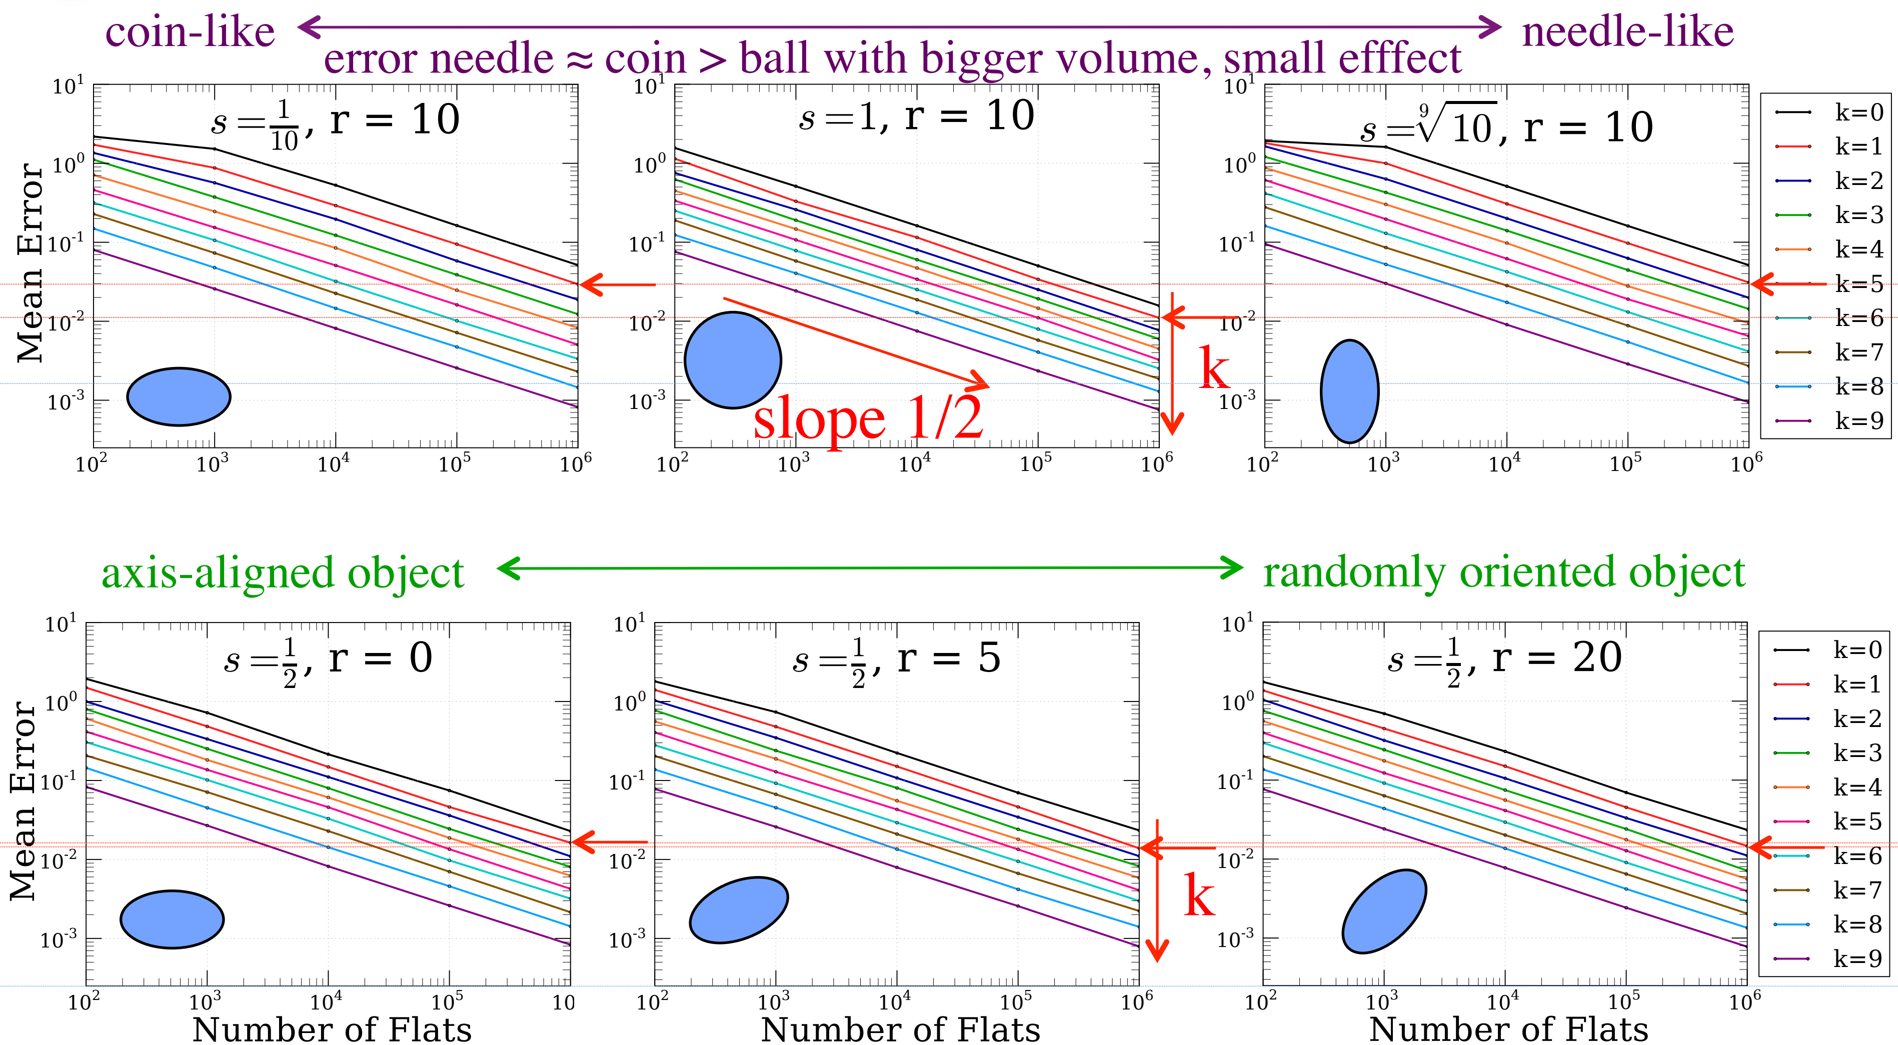

object orientation unaligned with axes helps a little, but not much

Estimate/true volume histograms for 1 million darts
Vary 10-d domain aspect ratio, orientation; k-dart-dimension
normal-like, sharper peaks for higher-k darts

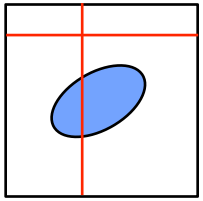

coin-like ⟷ needle-like

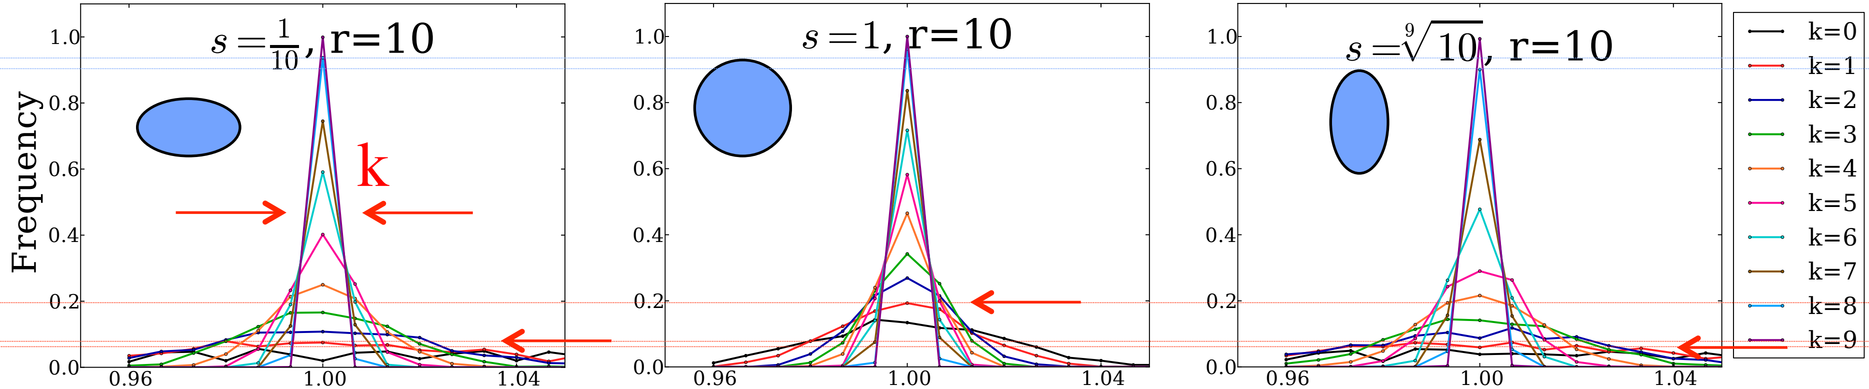

Squish matters a little bit, but volume matters much more.
We did 1-axis short, 1-axis long. Squish farther?

axis-aligned object ⟷ randomly oriented object

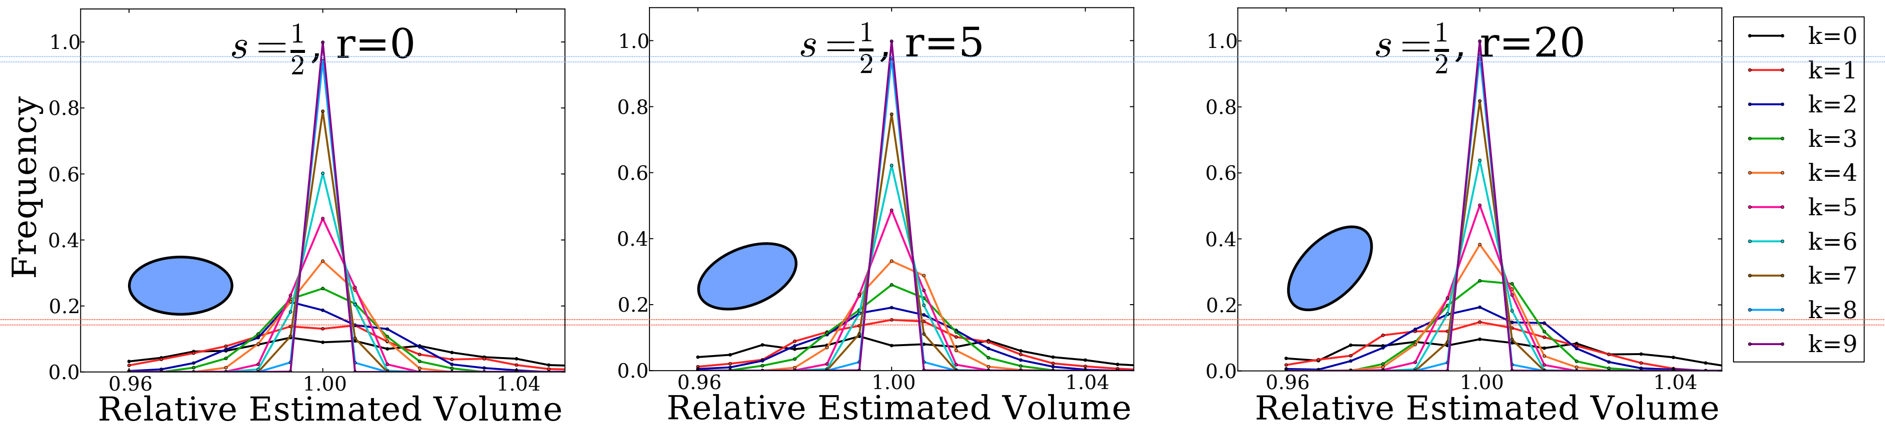

object orientation, doesn't matter very much

Trends by k

Conclusion

- **Higher k darts = less error, less variance**
  - **Because each dart gives more information**
- **Use a higher k if**
  - **You can compute its intersection with the object**
  - **And that computation is not too much slower**

Sample-Orientation Effects?
Axis-aligned flats just as accurate as randomly oriented darts...

...and faster and simpler. Axis-aligned best.

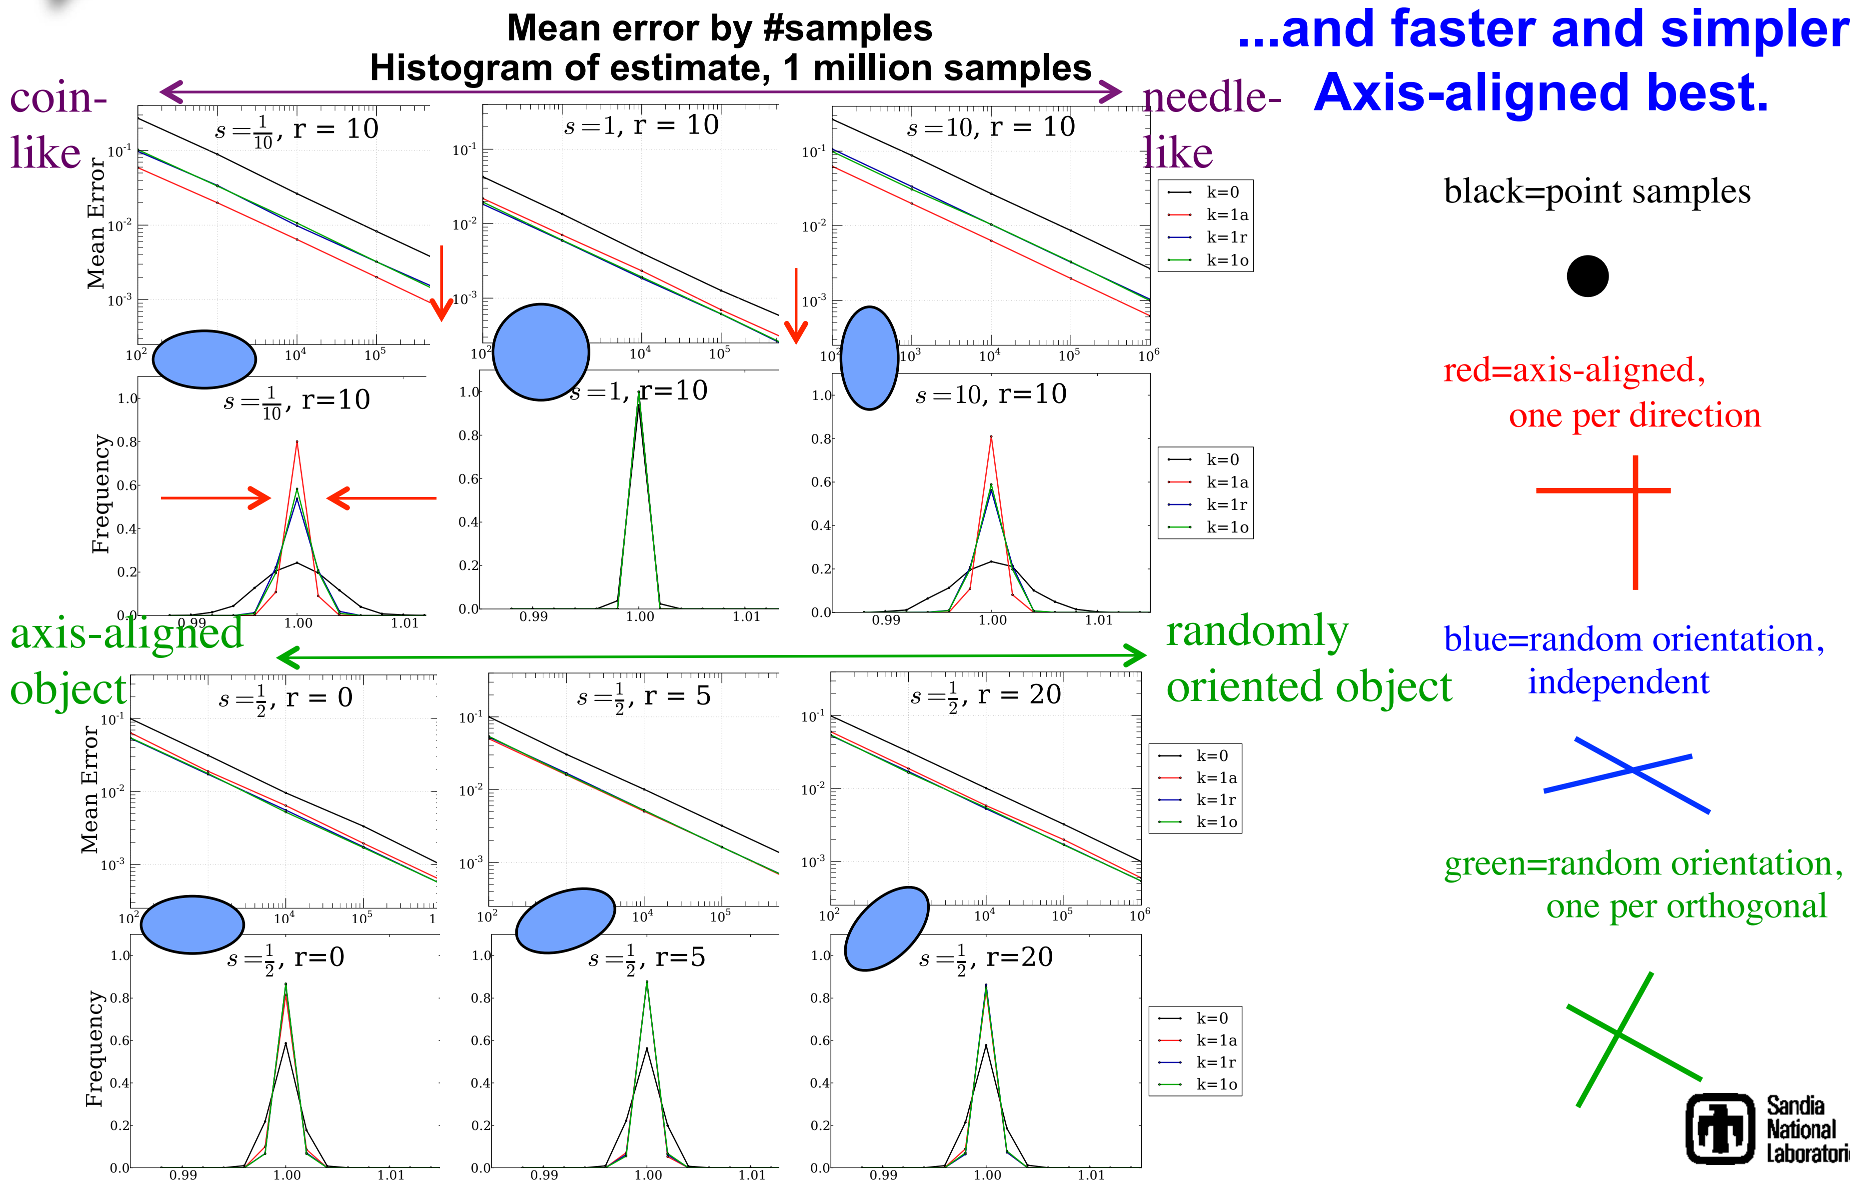

Dart orientation effects

Conclusion

- **(d choose k) orthogonal flats deterministically ≈ randomly**
  - **perhaps because we used so many samples.**
  - **random simpler?**
- **axis-aligned provides**
  - **good quality answers**
  - **simple, fast, through parameter substitution**

Use random-axis orientations, of independent flats

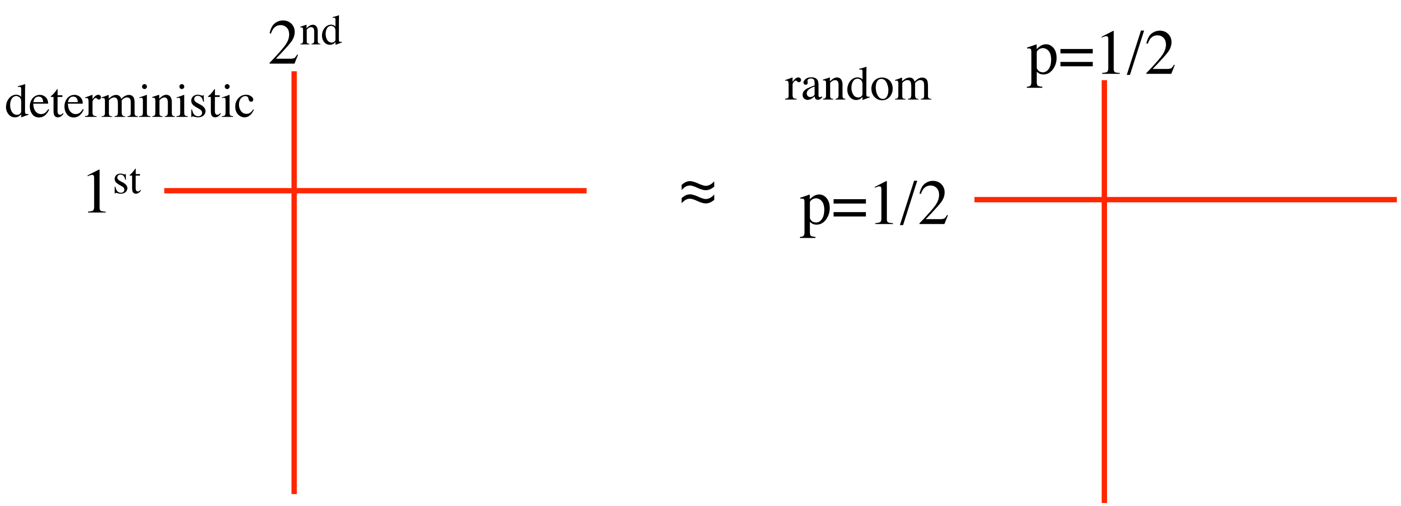

Application 1 of 3

Volume Estimation

LHS patterns

More interesting functions

Sandia National Laboratories

Volume Estimation Speedup
3-d ball, simple analytic

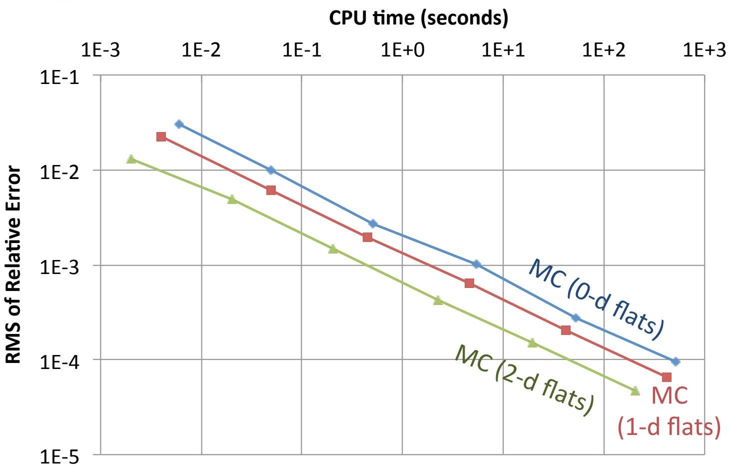

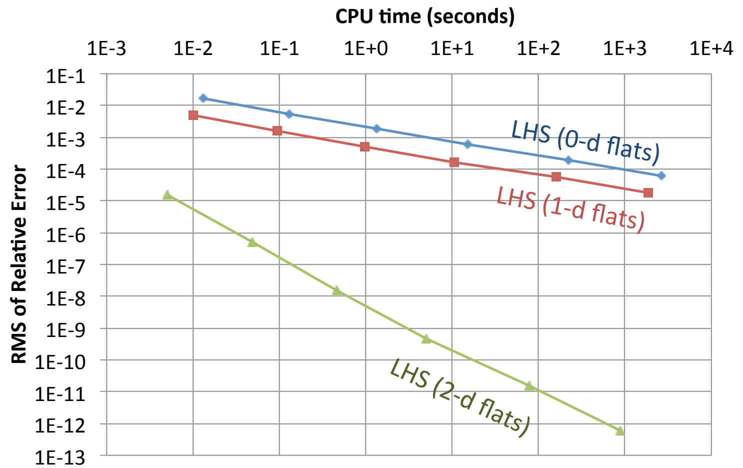

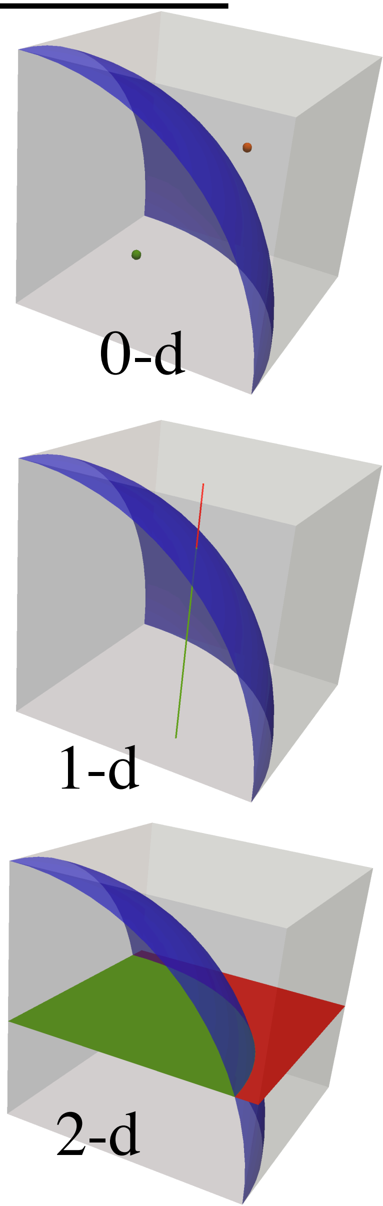

Algorithm: average lengths of lines (area of planes) inside sphere.

Why did this work so well?

—Evaluating f along k-d flats is cheap; in this case we exploited the analytic function of the ball.

—A k-d flat gives more information as k increases.

—A flat is cheap to generate. Each k-d flat requires $d - k$ random numbers; here $d = 3$.

—$(d-1)$-dimensional flats distributed in LHS fashion boosted the convergence rate from $O\left(\frac{1}{\sqrt{n}}\right)$ to $O\left(\frac{1}{n}\right)$.

Volume Estimation Speedup
Circular Parabola

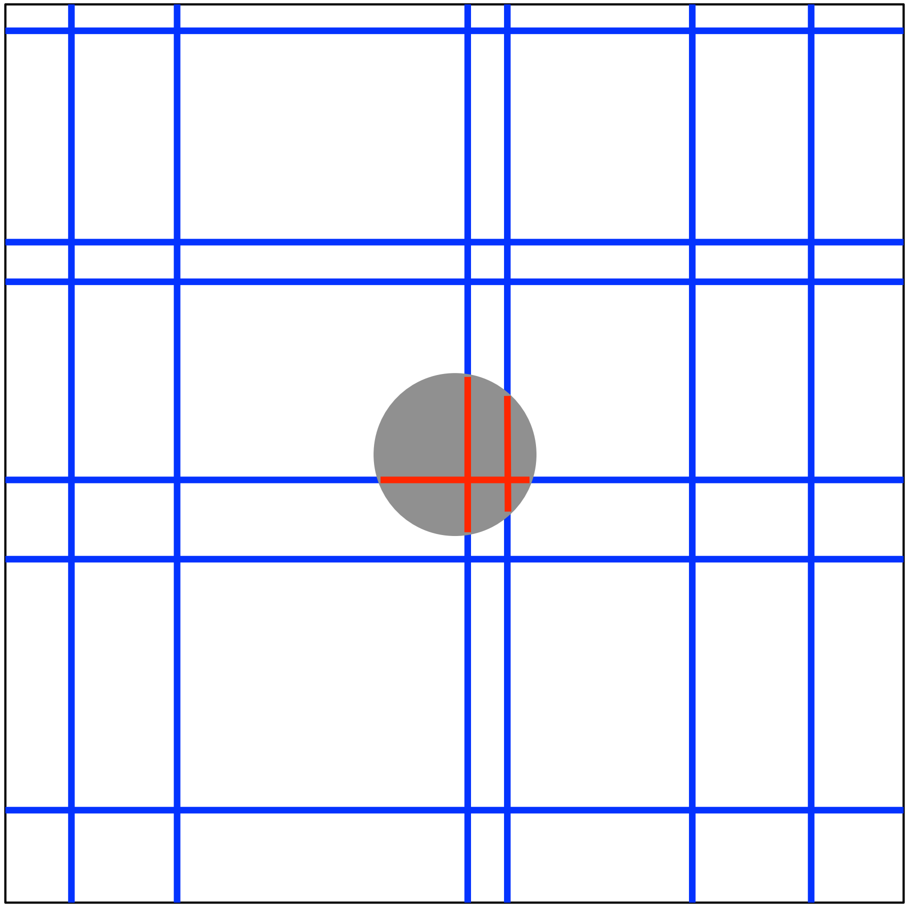

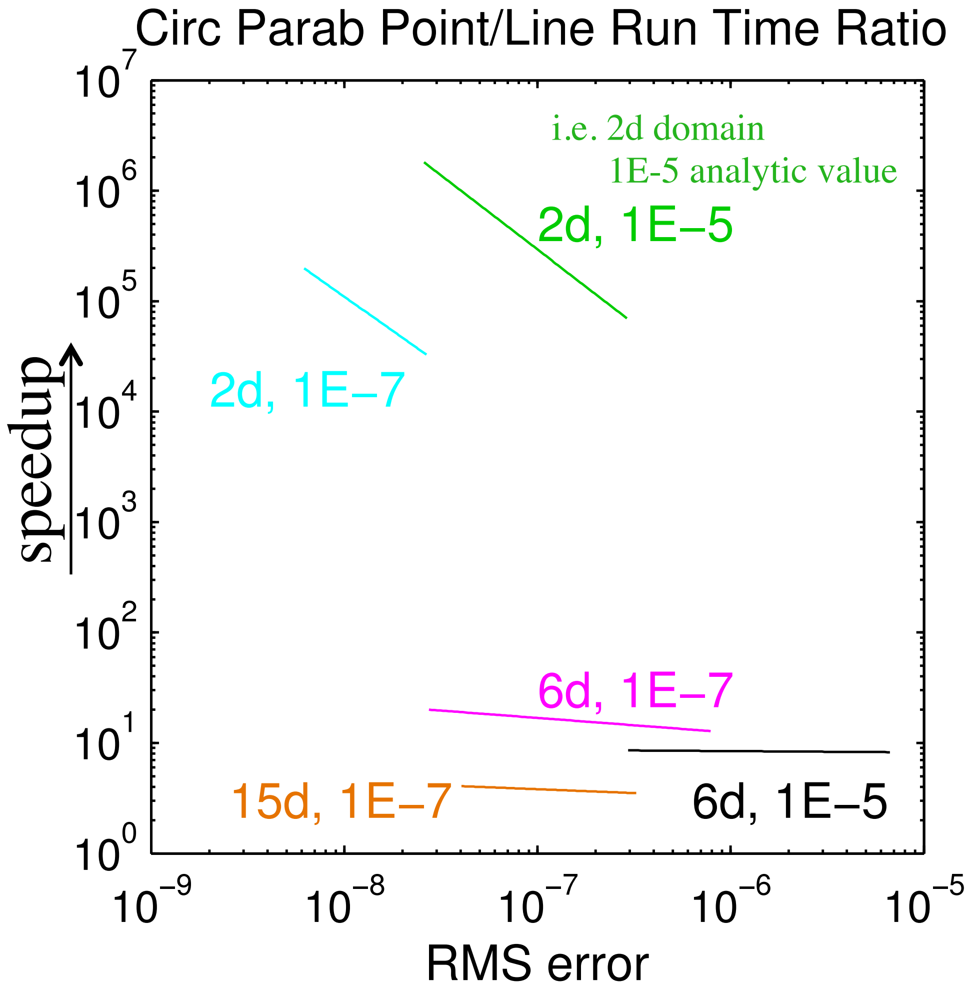

Volume Estimation Speedup
Planar Cross

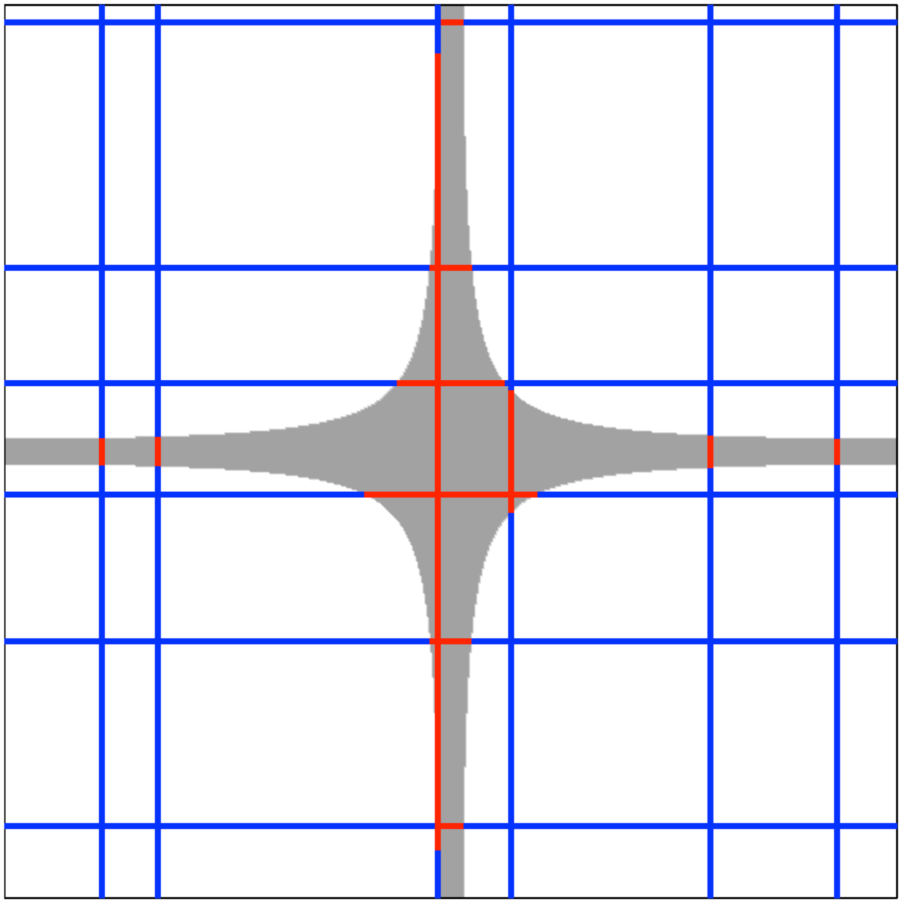

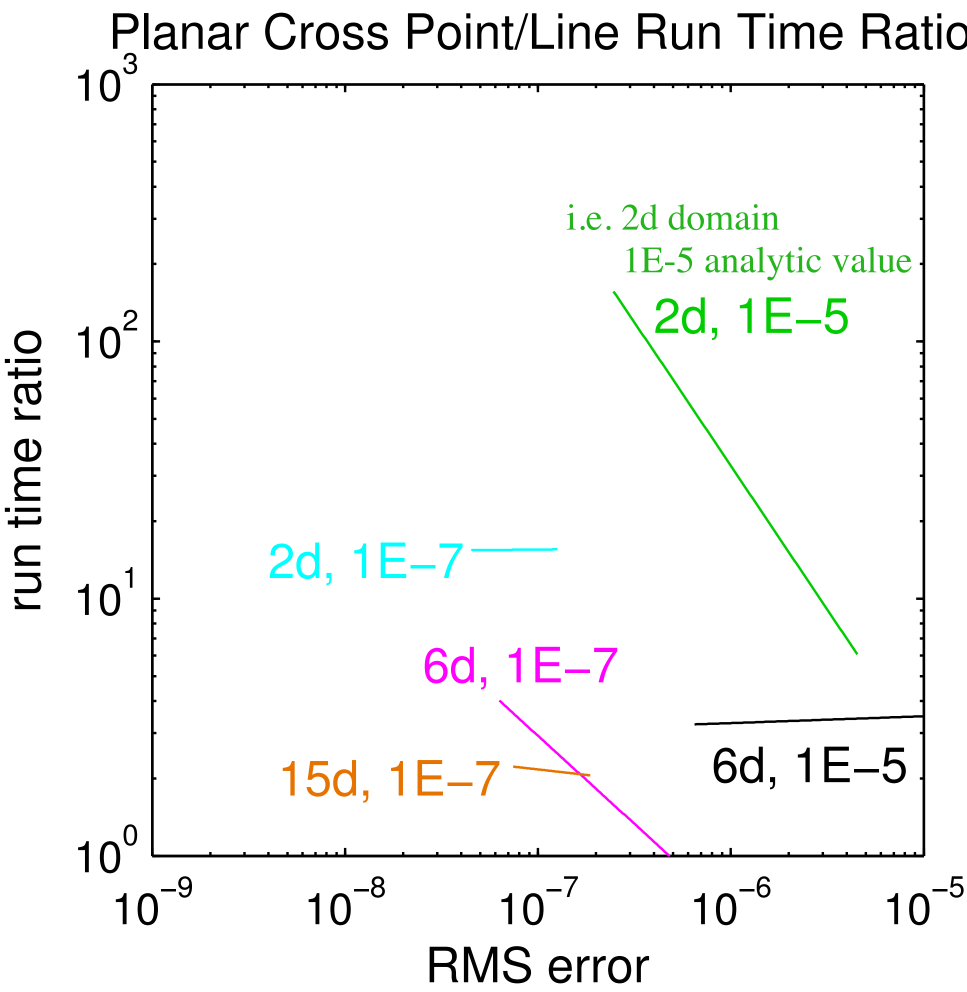

$$y(x) = \left[\prod_{i=1}^{d} \frac{1+\cos(2\pi x_i)}{2}\right]^{1/d}, \quad 0 < x_i < 1.$$ Estimate volume of y(x) < 0

Sandia National Laboratories

Application 2 of 3

Well-spaced points

Use line sampling to generate a point sampling, of a type popular in Graphics texture mapping

Maximal Poisson-Disk Sampling

- **Defined as the limit distribution of a statistical "dart-throwing" process**
  - **Random disks arive with Poisson-distributed arrival times, equivalent to random arrival order:**

Empty disk: $\forall x_i, x_j \in X, x_i \neq x_j : ||x_i - x_j|| \geq r$

Bias-free: $\forall x_i \in X, \forall \Omega \subset \mathcal{D}_{i-1} :$

$$P(x_i \in \Omega) = \frac{\text{Area}(\Omega)}{\text{Area}(\mathcal{D}_{i-1})}$$

Maximal: $\forall x \in \mathcal{D}, \exists x_i \in X : ||x - x_i|| < r$

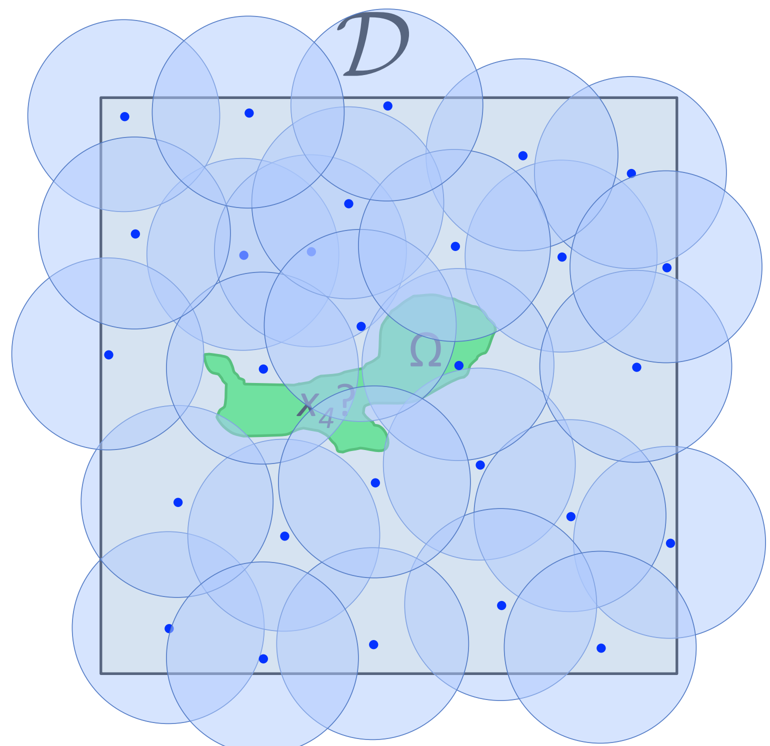

k-d Dart based Relaxed Poisson-Disk Sampling

- **Dart-based**
  - **search space using lines, planes, ...**
- **Relaxed**
  - **stop after many successive dart fails**
  - **expected uncovered volume is small**

Empty disk: $\forall x_i, x_j \in X, x_i \neq x_j : ||x_i - x_j|| \geq r$

~~Bias-free:~~ $\cancel{\forall x_i \in X, \forall \Omega \subset \mathcal{D}_{i-1} :}$

$$\cancel{P(x_i \in \Omega) = \frac{\mathrm{Area}(\Omega)}{\mathrm{Area}(\mathcal{D}_{i-1})}}$$

Maximal: $\cancel{\forall x \in \mathcal{D}, \exists x_i \in X : ||x - x_i|| < r}$

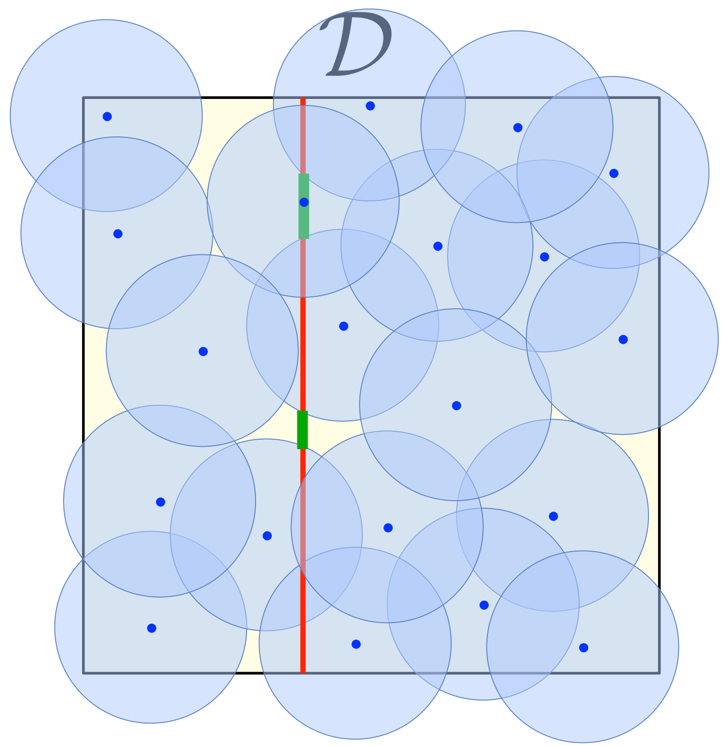

Spectral Quality Evaluation of Point Sets

- **White noise**

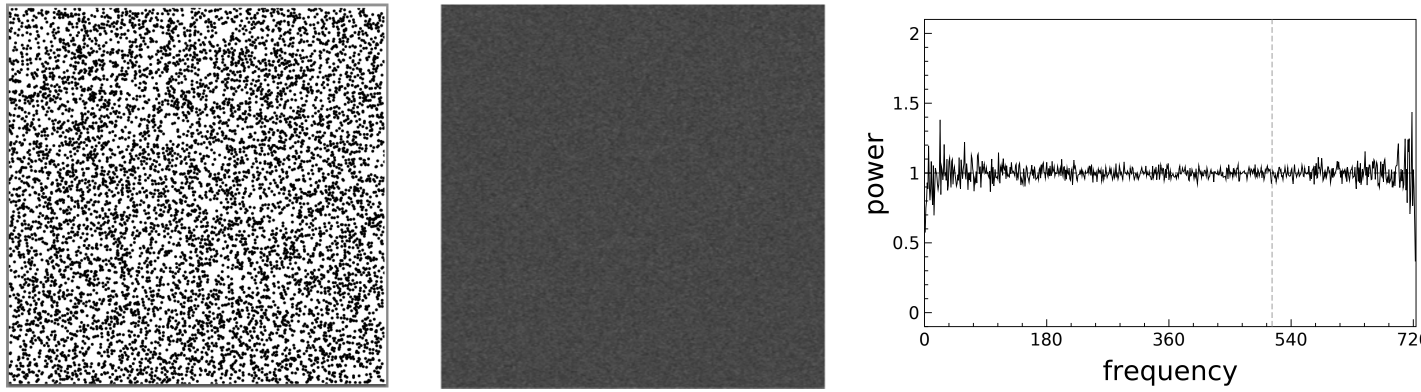

- **Maximal Poisson Disks**

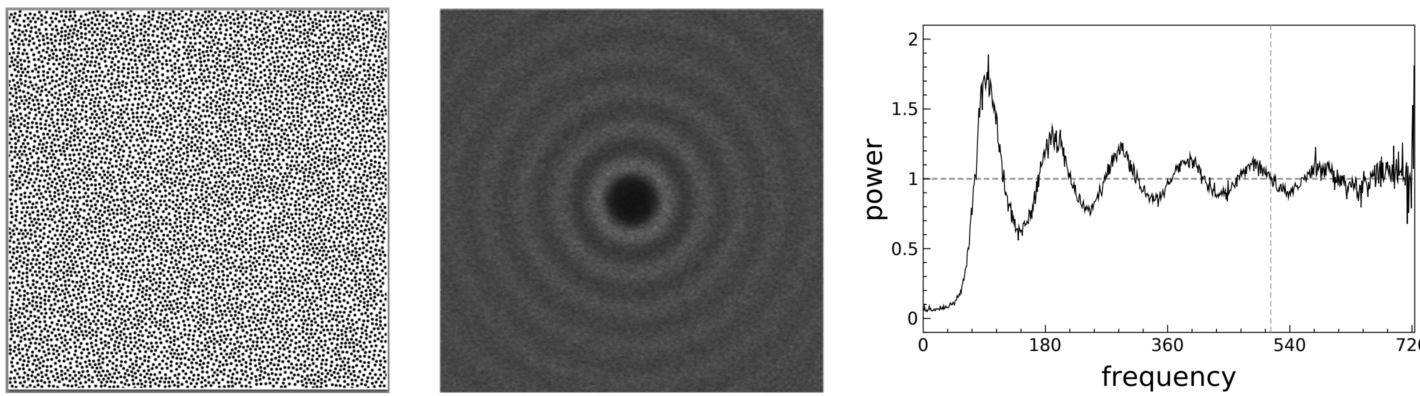

- **Maximal Correlated Disks**

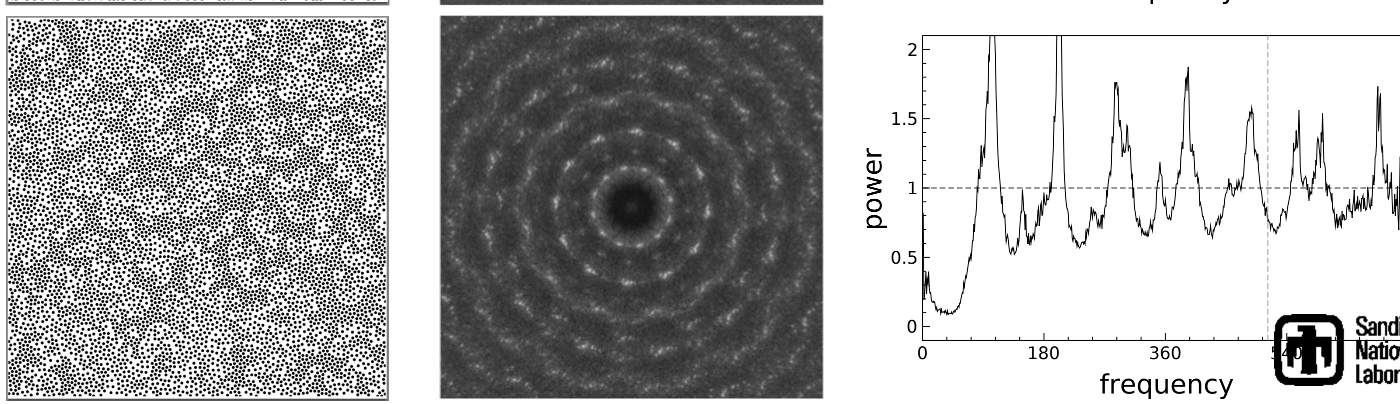

Sandia National Laboratories

MPS vs. line-dart vs. point-dart
Can you tell them apart?

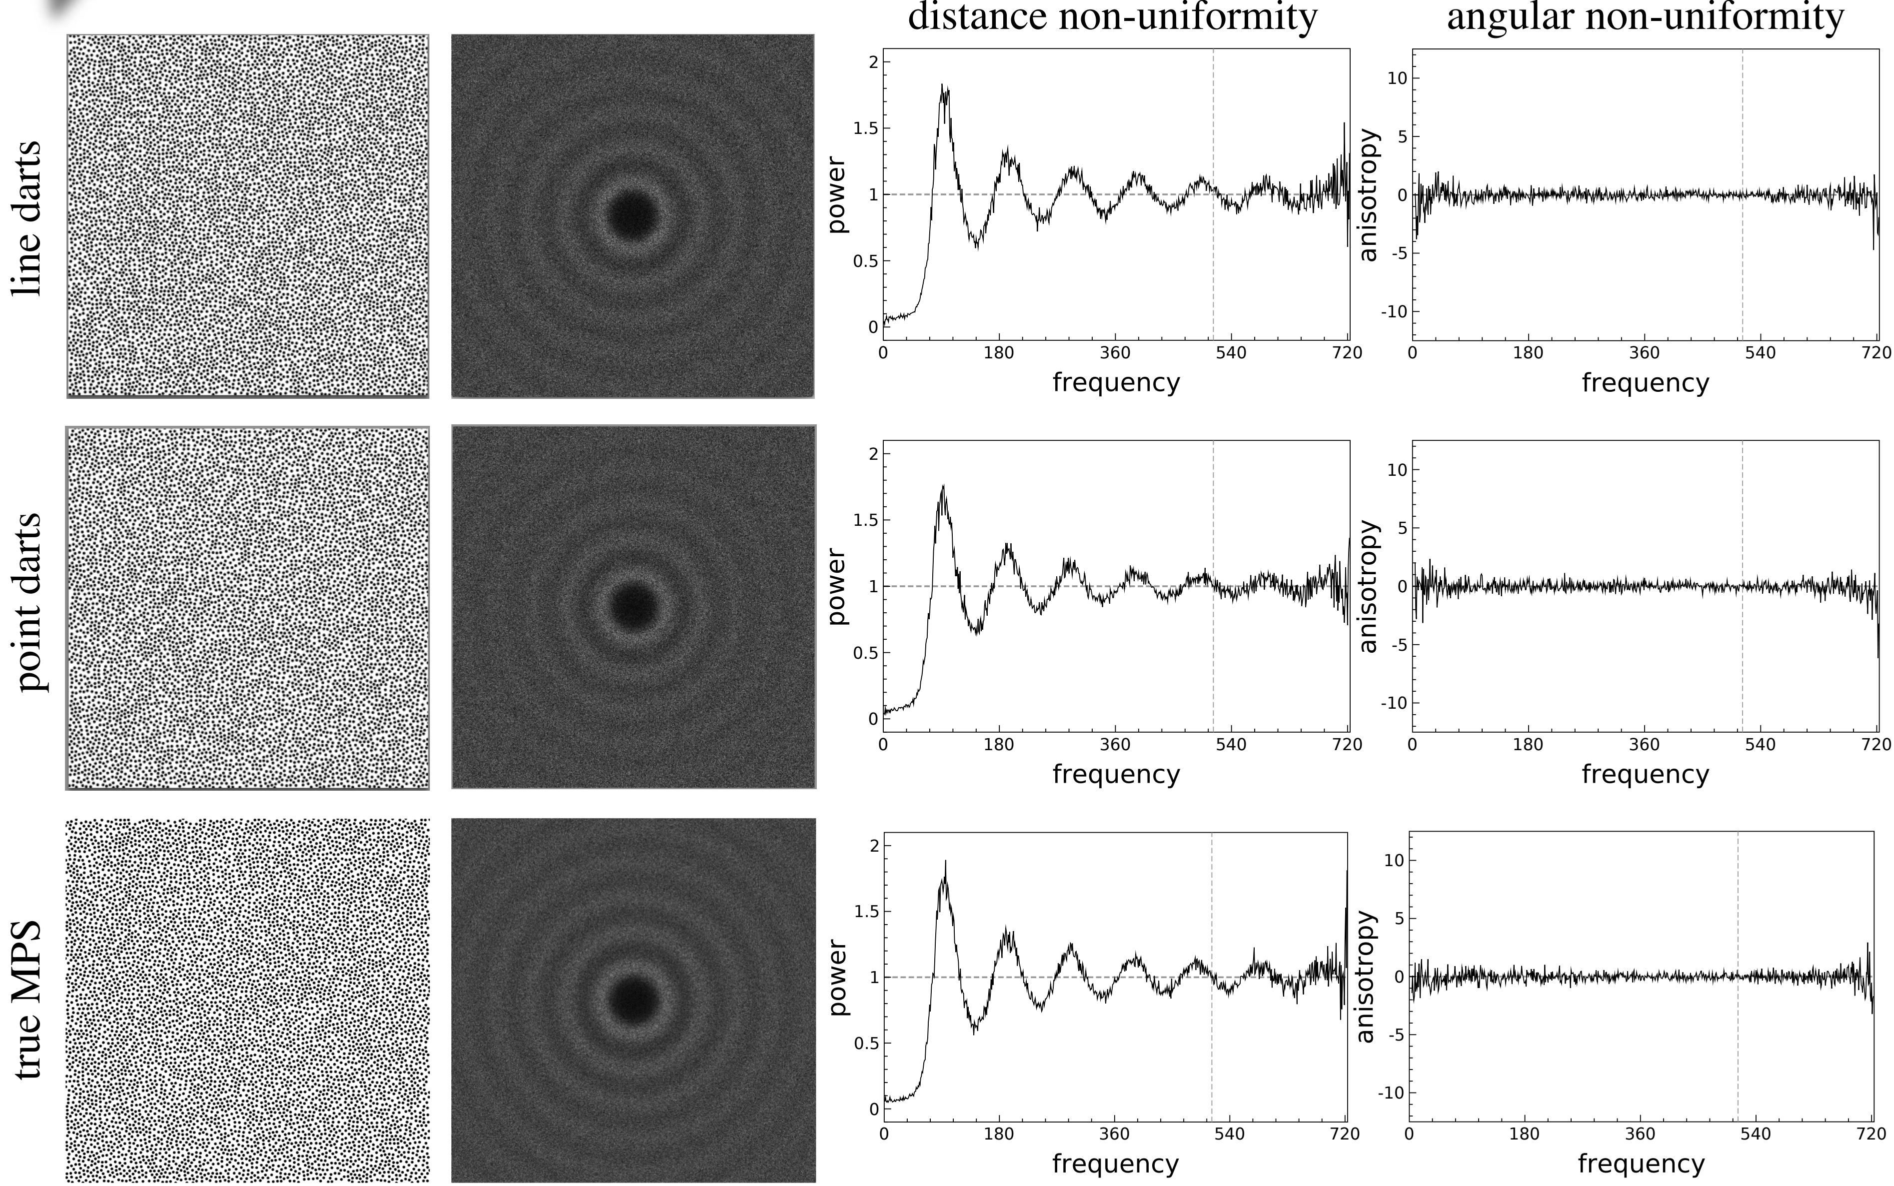

k-d Dart Relaxed MPS Properties

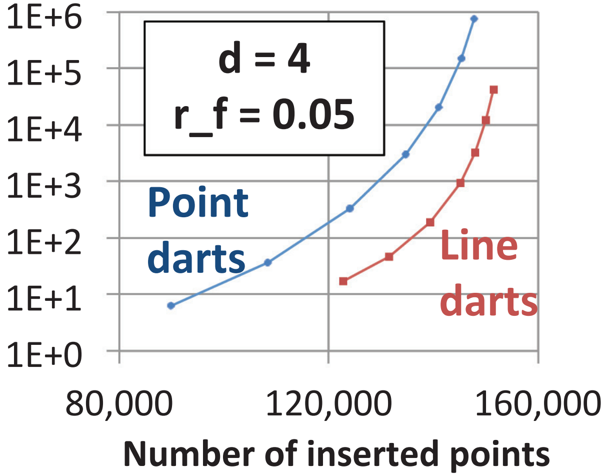

line darts add disks faster
than point darts

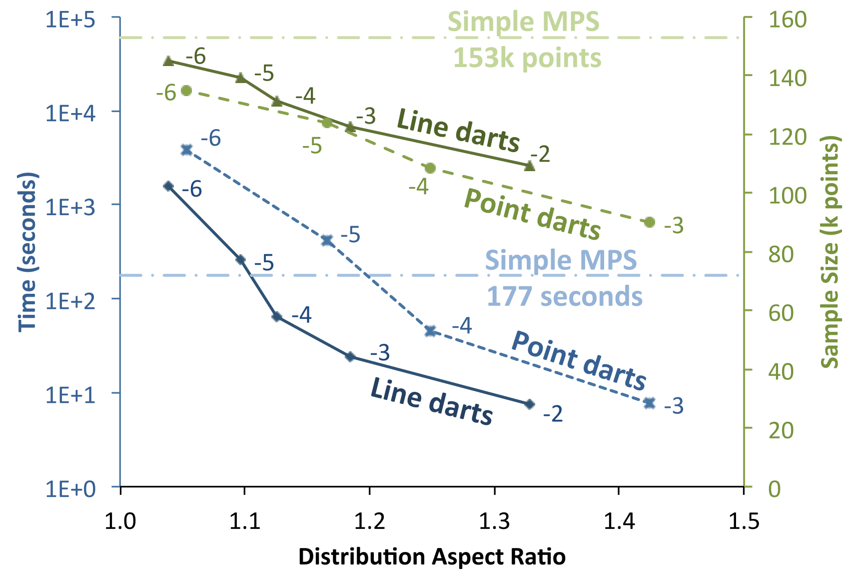

line darts are faster or slower than MPS in d=4,
depending on relaxed maximality

Simple MPS requires 2^d memory, intractable in d>6
but line-darts are linear memory

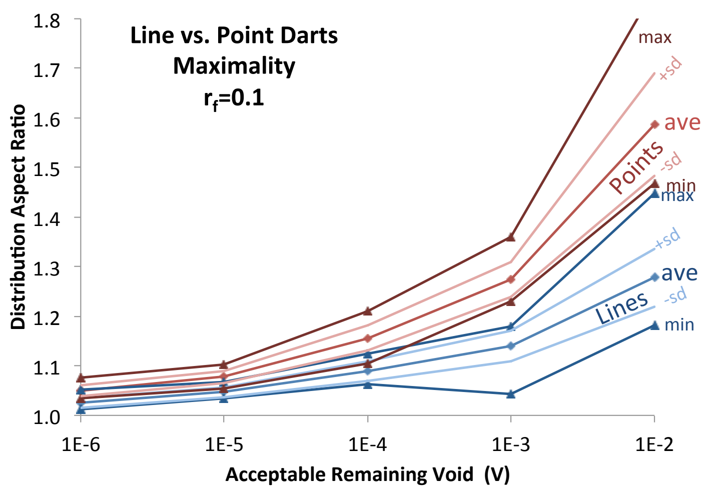

better quality
line darts produce fewer large gaps
than point darts

Application 3 of 3

- Graphics, depth of field blur
- Integration

Graphics Application
Depth of Field Sampling

Pixar's Toy Story 3

blurred
far from lens
focal plane

Depth-of-Field blur
requires many point samples
per pixel !

movie frame rendering is overnight,
not realtime,
as many simulation and UQ studies

Sandia
National
Laboratories

Line-sampling for Depth of Field Blur

- **Solution: point sampling**

  For every (a,b) pixel

  4-d: (u,v) lens plane, (x,y) scene space

  z scene space is dependent

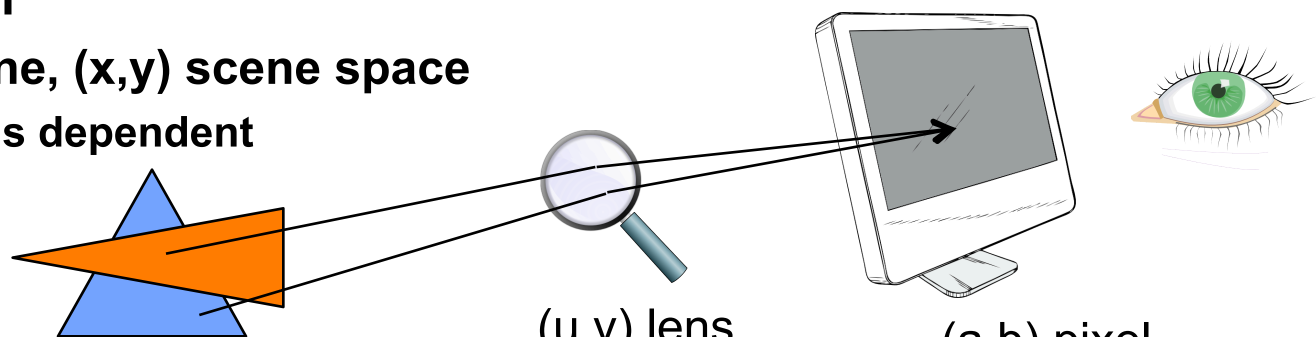

- **Solution: our algorithm**

  lines sample and compute

  occlusion depth (decision)

  color contribution (integration)

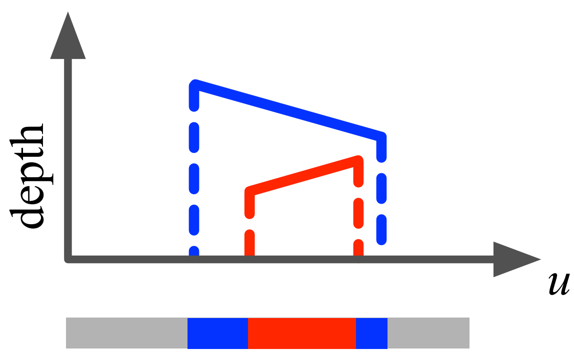

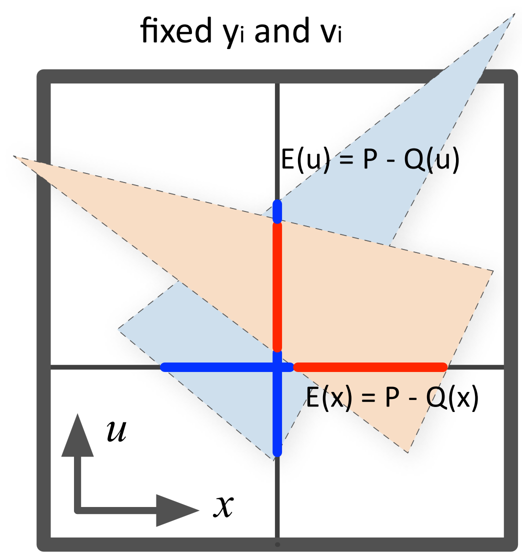

axis-aligned sample lines in (u,v,x,y)

e.g. pick x, y, v, let u vary

e.g. pick y, u, v, let x vary

Sandia National Laboratories

Blur, line-darts vs. point samples

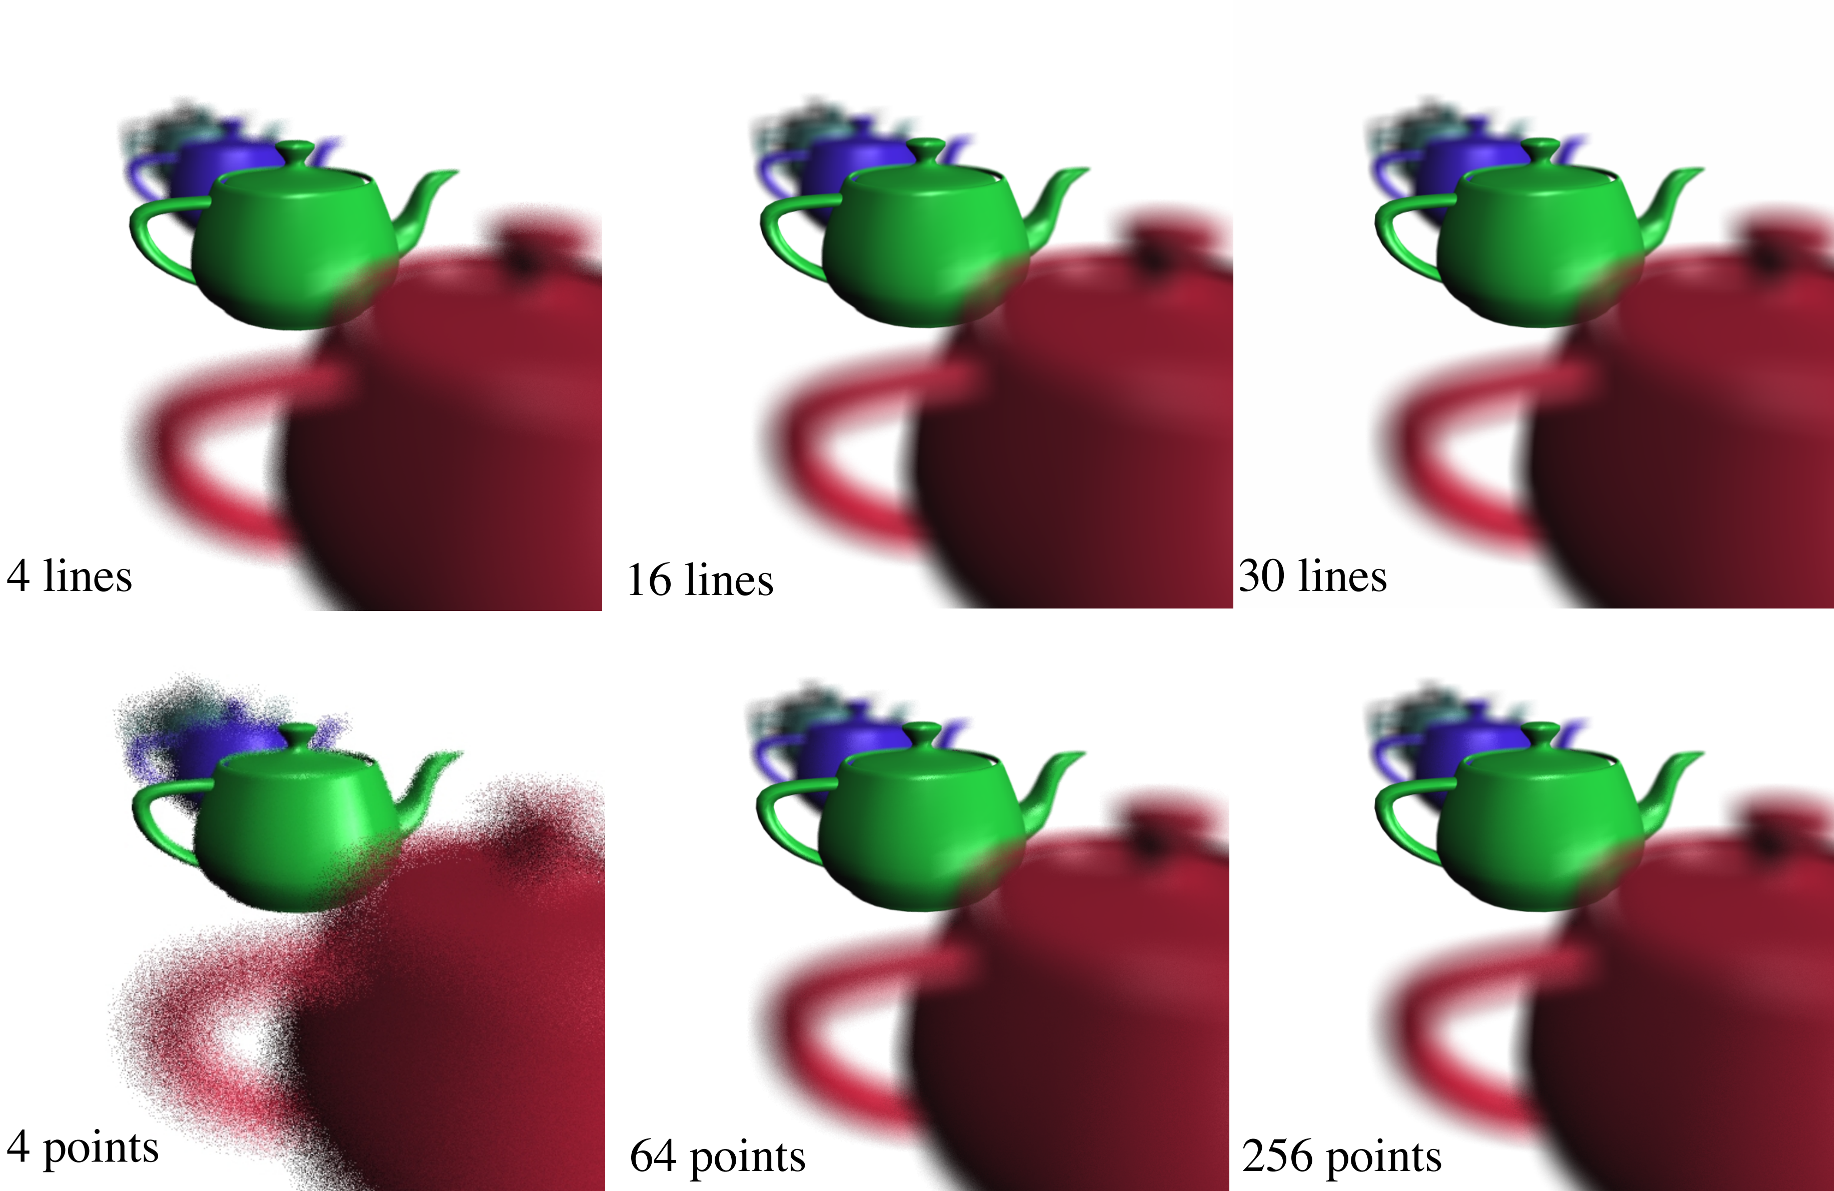

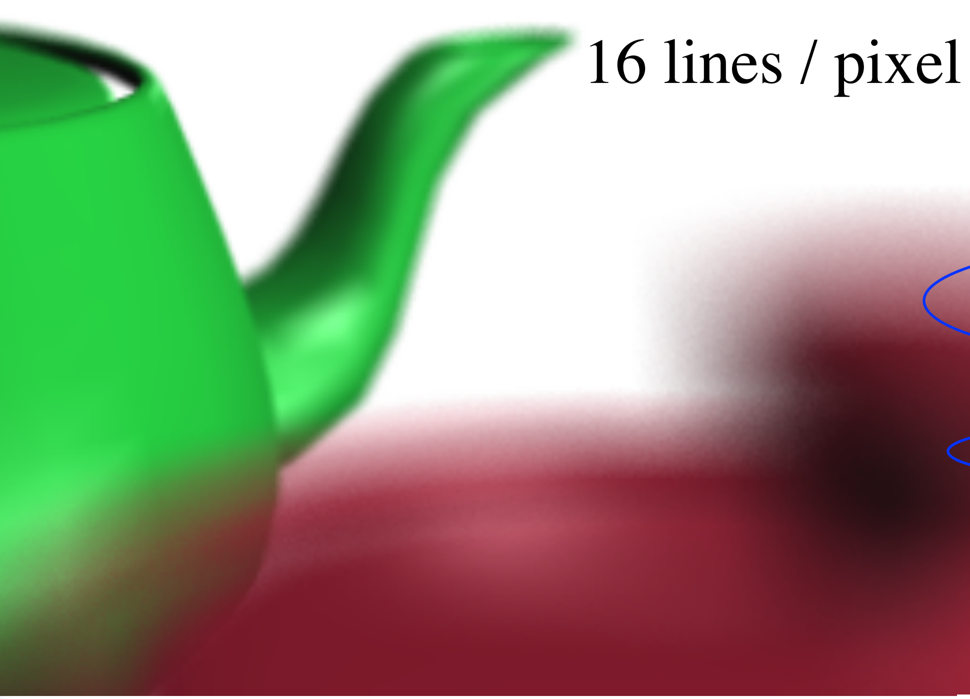

Performance of Our Point vs. Line Darts

		Rendering Time (s)	
Sample Type	Sample Count	Cessna	Teapot
	64	29.6	52.1
Points	256	116.7	198.6
	1024	453.0	792.1
	4	14.9	24.5
Line Darts	16	56.8	91.9
	30	105.1	169.4

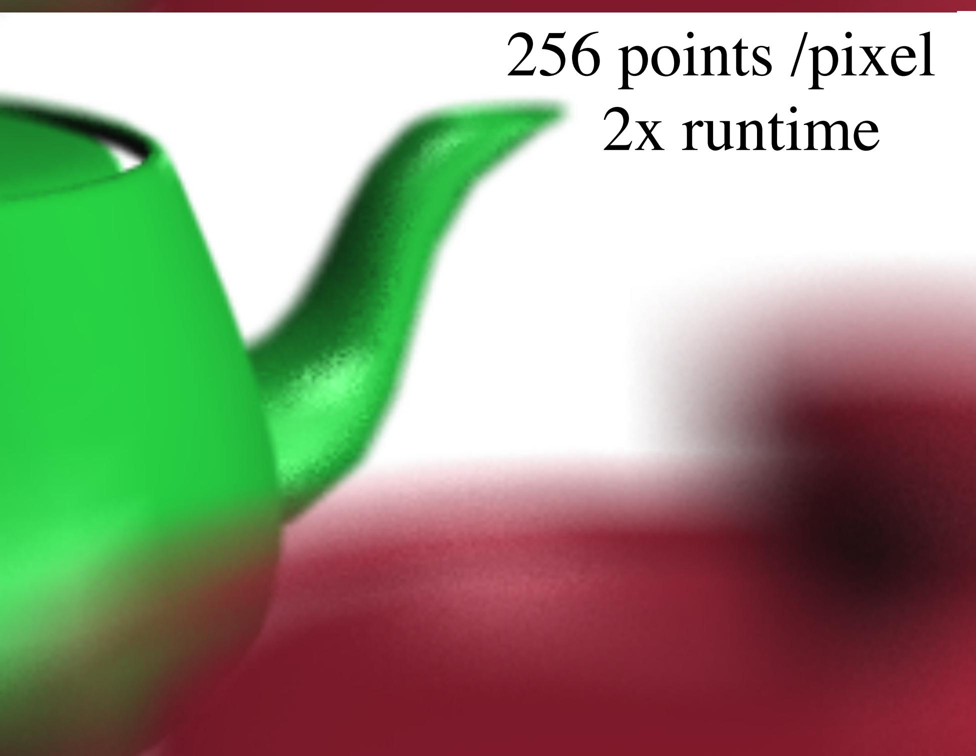

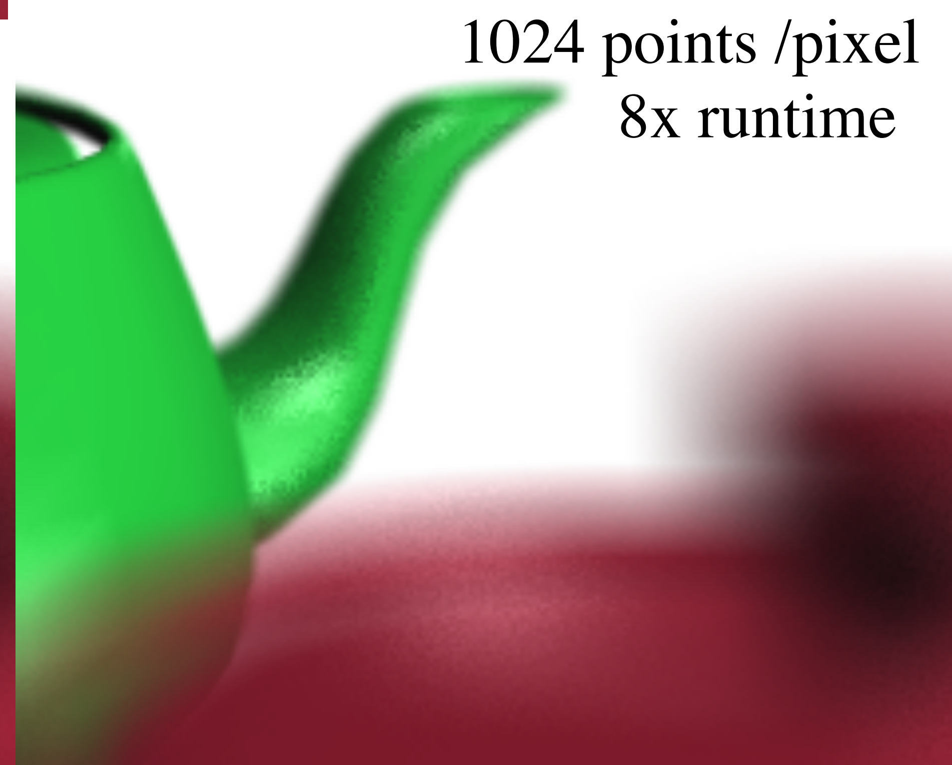

Summary

- Any point-sampling algorithm depending on function averages can be converted to a line-sampling algorithm
  - Including indicator functions e.g. volume estimation

$$u_{estimate} = \frac{\sum \text{WeightedVolume}(\mathit{flat})}{\sum \text{Volume}(\mathit{flat})} = \frac{\sum \text{Length}(\mathit{Line\ inside\ grey})}{\sum \text{Length}(\mathit{Line})}$$

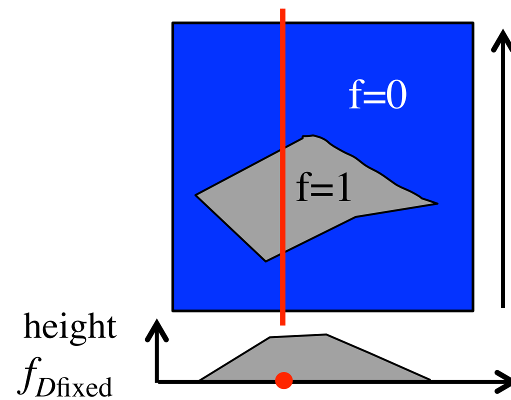

  - Need to evaluate function along a flat (line)
  - Efficiency depends on evaluation speed
    - This is the challenge for practical *k*-dimensional flats
  - Axis-aligned flats (lines)
    - efficient and random-enough
- Application variants
  - Generate a well-spaced disk packing
    - Although line samples are not uniform by area, effect on output distribution is unoticable.

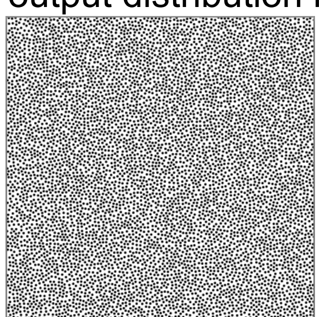

  - Depth of Field blur
    - Intrusive line-sampling
    - Efficient function integration without artifacts

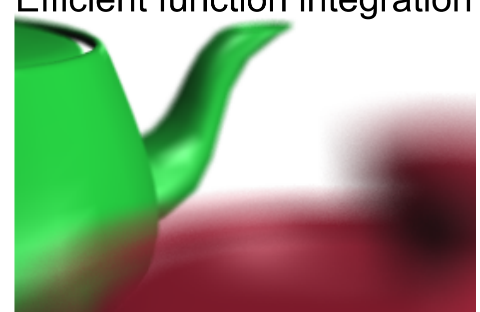

Sandia National Laboratories

Extra stuff

Heilmeier's Catechism

1. What is the problem, why is it hard?
   - Uncertainty quantification, small failure regions in vast spaces, expensive functions
2. How is it solved today?
   - Many sampling methods based on statistics and analysis
3. What is the new technical idea; why can we succeed now?
   - Borrowing Computational Geometry, Graphics concepts:
     - line searches
     - sample-neighborhoods, geometric balls
     - functional integration
4. What is the impact if successful?
   - Increased convergence rates, fewer parallel simulations

a computational geometer's view
Space for All Point Sampling Methods

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

Spatial Randomness

Fourier Spectrum, Power and Anisotropy

Pairwise Distances, Edge Orientations

Blue Noise

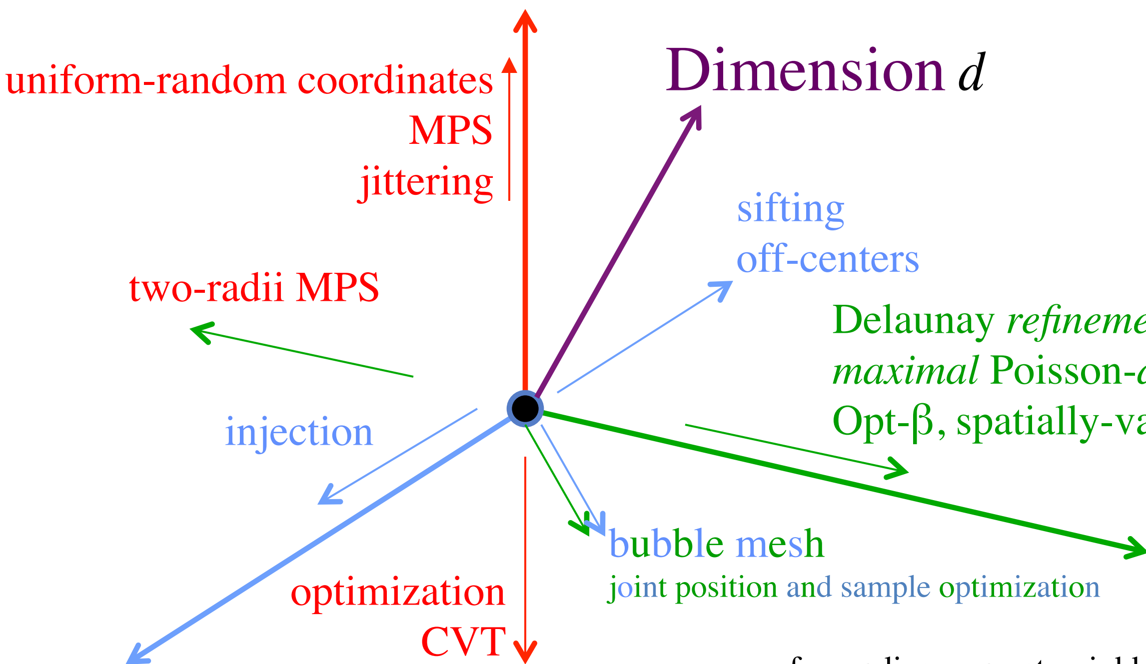

Geometric Density

r_f free radius, nearest-neighbor distance; Delaunay edge lengths

r_c coverage radius, Vornoi vertex distance

$\beta = r_c / r_f$ Distribution Aspect Ratio; DT angles,Vor cell aspect ratio

Lipschitz Conditions

Unique Coverage

Discrete Density

n number of samples

kissing number

number of neighbors, edges, cells,

Sandia National Laboratories

Main Challenge of Solving MPS in Higher Dimensions (Curse-Of-Dimensionality)

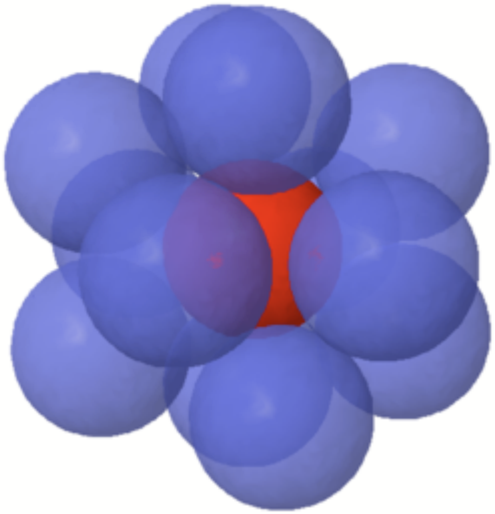

- Curse of dimensionality
  - Natural: Kissing Number grows exponentially with dimension
  - Artificial: Grid based methods (state of the art) to retrieve neighbors, and track remaining voids

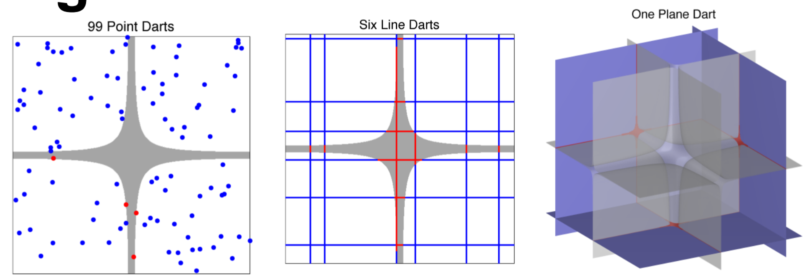

- Generalization of Sampling Entities
  - *k*-d darts: Random sampling using hyperplanes

→ void capturing, integration and UQ

k-d Darts for Solving MPS