SAND2018-8796C

Spoke-Darts for High-Dimensional Blue-Noise Sampling

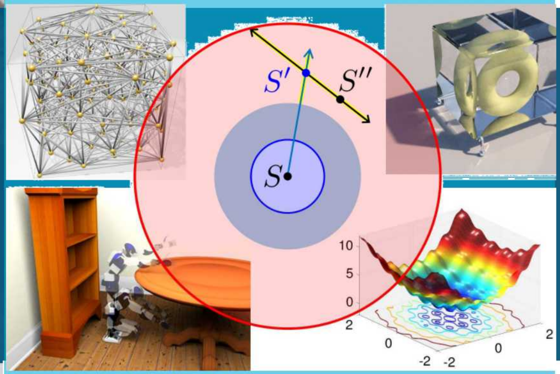

Center for Computing Research

Sandia National Laboratories

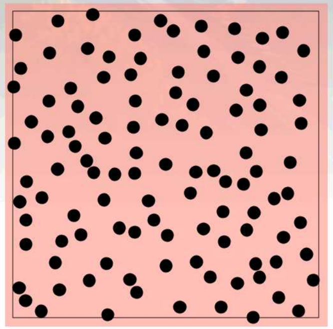

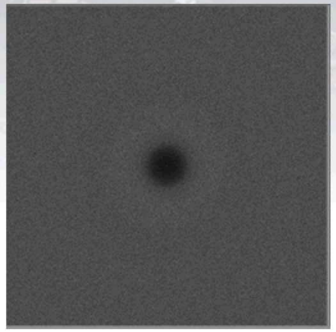

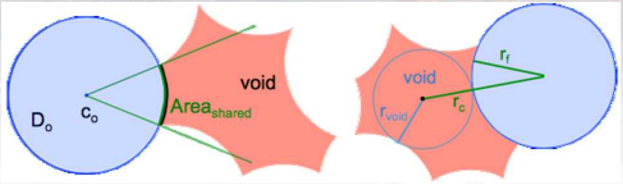

$$m = \lceil (-\ln \epsilon)(\beta^* - 1)^{1-d} \rceil \Leftrightarrow \beta^* = 1 + \left(\frac{-\ln \epsilon}{m} \right)^{1/(d-1)}$$

PRESENTED BY

Scott A. Mitchell

ACM TOG 37:2, May 2018

Scott A. Mitchell, Mohamed S. Ebeida, Muhammad A. Awad, Chonhyon Park, Anjul Patney, Ahmad A. Rushdi, Laura P. Swiler, Dinesh Manocha (now UMD), Li-Yi Wei

SIGGRAPH 2018
18 minutes talk
15 Aug 10:45 am session talk 4
West Building, Room 109-110

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

SAND2018-xxxx

Motivation

High-dimensional
blue-noise sampling is useful
but hard to generate

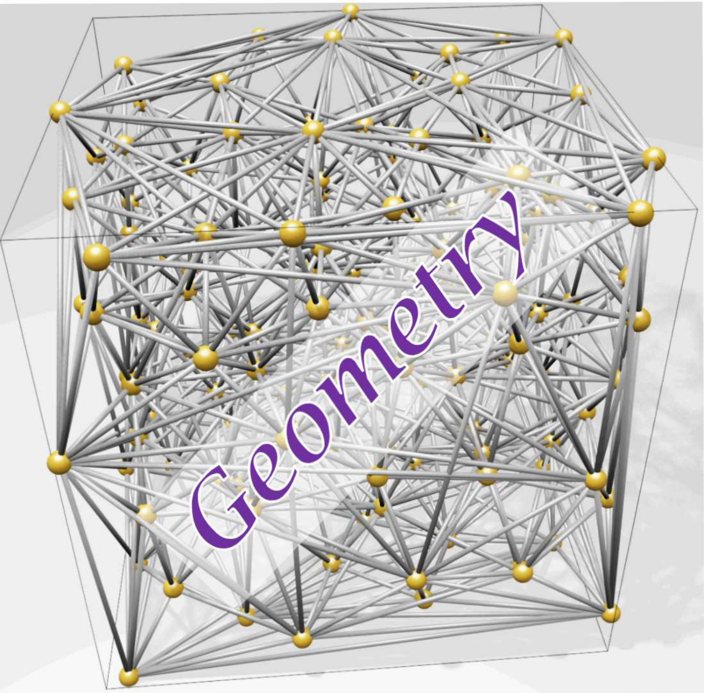

8-100D: approximate
Delaunay Graphs
(Voronoi neighbors)

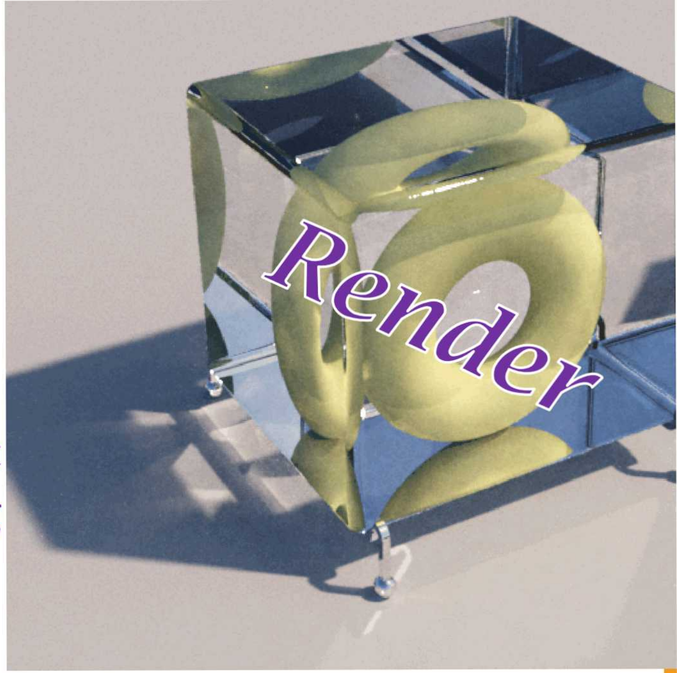

8D: proof-of-concept
path tracing

Social proof of utility:
Bridson SIGGRAPH 2007
cited 200 times

23D: find a path in
robot configuration space

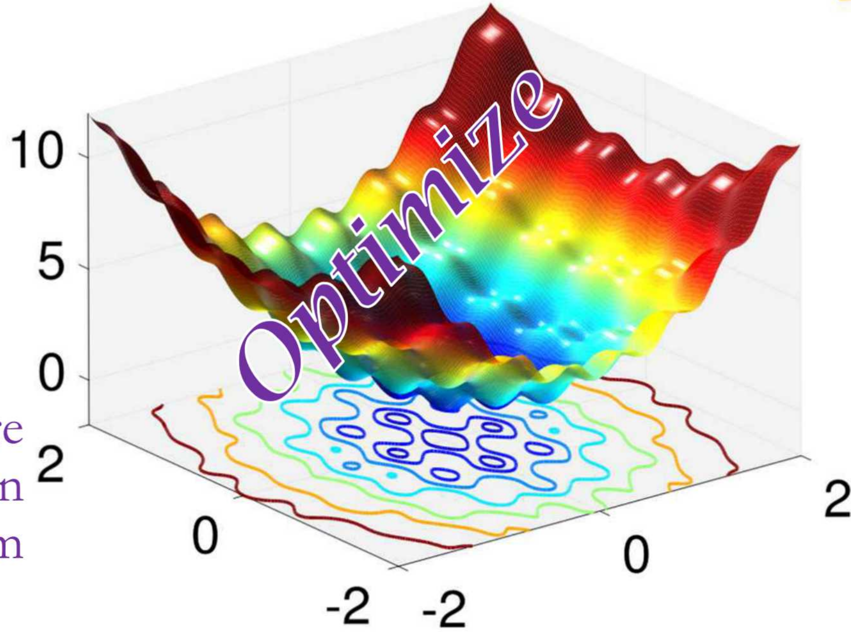

100D: adaptively explore
black-box function
to find global minimum

Motivation

Goal: algorithm to produce point distributions

Requirements

- ✓Output: blue-noise, uniform-random except
  - ✓No points too close
  - ✓No big gaps
- ✓Algorithm
  - ✓Memory & time scales to high dimensions (e.g. 20D)
  - ✓Locally adaptive, general domain shapes
- ✓Confidence: provable output quality

Maximum domain-to-sample distance

Minimum sample separation

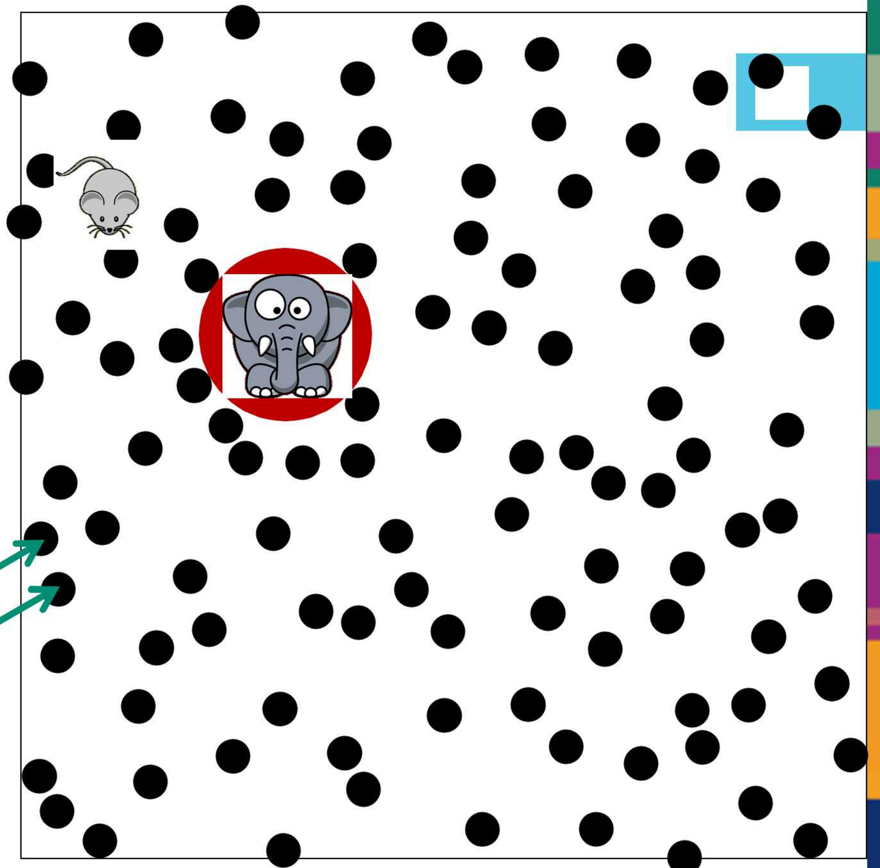

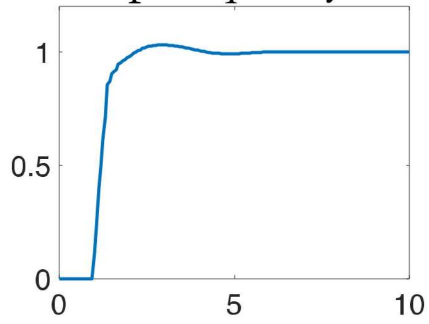

RDF Radial Distance Function

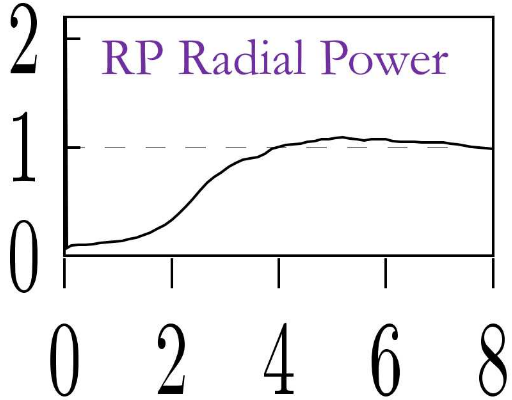

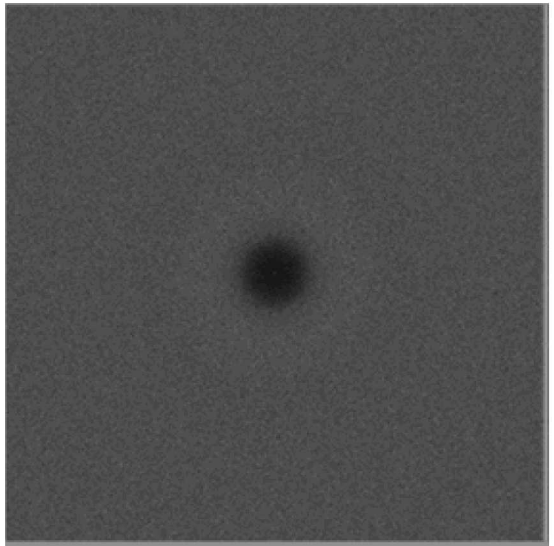

What's wrong with the algorithms we already have?

SimpleMPS guarantees saturation

$$\text{runtime \& memory } = O(n)$$

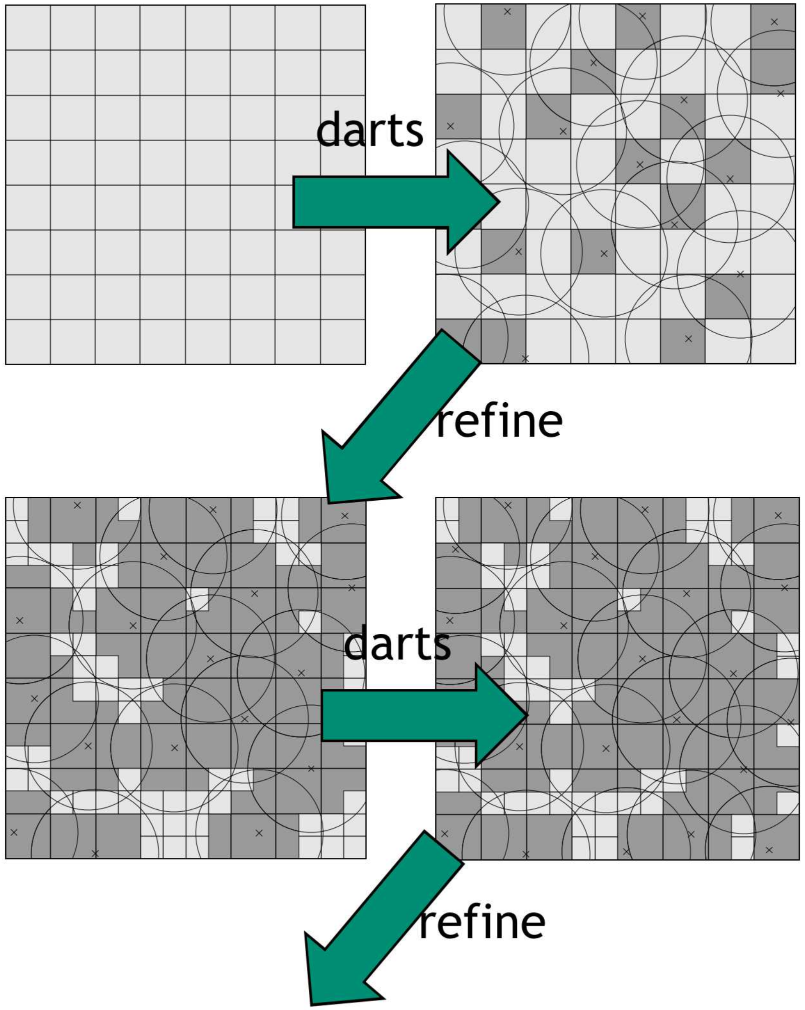


```
divide domain into cubes
while ∃cubes, and cube diagonal > machine precision
  do (A(d) × #cubes) times:
    pick a cube
    pick sample from cube
    if distance(sample, prior samples) > r
      accept sample
      discard cube
  refine cubes
  discard covered cubes
```

SimpleMPS, "A Simple Algorithm for Maximal Poisson-Disk Sampling in High Dimensions," Eurographics 2012.

What's wrong with the algorithms we already have?

SimpleMPS guarantees saturation … but doesn't scale by dimension, d. ☹

$$\text{runtime \& memory } = O(nd^{d/2})$$

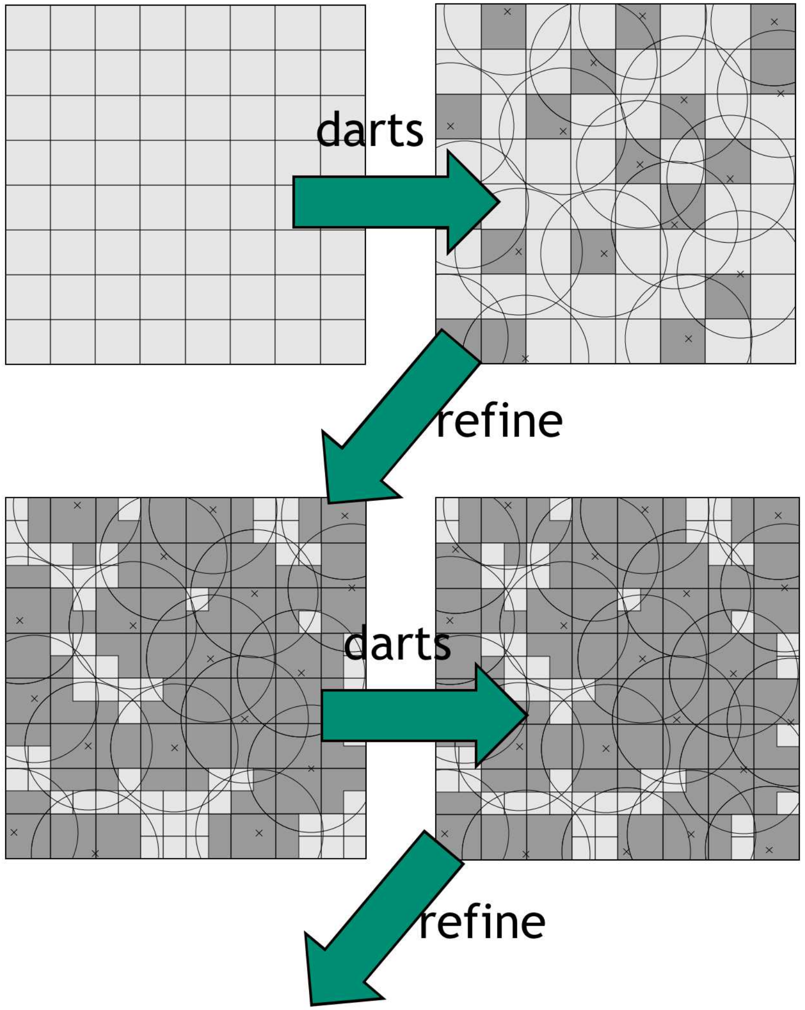

Problem: cube poorly approximates sphere
Exponentially worse as dimension increases
More empty boxes, memory limit d≈6

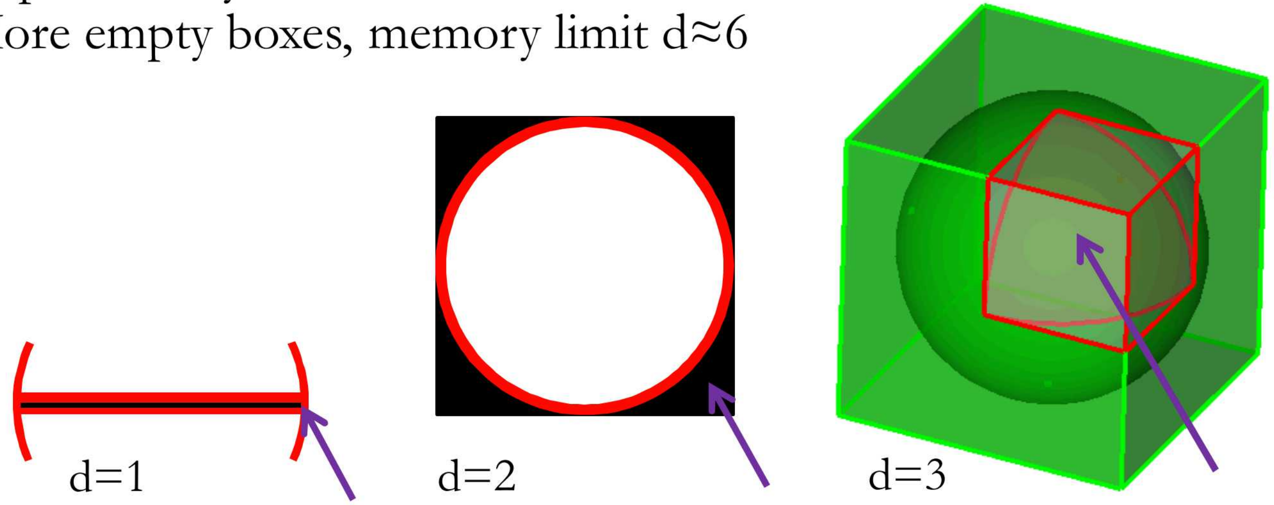

d=1	d=2	d=3	… d→∞
V(cube) = 1	V(cube) = 1	V(cube) = 1	V(c) = 1
V(sphere) = 1	V(sphere) = 0.785…	V(sphere) = 0.523…	V(s) → 0

Any algorithm that constructs (an approximation of) remaining sample space is doomed to be exponential-in-d

SimpleMPS, "A Simple Algorithm for Maximal Poisson-Disk Sampling in High Dimensions," Eurographics 2012.

What's wrong with the algorithms we already have?

Bridson 2007* scales

$$\text{runtime \& memory } = O(dn^2)$$

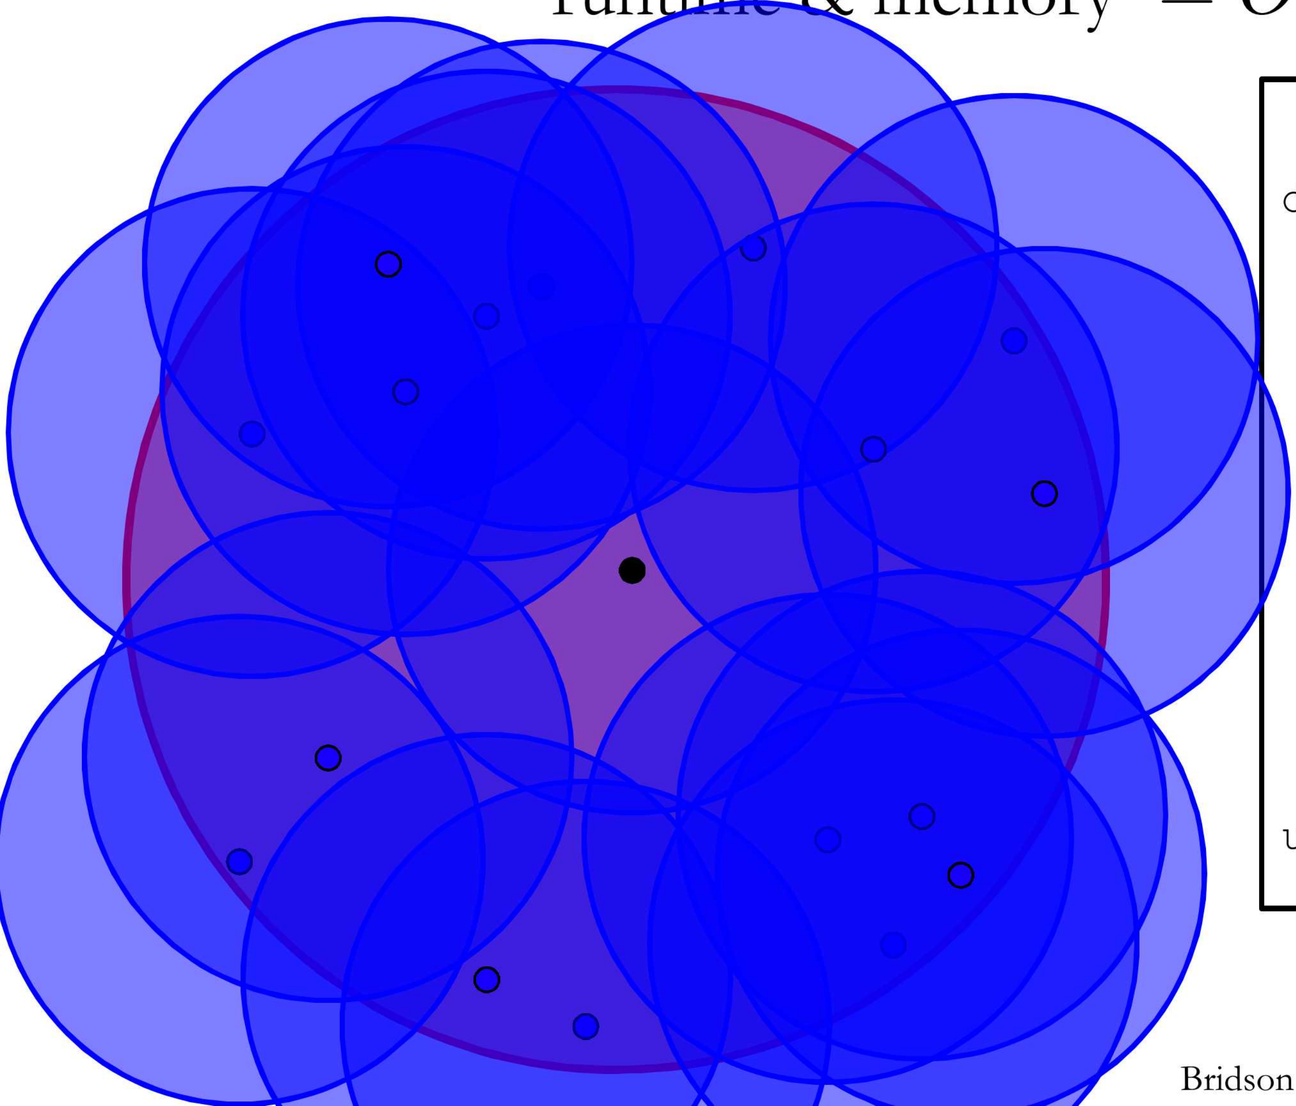

Bridson 2007* pseudocode

```
do
  prior = randomly pop Front
  do
    pick sample from volume of
      (r,2r) annulus of prior
    if distance(sample, all samples) > r
      accept sample
      add sample to Front
  until 30 consecutive rejections
until Front is empty
```

30+ depending on dimension

*Modified: no background grid, no exponential-in-d complexity.

Bridson 2007, "Fast Poisson disk sampling in arbitrary dimensions." SIGGRAPH '07 sketch.

What's wrong with the algorithms we already have?

Bridson 2007* scales … but doesn't guarantee saturation ☹

$$\text{runtime \& memory } = O(dn^2)$$

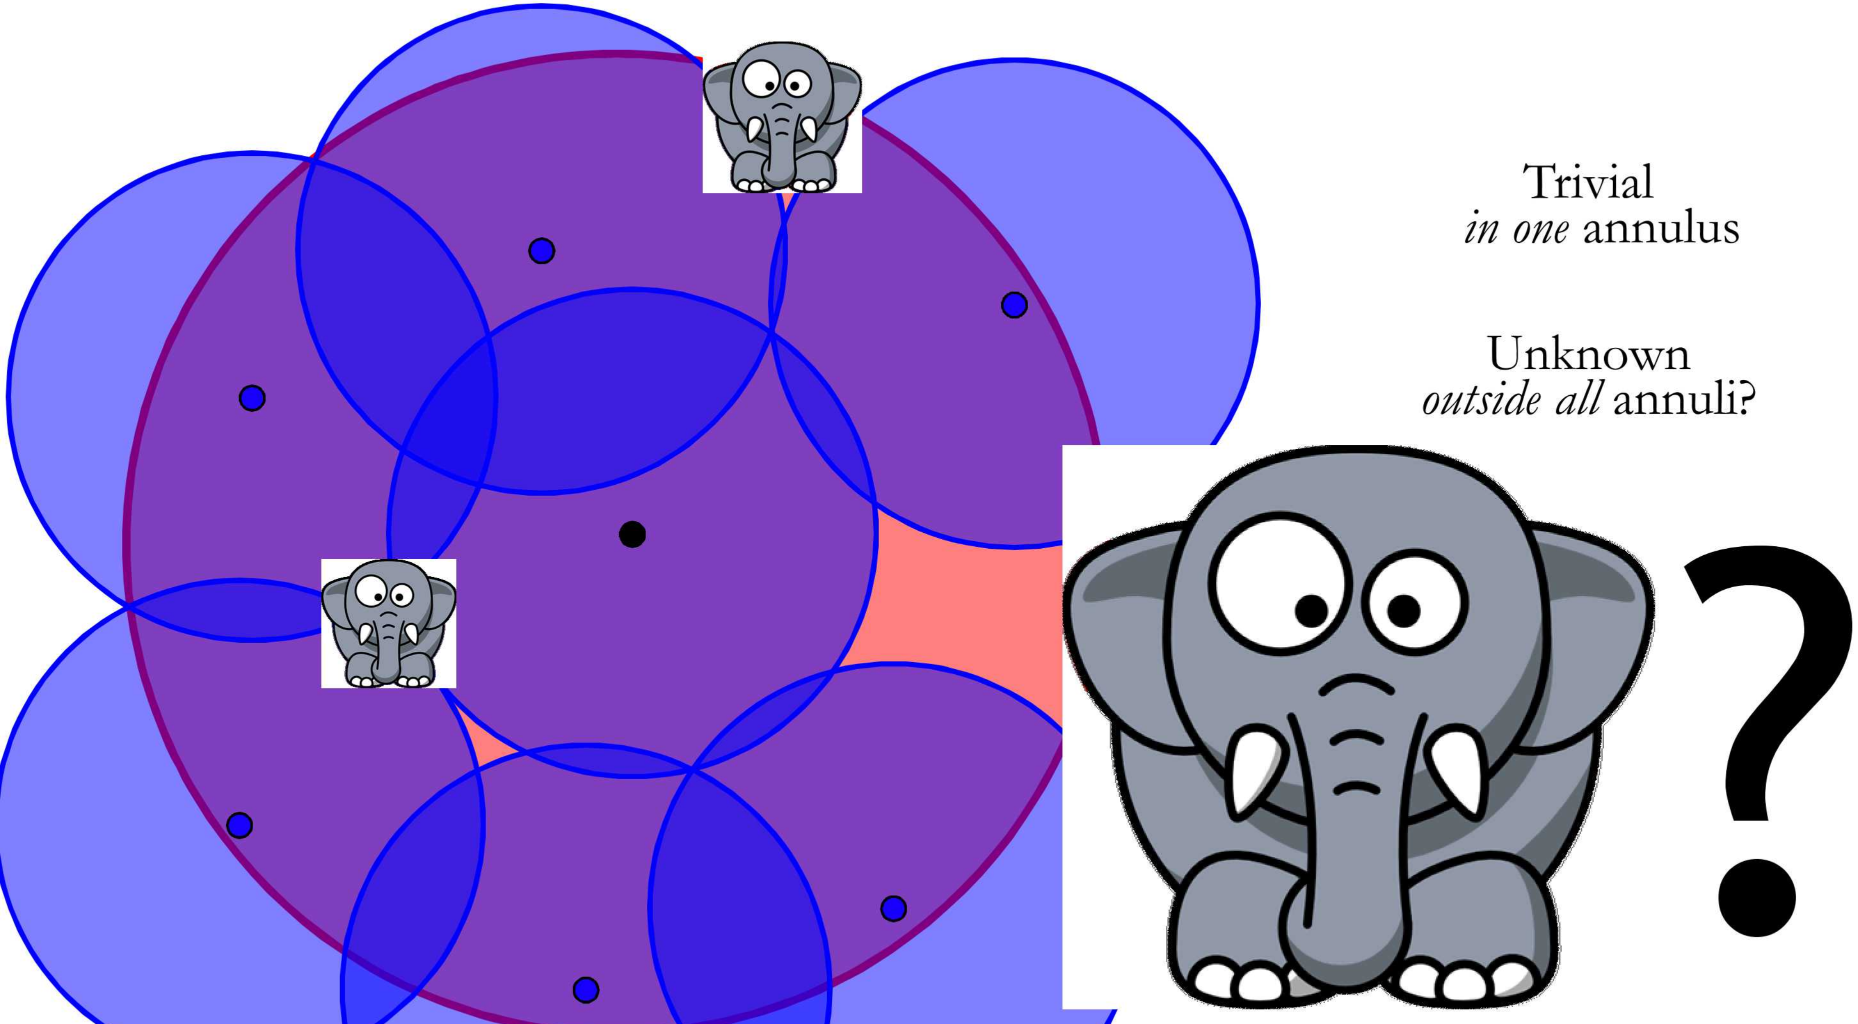

Bridson 2007, "Fast Poisson disk sampling in arbitrary dimensions." SIGGRAPH '07 sketch.

Proposed Spoke-Darts algorithm

Spoke-Darts scales

$$\text{runtime \& memory} = O(dn^2)$$

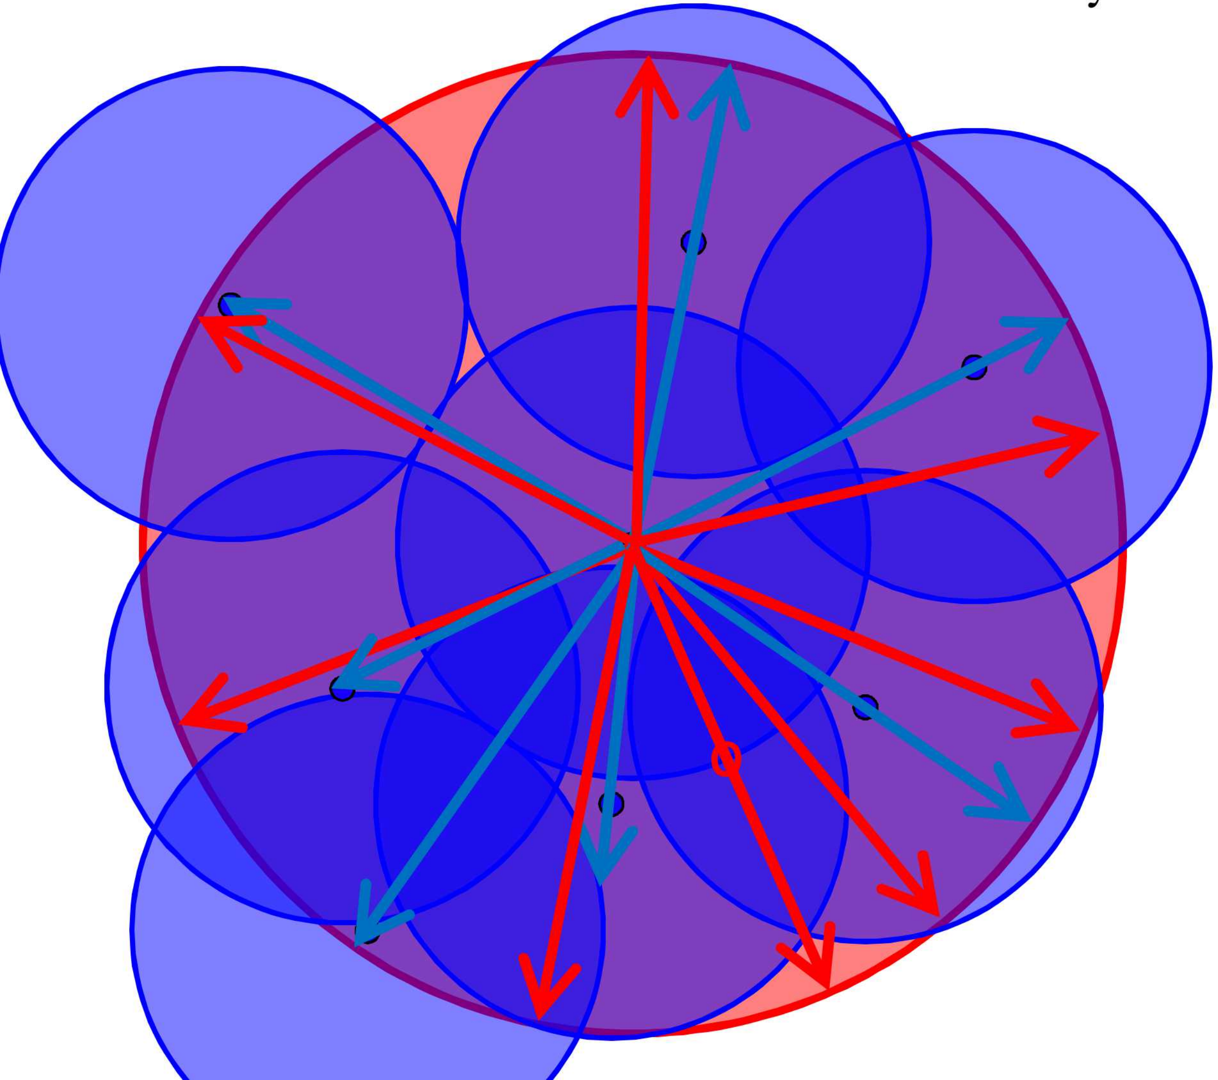

Spoke-Darts mod from Bridson

```
do
  prior = randomly pop Front
  do
    pick sample from radial line through
      (r,2r) annulus of prior
    if distance(sample, all samples) > r
      accept sample
      add sample to Front
  until 12 consecutive rejections
until Front is empty
```

12 same for all dimensions

Spoke-Darts for High-Dimensional Blue-Noise Sampling, TOG 72:2, SIGGRAPH 2018.

Proposed Spoke-Darts algorithm

Spoke-Darts scales … and guarantees saturation ☺

runtime & memory $= O(dn^2)$

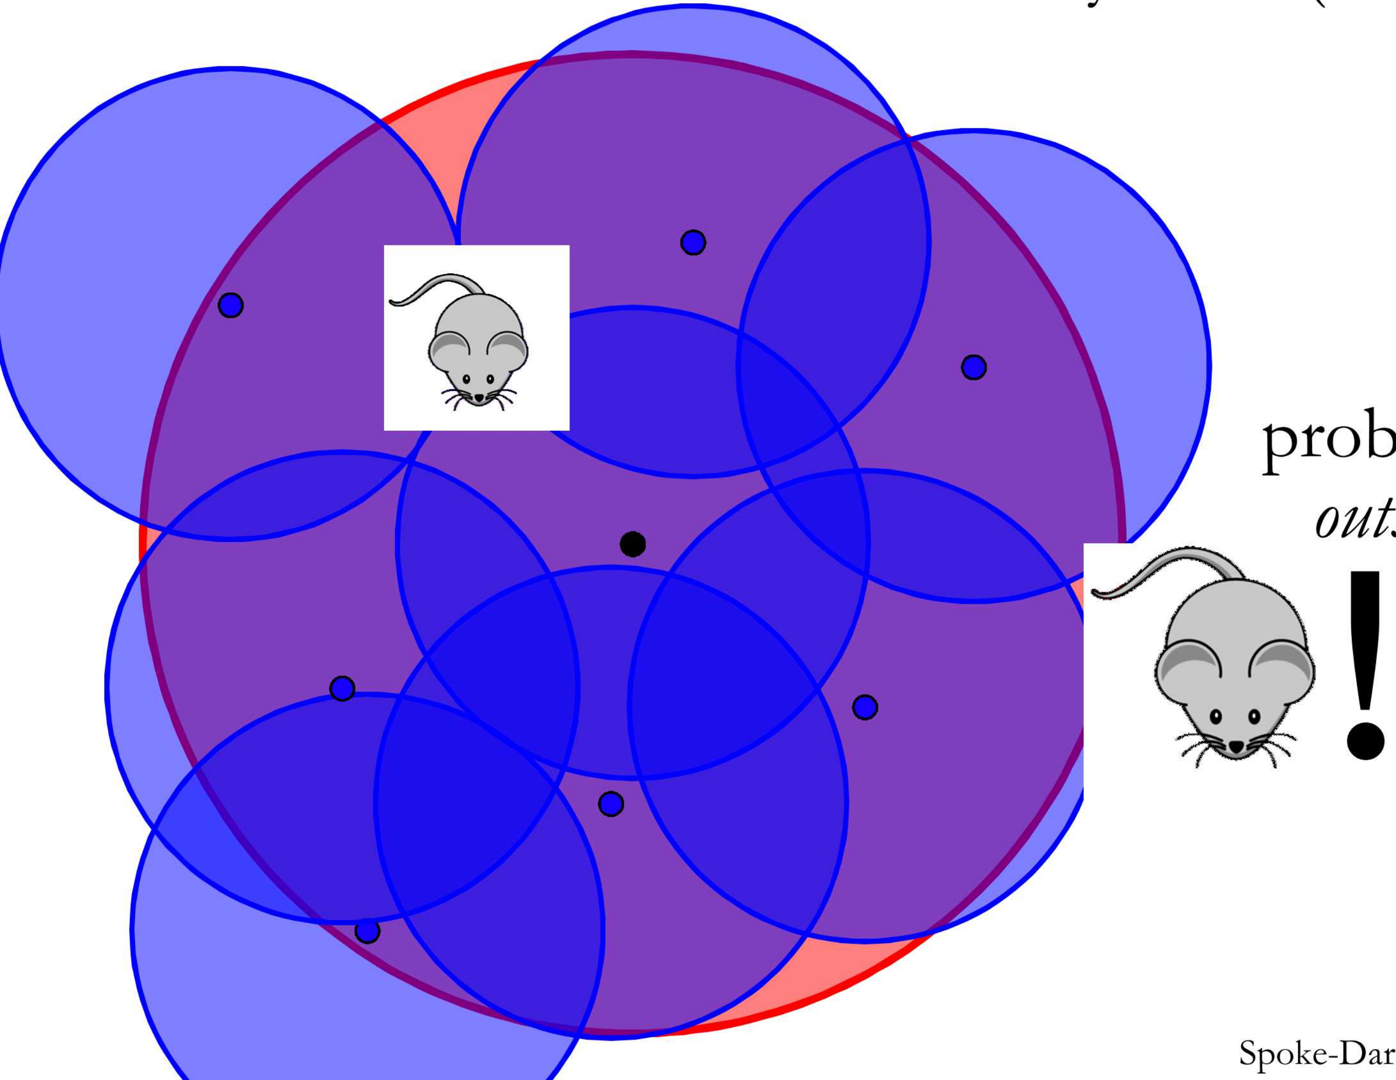

Proved probabilistic bound *outside all* annuli!

Proof of saturation

Spoke-Darts scales … and guarantees saturation

Suppose there exists a void

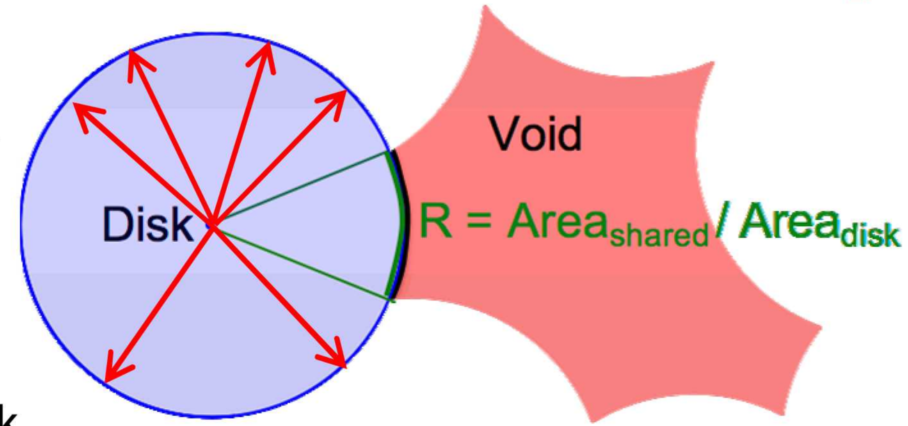

One disk

Missed 12 times, so area R is probably small,

Probability of 12 misses = $\Pi(1-R) = (1-R)^{12} < e^{(-12R)}$

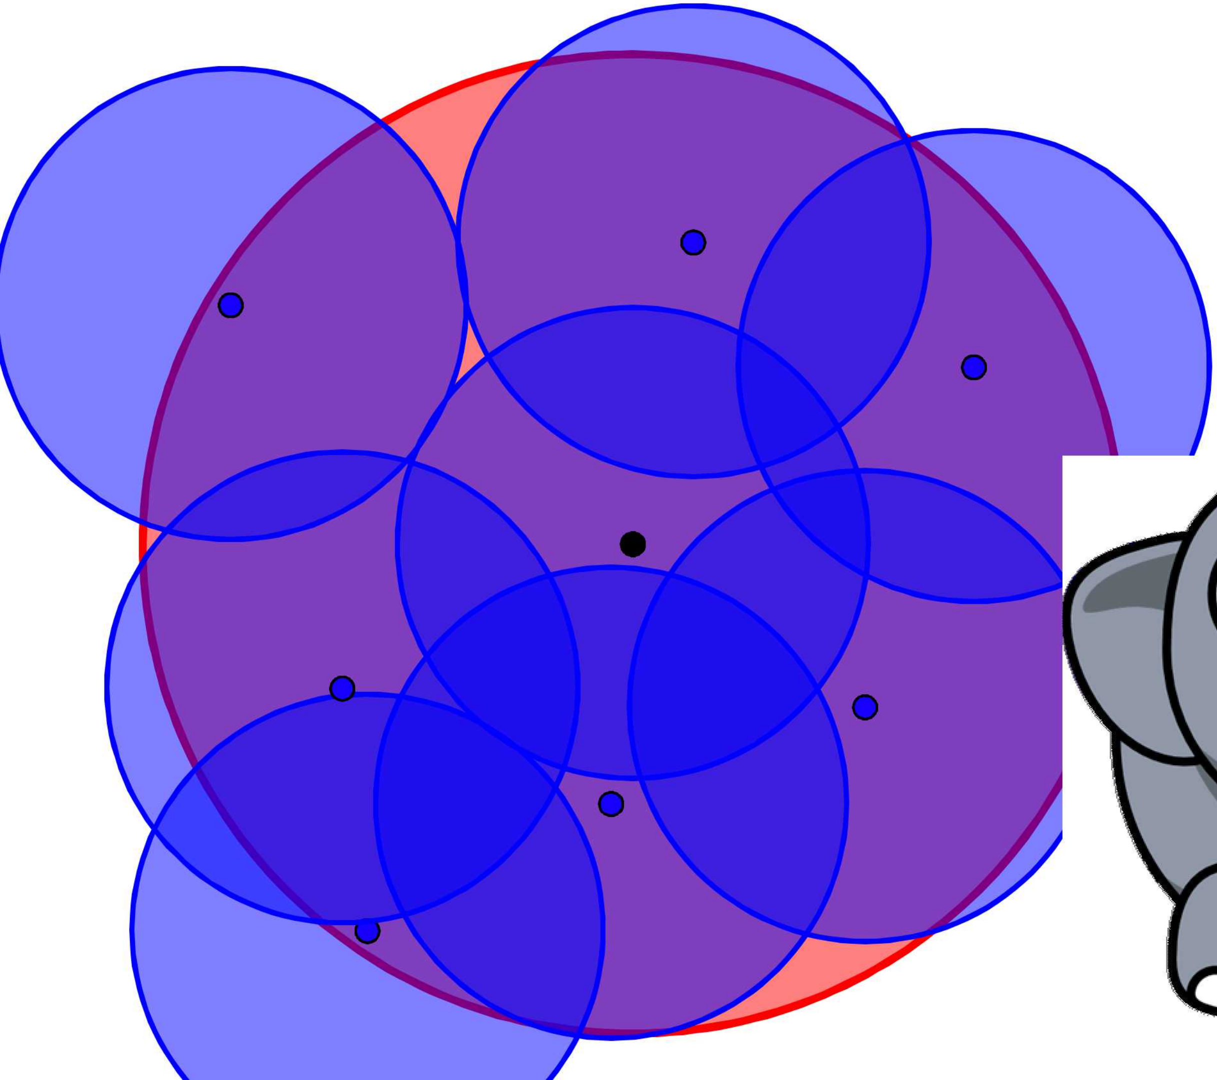

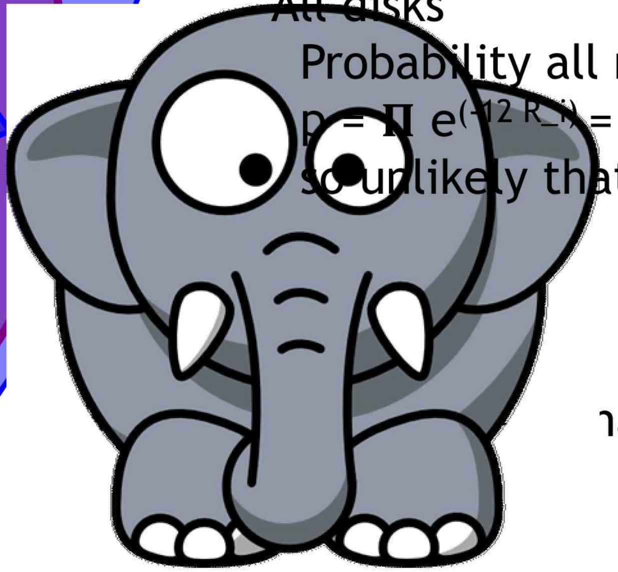

All disks

Probability all missed = product of probabilities

$p = \Pi\, e^{(-12\, R_i)} = e^{-12\, \Sigma\, R_i}$

so unlikely that the total surface area is large

nall void radius

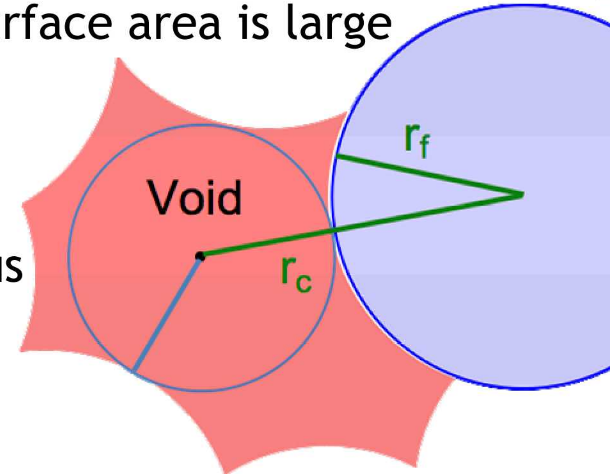

Structure of saturation guarantee

$$m = \lceil (-\ln \epsilon)(\beta^* - 1)^{1-d} \rceil$$

$$\Leftrightarrow$$

$$\beta^* = 1 + \left(\frac{-\ln \epsilon}{m} \right)^{1/(d-1)}$$

where

$$\beta^* = r_c / r_f$$

$\epsilon =$ chance beta exceeded

$m =$ #misses, failed darts

$d =$ dimension

pick any three

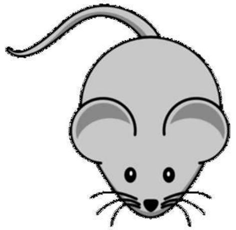

Magic values for dimensional independence

$$\frac{-\ln \epsilon}{m} = 1 \implies \beta^* = 2 \; \forall d$$

e.g.

$$m = 12$$
$$\implies \epsilon = 10^{-5}$$

probability $1 - 10^{-5}$
that $\beta^* < 2 \; \forall d$

e.g. $m = 14 \implies \epsilon = 10^{-6}$
$m = 30 \implies \epsilon = 10^{-13}$

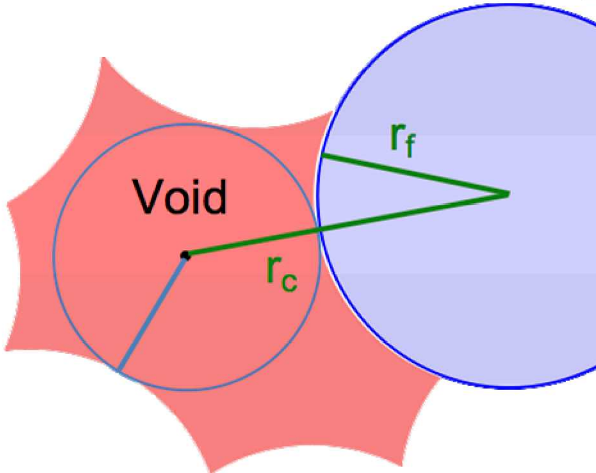

$\beta^* =$

Output Saturation in Practice = better than Bridson

one dimension "for free"

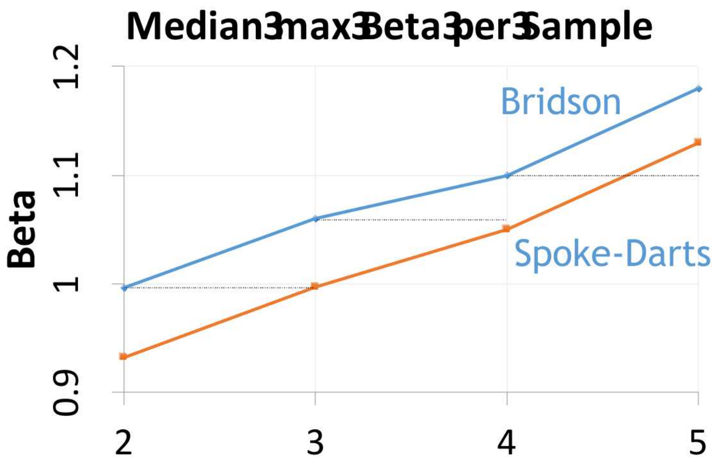

Theory: Beta = constant

Practice: Beta increases and narrows with d

- Narrowing consistent with theory

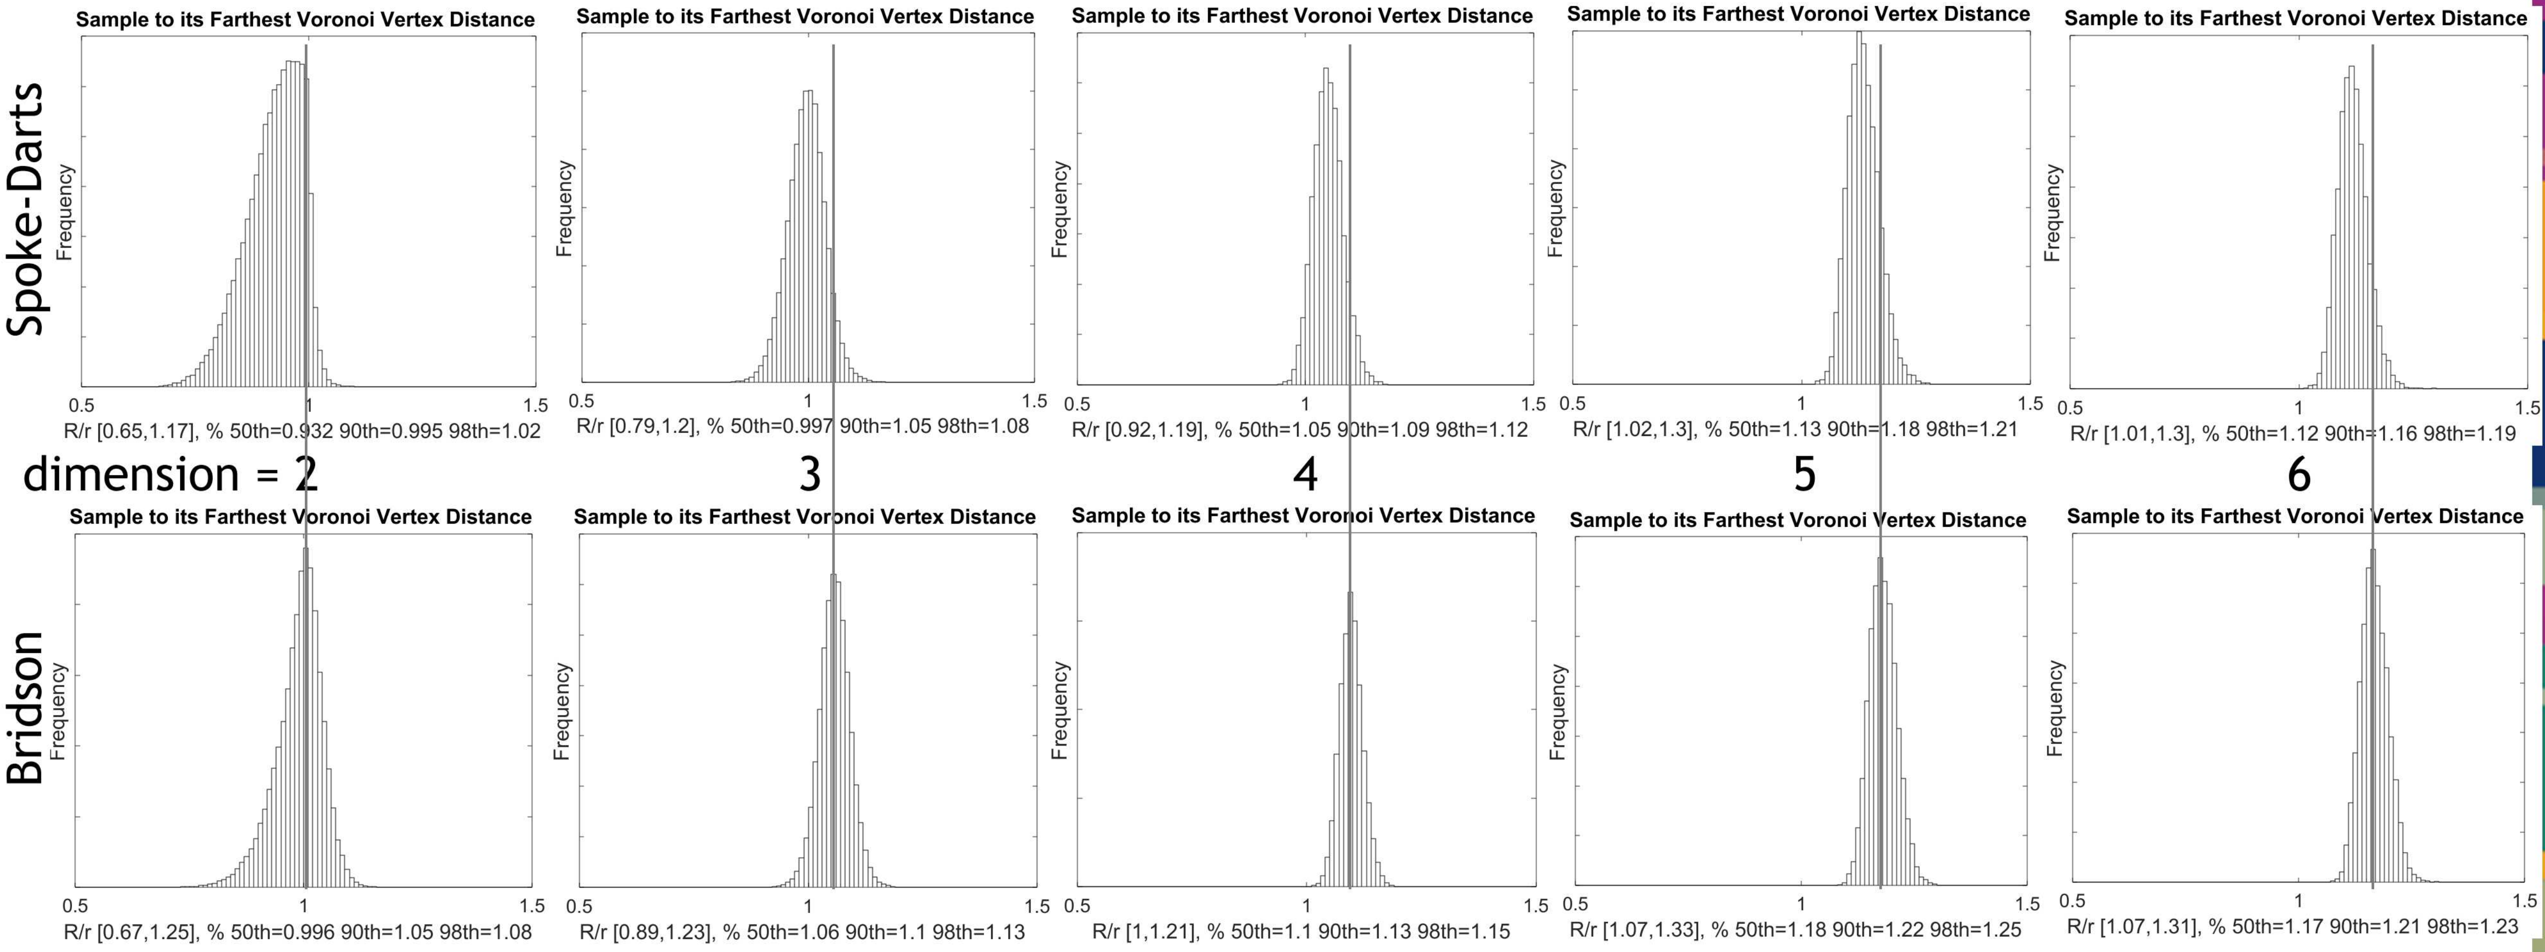

Output Randomness – Blue noise

Sharp spike at radius!
(annuli inner-radius)

d=2

Pairwise distances between samples

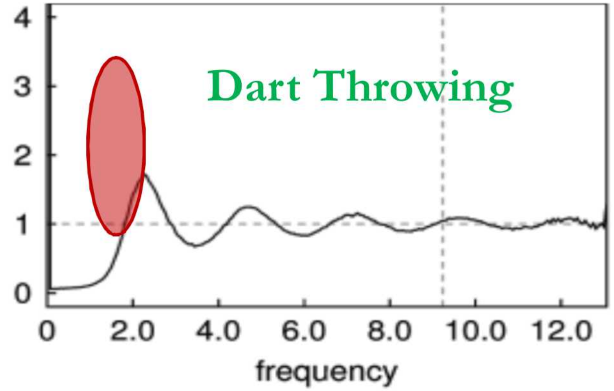

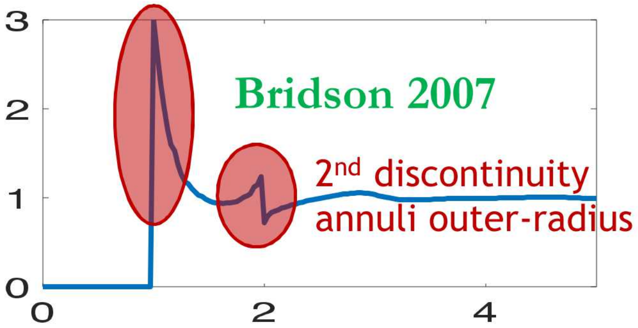

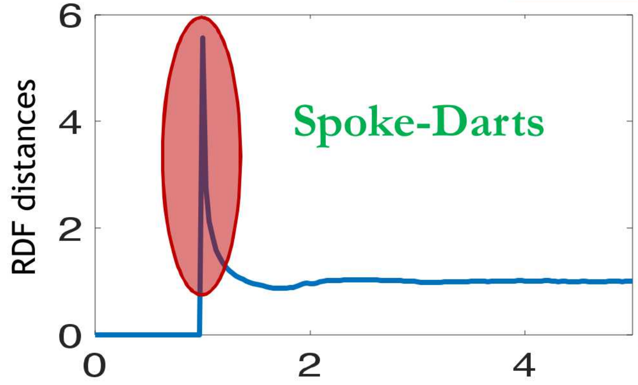

Fourier Transform, 1d
integral over angle by radius

Dart throwing images from
Blue Noise Sampling
with Controlled Aliasing,
TOG 32:3, 2013.

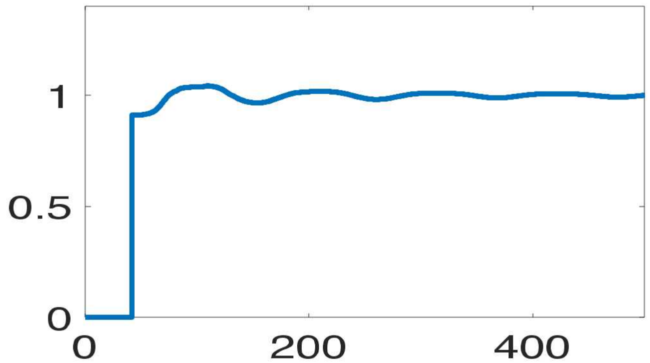

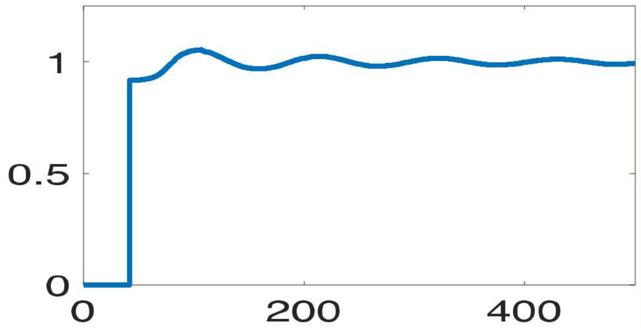

Fourier Transform

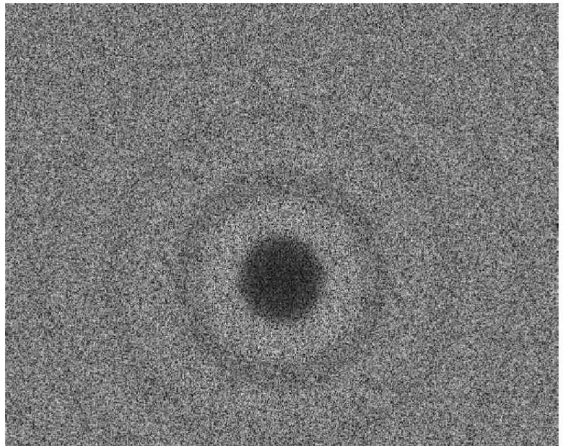

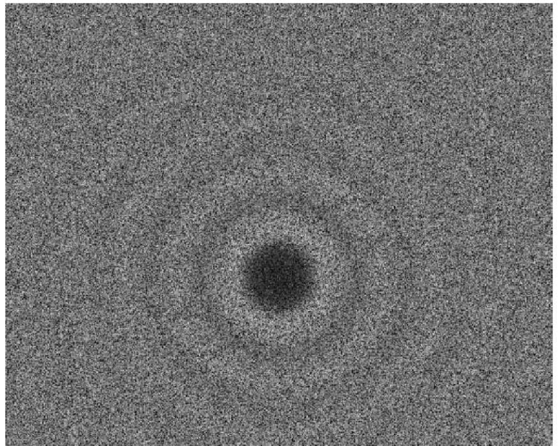

Output Randomness – Blue noise

Eliminate spike

- Post-processing optimization* = expensive, 2d
- Sample non-uniformly
  - By spoke-length? Uniform by swept volume? Worse!
  - Any sharp local rule → global discontinuity ☹

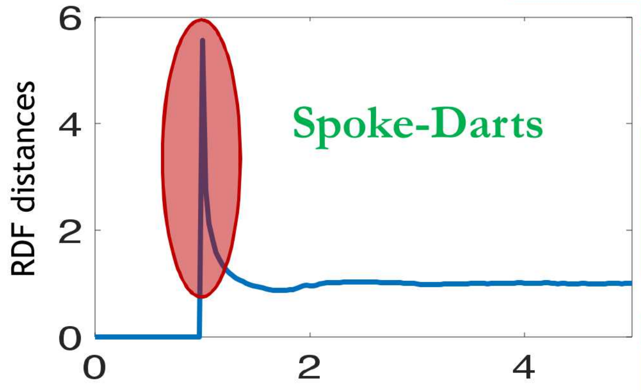

Good idea

- r_{cover} max domain-point-to-sample distance
- r_{free} min sample-to-sample distance

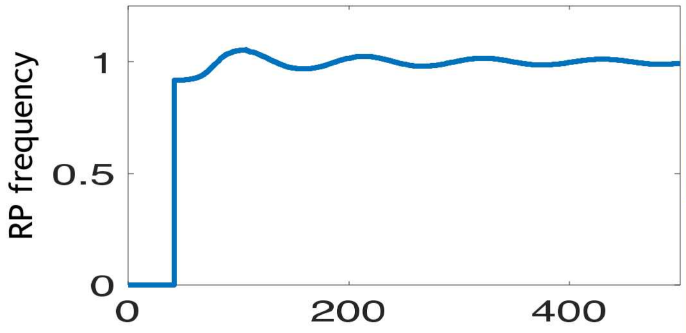

Bad idea

- $r_{free} = r_{cover} \iff \boldsymbol{B}=1$
- Traditional goal of Maximal Poisson-disk Sampling!
  - Serves no purpose

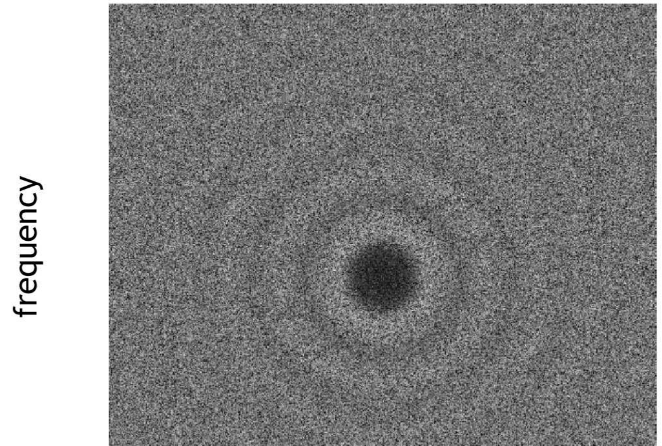

* Blue Noise Sampling with Controlled Aliasing, TOG 32:3, 2013.

Output Randomness – Blue noise

Eliminate spike – we solved in 2012 ☺

- † Two-radii : $r_{free} \neq r_{cover}$
  - Small blue r_{free} minimum sample separation
  - Large green r_{cover} maximum domain-to-sample
    - Unique coverage: accept only if covers white

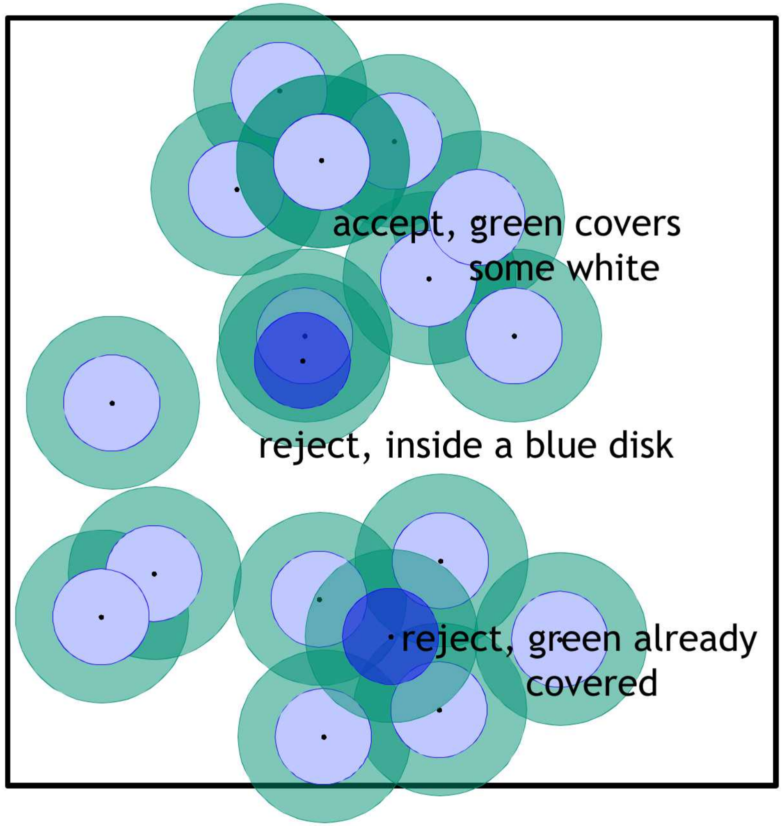

- **Two-spokes** mimic Two-radii
  - 1st spoke, trim by large green
  - 2nd spoke, trim by small blue

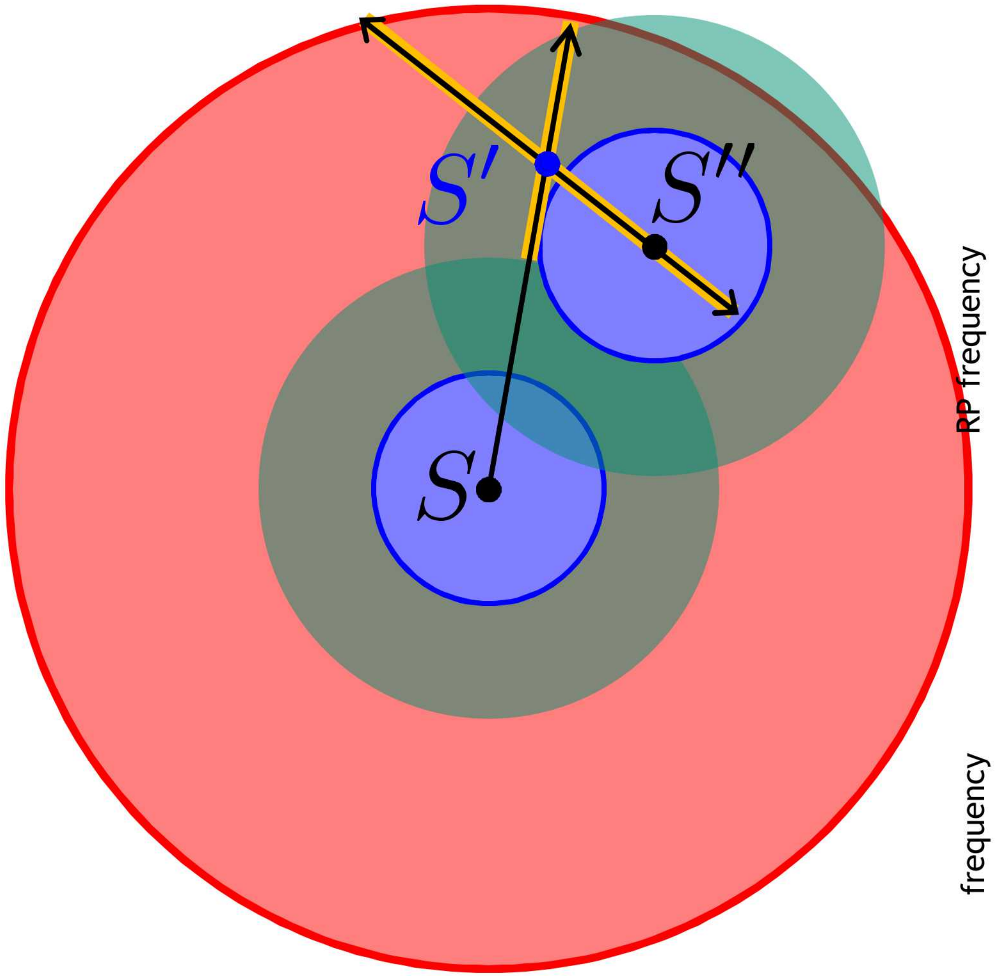

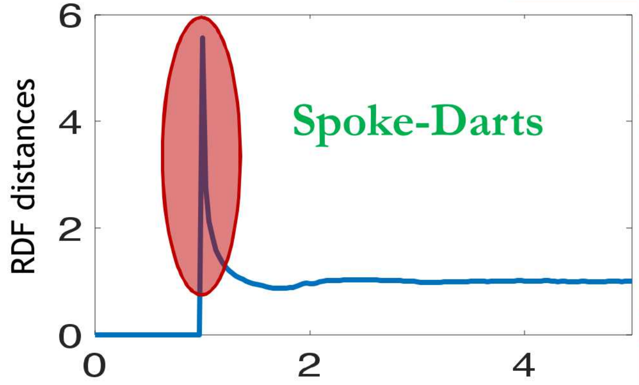

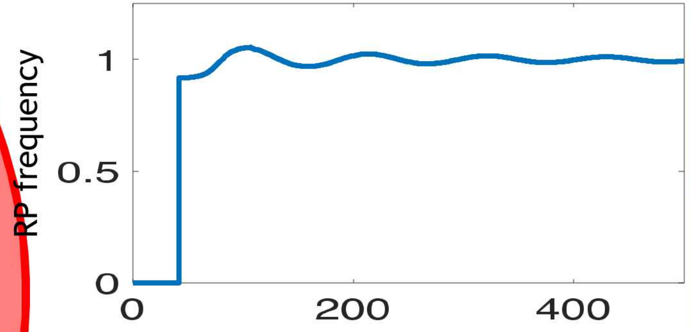

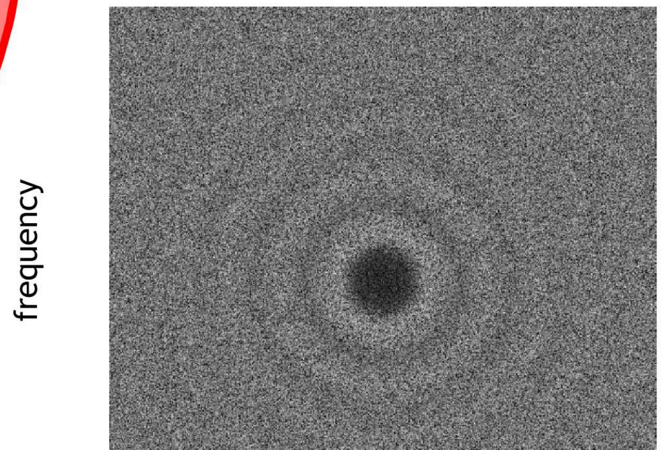

† Variable Radii Poisson-Disk Sampling, Canadian Conf. on Comp. Geometry (CCCG) 2012.

Output Randomness – Blue noise

Eliminate spike

- † Two-radii : $r_{free} \neq r_{cover}$
  - Small blue r_{free} minimum sample separation
  - Large green r_{cover} maximum domain-to-sample
    - Unique coverage: accept only if covers white

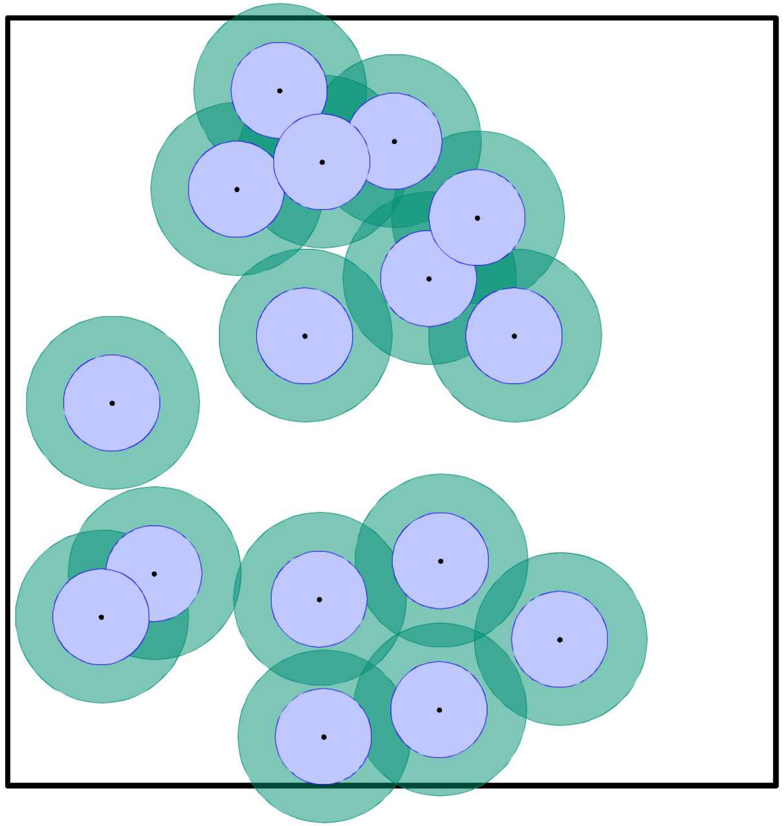

- **Two-spokes** mimic Two-radii
  - 1st spoke, trim by large green
  - 2nd spoke, trim by small blue

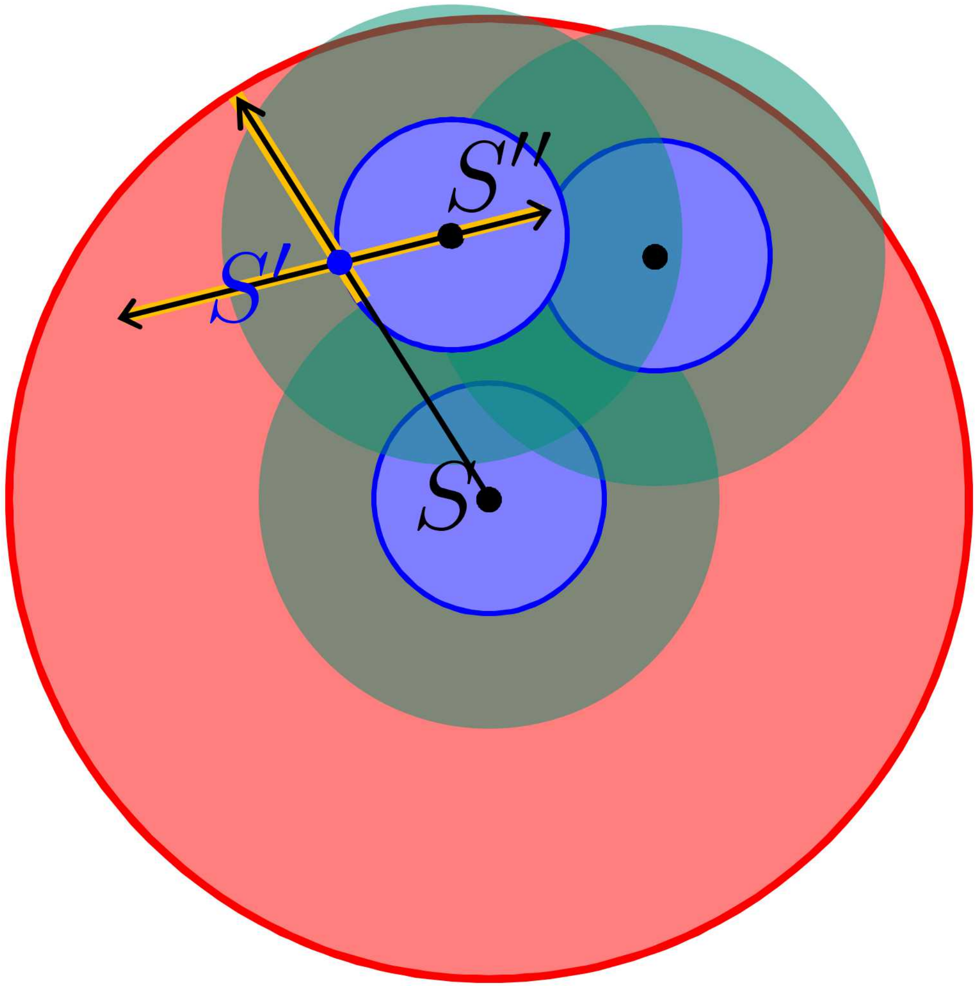

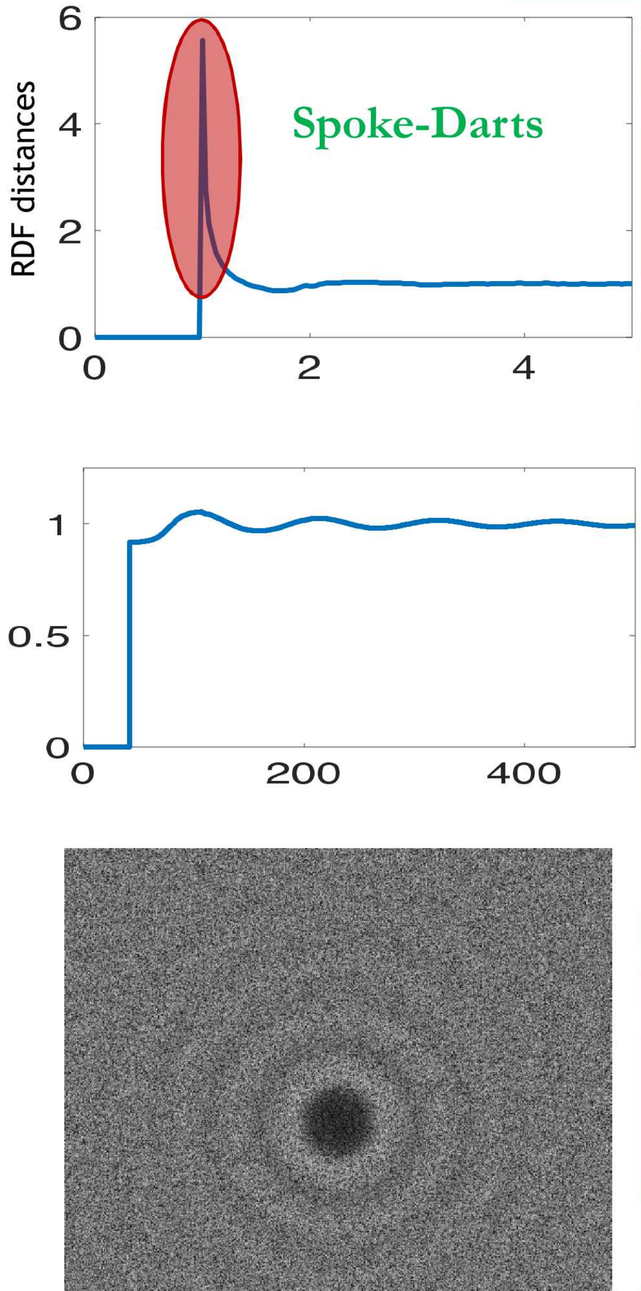

† Variable Radii Poisson-Disk Sampling, CCCG 2012.

Output Randomness – Blue noise

Eliminate spike

- † Two-radii : $r_{free} \neq r_{cover}$
  - Small blue r_{free} minimum sample separation
  - Large green r_{cover} maximum domain-to-sample
    - Unique coverage: accept only if covers white

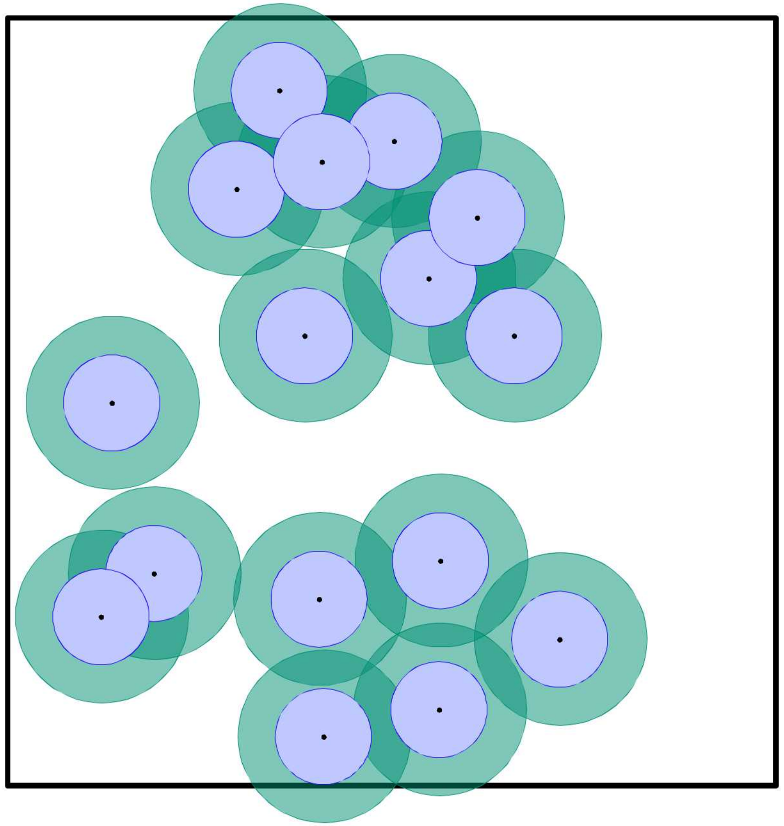

- **Two-spokes** mimic Two-radii
  - 1st spoke, trim by large green
  - 2nd spoke, trim by small blue

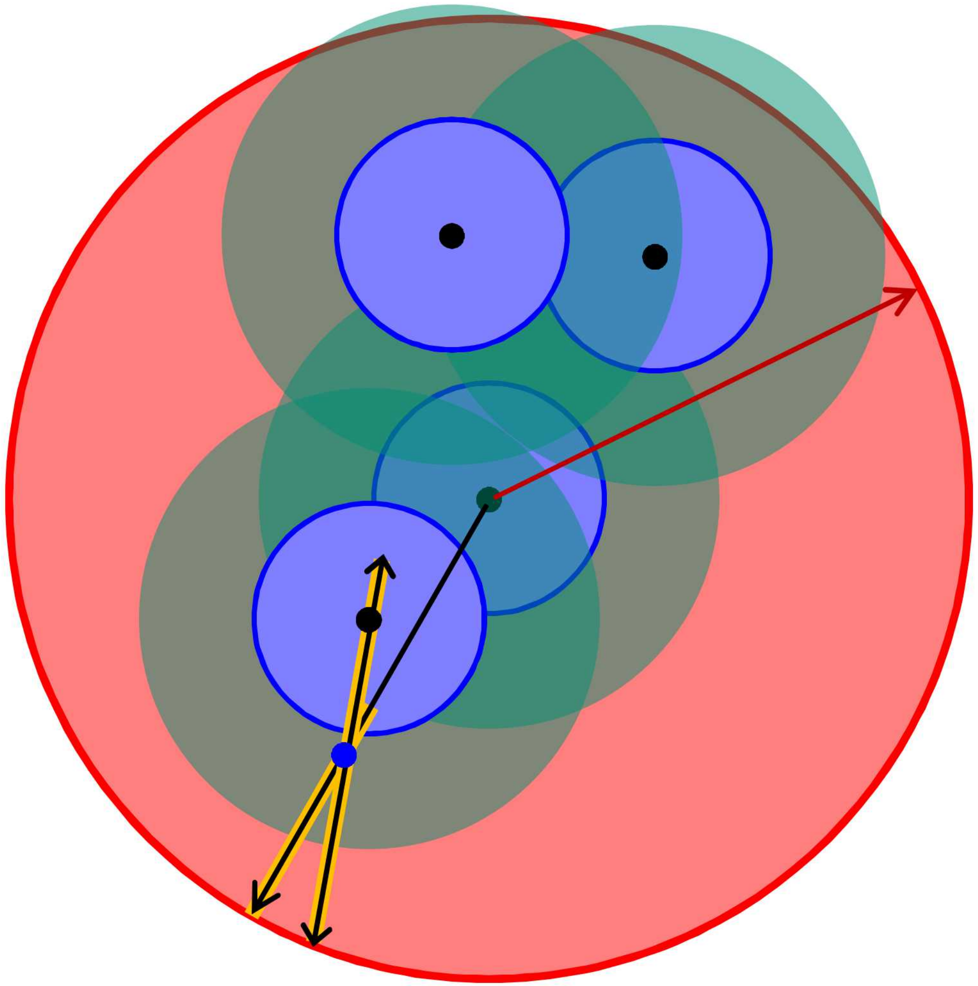

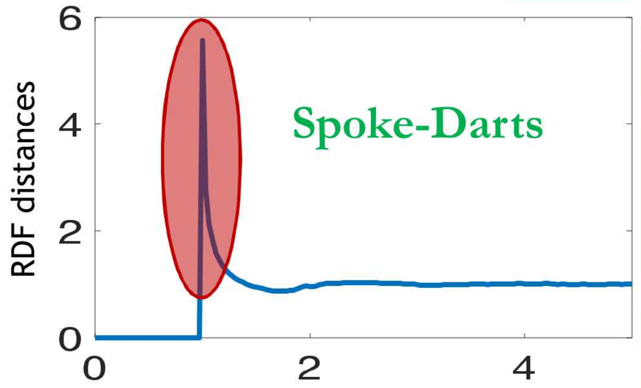

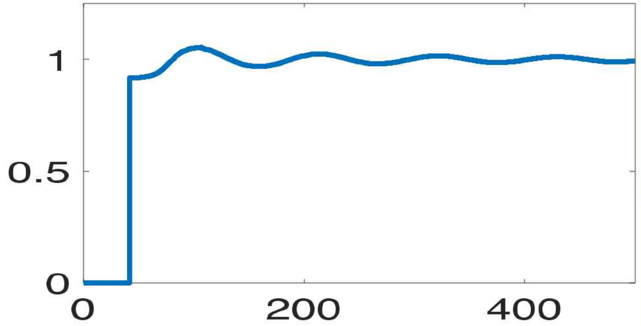

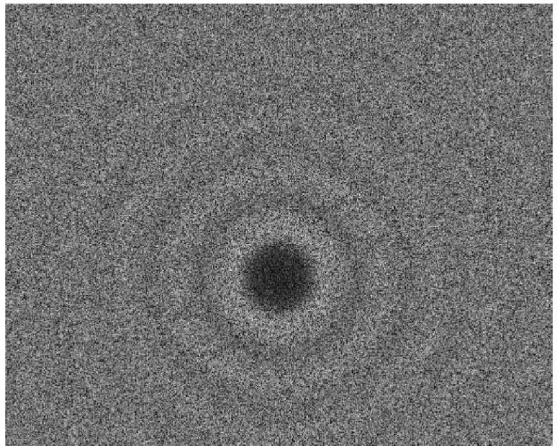

† Variable Radii Poisson-Disk Sampling, CCCG 2012.

Output Randomness – Blue noise

Eliminate spike

- † Two-radii : $r_{free} \neq r_{cover}$
  - Small blue r_{free} minimum sample separation
  - Large green r_{cover} maximum domain-to-sample
    - Unique coverage: accept only if covers white

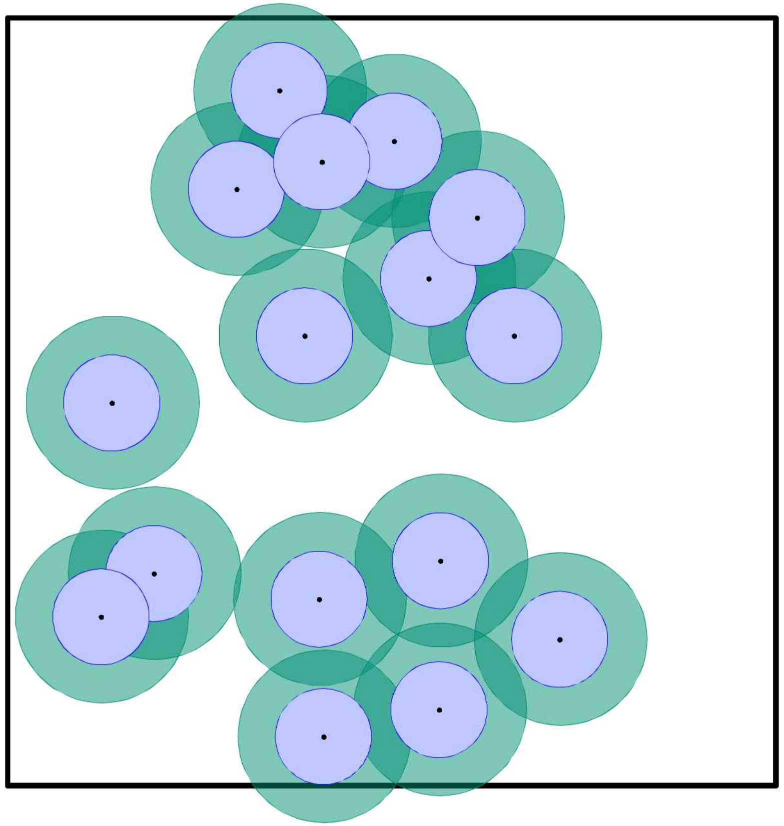

- **Two-spokes** mimic Two-radii
  - 1st spoke, trim by large green
  - 2nd spoke, trim by small blue

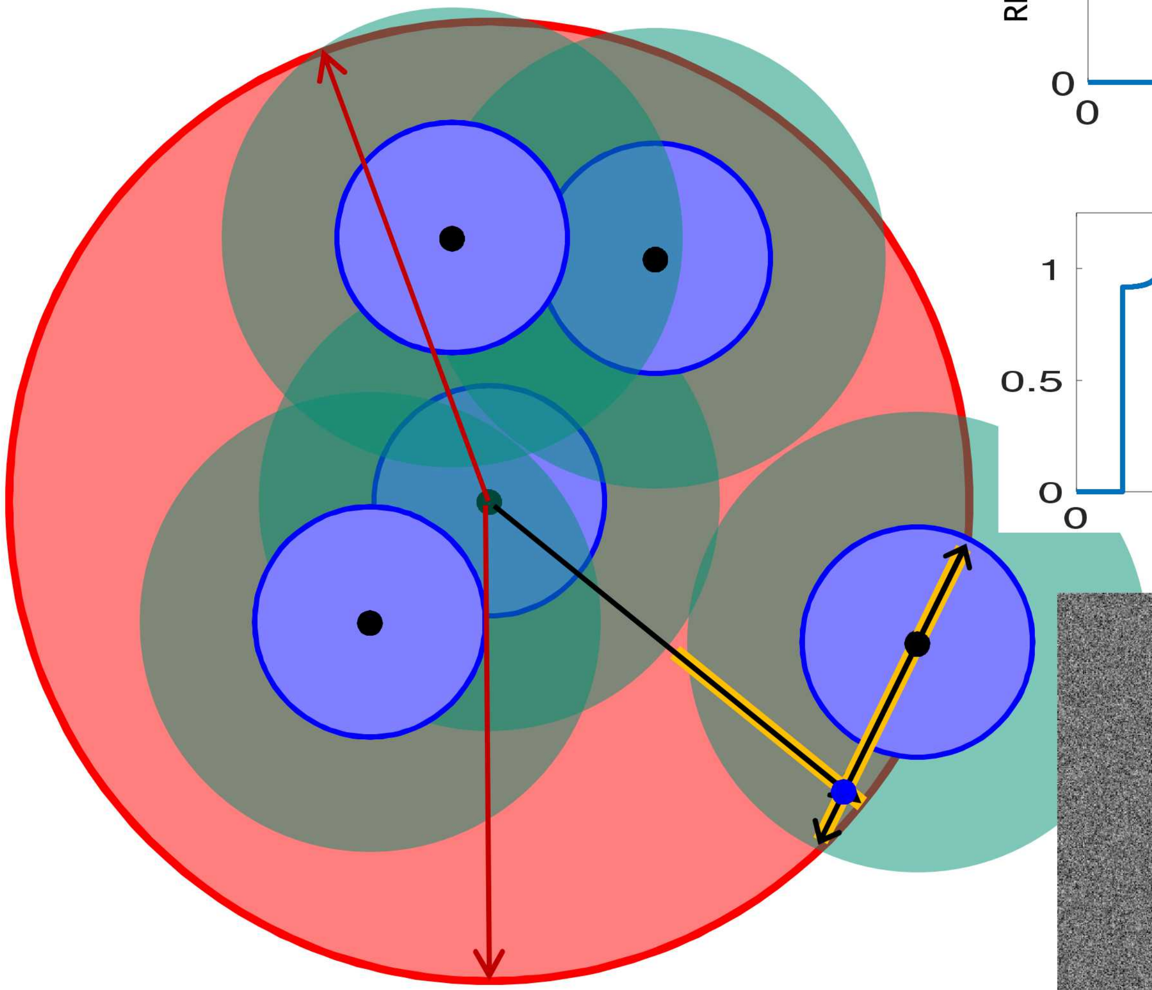

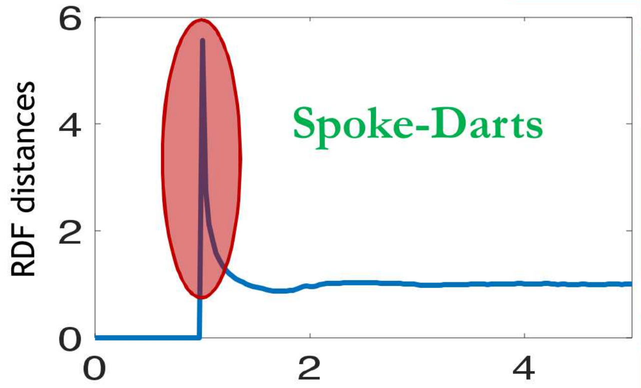

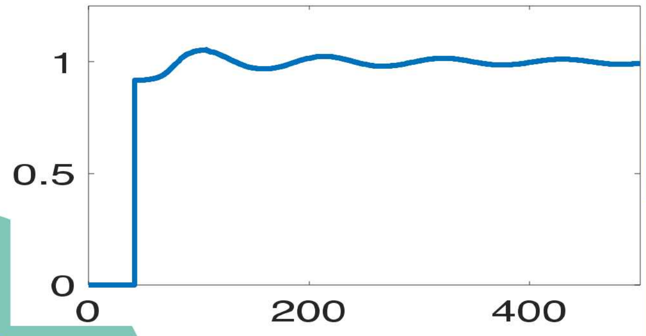

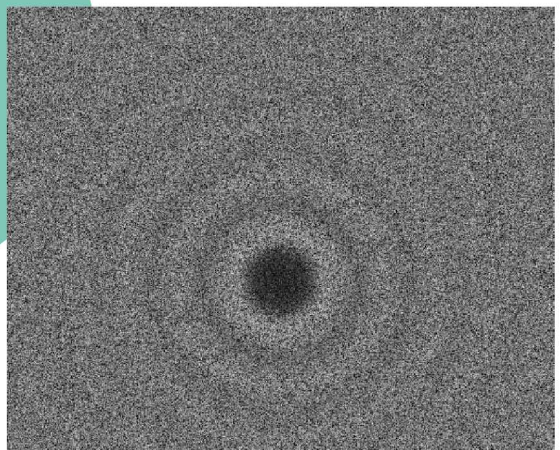

† Variable Radii Poisson-Disk Sampling, CCCG 2012.

Output Randomness – Blue noise

Eliminate spike

- †Two-radii : $r_{free} \neq r_{cover}$
  - Small blue r_{free} minimum sample separation
  - Large green r_{cover} maximum domain-to-sample
    - Unique coverage: accept only if covers white

- **Two-spokes** mimic Two-radii
  - 1st spoke, trim by large green
  - 2nd spoke, trim by small blue

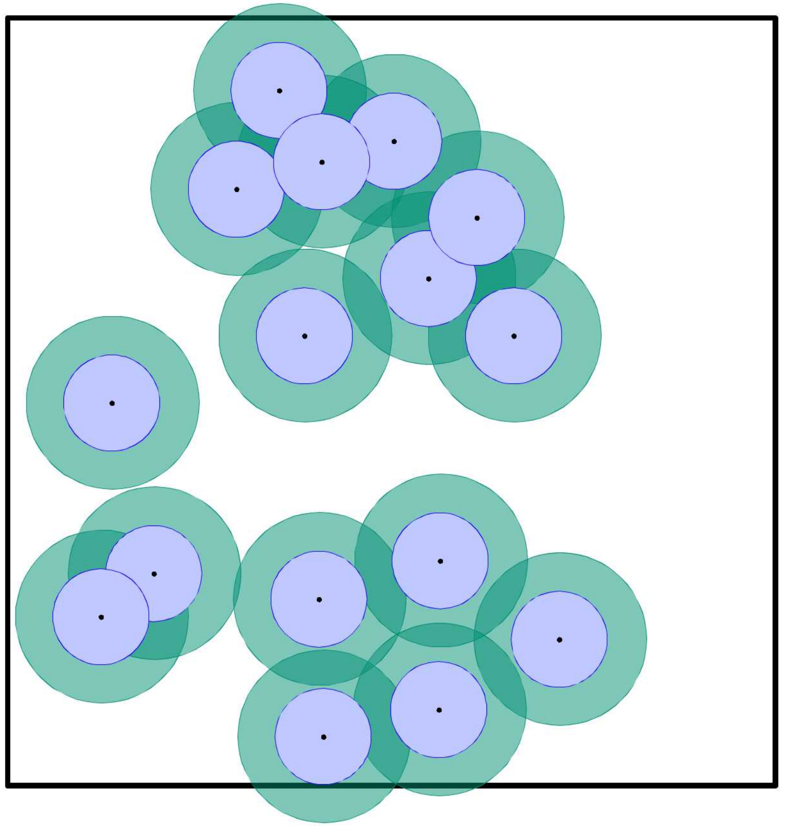

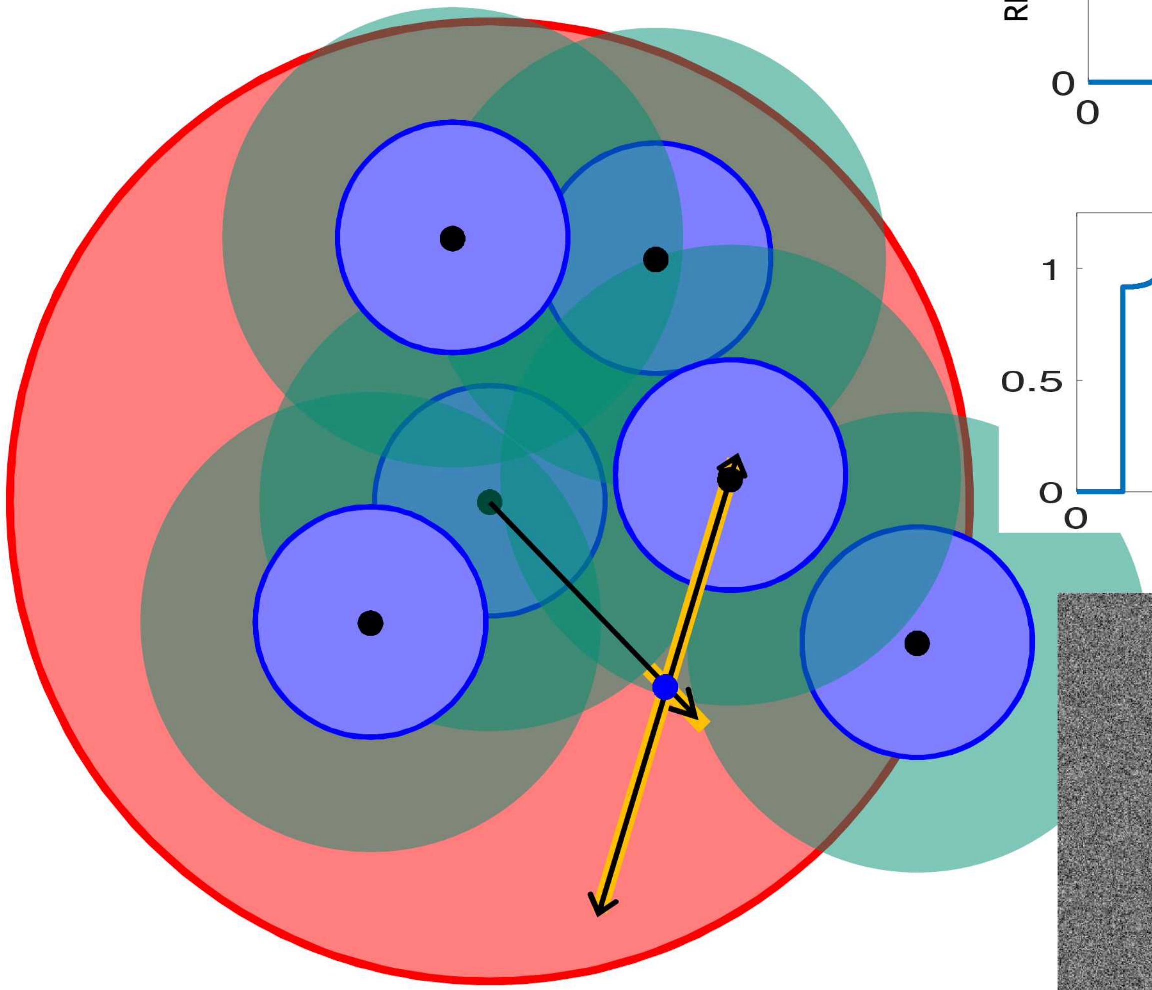

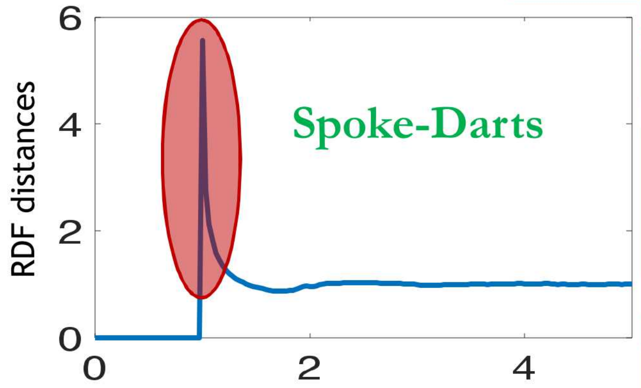

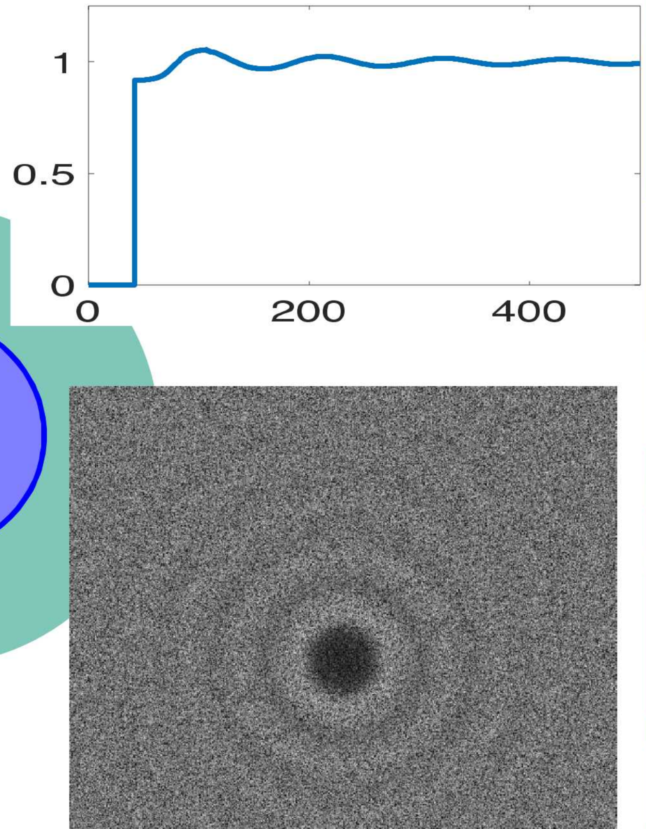

† Variable Radii Poisson-Disk Sampling, CCCG 2012.

Output Randomness – Blue noise

Eliminate spike

- † Two-radii : $r_{free} \neq r_{cover}$
  - Small blue r_{free} minimum sample separation
  - Large green r_{cover} maximum domain-to-sample
    - Unique coverage: accept only if covers white

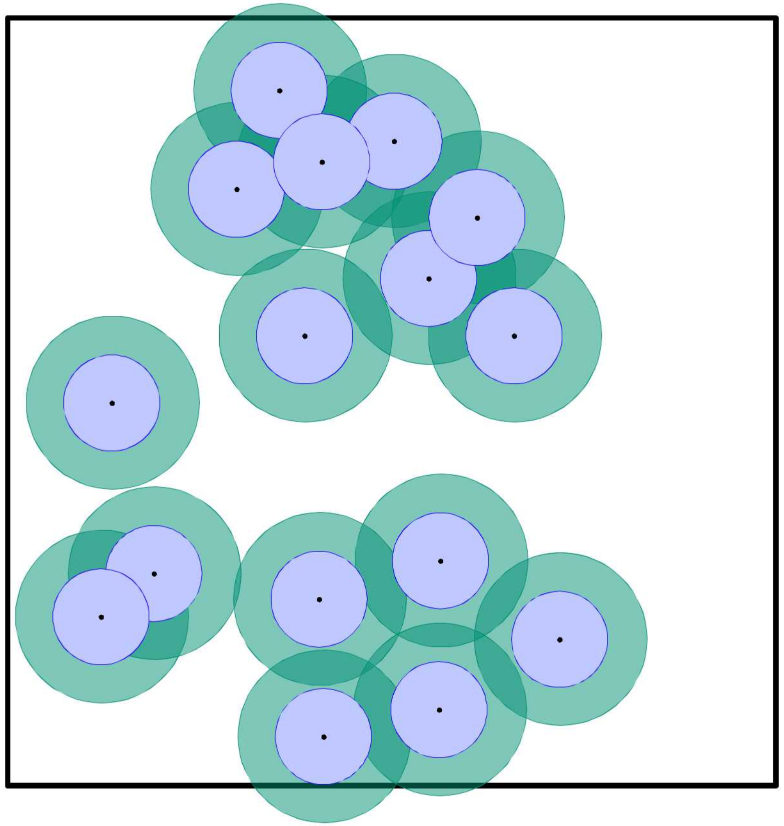

- **Two-spokes** mimic Two-radii
  - 1st spoke, trim by large green
  - 2nd spoke, trim by small blue

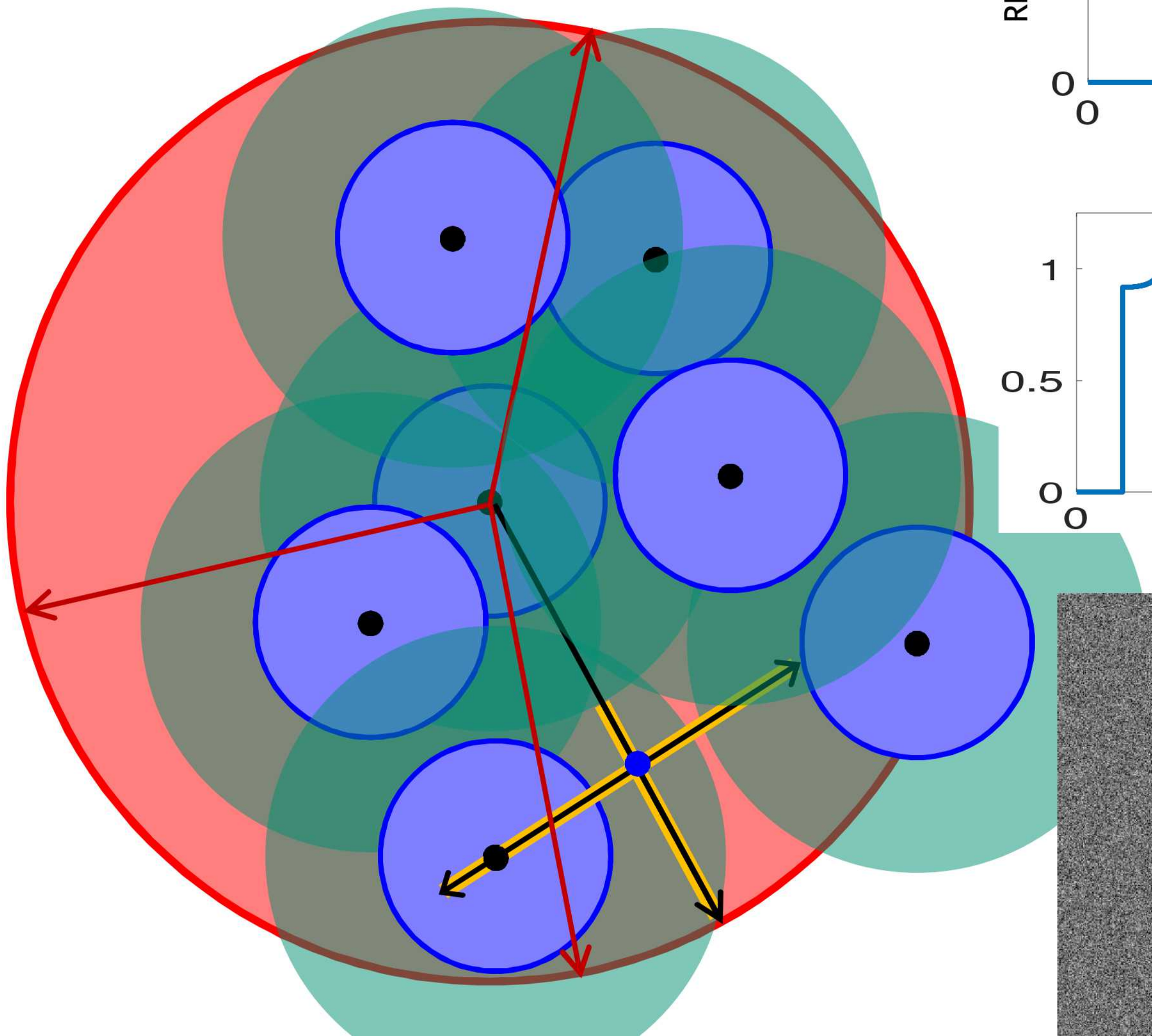

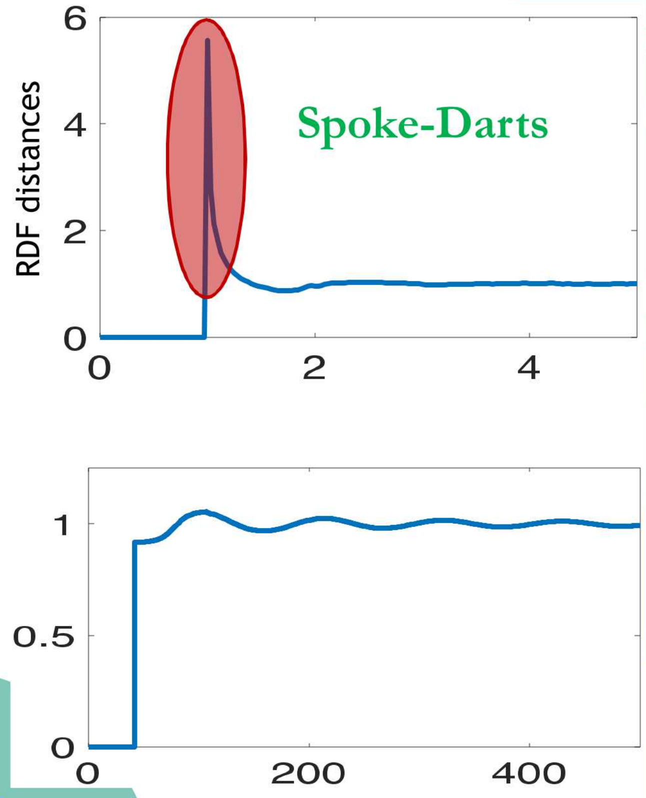

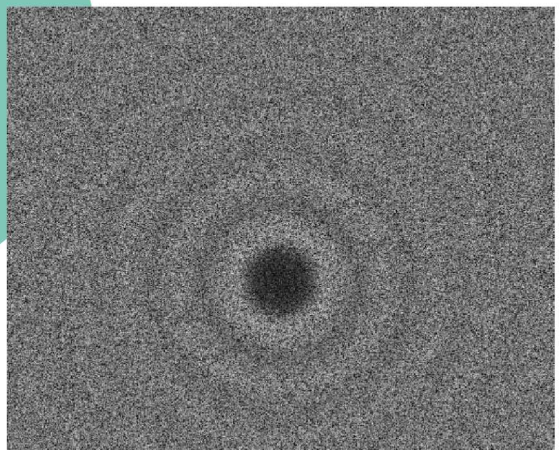

† Variable Radii Poisson-Disk Sampling, CCCG 2012.

Output Randomness – Blue noise

Eliminate spike

- † Two-radii : $r_{free} \neq r_{cover}$
  - Small blue r_{free} minimum sample separation
  - Large green r_{cover} maximum domain-to-sample
    - Unique coverage: accept only if covers white

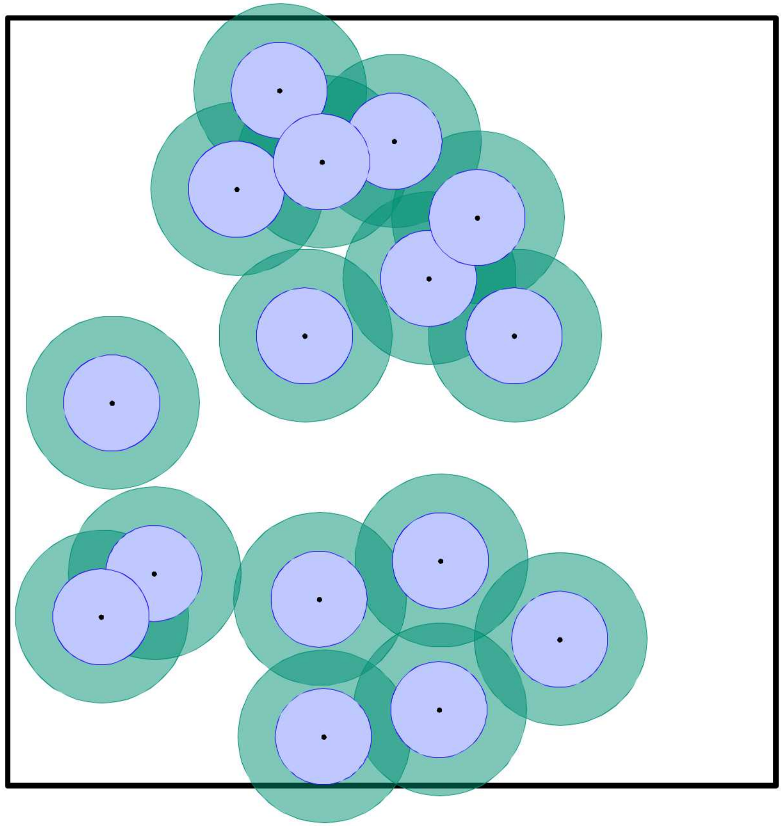

- **Two-spokes** mimic Two-radii
  - 1st spoke, trim by large green
  - 2nd spoke, trim by small blue

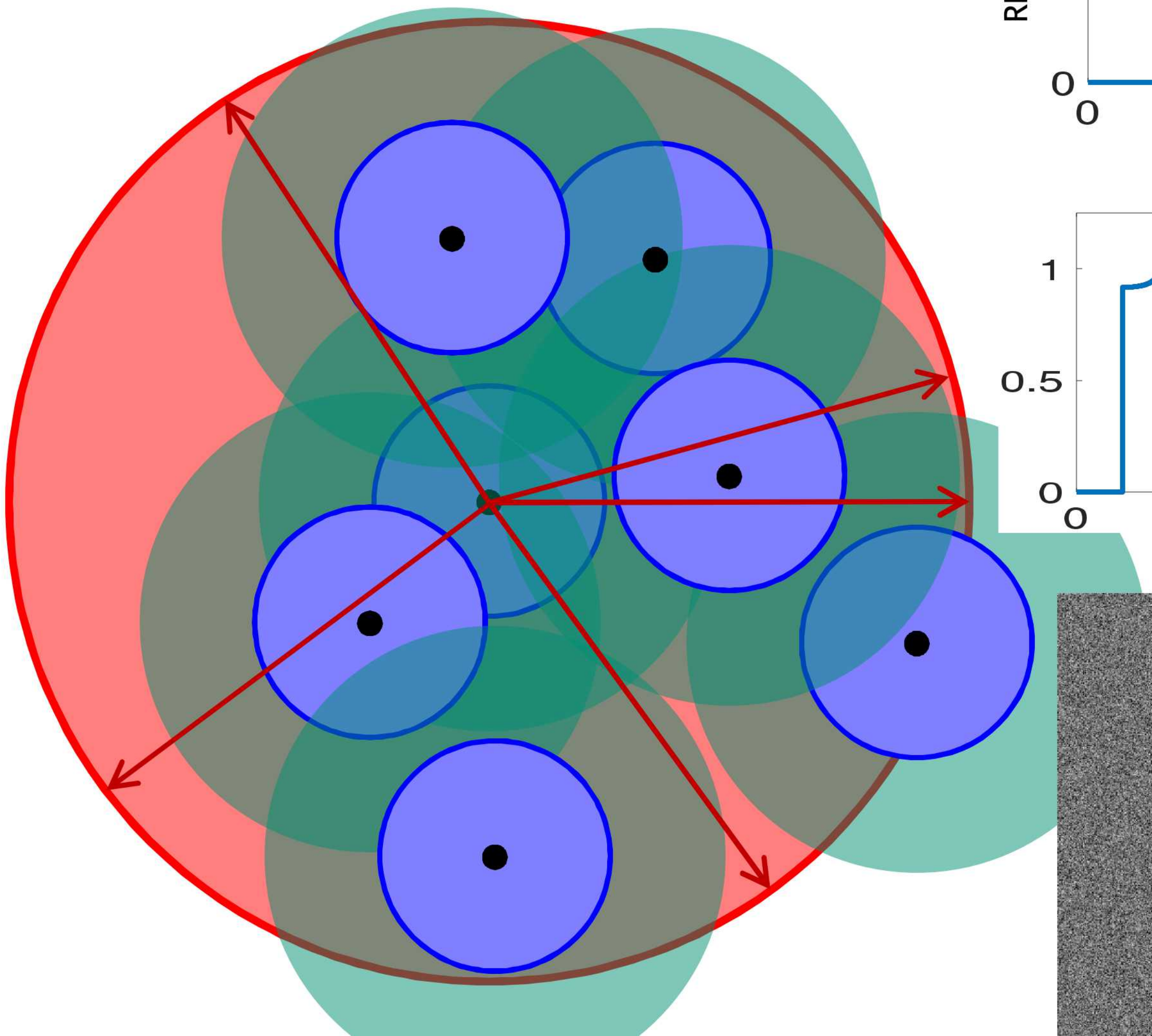

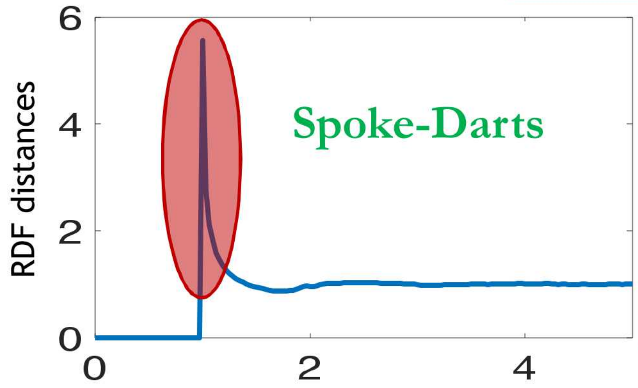

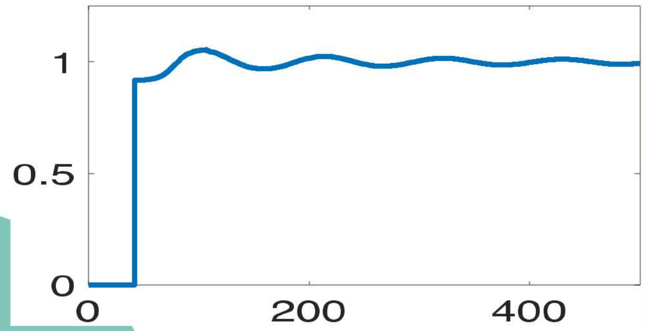

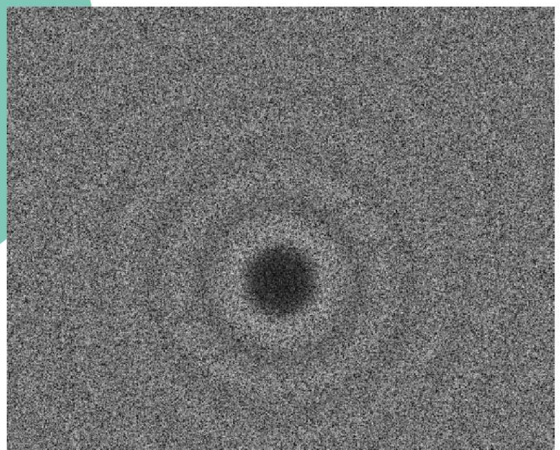

† Variable Radii Poisson-Disk Sampling, CCCG 2012.

Output Randomness – Blue noise

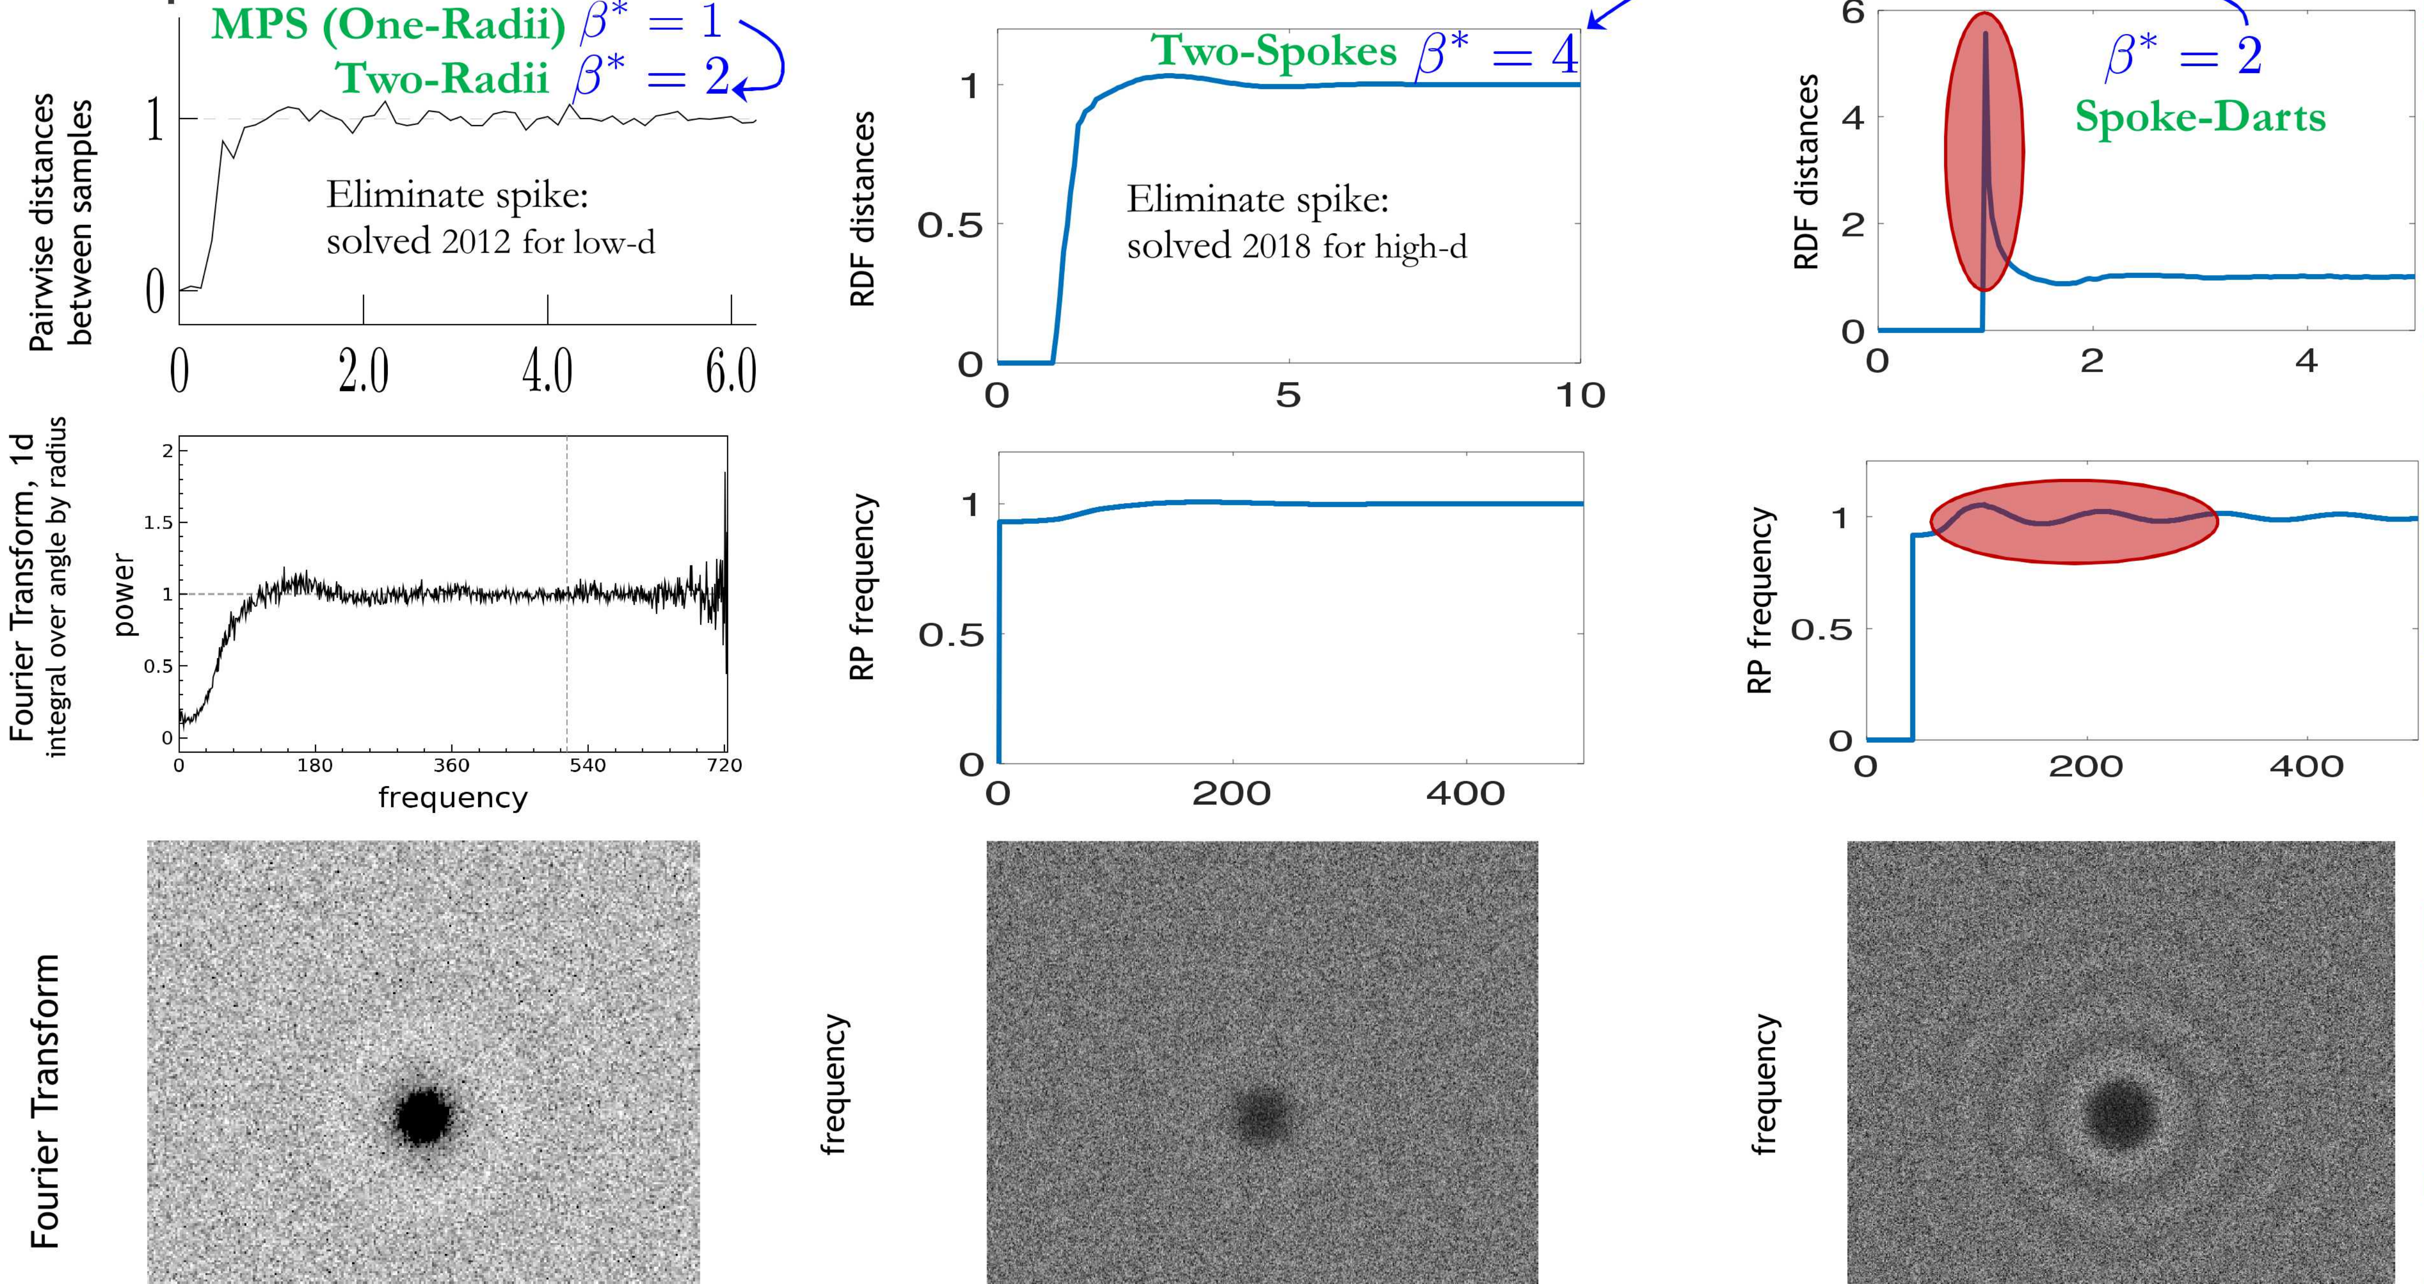

Application: Approximate Delaunay Graph

Trim line-spokes to find “significant” neighbors in Voronoi diagram

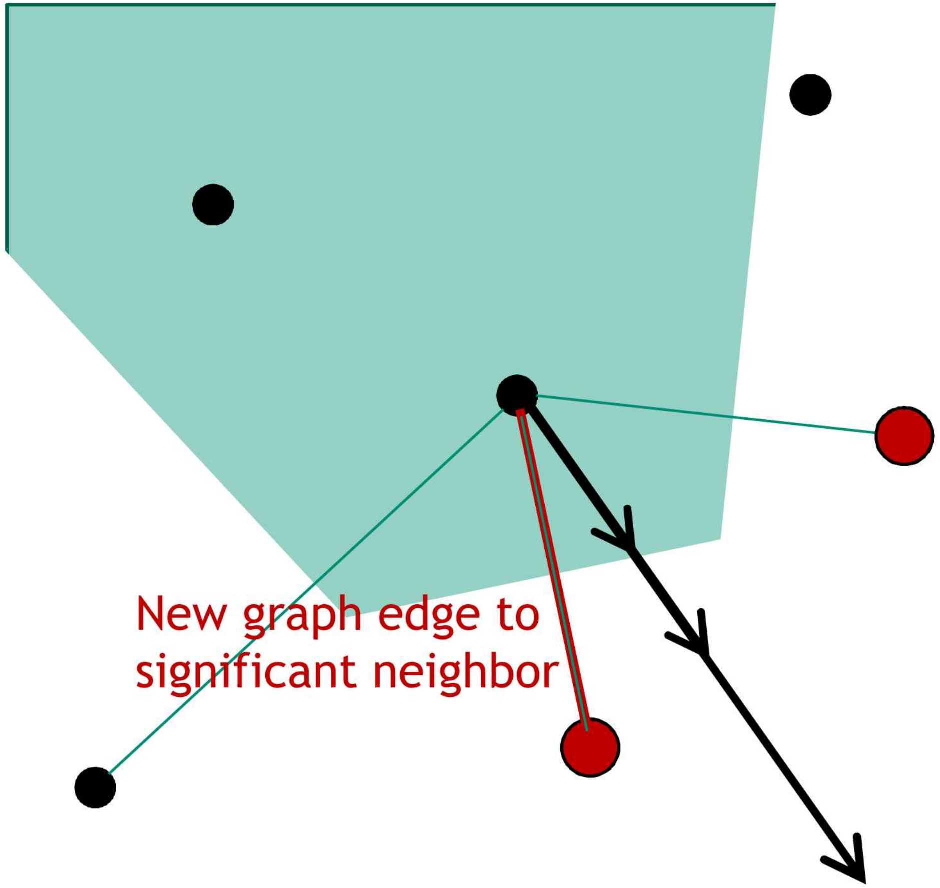

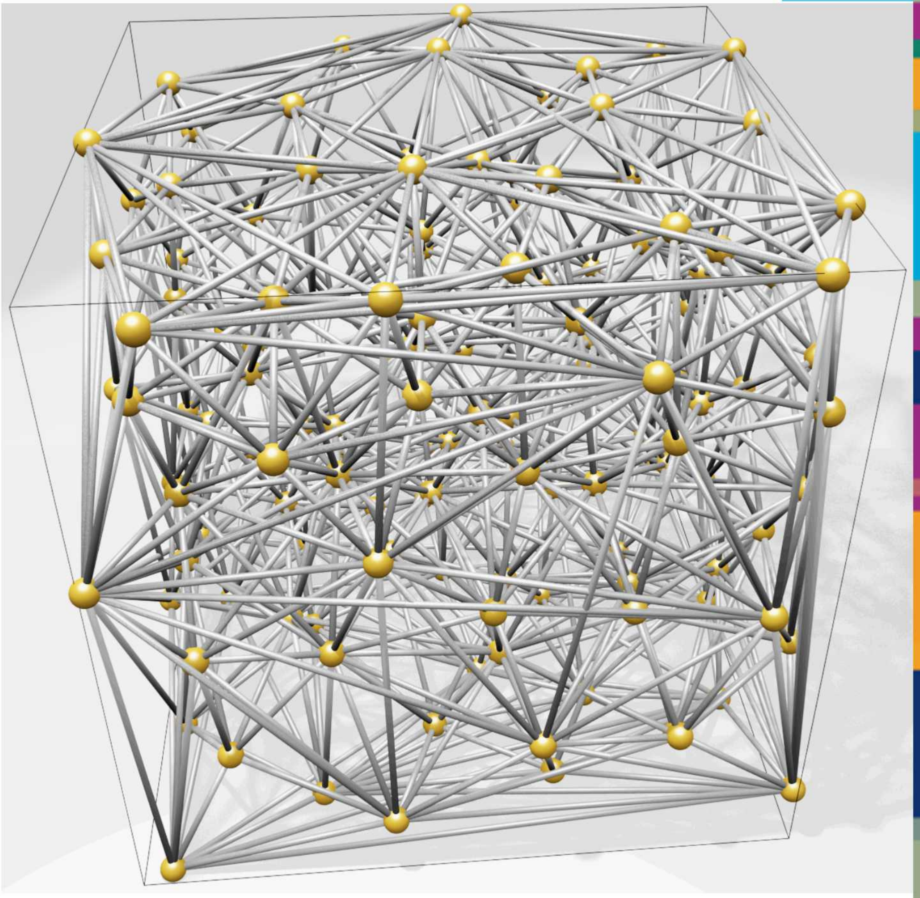

3d rendering of 8d graph

Application: Global Optimization

DIRECT = DIviding RECTangles
- Divide space by rectangles
- Refine large rectangles with small sample values

Opt-Darts
- Divide space by Voronoi cells (*implicit only*)
- Refine large cells with small sample values
  - Approx. Delaunay graph defines "large" cells
  - Spoke-darts selects refinement sample

Opt-Darts: 5-25× speedup over DIRECT
- Increasing speedup with d
  - Rectangle bad approximation of Voronoi cell
- More-random sample patterns

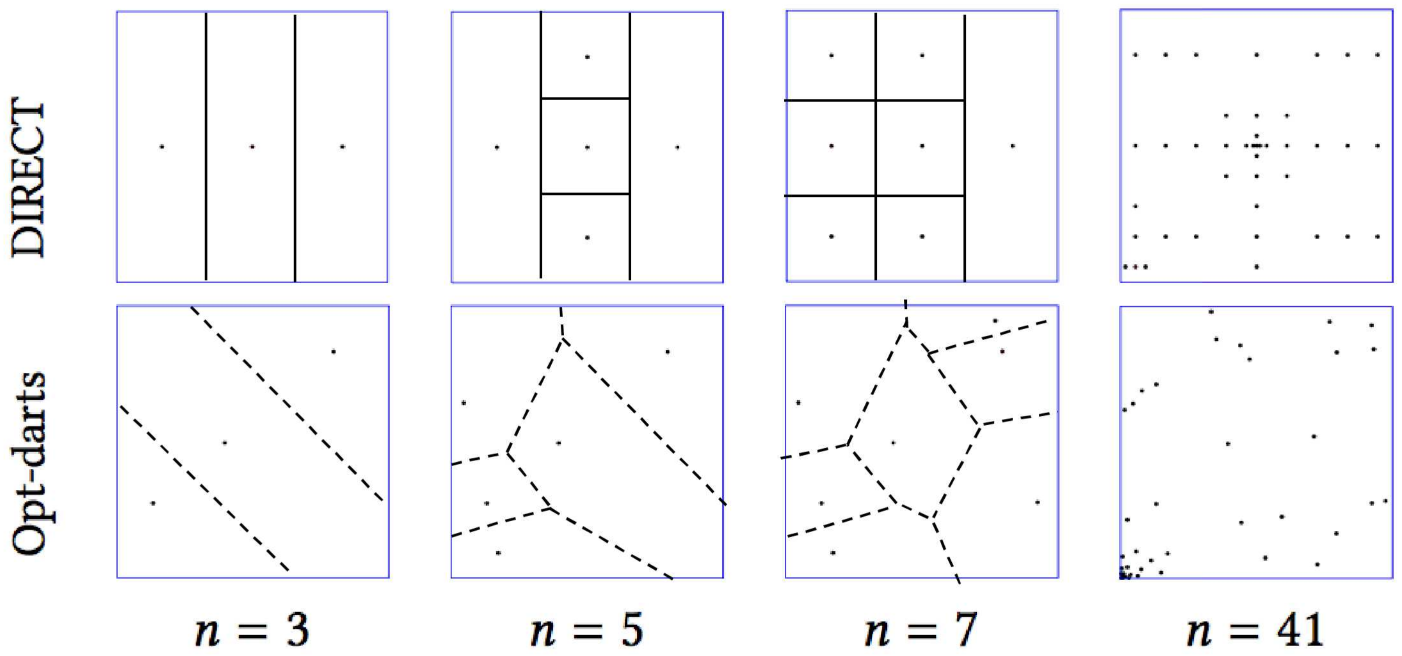

Application: Rendering

Mitsuba ray tracing

- 4D = 2 screen × 2 lens
- 16 samples per pixel

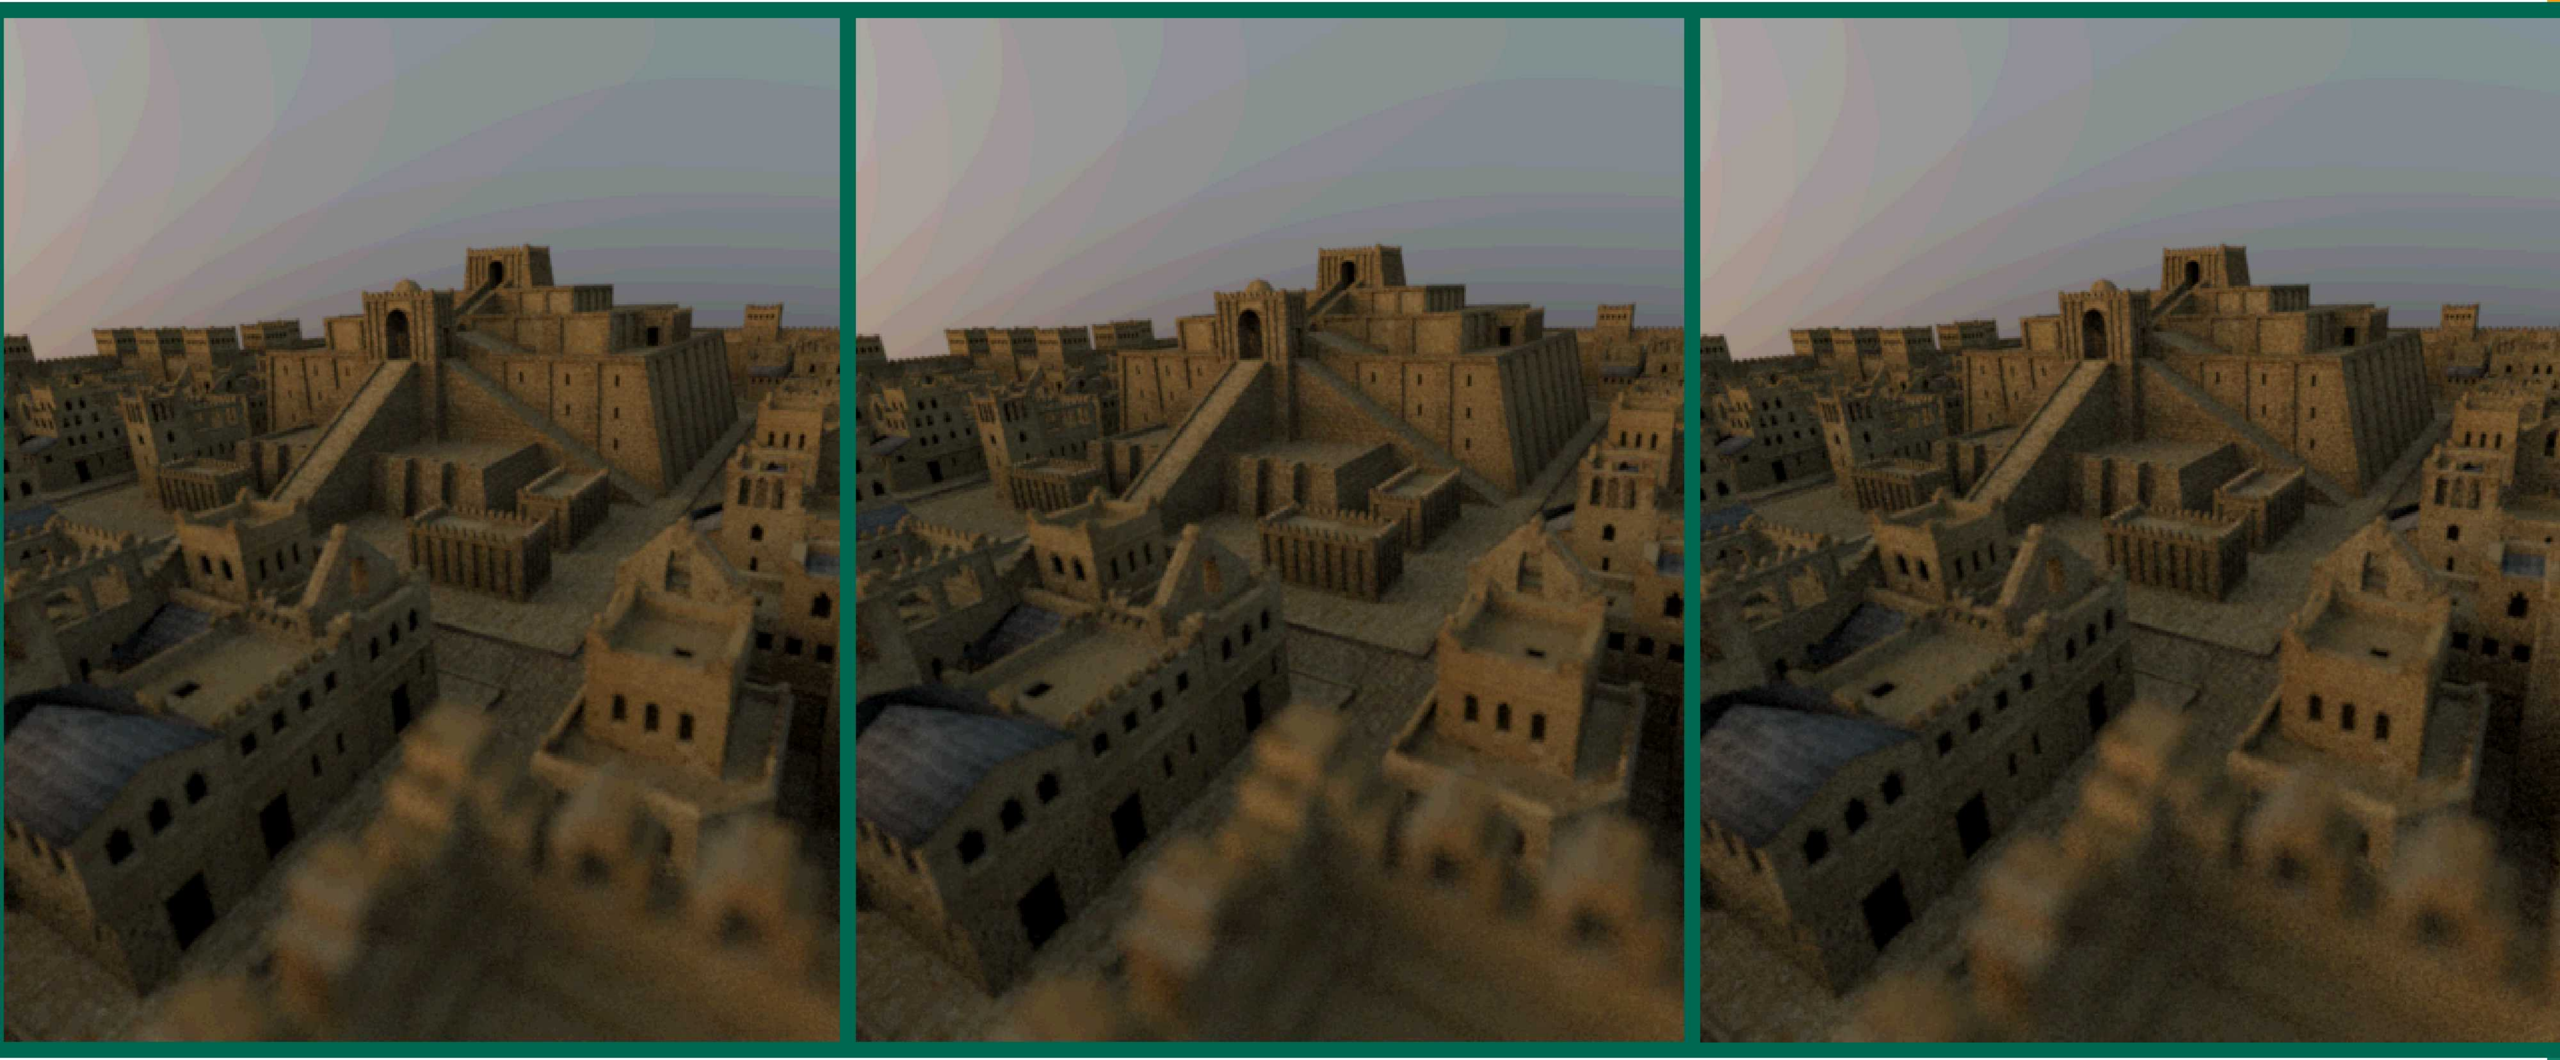

Stratified

Low discrepancy

Spoke-darts
comparable to alternatives

Application: Rendering

Mitsuba ray tracing

- 8D = 2 sky-emitter × 2 screen × 2 × 2 bounce
- 256 samples per pixel

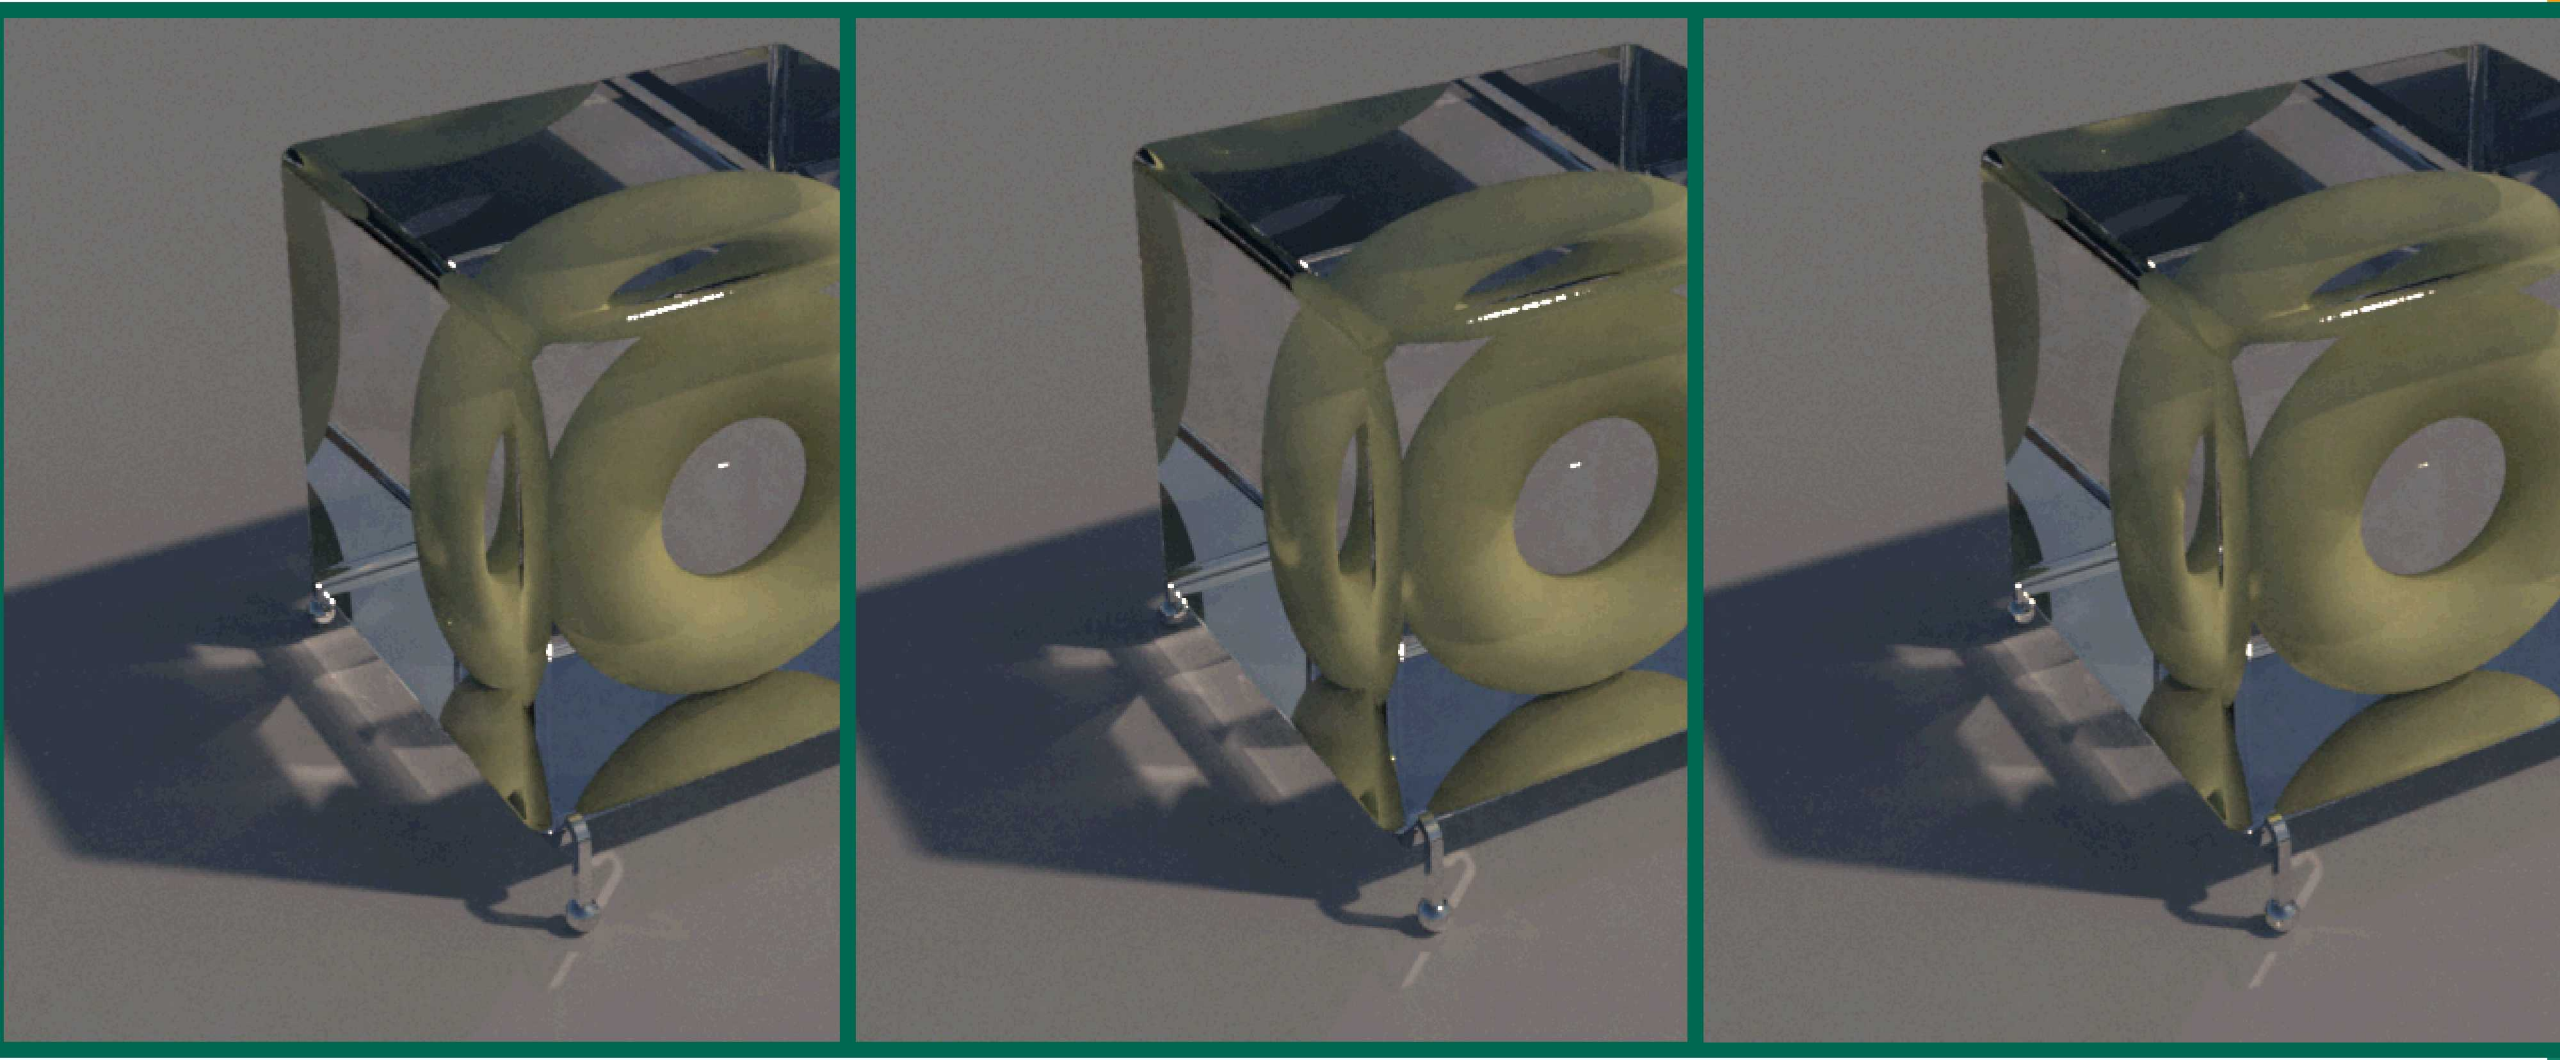

Stratified | Low discrepancy | Spoke-darts comparable to alternatives

Application: Motion Planning

Path = robot: under table to up
book: table to shelf

23d = degree of freedom robot
- Non-uniform: joint angles and room dimensions

Pre-sample configuration space with spoke-darts
- Explore paths in parallel, avoid redundant searches
- Order of magnitude speedup

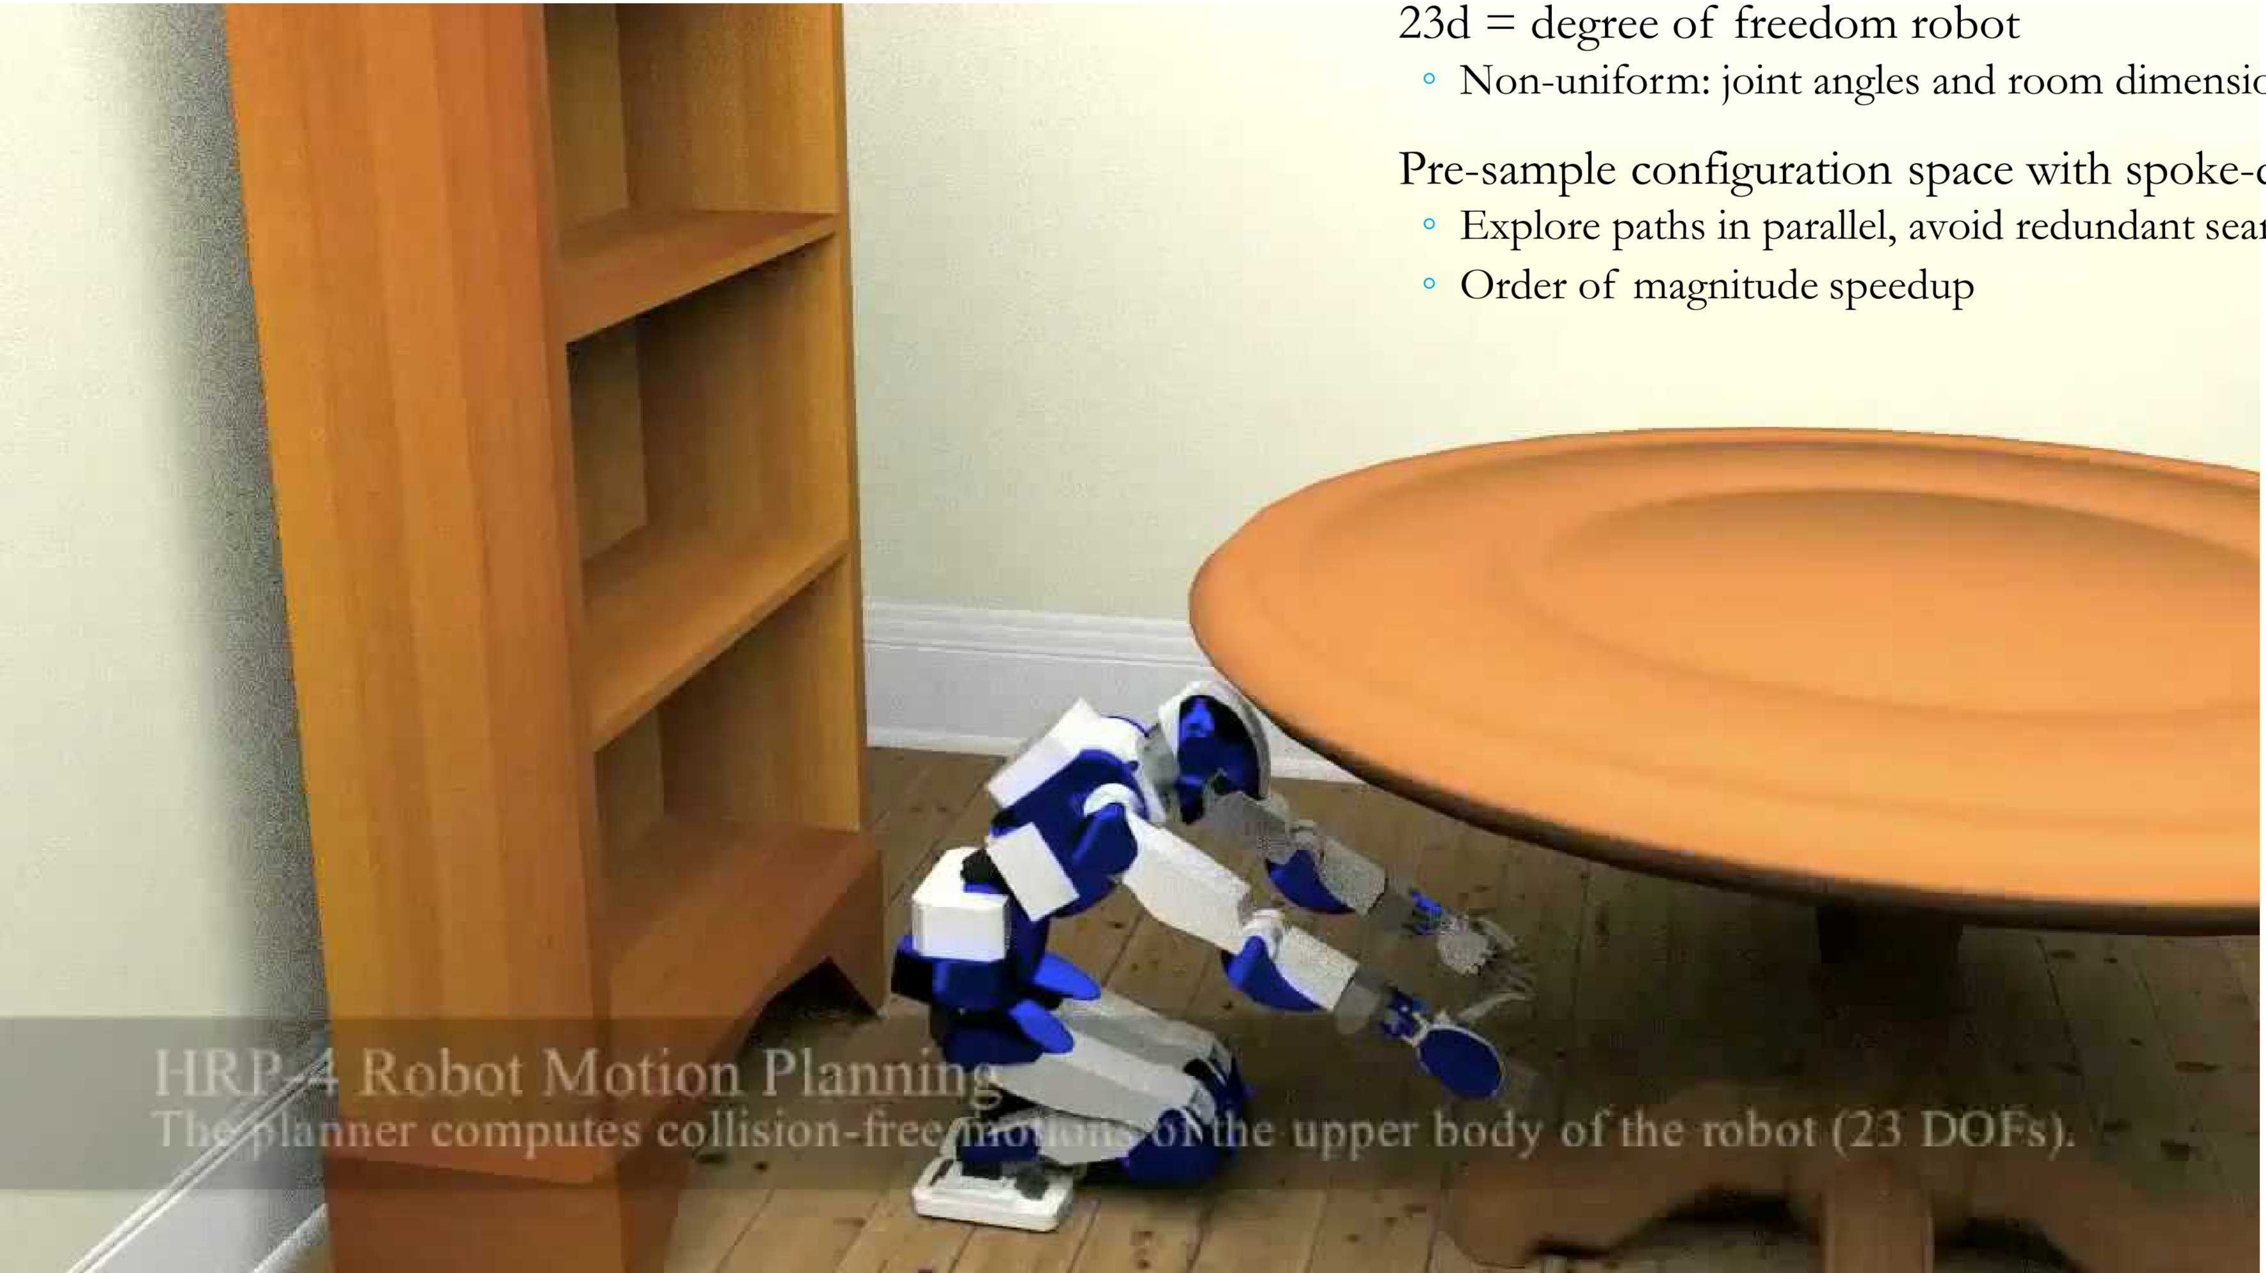

Conclusions

Center for Computing Research
Sandia National Laboratories

Free software

- SpokeDartsPublic on github
  https://github.com/samitch/SpokeDartsPublic

Recommendations

Two-radii

Quality blue-noise in $d < 5$

Spoke-Darts

Blue-noise $d \geq 5$
Linear scaling by dimension
Guaranteed saturation

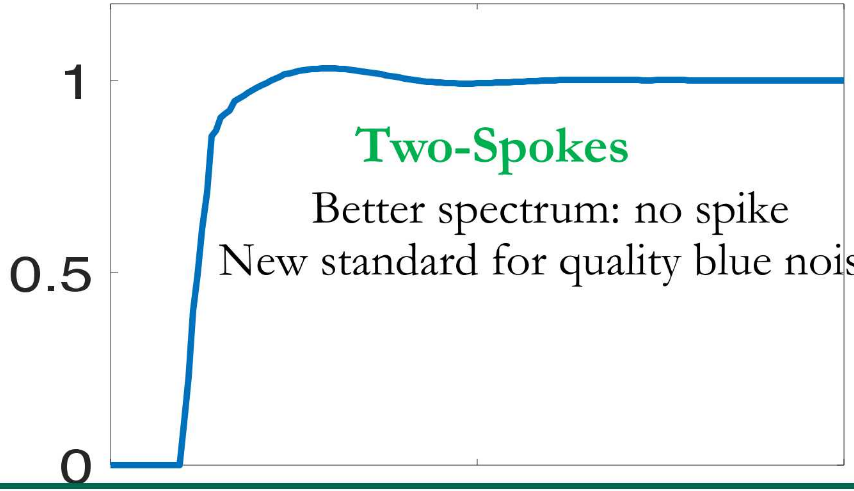

Two-Spokes

Better spectrum: no spike
New standard for quality blue noise

Open problems

- Faster than $O(d\ n^2)$ time?
- Our experiments had $r_{cover} = 2\ r_{free}$
  - Spectrum effects for other ratios?
- Prove non-trivial saturation-bound on Bridson 2007?
- Does saturation matter for high-d rendering, other "Graphics" applications ?

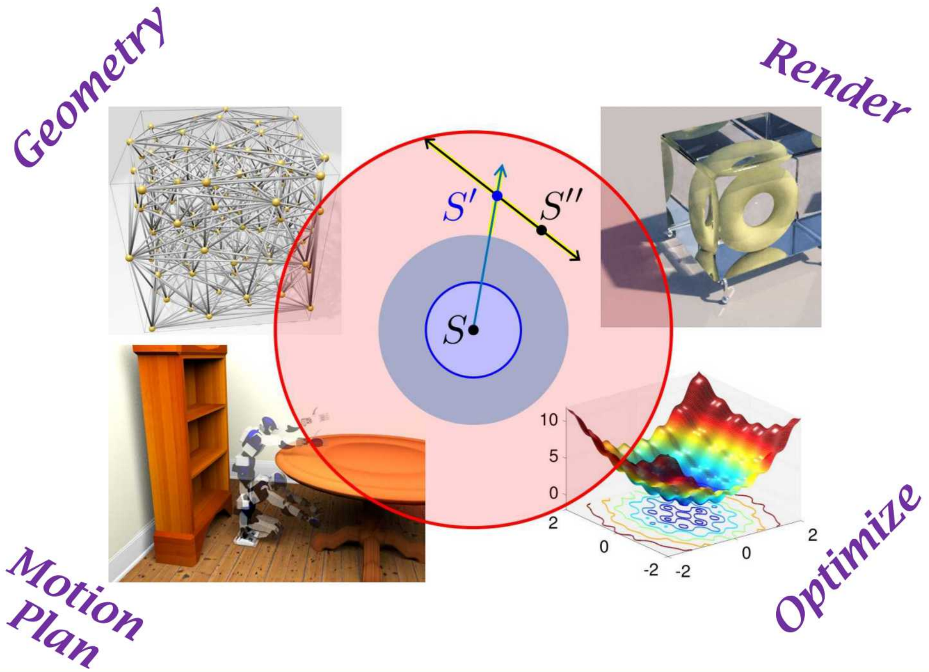

Backup slides

Spoke-Darts

Runtime $= O(dn^2)$, best one can hope for

d

- Linear scaling is perfect: primitives are O(d)

n^2

- Finding neighbors in high dimensions expensive
- Each of n samples, find neighbors
- Each of n may have n neighbors
  - Finding them quickly still won't prevent n^2

Spoke-Darts ≈ Bridson 2007 due to simple primitives

- Is point in sphere
- Trim segment by sphere

Runtime(Two-Spokes) > Runtime(One-Spoke)

- mostly due to 2 × radius of annulus, many more neighbor disks

Spectra by dimension