On Multiple Virtues of Blue Noise Sampling

Victor Ostromoukhov

University of Lyon/CNRS

Joint Work with David Coeurjolly, Adrien Pilleboue, Gurprit Singh, Helene Perrier, Abdalla Ahmed, Eric Heitz, Laurent Belcour, Matt Pharr, Michael Kazhdan, Jianwei Guo, Dongming Yan, Hui Huang, Oliver Deussen, Feng Xie, Pat Hanrahan

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

ABDALLA AHMED

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

MOTIVATION

- BN sampling is good for diminishing overall noise in MC integration
- BN sampling is good for improving visual appearance of synthetic images
- Advanced BN sampling can be efficiently implemented

OVERVIEW

- Theoretical foundation for BN sampling

    Based on Variance Analysis for Monte Carlo Integration, SIGGRAPH 2015

- Some efficient implementations

    Based on

    - Low-Discrepancy Blue Noise Sampling, SIGGRAPH-ASIA 2016
    - Sequences with Low-Discrepancy Blue-Noise 2-D Projections, EG2018
    - A Low-Discrepancy Sampler that Distributes Monte Carlo Errors as a Blue Noise in Screen Space, SIGGRAPH 2019 Talk

- Open Issues

BLUE NOISE

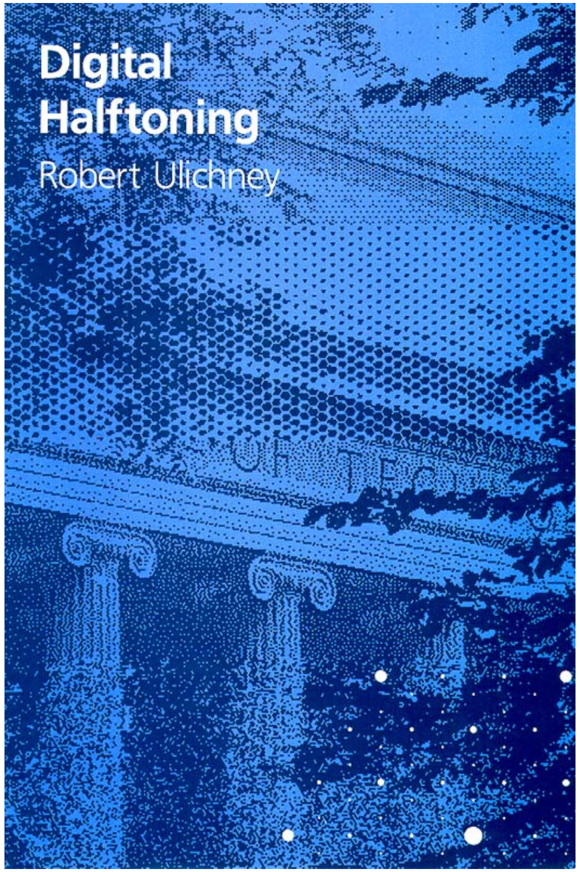

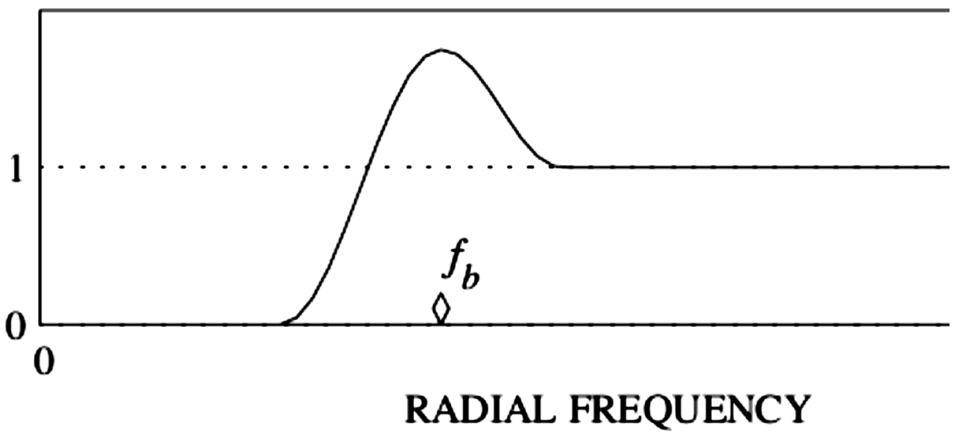

[Ulichney 1987]

thrive SIGGRAPH2019 LOS ANGELES • 28 JULY – 1 AUGUST

BN IN NATURE: COMPETITION FOR THE VITAL SPACE

BN (ORGANIC) VS. ARTIFICIAL (ORDERED) DISTRIBUTIONS

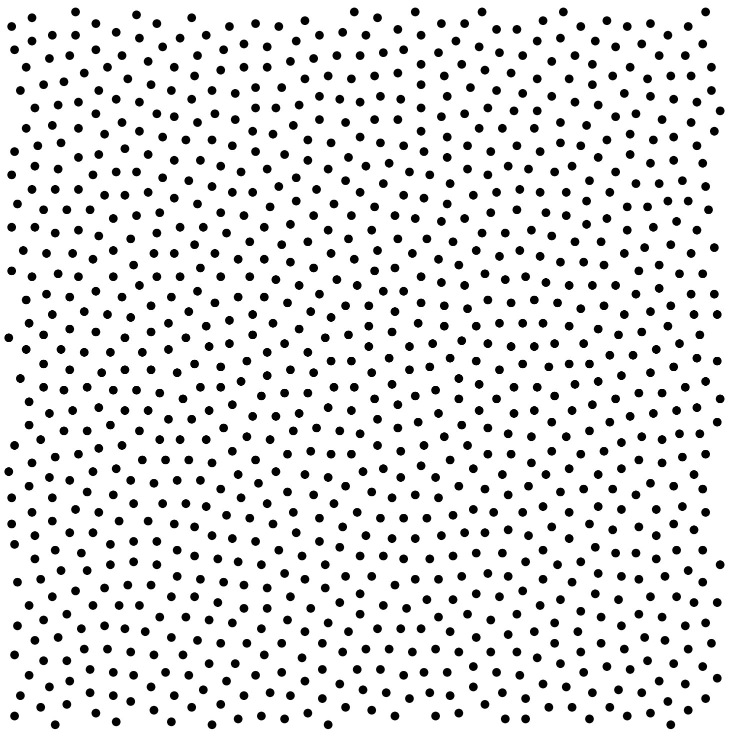

BN Points

[de Goes et al. 2012]

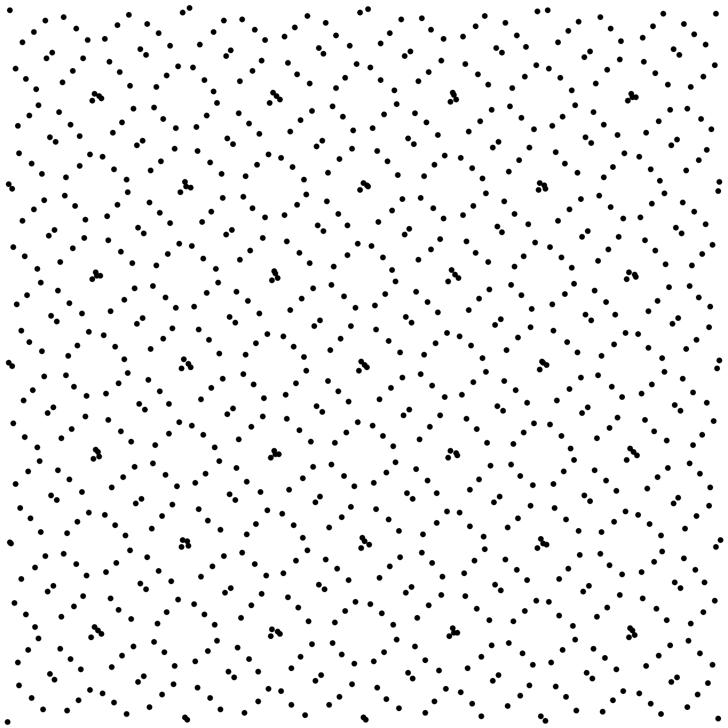

Sobol Points

BLUE NOISE POWER AND RADIAL SPECTRA IN 2D

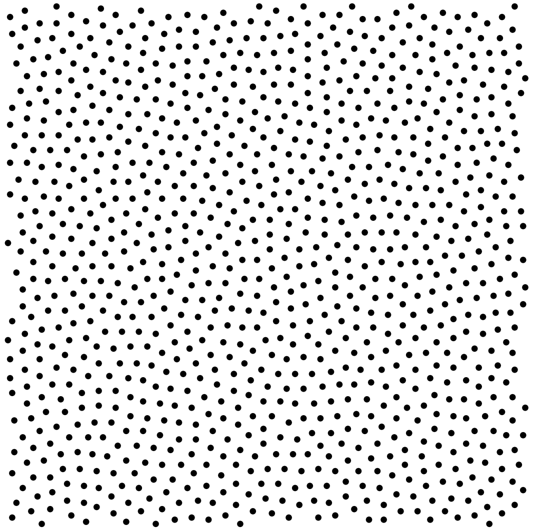

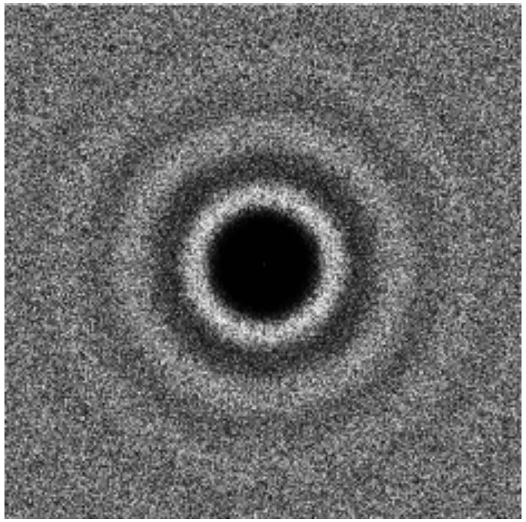

[de Goes et al. 2012]

BN: TARGET BEHAVIOR OF MSE IN INTEGRATION

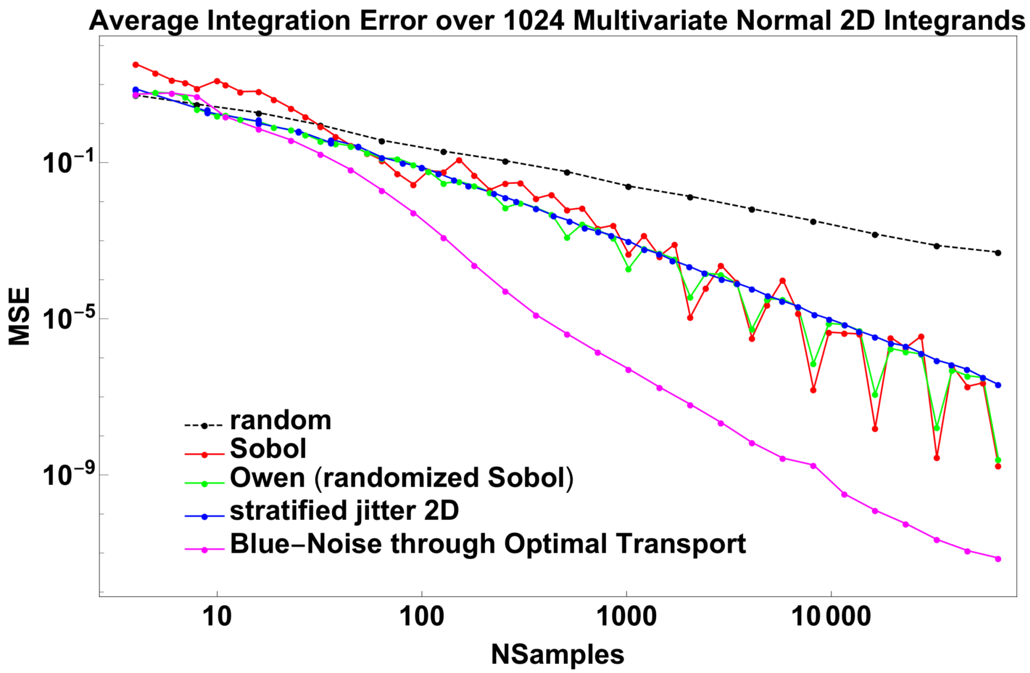

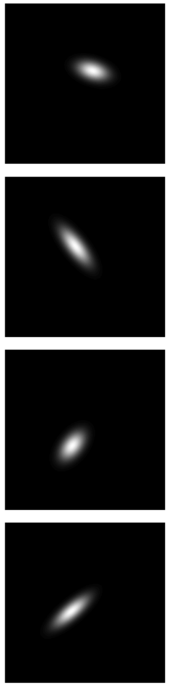

THEORETICAL FOUNDATION FOR BN SAMPLING

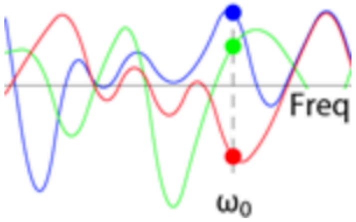

Variance Analysis for Monte Carlo Integration

Abstract. We propose a new spectral analysis of the variance in Monte Carlo integration, expressed in terms of the power spectra of the sampling pattern and the integrand involved. We build our framework in the Euclidean space using Fourier tools and on the sphere using spherical harmonics. We further provide a theoretical background that explains how our spherical framework can be extended to the hemispherical domain. We use our framework to estimate the variance convergence rate of different state-of-the-art sampling patterns in both the Euclidean and spherical domains, as the number of samples increases. Furthermore, we formulate design principles for constructing sampling methods that can be tailored according to available resources. We validate our theoretical framework by performing numerical integration over several integrands sampled using different sampling patterns.

Citation: Adrien Pilleboue, Gurprit Singh, David Coeurjolly, Michael Kazhdan, Victor Ostromoukhov, Variance Analysis for Monte Carlo Integration, SIGGRAPH 2015, ACM Trans. Graph. 34(4), pp. 124:1--124:14.

VARIANCE FORMULATION BASED ON FOURIER ANALYSIS

$$Var(I_N) = \frac{\mu(\mathcal{T}^d)\mu(\mathcal{S}^{d-1})}{N} \int_0^\infty \rho^{d-1} \breve{\mathcal{P}}_{\mathbf{S}}(\rho) \breve{\mathcal{P}}_{\mathbf{F}}(\rho) \mathrm{d}\rho$$

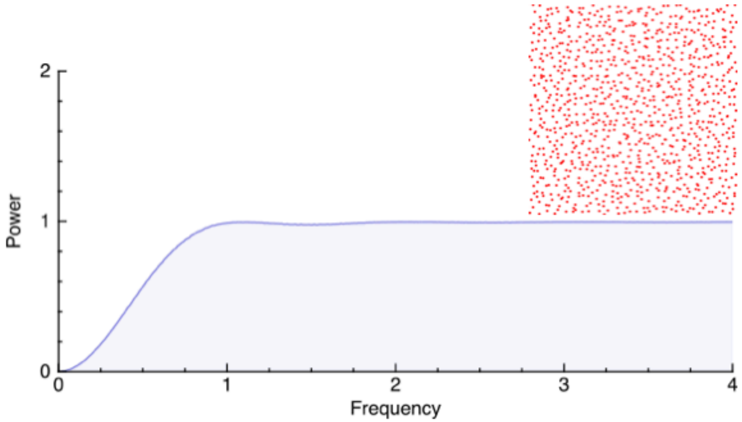

(Jittered Sampling Pattern)

VARIANCE FORMULATION BASED ON FOURIER ANALYSIS

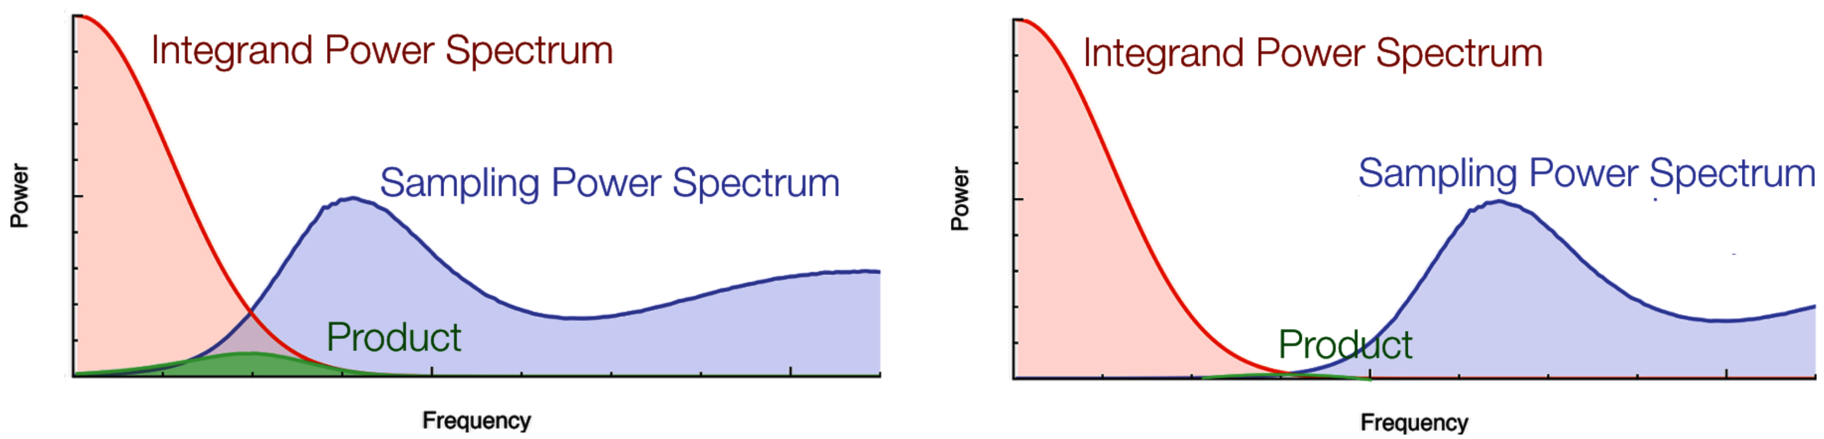

TAXONOMY OF CONVERGENCY CLASSES

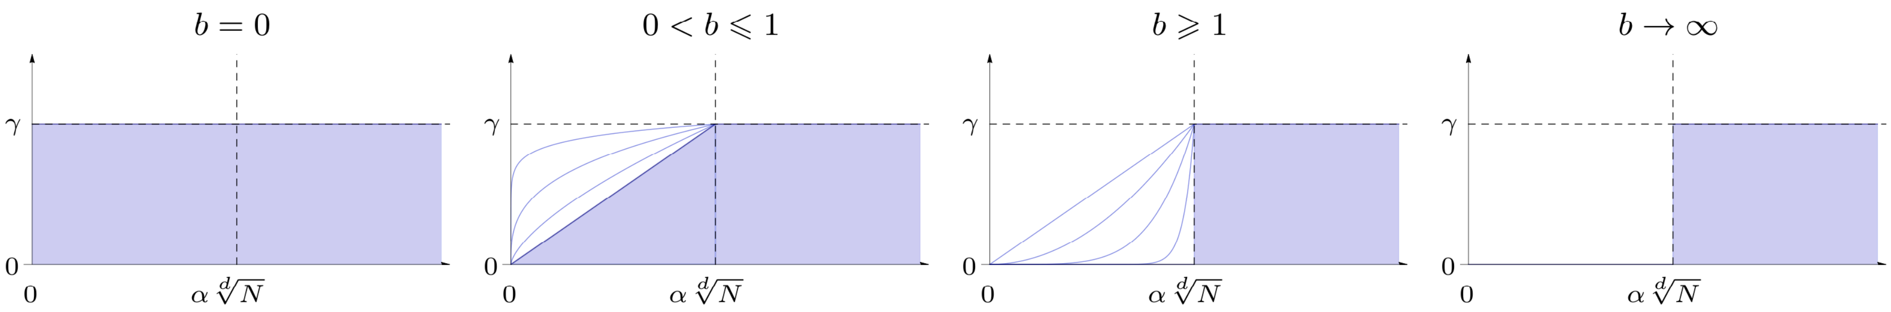

(a) (b) (c) (d)

(a)

Best-case: $O\left(\frac{1}{N}\right)$

Worst-case: $O\left(\frac{1}{N}\right)$

(b)

Best-case: $O\left(\frac{1}{N\sqrt[d]{N}^b}\right)$

Worst-case: $O\left(\frac{1}{N\sqrt[d]{N}^b}\right)$

(c)

Best-case: $O\left(\frac{1}{N\sqrt[d]{N}^b}\right)$

Worst-case: $O\left(\frac{1}{N\sqrt[d]{N}}\right)$

(d)

Best-case: 0

Worst-case: $O\left(\frac{1}{N\sqrt[d]{N}}\right)$

b: degree of the polynomial
d: dimensions
N: number of samples

LOW FREQUENCY REGION

Poisson Disk: $O(\frac{1}{N})$

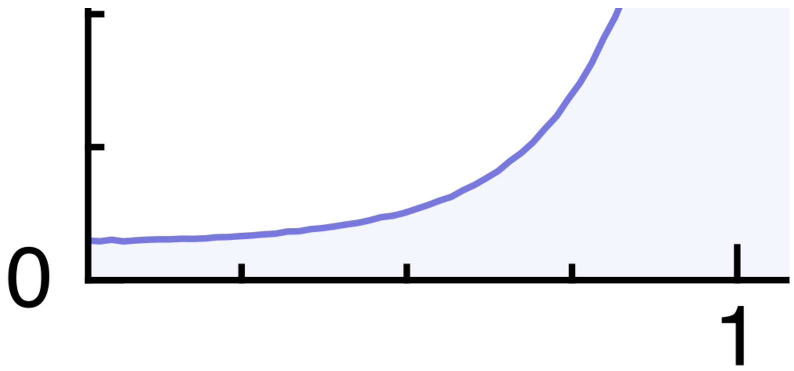

Jittered: $O(\frac{1}{N\sqrt{N}})$

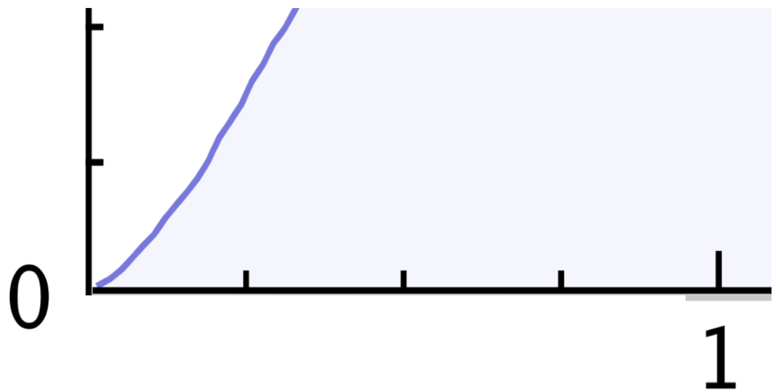

VERIFICATION OF THE THEORETICAL PREDICTION

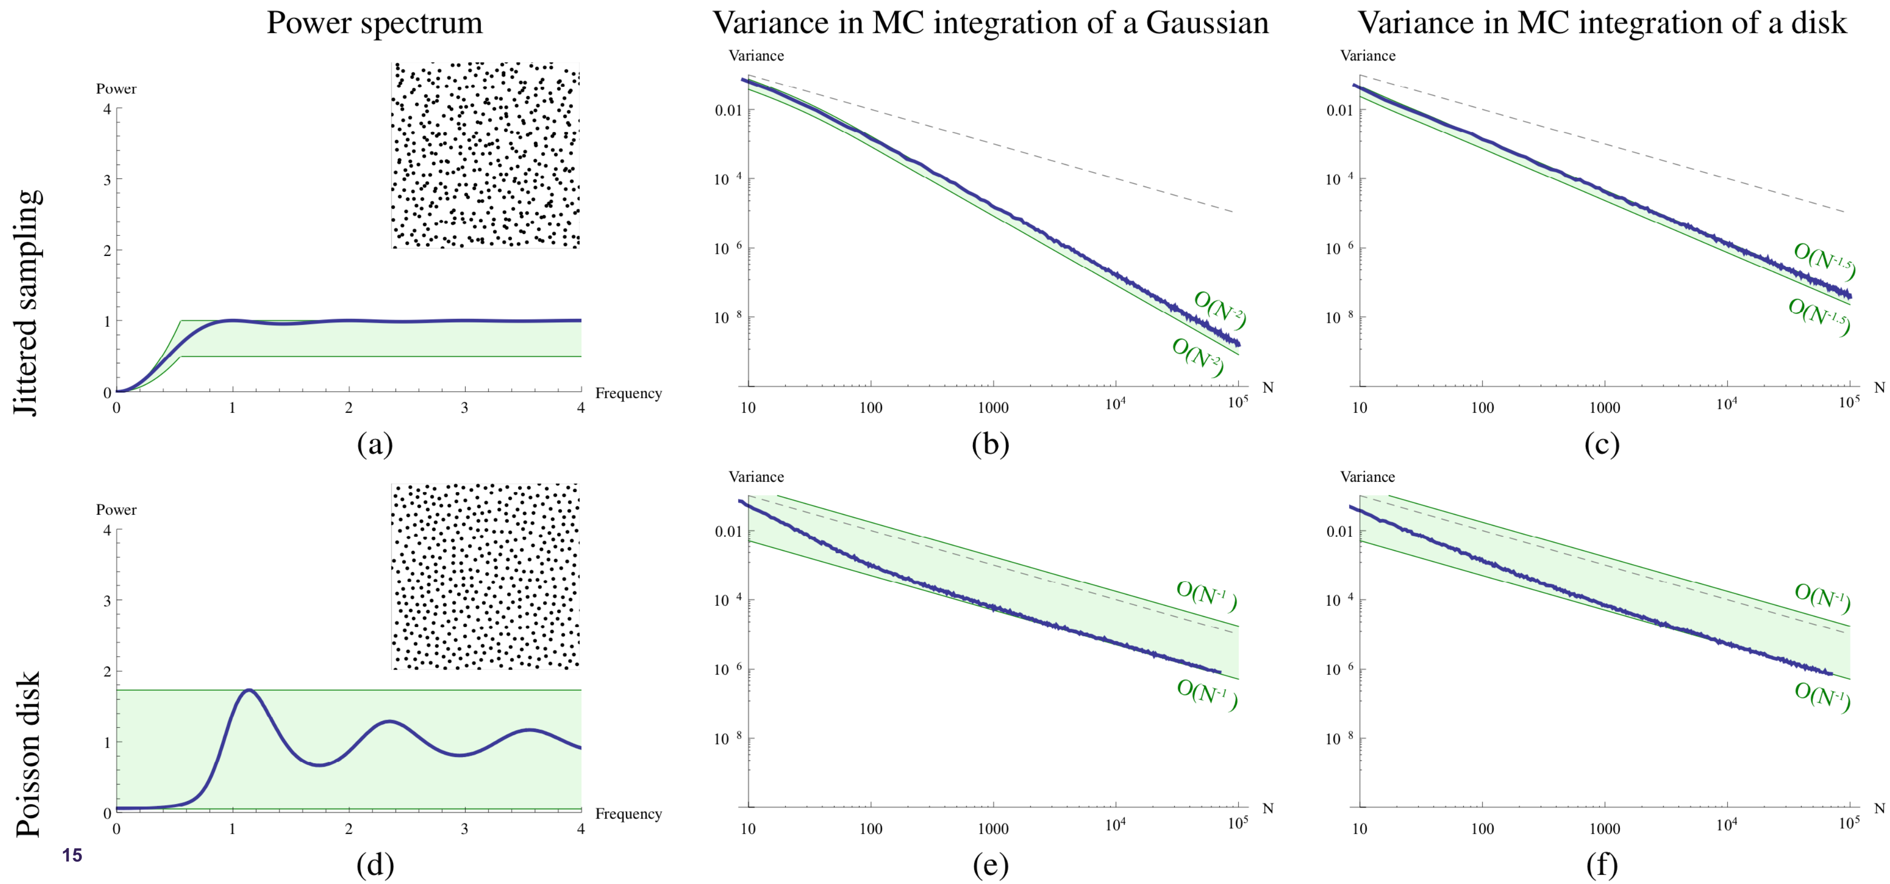

OVERVIEW

- Theoretical foundation for BN sampling
  - Based on Variance Analysis for Monte Carlo Integration, SIGGRAPH 2015
- Some efficient implementation
  - Based on
    - Low-Discrepancy Blue Noise Sampling, SIGGRAPH-ASIA 2016
    - Sequences with Low-Discrepancy Blue-Noise 2-D Projections, EG2018
    - A Low-Discrepancy Sampler that Distributes Monte Carlo Errors as a Blue Noise in Screen Space, SIGGRAPH 2019 Talk
- Open Issues

Low-Discrepancy Blue Noise Sampling

Abdalla G. M. Ahmed[1] Hélène Perrier[2] David Coeurjolly[2] Victor Ostromoukhov[2]
Jianwei Guo[3] Dongming Yan[3] Hui Huang[4,5] Oliver Deussen[1,5]

[1]University of Konstanz, Germany
[2]Université de Lyon, CNRS/LIRIS, France
[3]NLPR, Institute of Automation, CAS, China
[4]Shenzhen University, China
[5]SIAT, China

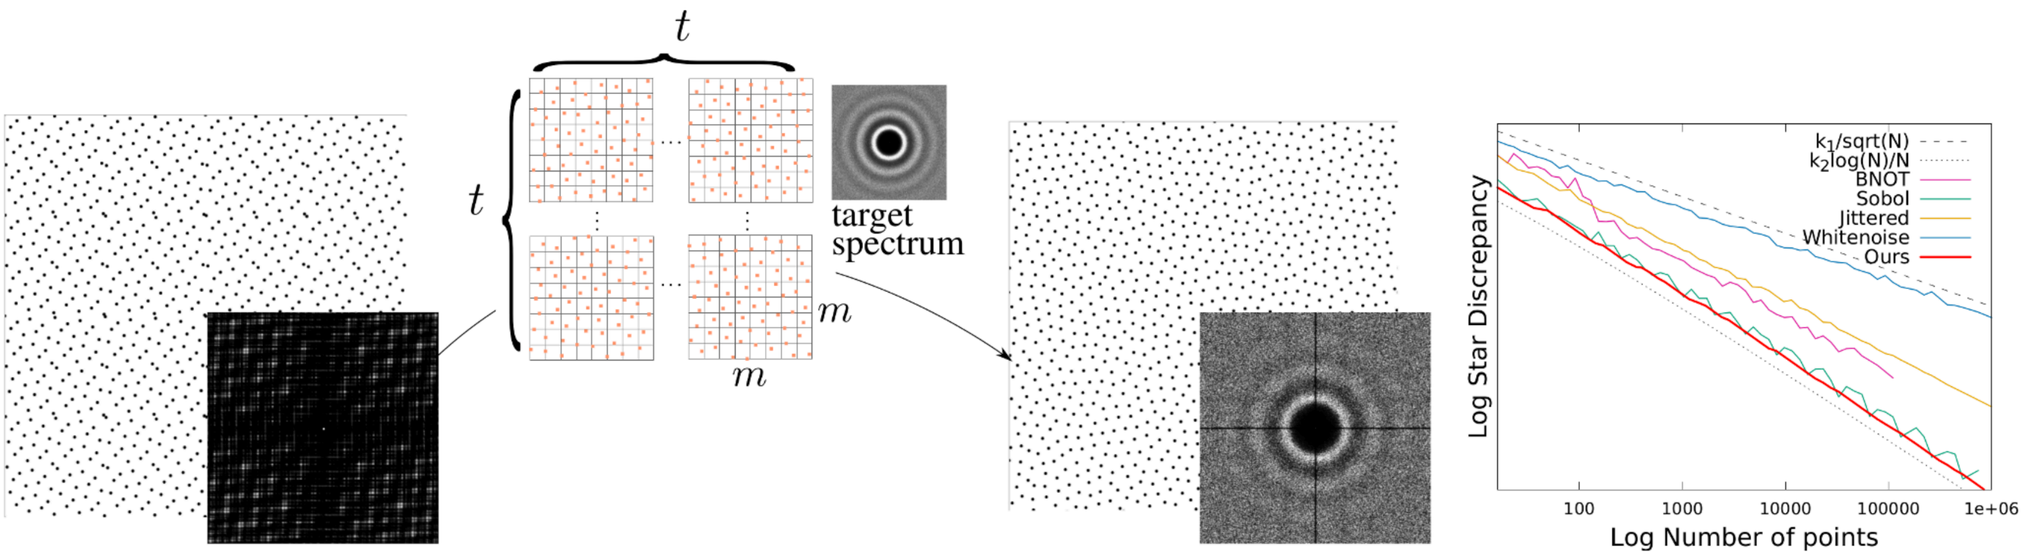

2D INDEXED LD SETS

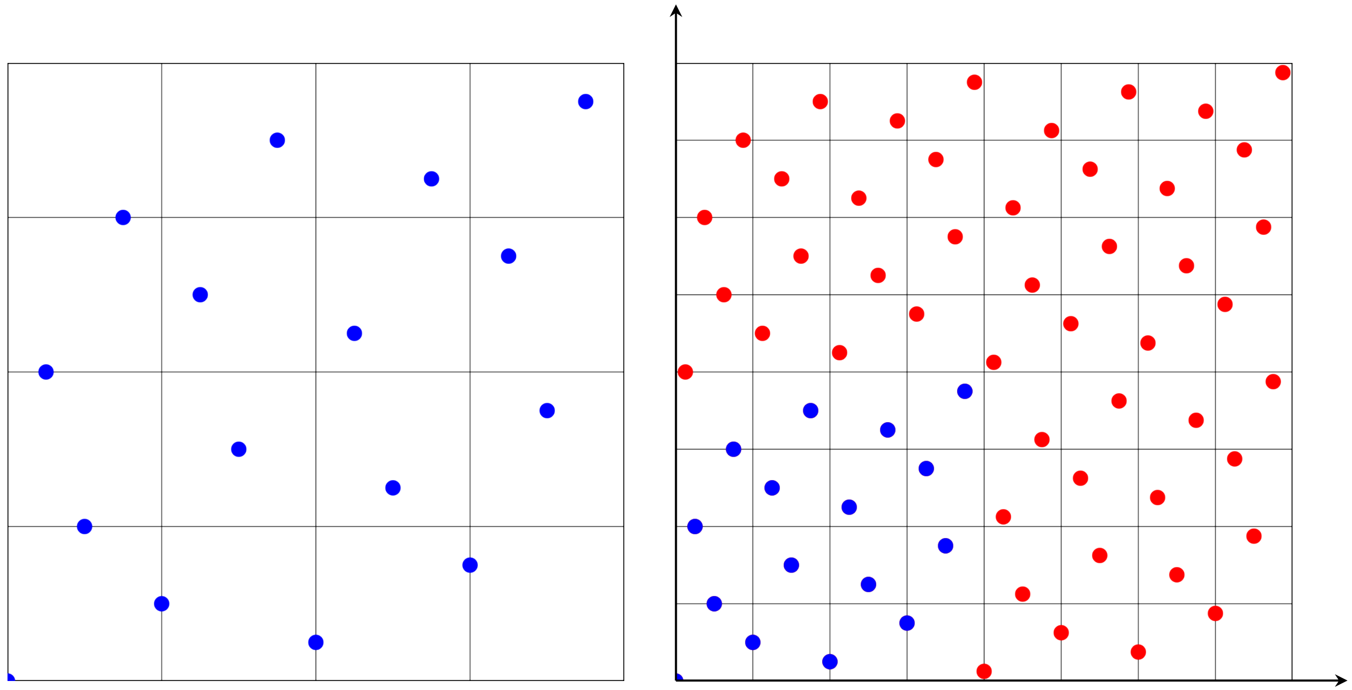

DISCREPANCY-PRESERVING REARRANGEMENT

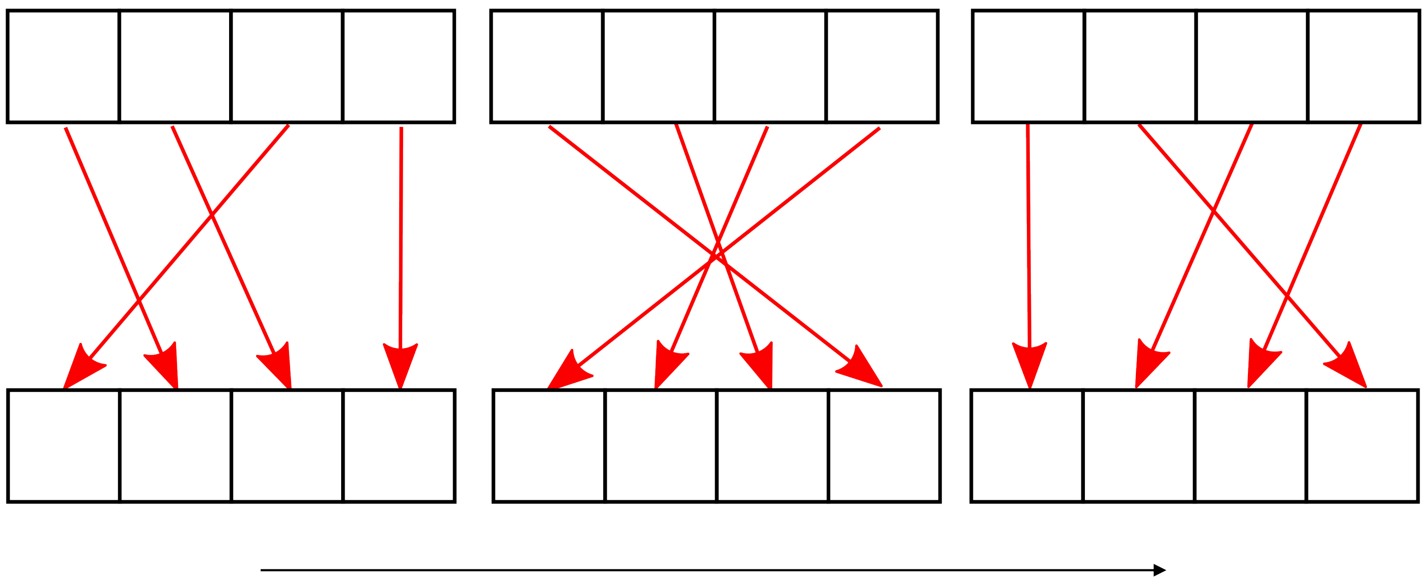

AXIS-WISE 2D REARRANGEMENT DEMO

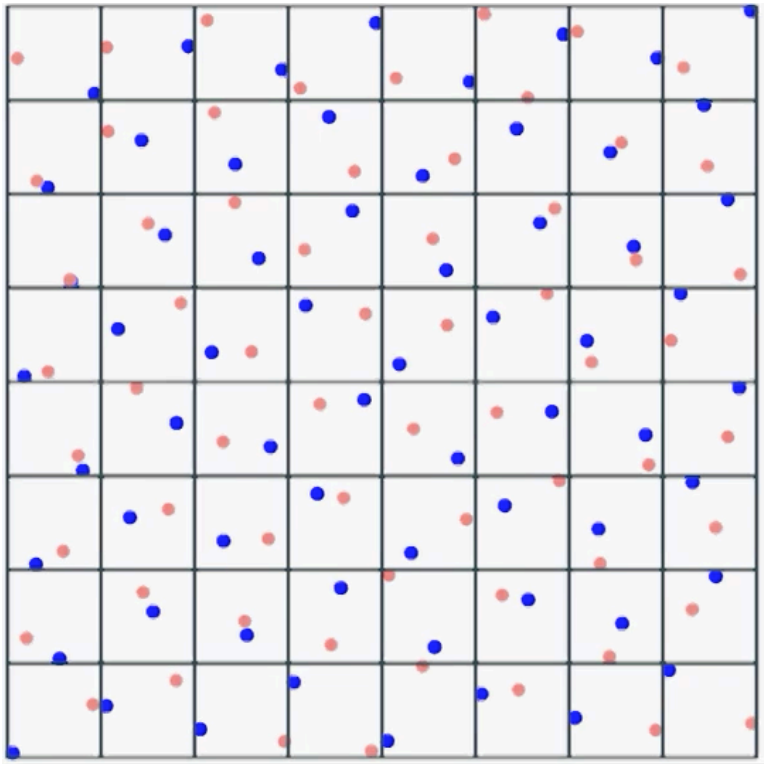

REFERENCE-MATCHING ALGORITHM

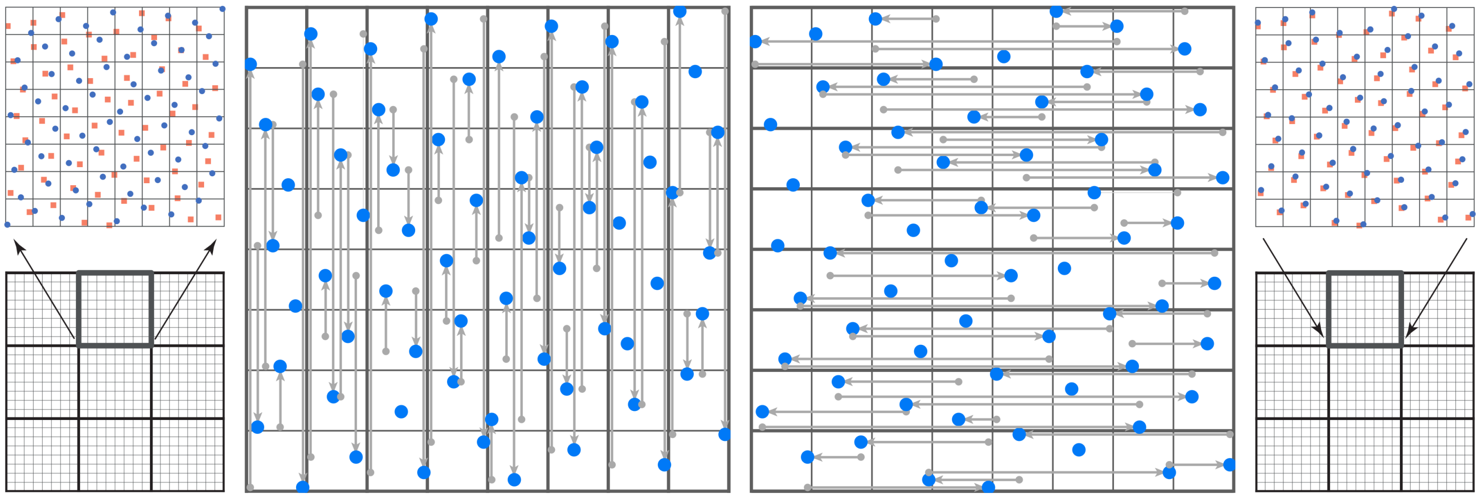

DEMO

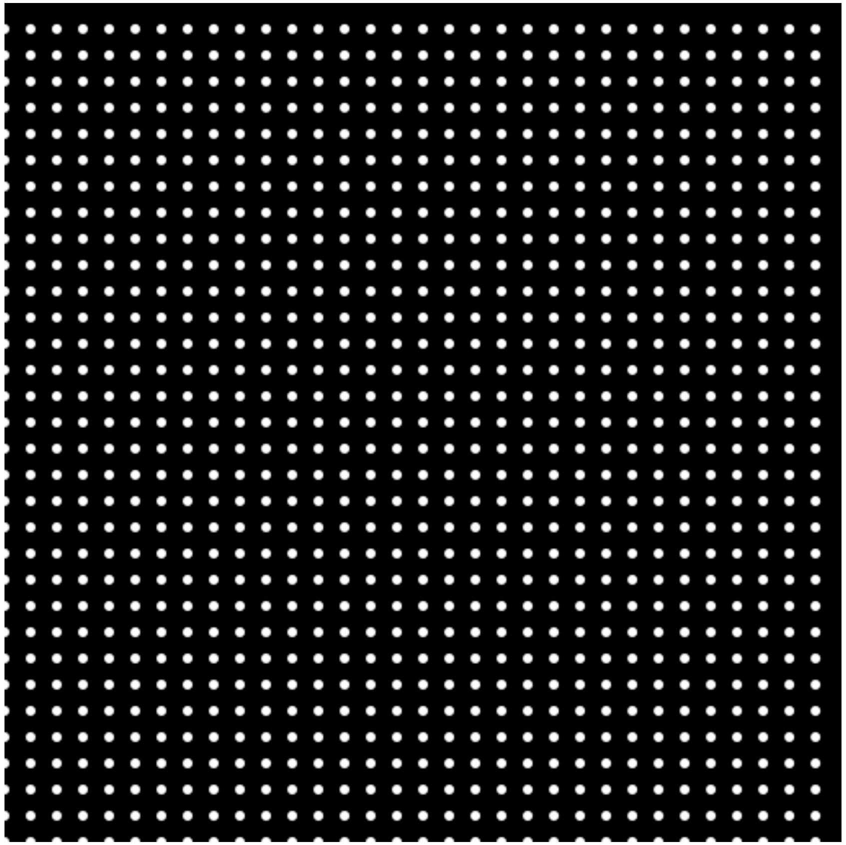

Sequences with Low-Discrepancy Blue-Noise 2-D Projections

Hélène Perrier[1] David Coeurjolly[1] Feng Xie[2] Matt Pharr[3] Pat Hanrahan[2] Victor Ostromoukhov[1]

[1]Université de Lyon, CNRS, LIRIS, France [2]Stanford, USA [3]Google, USA

In *Computer Graphics Forum (Proceedings of Eurographics)*, 2018

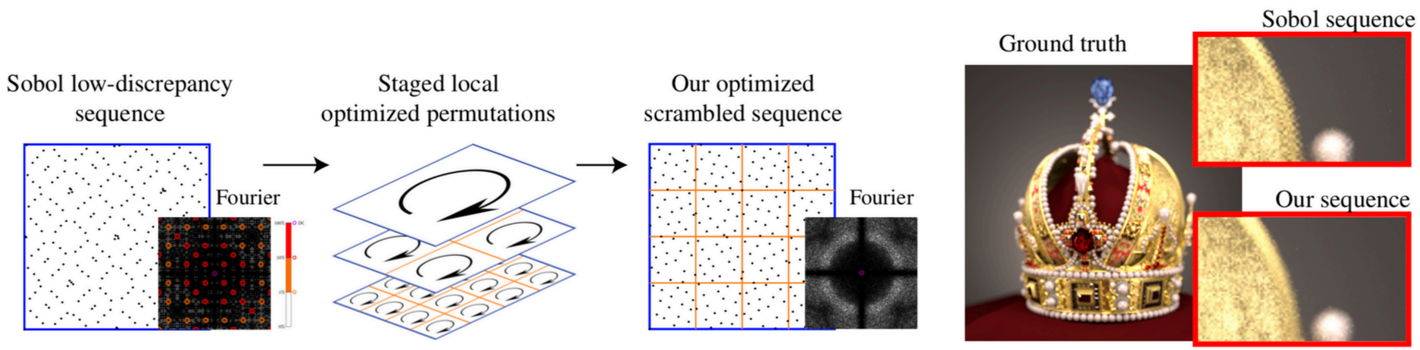

Left: our staged per-tile optimized scrambling, applied to a Sobol sequence of sampling points, produces a power spectrum close to Blue Noise. Right: Rendering of a challenging scene featuring depth of field and high specularity (jewels). The sampling was done with the Sobol sequence (top) and our sampler (bottom); both use 256 samples per pixel and 3 light bounces. Note the improvement in aliasing when using our method in comparison to the original Sobol sequence.

APH2019
28 JULY – 1 AUGUST

OWEN'S SCRAMBLING

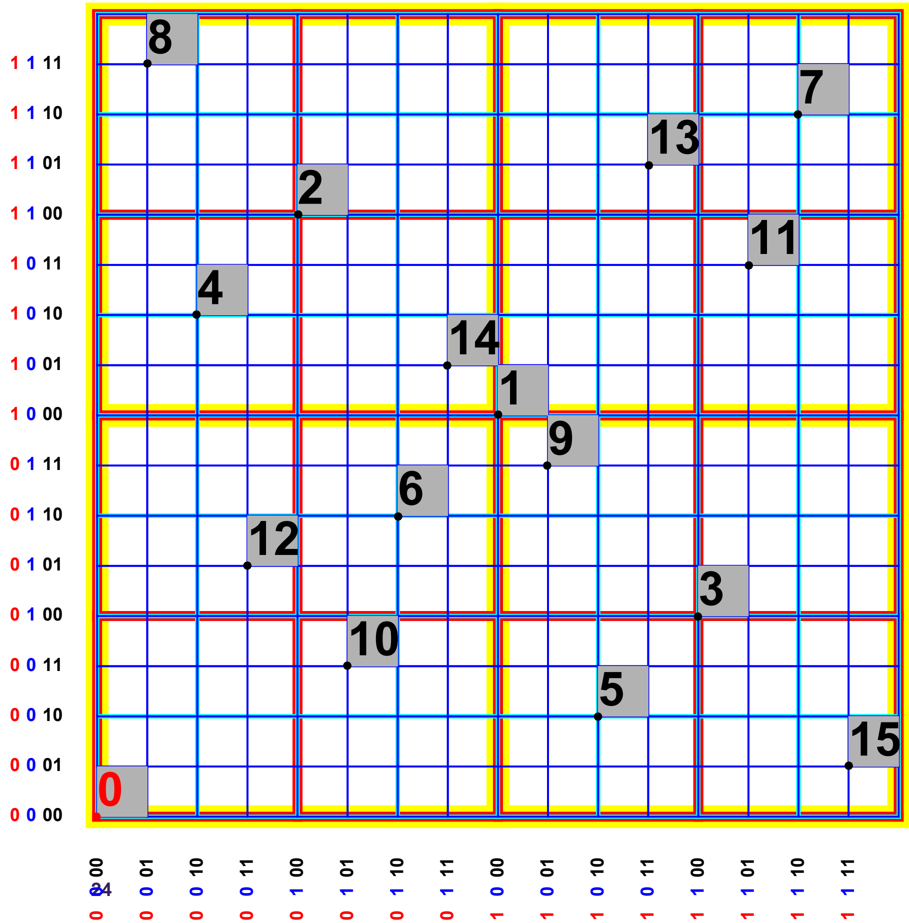

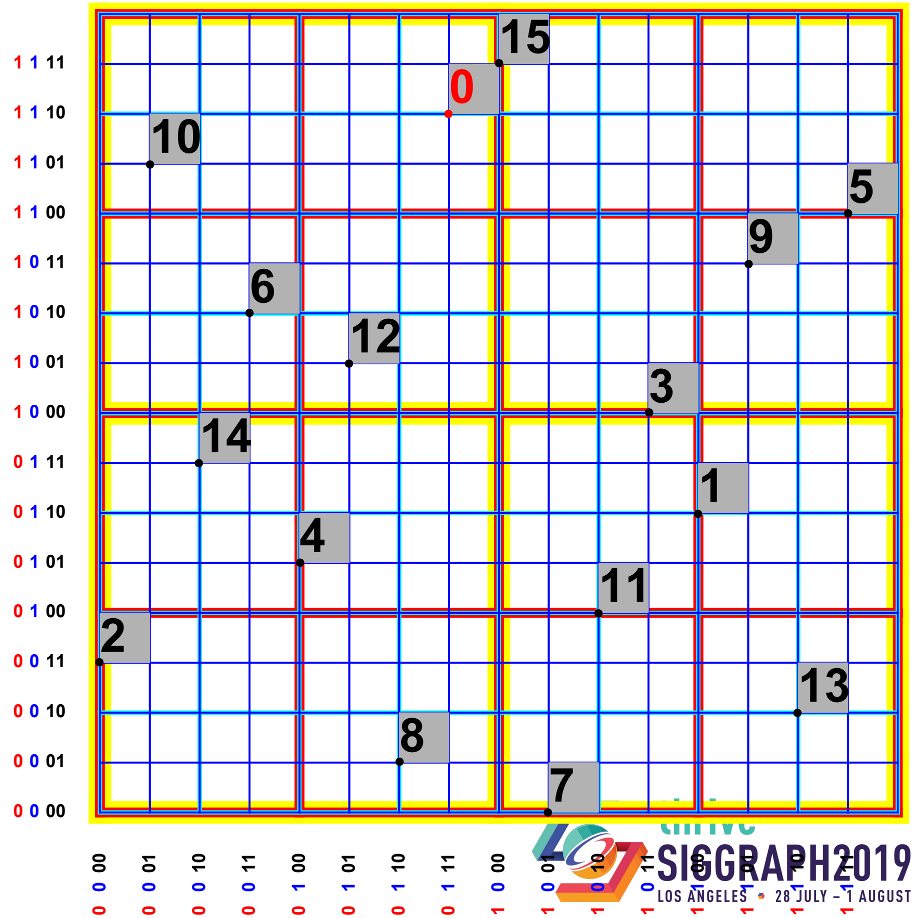

OWEN'S SCRAMBLING

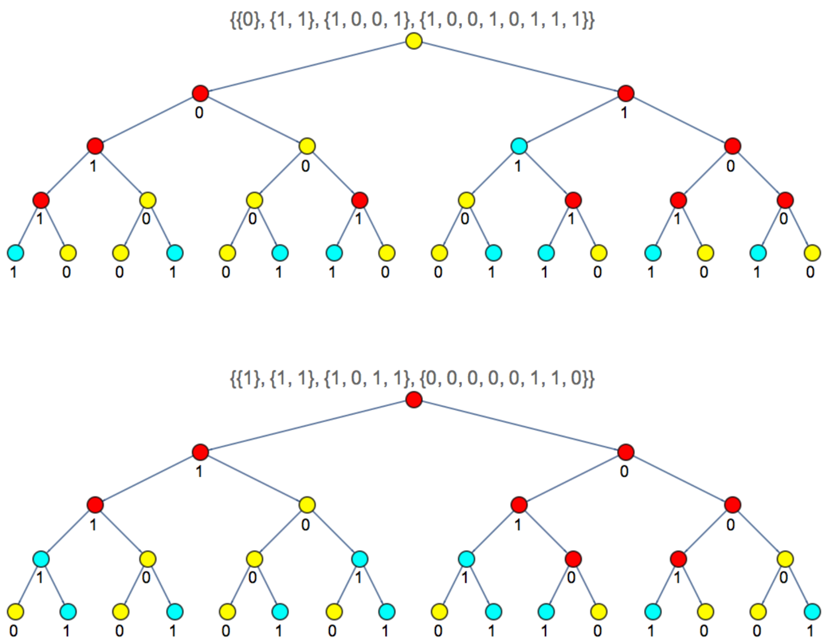

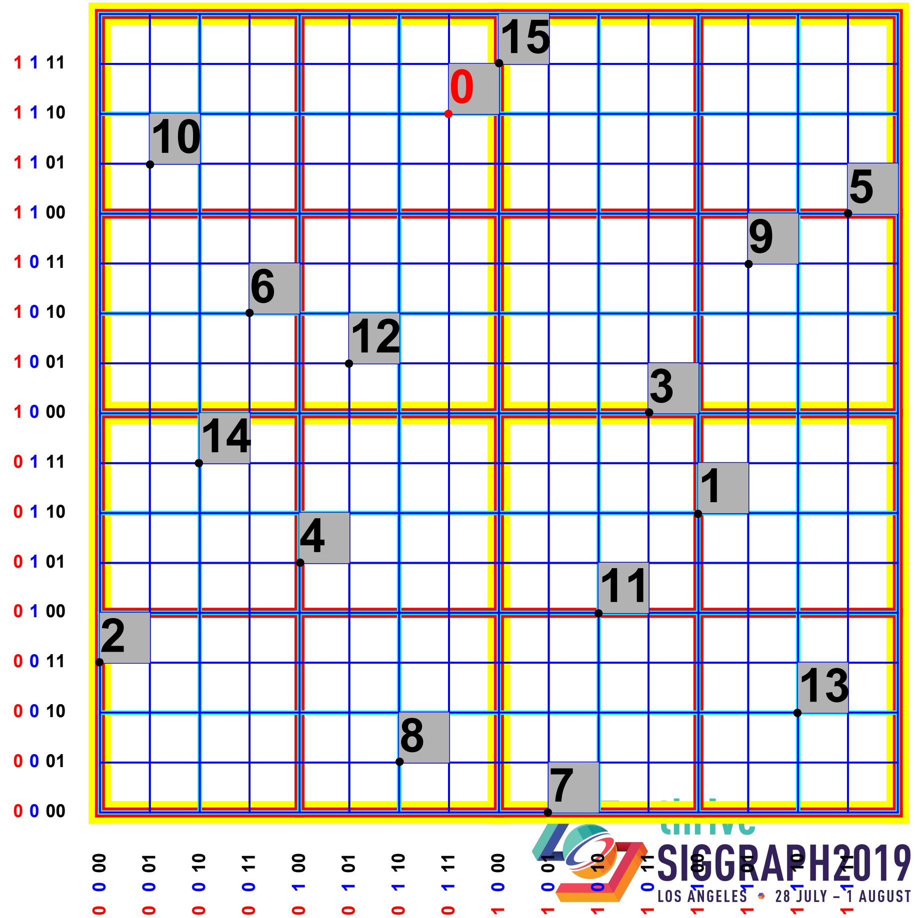

OWEN'S SCRAMBLING

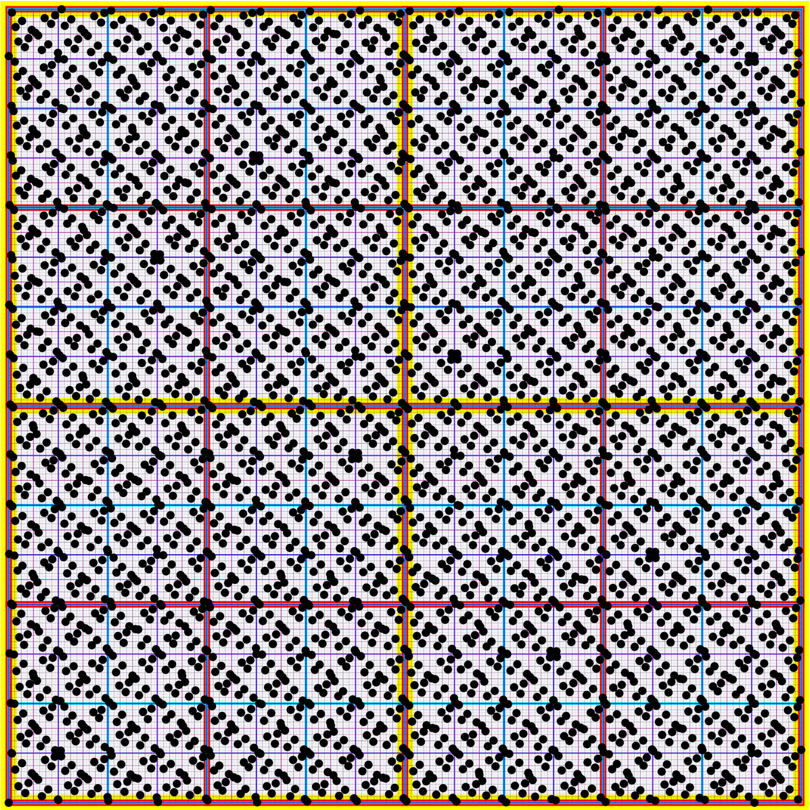

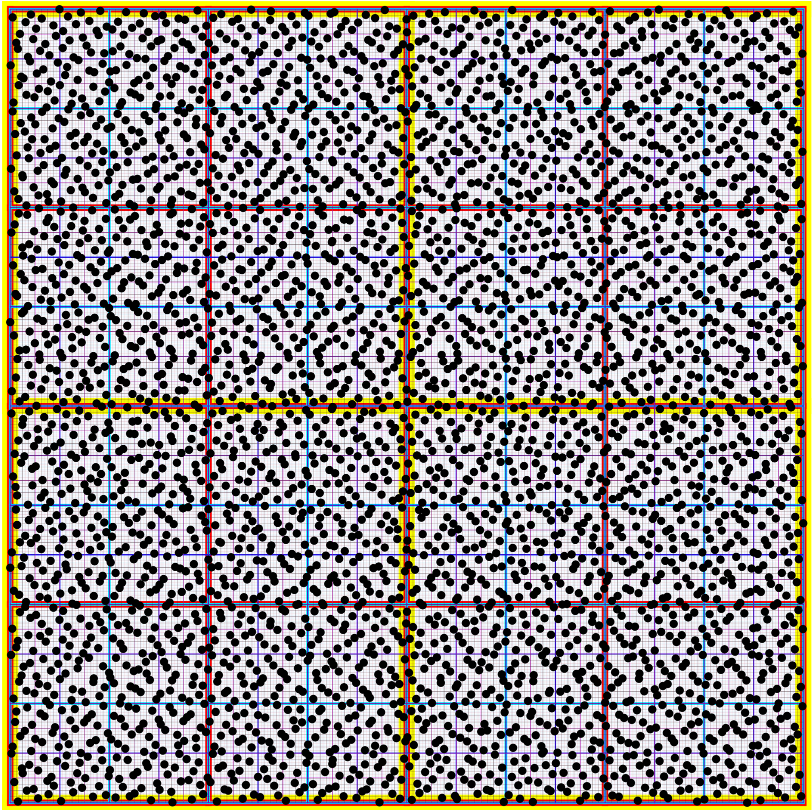

OWEN’S SCRAMBLING

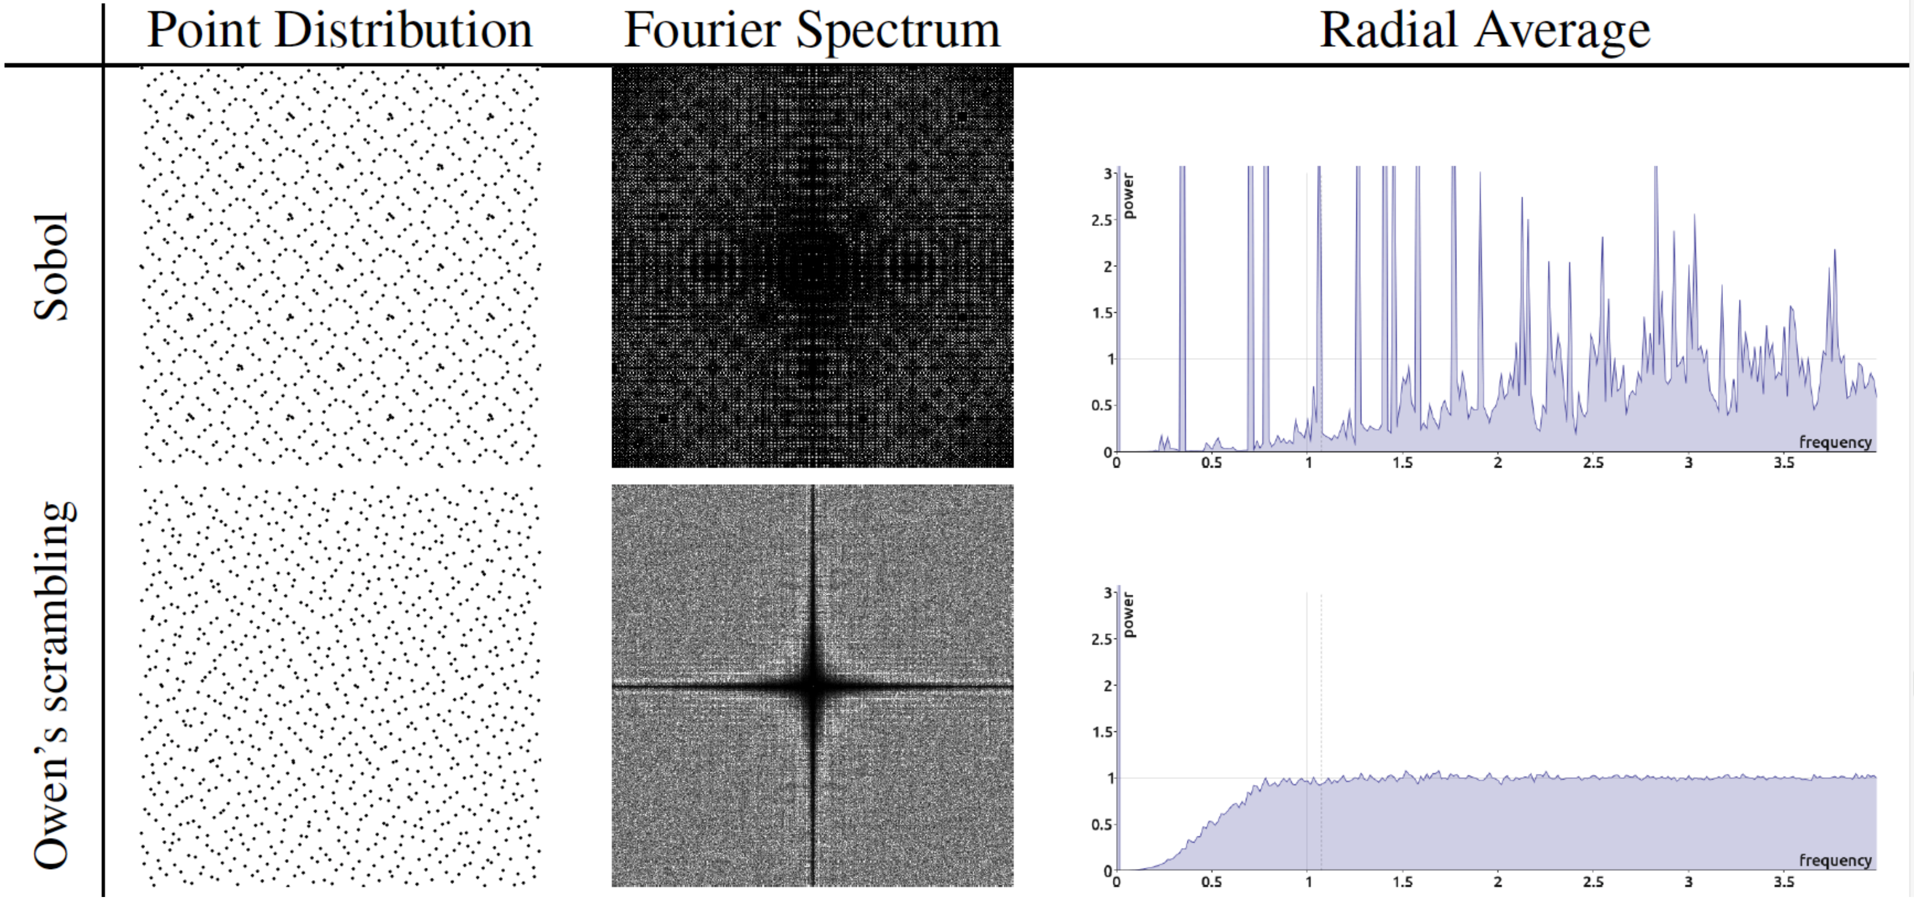

CONSTRUCTION IN [PERRIER ET AL. 2018]: THE KEY IDEA

Step 1: Identify all possible 16-point tiles in the original Sobol sets

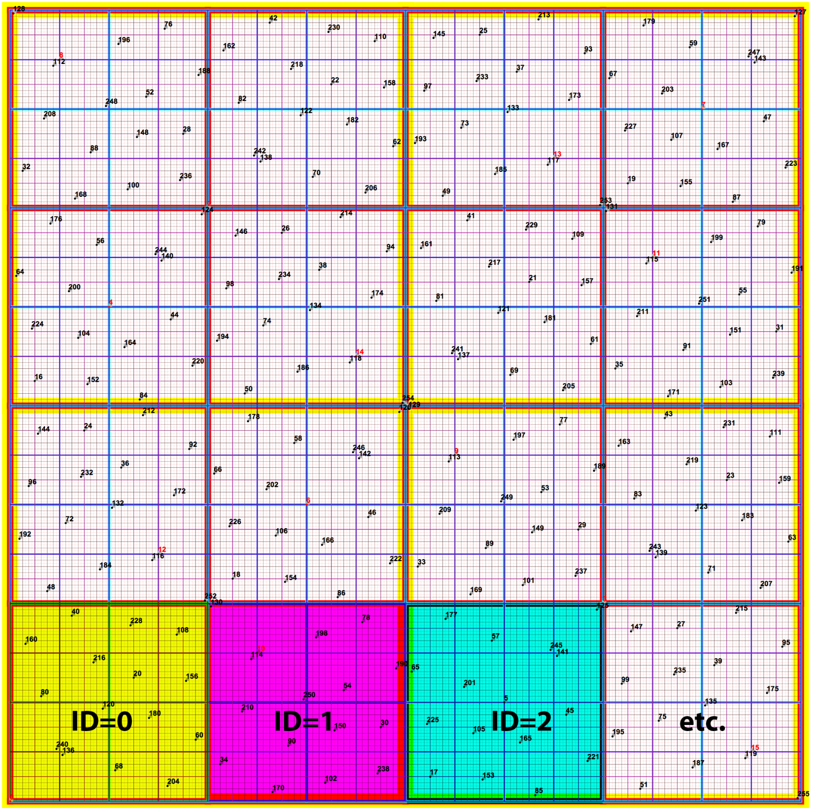

thrive
SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

CONSTRUCTION IN [PERRIER ET AL. 2018]: THE KEY IDEA

Step 2: For each ID, find Owen's permutations which maximize min dist

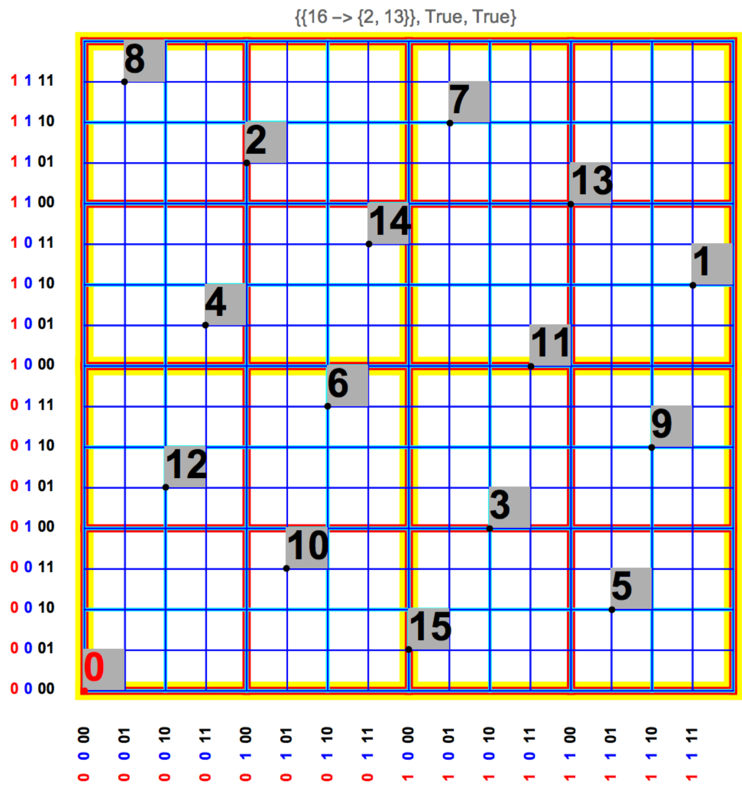

CONSTRUCTION IN [PERRIER ET AL. 2018]: THE KEY IDEA

Step 2: For each ID, find Owen’s permutations which maximize min dist

Dyadic Partitioning is preserved:

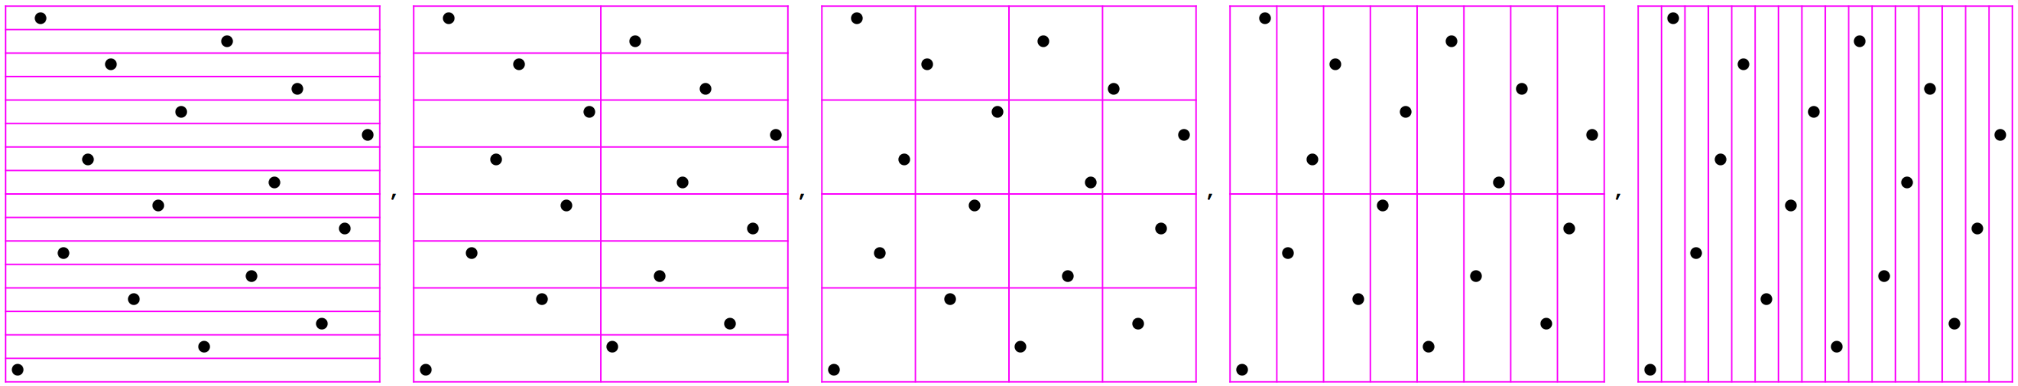

CONSTRUCTION IN [PERRIER ET AL. 2018]: THE KEY IDEA

Step 3: Store permuted patterns in a lookup table

Step 4: In runtime, take the pattern from the lookup table, and LSB bits from Sobol's codes:

GENERATED POINTS (4K): SOBOL

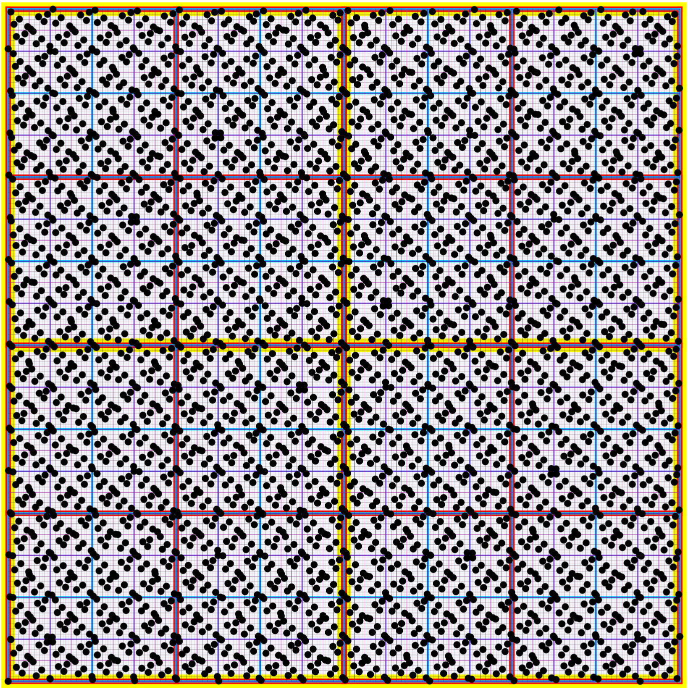

GENERATED POINTS (4K): OWEN'S SCRAMBLING

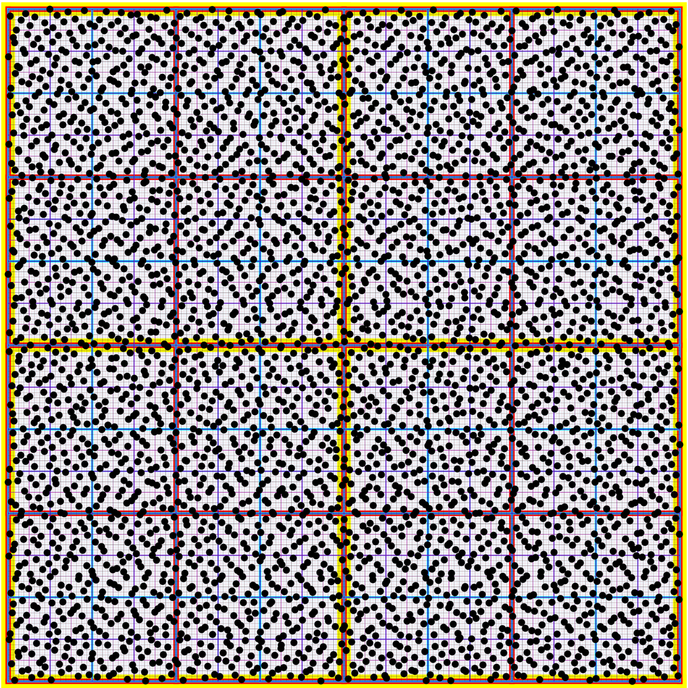

GENERATED POINTS (4K): [PERRIER ET AL. 2018]

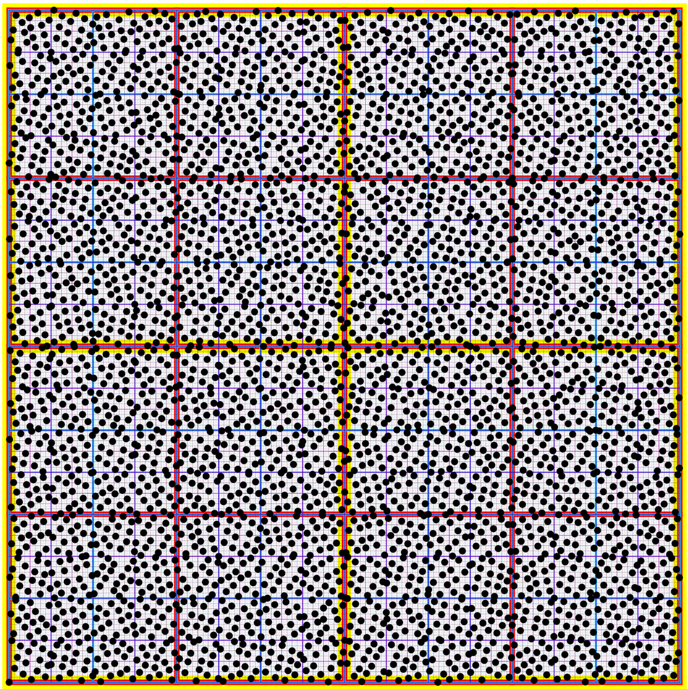

POWER SPECTRUM + RADIAL: OWEN VS. [PERRIER ET AL. 2018]

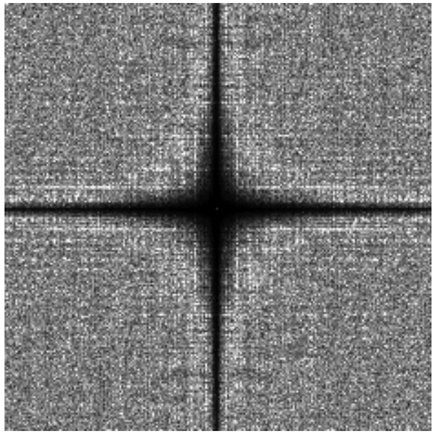

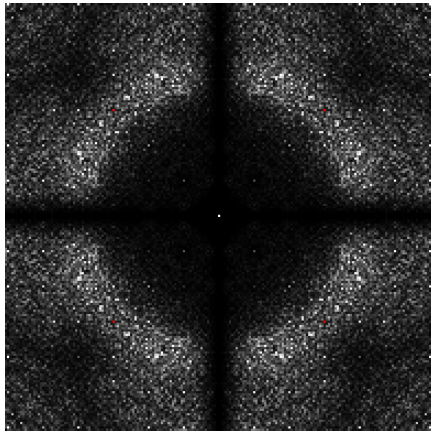

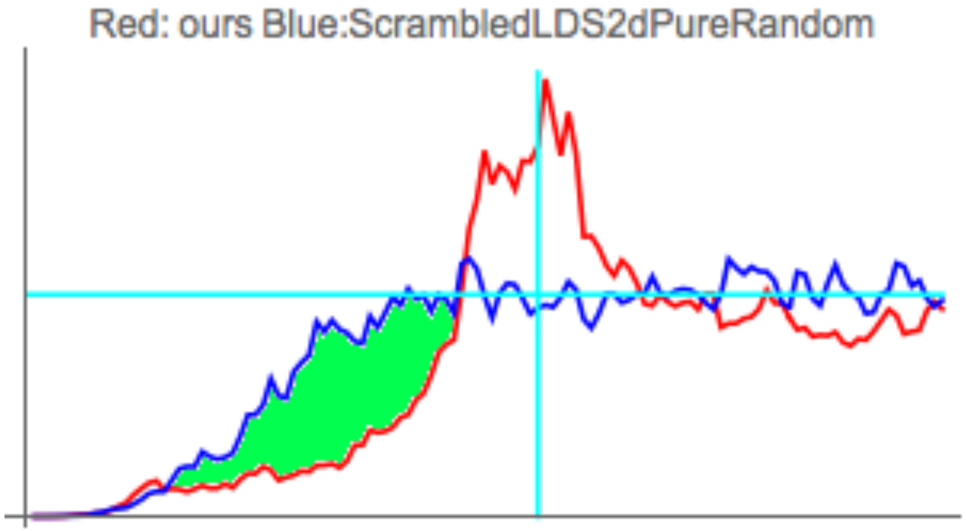

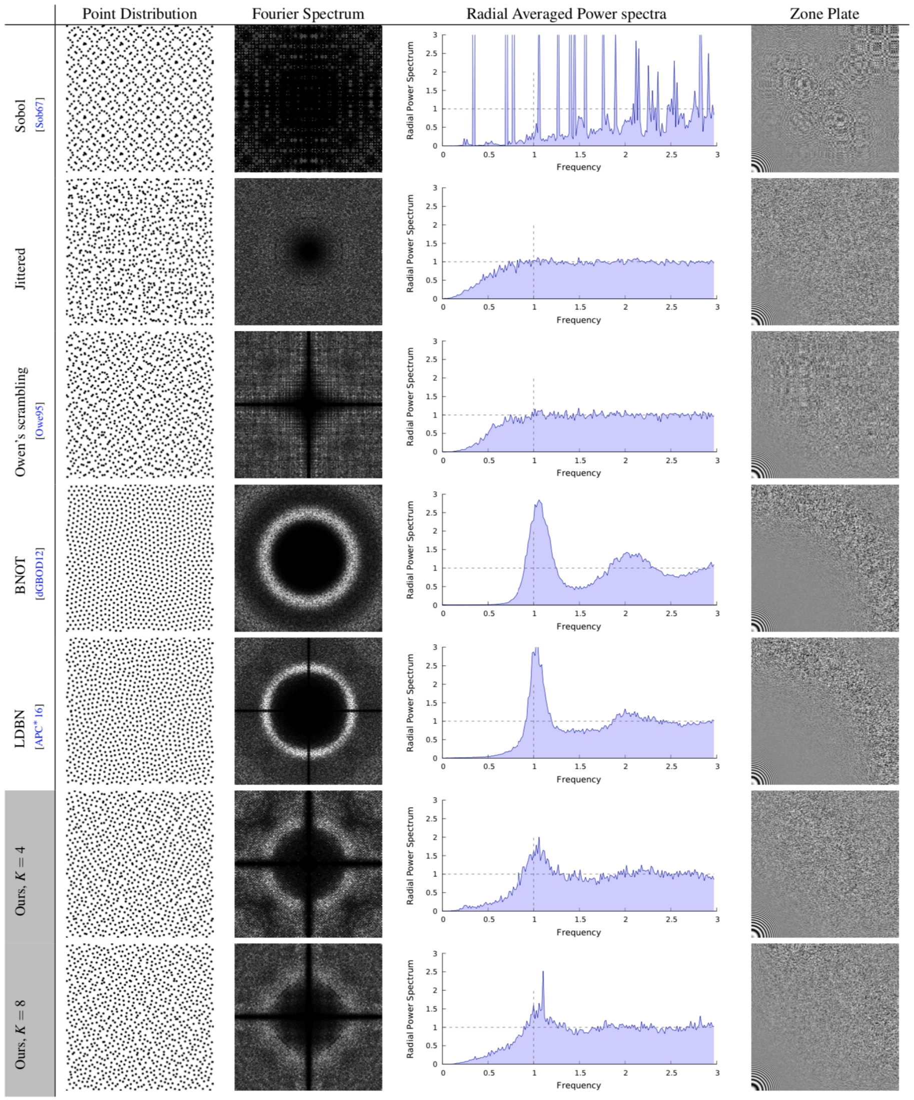

Point Distribution
Fourier Spectrum
Radial Averaged Power spectra
Zone Plate
Sobol [Sob67]
Jittered
Owen's scrambling [Owe95]
BNOT [dGBOD12]
LDBN [APC*16]
Ours, $K = 4$
Ours, $K = 8$
Radial Power Spectrum
Frequency

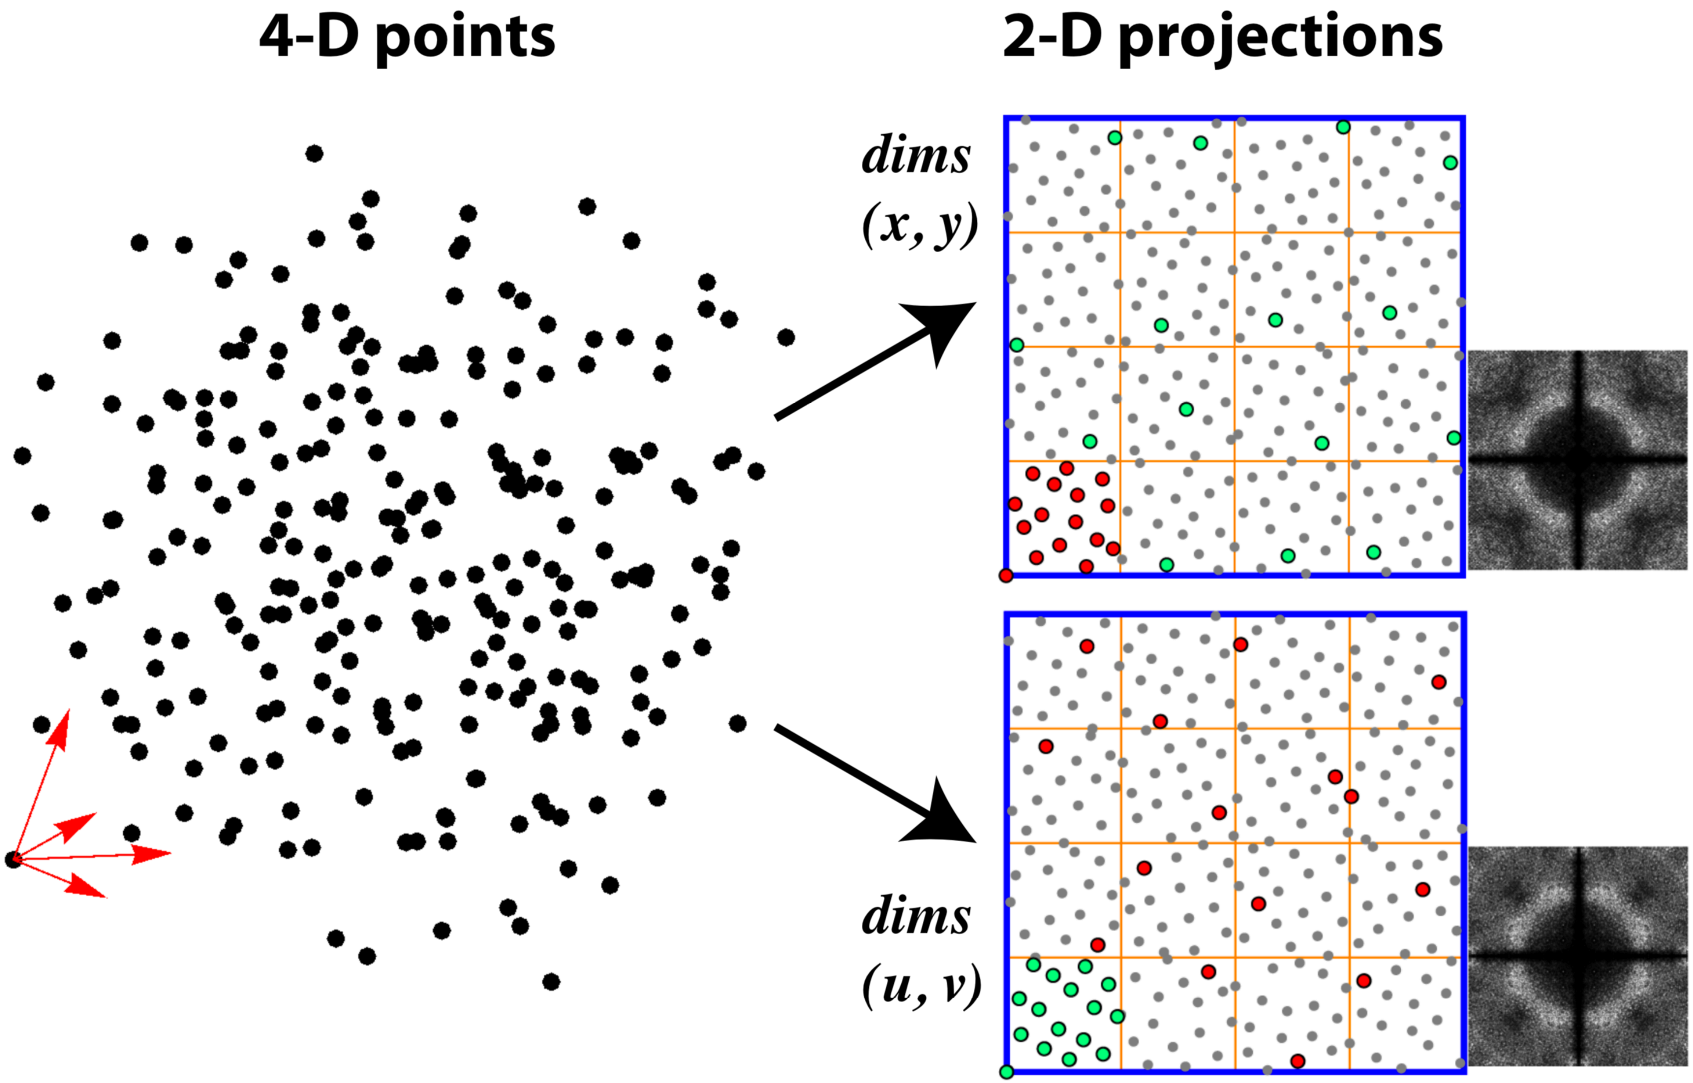

4-D points
2-D projections
dims
(x, y)
dims
(u, v)

[PERRIER ET AL. 2018]: CONCLUSIONS

What we Got:

- 2-D Low-Discrepancy *Sequences* (with Support for Progressive Sampling)
- Improved 2-D Fourier Spectra
- Extendable to 4-D and 6-D
- Supports Adaptive Sampling
- Fast, Low Memory Footprint
- Purely Deterministic, but Can Simulates Quasi-Randomness

Limitations:

- Hard to Get Higher Dimensions
- Power Spectra are “Blueish” rather then “Blue”

A Low-Discrepancy Sampler that Distributes Monte Carlo Errors as a Blue Noise in Screen Space

Eric Heitz[1] Laurent Belcour[1] Victor Ostromoukhov[2] David Coeurjolly[2] Jean-Claude Iehl[2]

[1]Unity Technologies [2]Université de Lyon, CNRS, LIRIS, France

In *ACM SIGGRAPH Talk*, 2019

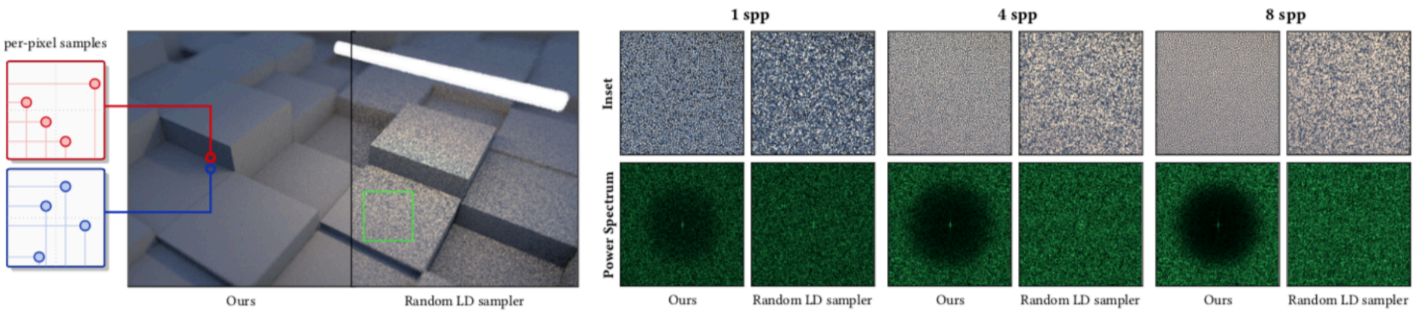

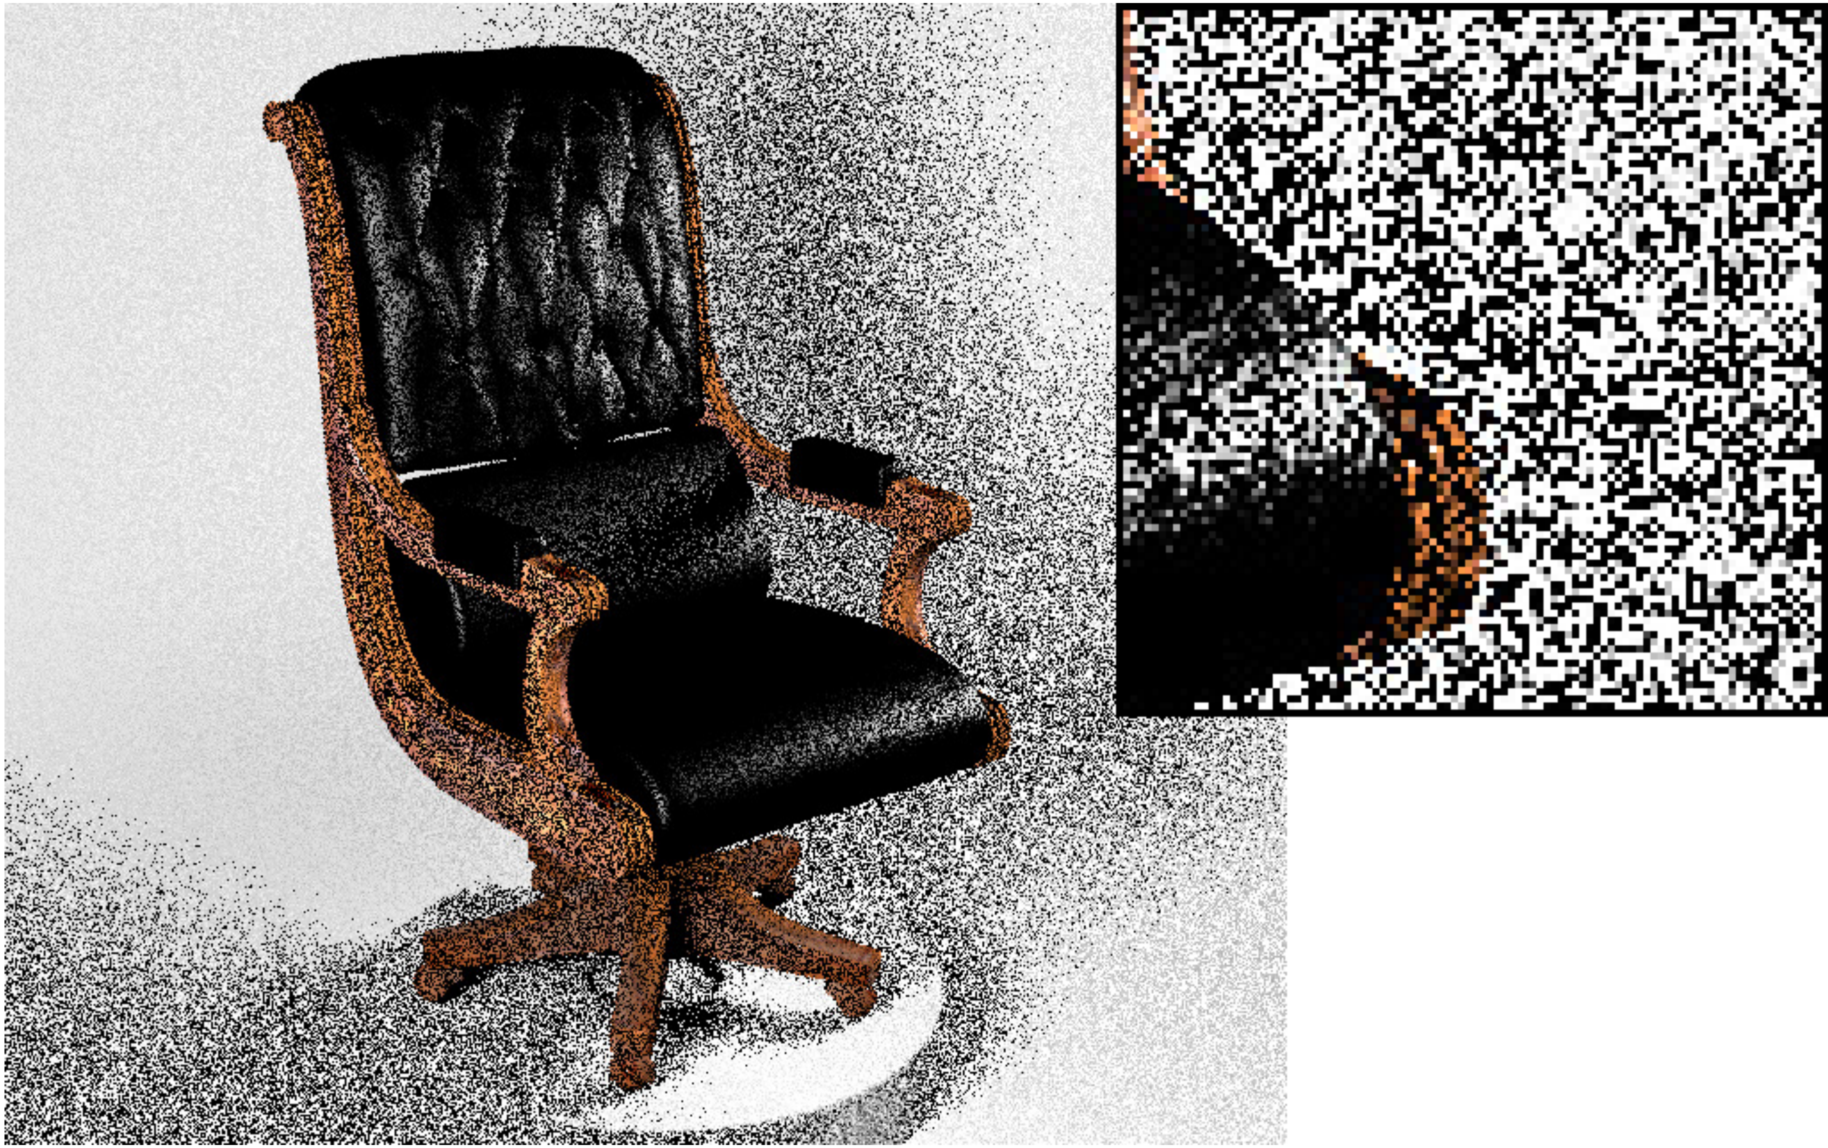

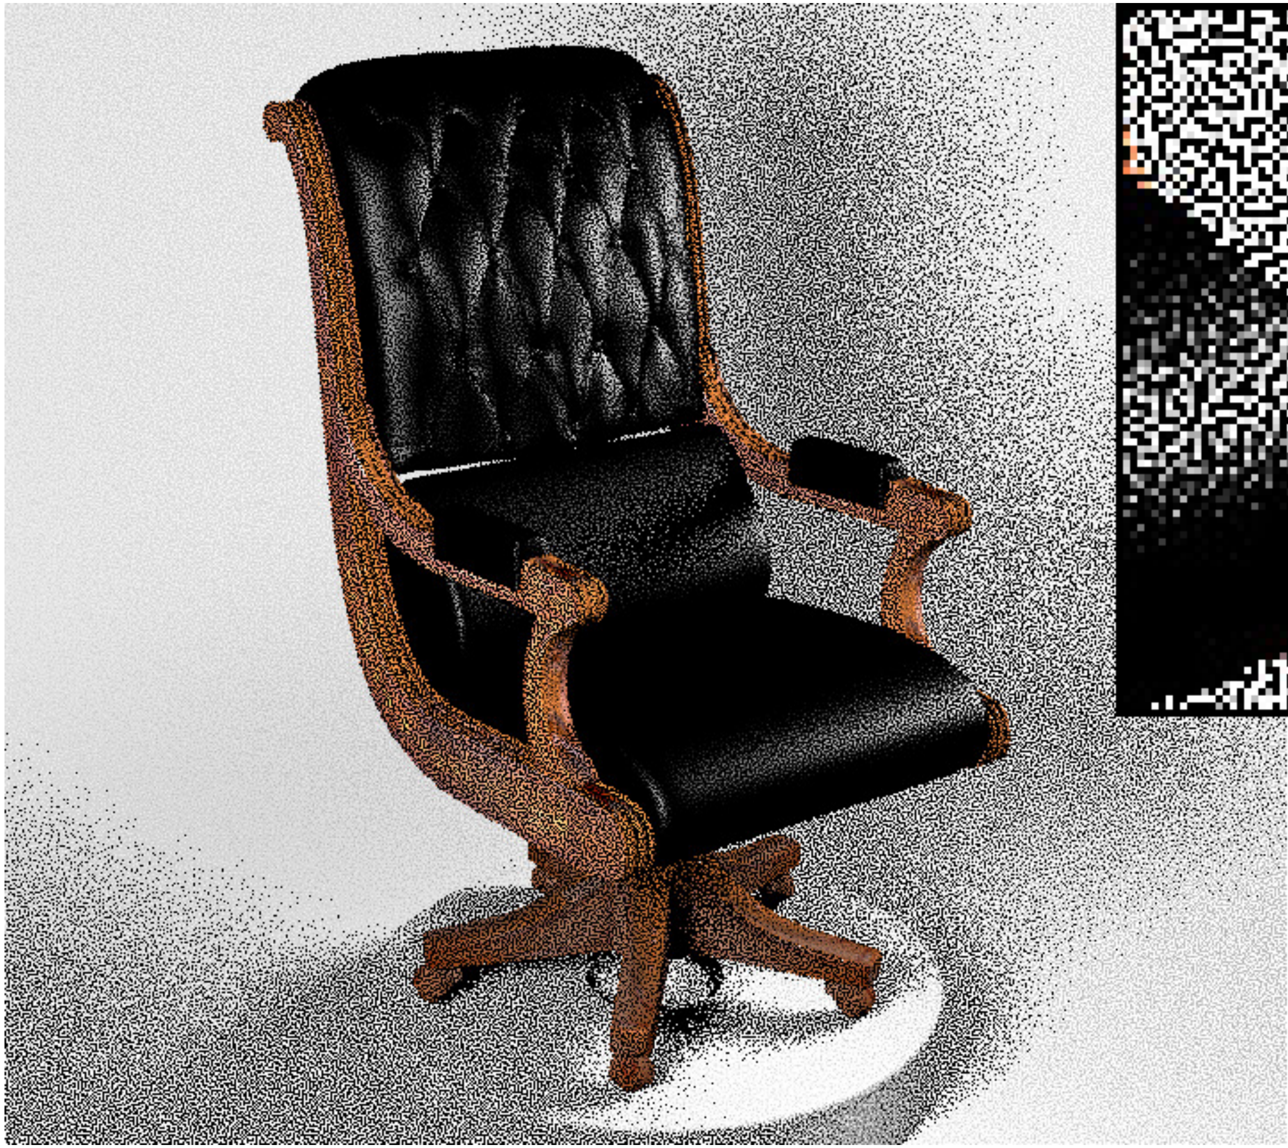


OVERVIEW

- Theoretical foundation for BN sampling
  - Based on Variance Analysis for Monte Carlo Integration, SIGGRAPH 2015
- Some efficient implementation
  - Based on
    - Low-Discrepancy Blue Noise Sampling, SIGGRAPH-ASIA 2016
    - Sequences with Low-Discrepancy Blue-Noise 2-D Projections, EG2018
    - A Low-Discrepancy Sampler that Distributes Monte Carlo Errors as a Blue Noise in Screen Space, SIGGRAPH 2019 Talk
- Open Issues

WHAT IS WRONG WITH DISCREPANCY AS MEASURE OF UNIFORMITY?

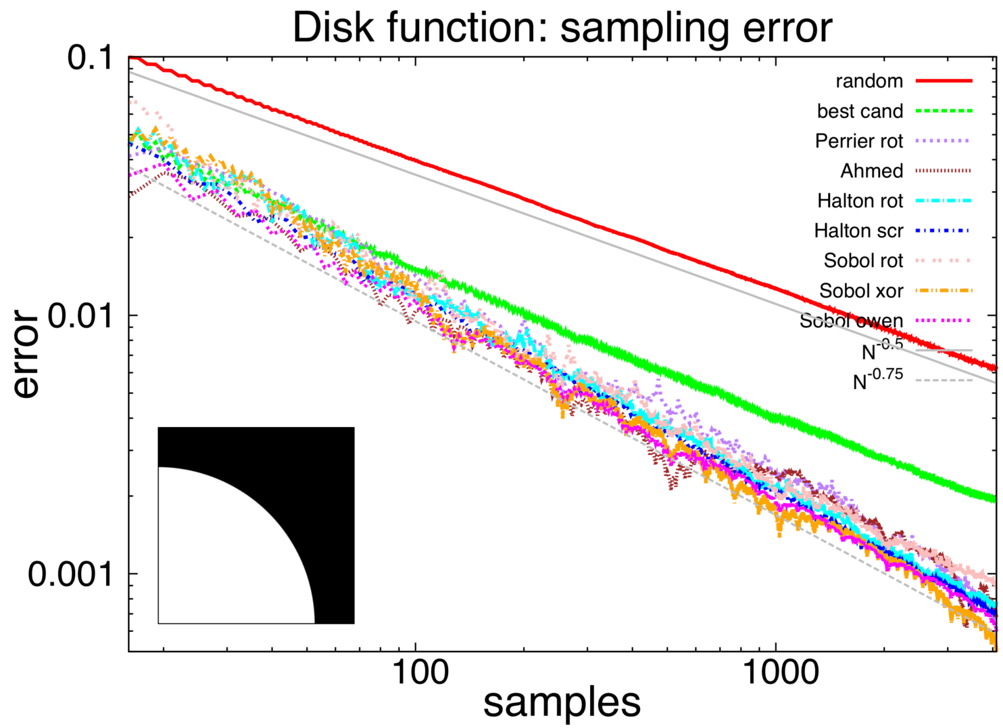

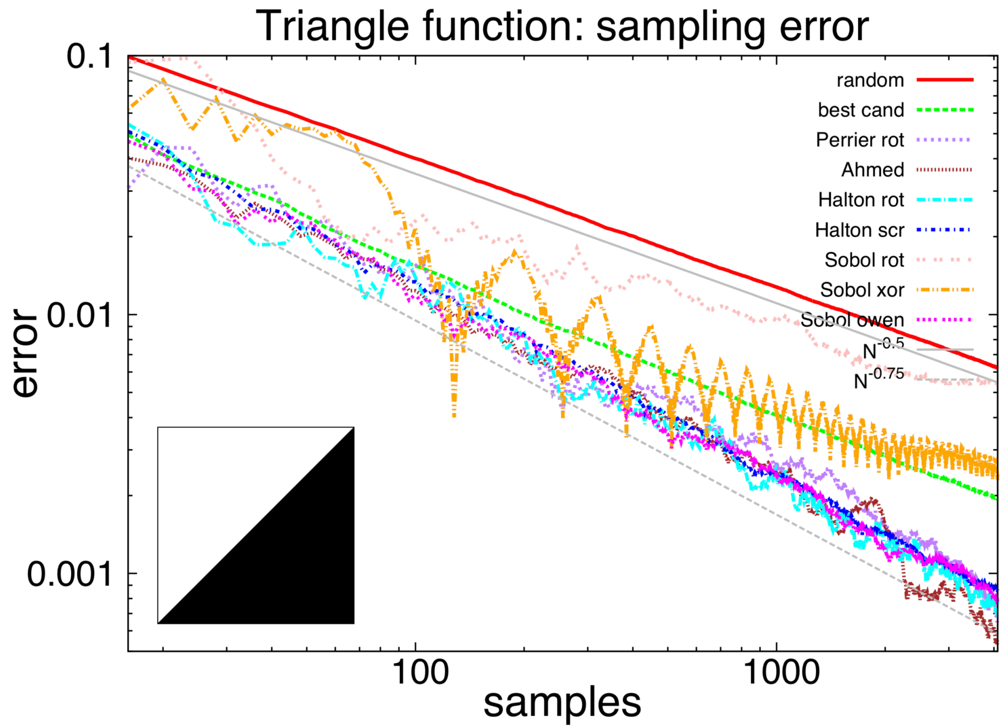

[Christensen et al. 2018]

WHAT IS WRONG WITH DISCREPANCY AS MEASURE OF UNIFORMITY?

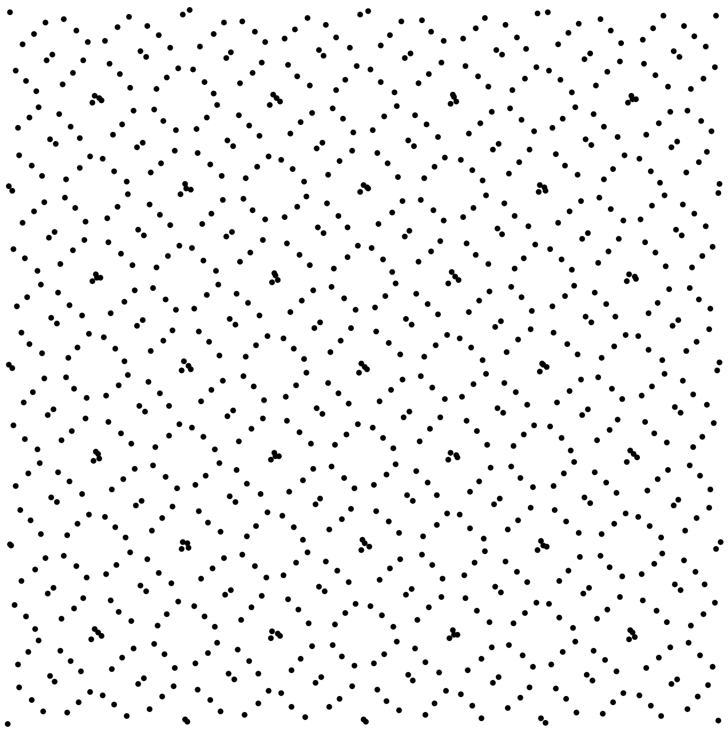

Sobol Points

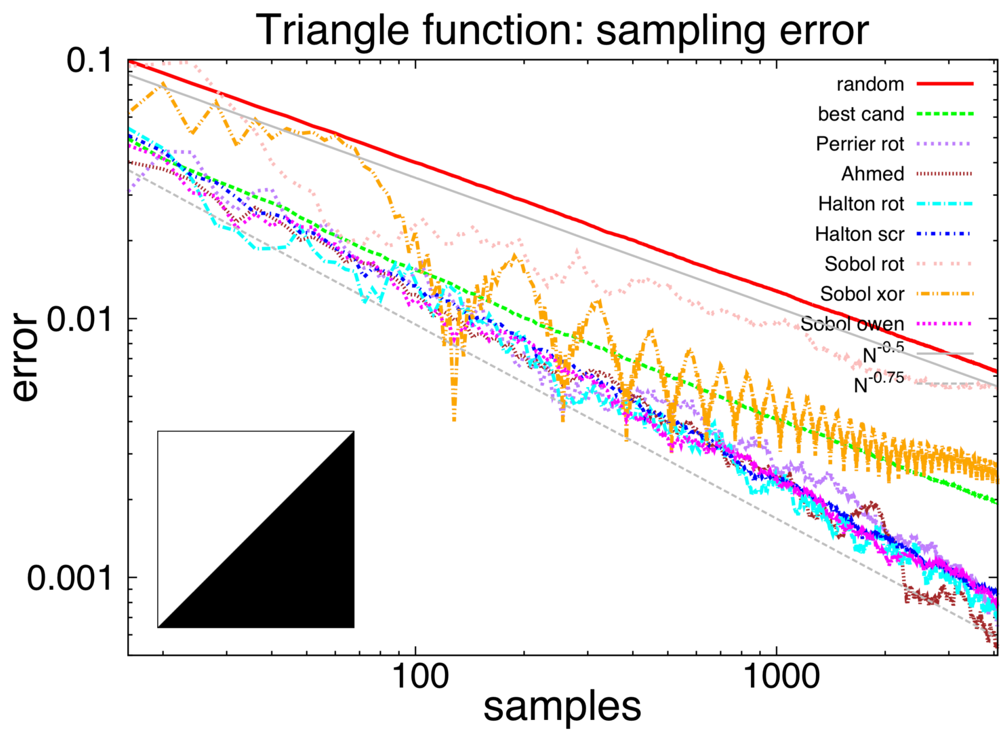

thrive SIGGRAPH2019
LOS ANGELES • 28 JULY – 1 AUGUST

WHAT IS WRONG WITH DISCREPANCY AS MEASURE OF UNIFORMITY?

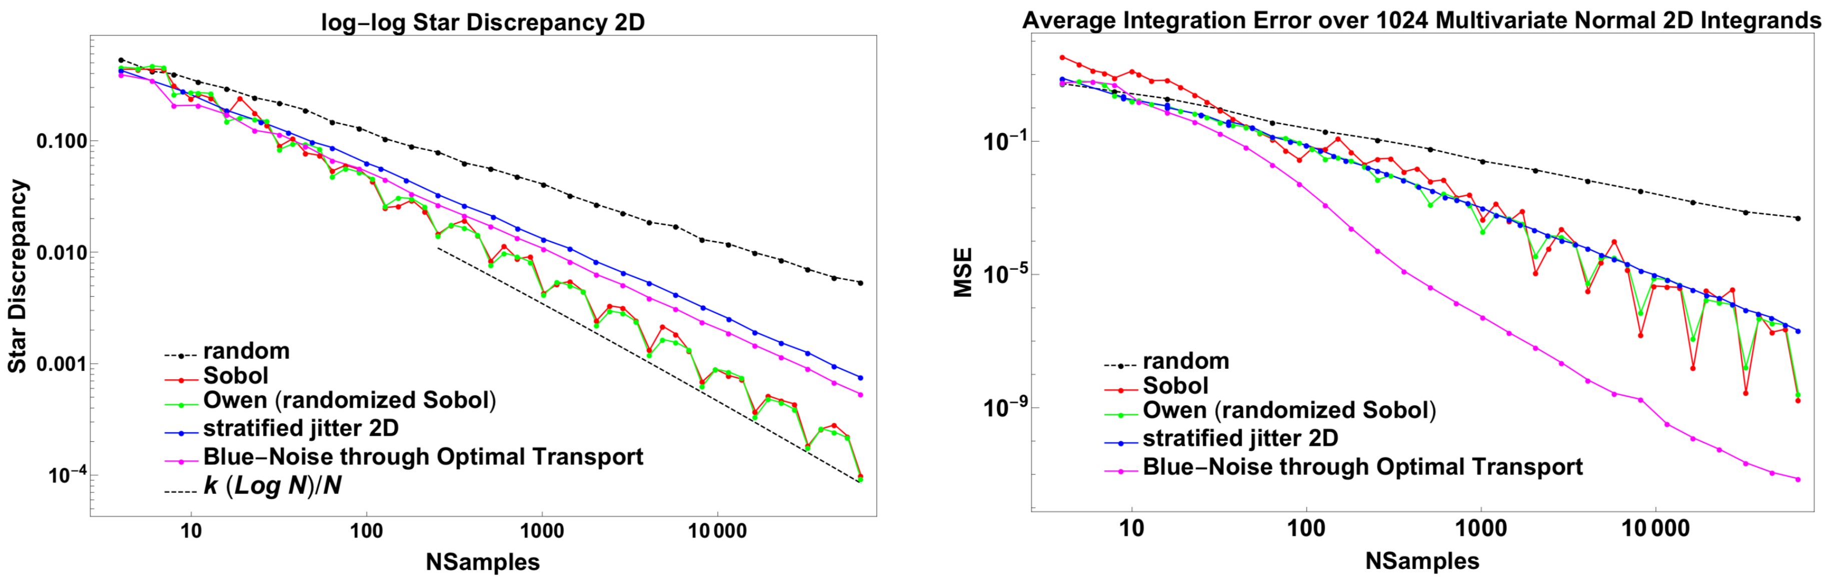

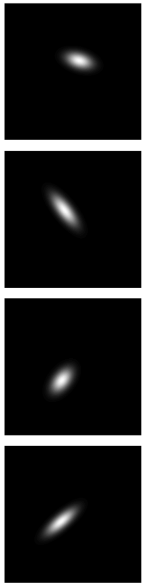

MY FAVORITE SAMPLER?

MY FAVORITE SAMPLER?

The future one!

MY FAVORITE SAMPLER?

The future one!

- Multi-dimensional
- Guarantees convergence to the true integral
- Prevents aliasing
- Minimizes noise
- Guarantees good frequency content of the noise
- Guarantees good computational efficiency

QUESTIONS?