

Lessons Learned and Improvements when Building Screen-Space Samplers with Blue-Noise Error Distribution

Laurent Belcour
Unity Technologies

Eric Heitz
Unity Technologies

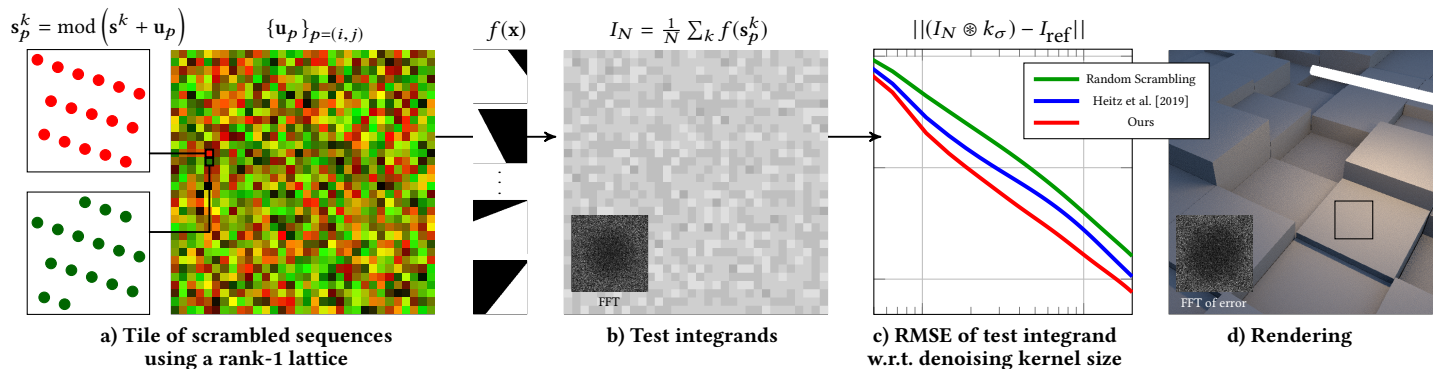


Figure 1: In this work, we optimize a tile of 2D points $\in [0, 1]$ (a-right) to shift the value of a 2D *rank-1 lattice* (a-left). We introduce a *GPU-optimizer* to achieve blue-noise when integrating test integrands (b). We found that our new rank-1 screen-space sampler produces lower error than our previous sampler [Heitz et al. 2019] after *denoising the test integrands tiles* (c). Finally, we apply our sampler using *dimension padding* in a path tracer to achieve blue-noise distribution of the error.

ABSTRACT

Recent work has shown that the error of Monte-Carlo rendering is visually more acceptable when distributed as blue-noise in screen-space. Despite recent efforts, building a screen-space sampler is still an open problem. In this talk, we present the lessons we learned while improving our previous screen-space sampler. Specifically: we advocate for a new criterion to assess the quality of such samplers; we introduce a new screen-space sampler based on rank-1 lattices; we provide a parallel optimization method that is compatible with a GPU implementation and that achieves better quality; we detail the pitfalls of using such samplers in renderers and how to cope with many dimensions; and we provide empirical proofs of the versatility of the optimization process.

CCS CONCEPTS

• Computing methodologies → Rendering.

ACM Reference Format:

Laurent Belcour and Eric Heitz. 2021. Lessons Learned and Improvements when Building Screen-Space Samplers with Blue-Noise Error Distribution. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH '21 Talks)*, August 09-13, 2021. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3450623.3464645>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGGRAPH '21 Talks, August 09-13, 2021, Virtual Event, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8373-8/21/08.
<https://doi.org/10.1145/3450623.3464645>

1 INTRODUCTION

Rendering via Monte Carlo (MC) integration is subject to numerical error that leads to perceivable artifacts such as aliasing and/or noise in the final image. Traditionally, these artifacts are understood to be linked to the *amplitude* of the integration error, and can thus be attenuated via variance-reduction techniques such as importance sampling. Georgiev and Fajardo [2016] showed that the *screen-space distribution* of this error also has a direct impact on the perceptual quality of the image.

They introduced an optimization technique to build a screen-space sampler, a sampler where each pixel of a rendered image uses a correlated sequence of random numbers. Their showed that the perceptual quality of images improved for one-sample-per-pixel renderings. In a previous talk [Heitz et al. 2019], we built on this idea to support the many-samples-per-pixel case and preserve convergence properties. We relied on the Sobol sequence and scrambled it with XORs to generate a different sequence per pixel.

In this talk, we continue our research effort on this topic and provide different improvements to our previous method. First, we found it tedious to judge rendering quality as it relies on perception and might be prone to bias. We developed a *quantifiable criterion* to assess the quality of screen-space sampler. From this criterion, we changed the *sequence and its scrambling mechanism*. Furthermore, we improve the optimization process by *modifying the cost function* (see our supplemental material) and *parallelizing the optimizer* to benefit from GPU's performances. We found how optimizing using discontinuous test integrands could also *benefit other classes of integrands*. Finally, we render images using *dimension padding* for which some care is needed.

2 OUR PREVIOUS METHOD

A screen-space sampler is a 2D map of quasi-random numbers we defined as $\mathcal{S} = \{\mathbf{s}_p^k : p, k\}$, where p is the pixel index $p = (i, j)$ and k is the sample index, \mathbf{s}_p^k is a n -D vector $\in [0, 1]^n$. Following Georgiev and Fajardo [2016], individual sequences of quasi-random numbers are generated by scrambling a *main sequence* \mathbf{s}^k , with a *scrambling function* $\mathcal{K}(\mathbf{s}^k, \mathbf{u}_p)$ where \mathbf{u}_p is a random scramble value. The screen-space sampler is then defined by the main sequence and the 2D *tile* storing the scrambling parameters $\mathcal{U} = \{\mathbf{u}_p\}_p$. For example, in our previous implementation, we used a Owen scrambled Sobol sequence as the *main sequence* and scrambled its values with XORs at each pixel [Heitz et al. 2019]. To distribute the error as a blue-noise, prior to rendering, we optimized the tile by permuting its pixels to reduce the following loss:

$$\mathcal{L} = \sum_p \sum_{q \neq p} \exp\left(-\frac{\|(\mathbf{p} - \mathbf{q})\|^2}{2.1^2}\right) \|\mathcal{I}_p - \mathcal{I}_q\|^2, \quad (1)$$

where $\mathcal{I}_p = \{1/N \sum_k f_i(\mathbf{s}_p^k)\}$ is the 2D map of *test integrands* (random oriented Heavisides). Each pixel contains the vector of Monte-Carlo integrals of random Heavisides function (Fig. 1(b)). Last, the order of the sequence in each pixel is also altered using a different ranking to ensure that blue-noise is achieved for different sample counts. In this talk we will omit this part as we did not improve it.

3 IMPROVING A SCREEN-SPACE SAMPLER

3.1 Padding Dimensions

The first improvement we made to the screen-space sampler was to restrict the main sequence and scramble map to 2D values and padded this sequence to integrate higher dimensions. Since blue-noise is only optimized for pairs of dimensions, dimensions must be consumed by the renderer in pairs. That is, to integrate motion blur (1D), two dimensions must be allocated for this effect and one of them is discarded. Because, to achieve blue-noise, neighboring pixels must use samples dimensions for the same dimensions of the integrand. Thus, a screen-space sampler (either this improvement or the previous one) should not be used with bidirectional path tracing (as done by Ahmed and Wonka [2020]).

3.2 Rank-1 Lattice

After many tests, we opted for a different main sequence and scrambling function. We choose the Rank-1 Lattice sequence: $\mathbf{s}^k = \text{mod}(\Phi(k) \times \mathbf{d}, 1)$, where $\Phi(k)$ is the van der Corput sequence, and \mathbf{d} is the *direction vector* of the sequence. In our implementation, we used the direction vectors of Hickernell et al. [2012]. A natural scrambling mechanism for such a sequence is to offset it modulo one: $\mathcal{K}(\mathbf{s}^k, \mathbf{u}_p) = \text{mod}(\mathbf{s}^k + \mathbf{u}_p, 1)$ (see Fig. 1(a)). Although Sobol is theoretically superior w.r.t. convergence, we found that rank-1 lattices offered the same convergence for rendering integrands while improving upon the blue-noise distribution of rendering error (see our supp. mat.).

3.3 Evaluating a Screen-Space Sampler

In this work, we advocate that the correct method to evaluate a screen-space sampler is *how much it reduces variance after denoising*.

Indeed, a Monte-Carlo renderer image is most of the time denoised. Denoising consists of applying an edge-avoiding, low pass filter in screen-space. Blue-noise error distribution does not contain energy in the low frequency region, this permits to reduce the magnitude of the error after denoising. A denoising kernel is driven by additional buffers (z-buffer, normal buffer, ...) and varies in size. Thus, we apply the criterion on a range of filter size. We use gaussian filters of varying standard deviation (we took $\sigma \in [0.25, 20]$ in our examples). To ensure that our criterion is scene-independent, we leverage again the test integrands and apply the kernel on the tile of QMC estimates (see Figure 1 (c)). In this setting, we found that our new screen-space sampler consistently performed better than our previous XORed Owen sequence. We found the same difference when applying blurring to rendering images (see supp. mat.). This idea bears similarity with the work of Chizhov et al. [2020]. However, they use screen-space blur to drive the optimization, and not to evaluate the final sampler.

3.4 Efficient Optimization

We parallelize the optimization process by testing multiple couples of pixels for swapping at the same time. To avoid collision, we ensure that a pixel in a swap couple is not used in another couple. We evaluate the cost function with no modification and perform the mutation as in the sequential optimisation. Of course, this method can produce a tile with a higher cost function due to concurrent swappings of neighboring pixels. However, if we restrict the number of pixels for swapping to $N/4$, we observe a behavior comparable to simulated annealing: at the beginning of the optimisation, many concurrent swappings occur and non optimal configuration are accepted. However, as the number of mutation increases and the number of valid swapping candidates reduces, the algorithm accepts only a few valid mutations per cycle. We implemented our GPU-optimizer using compute shaders. To generate the list of swap candidates, we precompute a permutation of pixel indices that we store in a linear array. We scramble the pixels indices using an XOR with a different seed per compute pass to avoid testing the same pixels again and computing another permutation. This optimizer allowed us to generate tiles in under a minute.

3.5 Versatility of Test Integrands

Despite being optimized for discontinuous integrands, the screen-space sampler achieves blue-noise distribution of integrand noise even for low frequency integrands. We found that the vector space of test integrands shed light on this behavior (see our supp. mat.).

REFERENCES

- Abdalla GM Ahmed and Peter Wonka. 2020. Screen-space blue-noise diffusion of Monte Carlo sampling error via hierarchical ordering of pixels. *ACM Transactions on Graphics (TOG)* 39, 6 (2020), 1–15.
- Vassilen Chizhov, Iliyan Georgiev, Karol Myszkowski, and Gurprit Singh. 2020. Perceptual error optimization for Monte Carlo rendering. *arXiv preprint arXiv:2012.02344* (2020).
- Iliyan Georgiev and Marcos Fajardo. 2016. Blue-noise dithered sampling. In *ACM SIGGRAPH 2016 Talks*. ACM, 35.
- Eric Heitz, Laurent Belcour, Victor Ostromoukhov, David Coeurjolly, and Jean-Claude Iehl. 2019. A low-discrepancy sampler that distributes Monte Carlo errors as a blue noise in screen space. In *ACM SIGGRAPH 2019 Talks*. 1–2.
- Fred J Hickernell, Peter Kritzer, Frances Y Kuo, and Dirk Nuyens. 2012. Weighted compound integration rules with higher order convergence for all N. *Numerical Algorithms* 59, 2 (2012), 161–183.