Blue-noise Dithered Sampling

Iliyan Georgiev
Solid Angle

Marcos Fajardo
Solid Angle

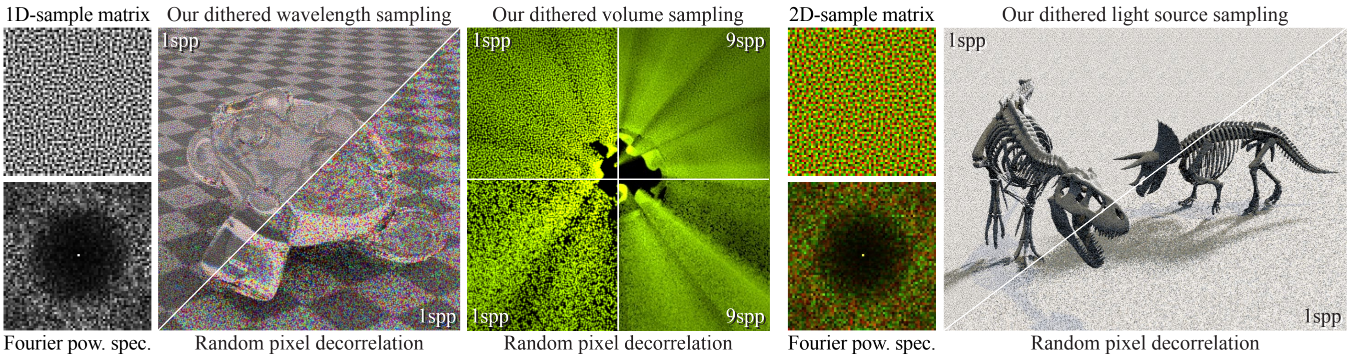

Figure 1: *Our method uses specially constructed blue-noise matrices to correlate the samples between image pixels and minimize the low-frequency content in the output error distribution. Without actually reducing the numerical approximation error, our method can produce images with higher visual fidelity than traditional random pixel decorrelation, especially with a small number of samples per pixel (spp).*

Keywords: Monte Carlo, sampling, blue noise, dithering

Concepts: •**Computing methodologies** → *Rendering;*

The visual fidelity of a Monte Carlo rendered image depends not only on the magnitude of the pixel estimation error but also on its distribution over the image. To this end, state-of-the-art methods use high-quality stratified sampling patterns, which are randomly scrambled or shifted to decorrelate the individual pixel estimates. While random pixel decorrelation yields an eye-pleasing white-noise image error distribution, it is far from perceptually optimal. We show that visual fidelity can be substantially improved by instead *correlating* the samples among pixels in a way that minimizes the low-frequency content in the output noise. Inspired by digital halftoning, our *blue-noise dithered sampling* method can produce significantly more faithful images, especially at low sampling rates.

Blue-noise dithering. In digital halftoning, dithering is the intentional application of noise to randomize the error from quantizing a continuous-tone image [Lau and Arce 2008]. An efficient approach is to threshold the pixels using a blue-noise dither mask, which is a 2D matrix of scalar values arranged such that the matrix Fourier power spectrum is isotropic and devoid of low frequencies. That is, neighboring matrix elements have very different values, and similar values are assigned to elements far away from each other.

Dithered sampling. Our idea is to apply the concept of dithering to correlate pixel estimates in Monte Carlo distribution ray tracing. Given a d-dimensional sampling pattern, we toroidally shift it for every pixel, but rather than choosing the offset randomly, as done traditionally, we look it up in a *blue-noise sample matrix* tiled over the image. The value of every element in such a pre-computed matrix is a d-dimensional vector, and for $d = 1$ the matrix is identical to a dither mask. In this setting, traditional random-offset pixel decorrelation is equivalent to using a white-noise sample matrix.

When using a blue-noise matrix, neighboring pixels evaluate different locations in the sampling domain, yielding a high-frequency noise distribution in the rendered image. Mitchell [1991] also observed that this is a desirable property and proposed a sample distribution optimality condition along with a simple algorithm to approximate it. Our approach aims to efficiently meet this condition and improve over Mitchell's algorithm.

Sample matrix construction. Our sampling matrices are scene-independent and can be built offline. We use simulated annealing to construct a matrix M by minimizing the following energy function:

$$E(M) = \sum_{i \neq j} E(i,j) = \sum_{i \neq j} \exp\left(-\frac{\|c_i - c_j\|^2}{2\sigma^2}\right) \frac{1}{\sqrt{\|v_i - v_j\|^d}},$$

where i and j are matrix elements, c_i and c_j are their 2D (integer) matrix coordinates, v_i and v_j are their d-dimensional vector values, and σ is a parameter (we use an empirically determined $\sigma = 1.1$).

The exponential term above measures the distance between matrix elements, and is borrowed from Ulichney's [1993] void-and-cluster dither mask construction method. The element distance $\|c_i - c_j\|$ is computed by wrapping the boundaries of the matrix toroidally so it can be tiled seamlessly over the image. The second term measures the distance between the element values (vectors), and the two terms work together to space similar values apart in the matrix.

Results. In Figure 1 we show our 1D- and 2D-sample matrices and their corresponding Fourier power spectra. We also show how these matrices compare against random pixel decorrelation for offsetting patterns with 1 and 9 samples. Note that we plot matrices of size 64^2, but we used size 128^2 to render our images. Our approach can produce images with blue-noise error distribution that appear less noisy, even though their numerical error is roughly the same as that of traditional white-noise decorrelation.

References

LAU, D., AND ARCE, G. 2008. *Modern Digital Halftoning, Second Edition*. CRC Press.

MITCHELL, D. P. 1991. Spectrally optimal sampling for distribution ray tracing. *ACM SIGGRAPH Comput. Graph. 25*, 4.

ULICHNEY, R. A. 1993. The void-and-cluster method for dither array generation. *Proc. SPIE 1913*, 332–343.

SIGGRAPH '16, July 24-28, 2016, Anaheim, CA,
ISBN: 978-1-4503-4282-7/16/07
DOI: http://dx.doi.org/10.1145/2897839.2927430