


Line Segment Sampling with Blue-Noise Properties

Xin Sun¹ Kun Zhou² Jie Guo³ Guofu Xie^{4,5} Jingui Pan³ Wencheng Wang⁴ Baining Guo¹

¹Microsoft Research Asia ²State Key Lab of CAD & CG, Zhejiang University

³State Key Lab for Novel Software Technology, Nanjing University

⁴State Key Laboratory of Computer Science, ISCAS ⁵GUCAS & UCAS



(a) 1 spp in 5D. (b) 64 spp in 5D, (c) Refocused (a). (d) Refocused (b). (e) 1 spp in 2D. (f) Reference. drop to 1 spp in 2D.

Figure 1: Comparison of temporal light field reconstruction [Lehtinen et al. 2011] with different schemes for low discrepancy sampling. The samples in (a) have blue-noise properties in 5D light field space, and the properties are greatly diminished in 2D image space. The samples in (b) and (e) maintain blue-noise properties in 2D image space, which greatly enhances the rendering quality. The quality of refocus rendering in (c) and (d) is consistent to the blue-noise properties in 2D image space. Please refer to Section 4.4 for a detailed description.

Abstract

Line segment sampling has recently been adopted in many rendering algorithms for better handling of a wide range of effects such as motion blur, defocus blur and scattering media. A question naturally raised is how to generate line segment samples with good properties that can effectively reduce variance and aliasing artifacts observed in the rendering results. This paper studies this problem and presents a frequency analysis of line segment sampling. The analysis shows that the frequency content of a line segment sample is equivalent to the weighted frequency content of a point sample. The weight introduces anisotropy that smoothly changes among point samples, line segment samples and line samples according to the lengths of the samples. Line segment sampling thus makes it possible to achieve a balance between noise (point sampling) and aliasing (line sampling) under the same sampling rate. Based on the analysis, we propose a line segment sampling scheme to preserve blue-noise properties of samples which can significantly reduce noise and aliasing artifacts in reconstruction results. We demonstrate that our sampling scheme improves the quality of depth-of-field rendering, motion blur rendering, and temporal light field reconstruction.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Antialiasing I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture;

Keywords: Sampling, aliasing, blue noise, Poisson disk sampling

Links:  DL  PDF

1 Introduction

Recently, line segment sampling has demonstrated substantially higher efficiency than point sampling in terms of both storage and computation in many rendering applications, such as motion blur [Gribel et al. 2011], depth of field [Tzeng et al. 2012] and scattering media [Sun et al. 2010; Novák et al. 2012b]. A line segment samples a function continuously at an infinite number of points along the line segment. Evaluation within the range of the sample can be analytically or semi-analytically performed, as each point along the line is naturally parameterized. Although the computational cost of a line segment sample is often more expensive than that of a point sample, the sampling rate of line segments can be much lower because the evaluation result of a line segment sample is equivalent to that of numerous discretized point samples.

A question consequently raised is how to generate line segment samples with good properties. The research on sampling [Cook 1986] has proven that samples with blue-noise properties perform excellently in applications because of their low discrepancy and randomness. The low discrepancy reduces variance while the randomness removes aliasing. Although many algorithms have been proposed over the past twenty years for generating blue-noise point samples, to our knowledge there does not exist much work on analysis and sampling schemes for line segment samples. Simple uniform sampling, random sampling or specific sampling schemes extended from previous point sampling methods have been tried in existing works, which sometimes work well but sometimes not. It is thus necessary and important to perform an analysis on the properties of line segment samples and design appropriate sampling schemes based on the analysis.

This paper conducts a frequency analysis of line segment sampling. One important conclusion from this analysis is that the frequency content of a line segment sample is equivalent to the weighted frequency content of a point sample. The weight introduces anisotropy that smoothly changes among point samples, line segment samples and line samples according to the lengths of the samples. Line segment sampling thus makes it possible to achieve a balance between noise (point sampling) and aliasing (line sampling) under the same sampling rate. These conclusions are drawn in the 2D space, but can be generalized to high dimensional spaces and samples of arbitrary non-point shapes.

Based on the frequency analysis, we propose a line segment sampling scheme that best preserves the blue-noise properties of samples and minimizes anisotropy. We experiment with the sampling scheme in several applications, namely image reconstruction, motion blur rendering, depth of field rendering, and temporal light field reconstruction. The line segment samples generated by our scheme perform better with much less variance and aliasing than other sampling schemes such as uniform sampling, random sampling or some straightforward extensions of blue-noise point sampling.

In the rest of the paper, we first review some related work. In Section 3, we present the frequency analysis and sampling schemes for line segments, including line sampling in Section 3.1, line segment sampling in Section 3.2, line segment jittering in Section 3.3, and extensions to high dimensional spaces in Section 3.4 and other non-point samples with arbitrary shapes in Section 3.5. We demonstrate applications of our sampling method in Section 4 and conclude the paper in Section 5.

2 Related Work

Blue-noise properties of Poisson disk sampling are critical to a wide range of computer graphics applications including texture mapping [Lagae and Dutré 2005], geometry processing [Nehab and Shilane 2004; Öztireli et al. 2010] and photorealistic rendering [Jensen and Buhler 2002; Cheslack-Postava et al. 2008]. The importance of sampling with blue-noise properties was highlighted in the pioneering work of Lloyd [1983] and Cook [1986], who also introduced the classic schemes of relaxation and dart throwing for Poisson disk sampling. Following the same idea of Poisson disk sampling that each sample should be away from its closest neighbor as much as possible while maintaining a constant sampling rate, many works have focused on efficient blue-noise sampling on the fly [Dunbar and Humphreys 2006; Bridson 2007], or via precomputation [Cohen et al. 2003; Ostromoukhov et al. 2004; Lagae and Dutré 2005; Kopf et al. 2006; Ostromoukhov 2007], spatial hierarchies [Mitchell 1987; McCool and Fiume 1992; White et al. 2007] and parallelism [Wei 2008; Bowers et al. 2010; Ebeida et al. 2011; Ebeida et al. 2012]. We refer readers to the excellent survey by Lagae and Dutré [2008]. There are various directions of research related to blue-noise point sampling, including adaptive sampling schemes and anisotropic reconstruction [Hachisuka et al. 2008], and statistical mechanics to enhance blue noise spectral properties [Fattal 2011]. Blue-noise point sampling can be generalized to a manifold with arbitrary dimensions [Öztireli et al. 2010].

Anisotropic blue-noise sampling is also useful in many applications such as importance sampling [Ostromoukhov et al. 2004], stippling [Balzer et al. 2009] and half-toning [Pang et al. 2008]. The algorithm introduced by Li et al. [2010] can introduce orientation and even place non-point shapes at sampled positions to enrich visual effects, but the samples are still evaluated by points, which is fundamentally different from our line segment sampling. Another form of heterogeneous blue-noise sampling is multi-class blue-noise sampling [Wei 2010]. The samples are divided into a number of classes, and each class of samples is blue-noise sampled while their union also has blue-noise properties.

The blue-noise properties of point samples are often achieved by preventing any pair of points from being placed too closely together. Extending this to continuous samples like line segments is difficult because every point on a continuous line segment is counted. Our work is inspired by the quantitative analysis of Poisson disk sampling [Wei and Wang 2011; Zhou et al. 2012], whose explicit and analytic frequency formulation opens the door to analyzing the blue-noise properties of continuous samples. The concurrent work by Öztireli and Gross [2012] analyzes blue-noise properties based on pair correlations to incorporate an arbitrary metric space.

Line segment sampling is considered to have advantages in anti-aliasing due to the analytical evaluation along its direction [Jones and Perry 2000]. Recently, it has been incorporated into stochastic rasterization [Akenine-Möller et al. 2007], which leads to high quality and high performance in motion blur rendering [Gribel et al. 2010; Gribel et al. 2011] with the accompanying semi-analytical visibility test. The rendering of defocus blur also benefits from a similar idea by placing line samples on the camera lens [Tzeng et al. 2012]. Global illumination rendering and volumetric scattering rendering are notorious for their computational expense. To address this issue, Havran et al. [2005] introduced line segments to sample the lighting. Later, Jarosz et al. [2008] used line segment samples to represent viewing rays. In recent years, many efficient rendering algorithms have benefited from line segments for sampling both lighting and viewing rays [Sun et al. 2010; Jarosz et al. 2011a; Jarosz et al. 2011b; Novák et al. 2012b; Novák et al. 2012a]. Researchers have also begun to use line samples for hair rendering because of their efficiency in the computationally intensive and highly detailed visibility test [Barringer et al. 2012]. In these algorithms, the computational cost of a single line segment sample is usually much greater than that for a point sample. But its evaluation result is equivalent to that of numerous discretized point samples, which is why line segment sampling has significant advantages over point sampling in general.

Frequency analysis of light transport has yielded much insight on how to control sampling rates adaptively for rendering [Durand et al. 2005; Ramamoorthi et al. 2007; Soler et al. 2009]. Another interesting result from frequency analysis is to perform sheared reconstruction from high dimensional samples with a sheared filter instead of isotropic filters, and this has proven to be very efficient for rendering distribution effects such as motion blur and depth of field [Egan et al. 2009; Lehtinen et al. 2011; Egan et al. 2011; Mehta et al. 2012]. However, it is not clear what is a good sampling scheme for sheared reconstruction as previous sampling algorithms are built on the assumption that the reconstruction kernel is isotropic. Our analysis shows that sheared reconstruction of point samples is equivalent to the traditional isotropic reconstruction of samples with non-point shapes and thus solves the sampling problem of sheared reconstruction from another perspective.

3 Line Segment Sampling

Our goal is to generate line segment samples with blue-noise properties. This means that the low frequency content of the samples should be near zero while the high frequency content should be like white noise. In this section, we first analyze the frequency content of line samples and line segment samples in $2D$ space. Then we extend the analysis to high dimensional spaces and non-point samples with arbitrary shapes.

To evaluate different sampling schemes, we use image reconstruction as an example application. The benchmark image is shown in Fig. 2 (d) accompanied by the reconstruction results of different point sampling schemes. Some preliminaries of blue-noise properties are introduced in the appendix and in Fig. 2. We would like to generate line segment samples with frequency content similar to Fig. 2 (c) but without the limitations of uniform sampling in Fig. 2 (a) and random sampling in Fig. 2 (b). All images in this paper are of resolution 512×512 unless mentioned otherwise.

3.1 Line Samples of Infinite Length

We first consider line samples of infinite length, which is a special case of line segment samples. A line sample traverses the




Figure 2: Point sampling and reconstruction. For subfigures (a) (b) (c): top-left is the samples, top-right is the frequency content, bottom-left is the power spectrum and anisotropy in blue and red respectively, and bottom-right is the reconstructed image.

whole domain instead of just a single point, producing very different frequency content from a point sample. In the following, we will describe the frequency content of a single line sample.

Given a line in $2D$ space, each point $\mathbf{x} = (x, y)^T$ in it satisfies

$$\mathbf{L}_{P,Q,R}(\mathbf{x}) = \{\mathbf{x} | Px + Qy + R = 0\}. \quad (1)$$

Without loss of generality, we assume that the coefficients are normalized, i.e., $P^2 + Q^2 = 1$. The direction of the line is $\mathbf{d}_{P,Q}^\perp = (-Q, P)$ and its perpendicular direction is $\mathbf{d}_{P,Q}^\parallel = (P, Q)$. R is the signed distance from the origin to the line.

A line sample can be represented as a line impulse:

$$\delta_{P,Q,R}(\mathbf{x}) = \delta(Px + Qy + R) = \delta(\mathbf{d}_{P,Q}^\perp \cdot \mathbf{x} + R), \quad (2)$$

with δ denoting the impulse function. It can easily be derived that the frequency content of Eq. (2) is

$$\begin{aligned} \mathcal{F}(\delta_{P,Q,R}(\mathbf{x})) &= e^{2\pi i R(\mathbf{d}_{P,Q}^\perp \cdot \boldsymbol{\omega})} \delta_{-Q,P,0}(\boldsymbol{\omega}) \\ &= e^{2\pi i R \boldsymbol{\omega}^\perp} \delta(\boldsymbol{\omega}^\parallel), \end{aligned} \quad (3)$$

where \mathcal{F} represents the Fourier transform operator, $\boldsymbol{\omega} = (\omega_x, \omega_y)^T$ represents the frequency domain, $\boldsymbol{\omega}^\perp = \mathbf{d}_{P,Q}^\perp \boldsymbol{\omega}$ and $\boldsymbol{\omega}^\parallel = \mathbf{d}_{P,Q}^\parallel \boldsymbol{\omega}$.

A line is a union of all points the line covers. So Eq. (3) is determined by the integral of point samples along the line. As shown in Eq. (3), the frequency content of a line sample is the product of two parts. The form of the first part, $e^{2\pi i R \boldsymbol{\omega}^\perp}$, is very similar to the frequency content of a point sample (Eq. (27)). The second part, $\delta(\boldsymbol{\omega}^\parallel)$, results from the line (Eq. (1)) and determines the magnitude of the spectrum throughout the frequency domain.

Frequency content of line impulses. Because of the line impulse $\delta(\boldsymbol{\omega}^\parallel)$ in Eq. (3), the frequency content of a line sample in the spatial domain also lies on a line in the frequency domain, perpendicular to the original line (Eq. (1)) and passing through the origin. So in the frequency domain, line samples with different directions can only meet at the origin, while parallel line samples will accumulate on the same line impulse. In other words, the frequency content at any point $\boldsymbol{\omega}$ (except the origin) is only determined by line samples with the same direction $\mathbf{d}_{P,Q}^\parallel$ that satisfies $\boldsymbol{\omega}^\parallel = 0$.

According to Eq. (3), a line sample's frequency value along the line impulse $\delta_{-Q,P,0}(\boldsymbol{\omega})$ is a Fourier basis function with parameter $\boldsymbol{\omega}^\perp$ and frequency R . According to the frequency analysis for $1D$ point sampling in the appendix¹, we can draw the conclusion that *the frequency content of parallel line samples is equivalent to that of point samples in the $1D$ space with parameter $\boldsymbol{\omega}^\perp$, and the position of a point sample is given by $-R$.*

Consequently, the frequency properties of parallel line samples are determined by the distribution of $-R$ only. If $-R$ is sampled uniformly, there are periodic impulses along $\delta_{-Q,P,0}(\boldsymbol{\omega})$, producing aliasing artifacts as in Fig. 3 (a). The randomly sampled $-R$ generates white noise and large variations in reconstruction as shown in Fig. 3 (b). To achieve blue-noise properties along the line impulse $\delta_{-Q,P,0}(\boldsymbol{\omega})$, we need to generate the distribution of $-R$ as $1D$ Poisson disk samples as shown in Fig. 3 (c). These observations are in accord with common knowledge about point sampling.

Anisotropy. A problem with line sampling is its strong anisotropy. The frequency content of parallel line samples lies on a single line impulse. The low frequency noise in the reconstruction also lies on the same line impulse and remains zero everywhere away from the line, which results in angular aliasing as shown in the right image of Fig. 3 (c). The best way to eliminate angular aliasing is to distribute the low frequency noise along all directions in the frequency domain, which is unfortunately difficult because a finite number of line samples cannot continuously cover all directions. An approximate solution is to sample a few groups of lines, each of which has a different direction. Then the angular aliasing can be greatly reduced because the low frequency noise will be distributed among more than one direction in the power spectrum as shown in Fig. 3 (f). Sampling of multiple directions can help uniform sampling (Fig. 3 (d)) and random sampling (Fig. 3 (e)) of $-R$.

Line sampling scheme. We note that a key factor for line sampling with blue-noise properties is perfect directional alignment. We first sample line directions $\{P_i, Q_i\}$. For each sampled direction, we generate a group of parallel line samples. The sampling rate of each group should be high enough to achieve blue-noise properties along the corresponding line impulse $\delta_{-Q_i,P_i,0}(\boldsymbol{\omega})$. Within a group of line samples, their direction should be exactly the same without any jittering or perturbation. Otherwise the frequency content becomes close to white noise (Fig. 3 (g)). Multiple direction sampling would also not help in this case (Fig. 3 (h)), since its frequency properties are similar to randomly sampled lines (Fig. 3 (i)) with very minor improvements.

¹The readers can refer to the appendix and the supplementary material for preliminaries on point sampling, including aliasing artifacts and noise, and the frequency content of a point sample. Conclusions about Poisson disk sampling are also presented, such as the minimum distance r (Poisson disk distance) between two point samples and the corresponding Nyquist limit $\nu_N = \frac{1}{r}$, the ideal Poisson disk sampling (Eq. (28)) and the resulting quantitative measurements (Eq. (29)).

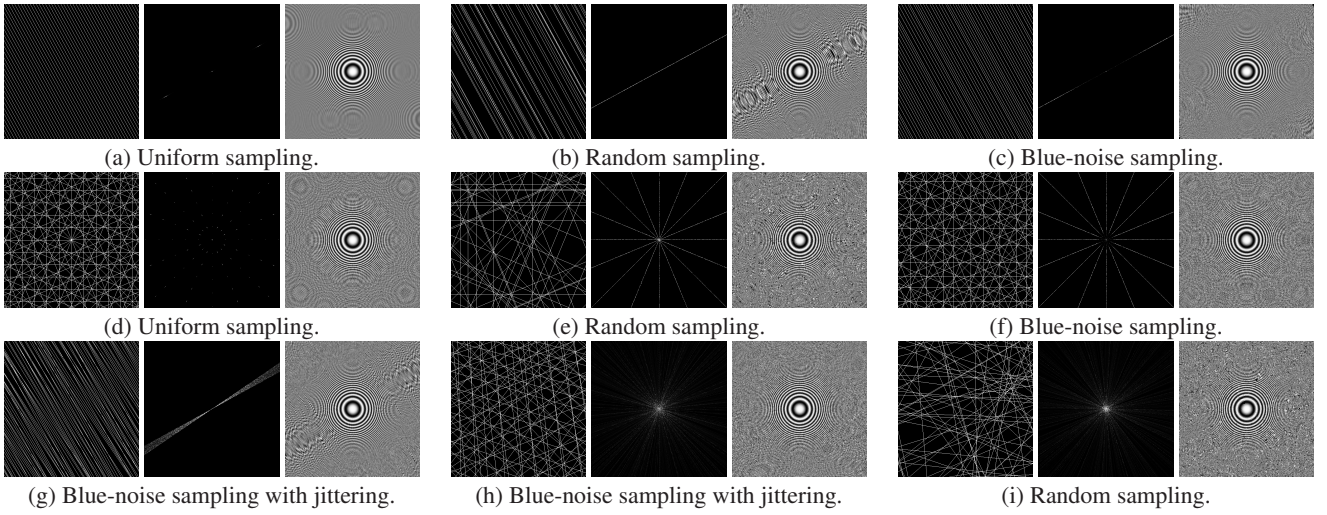


Figure 3: Different line sampling schemes with respect to R . Top row and (g) sample one direction. Middle row and (h) sample eight evenly distributed directions. (i) randomly sampled directions. For each scheme, 800 samples are used and we show the samples, power spectrum and reconstructed image from left to right. For blue-noise sampling, we set the Poisson disk radius r to 0.001 and 0.005 for single and multiple directions, respectively. The sample images show only about 10% of the samples for better visualization.

3.2 Line Segment Samples

A line segment sample is a part of a line sample. So the frequency content of a line sample introduced in Section 3.1 should be modified by the position and length of the line segment.

Consider a line segment sample with a center position $\mathbf{x}_c = (x_c, y_c)^T$ and finite length l :

$$\delta_{P,Q,R,\mathbf{x}_c,l}(\mathbf{x}) = \delta_{P,Q,R}(\mathbf{x}) \text{H} \left(\frac{l}{2} - \left| \mathbf{d}_{P,Q}^{\parallel}(\mathbf{x} - \mathbf{x}_c) \right| \right) \quad (4)$$

where H is the Heaviside step function. The corresponding frequency content of Eq. (4) exists throughout the frequency domain and is not an impulse. It can be thought of as the weighted frequency content of point samples according to Eq. (27):

$$\begin{aligned} \mathcal{F}(\delta_{P,Q,R,\mathbf{x}_c,l}(\mathbf{x})) &= \mathcal{F}(\delta(\mathbf{x} - \mathbf{x}_c)) w(l, \omega^{\parallel}) \\ w(l, \omega^{\parallel}) &= l \text{sinc}(l\omega^{\parallel}) \end{aligned} \quad (5)$$

where $\text{sinc}(x) = \sin(\pi x)/(\pi x)$.

Similar to Eq. (3), there are two parts of Eq. (5). The first part, $\mathcal{F}(\delta(\mathbf{x} - \mathbf{x}_c))$, is exactly the frequency content of a point sample. The key difference comes from the second part, $w(l, \omega^{\parallel})$, which is not an impulse anymore. And the length of the line segment also determines the frequency content. Consequently, the frequency content of a line segment sample is more complex than a line sample.

Weighted frequency content of point samples. According to Eq. (5), the magnitude of the frequency content of line segment samples is mainly affected by the weight $w(l, \omega^{\parallel})$, which is invariant along the direction $\mathbf{d}_{P,Q}^{\perp}$. The weight is of largest magnitude on the line $\mathbf{L}_{-Q,P,0}(\omega)$ and decreases from l to 0 with an increasing distance from the line (i.e., $|\omega^{\parallel}|$).

Note that the weight varies throughout the frequency domain, but with more than 90% of the power concentrated within the range

$|\omega^{\parallel}| < \frac{1}{l}$. Thus the line segment sample's frequency content mainly lies within a band. The position and direction of the band is the same as that of the line sample (Eq. (3)). The difference is that the band has a non-zero width of $\frac{2}{l}$.

Within the frequency band, the variation of frequency content along the direction $\mathbf{d}_{P,Q}^{\perp}$ is mainly determined by the center positions \mathbf{x}_c of line segment samples. Given a group of parallel line segment samples with the same length l , the distribution of sample centers \mathbf{x}_c determines the frequency properties within the frequency band according to Eq. (5). Without exception, uniform sampling (Fig. 4 (a)) and random sampling (Fig. 4 (b)) of \mathbf{x}_c show impulses and white noise in the frequency content respectively, and produce obvious aliasing and noise in the reconstructions. If the center position \mathbf{x}_c is a Poisson disk sampled according to Eq. (28), the power spectrum is the 2D version of Eq. (29) with a weight of $w(l, \omega^{\parallel})$, which has blue-noise properties on the band along $\mathbf{d}_{P,Q}^{\perp}$ (Fig. 4 (c)).

The width of the frequency band is determined by the line segment length l , which controls not only the value range but also the level of anisotropy. For a small segment length l , the weight changes smoothly as shown in Fig. 4 (h), while a large l leads to sharp falloff as shown in Fig. 4 (g). We can see the different levels of aliasing and noise in the reconstructions in the right images of Fig. 4 (g)(h).

Relationships among point, line and line segment sampling.

The frequency content of point samples, line samples and line segment samples are different only in the magnitude and anisotropy of the weight $w(l, \omega^{\parallel})$, which is determined by the directions and lengths of samples (Eq. (5)). A line sample is a special case of a line segment sample with infinite length. Intuitively, increasing l leads to steeper variation and a higher peak value in the weight $w(l, \omega^{\parallel})$, and the band becomes narrowed and converges to a line impulse. Formally, the weak limit of the weight is the line impulse for an infinitely large length l . The frequency content on the line impulse is also consistent with the line sample:

$$\lim_{l \rightarrow \infty} w(l, \omega^{\parallel}) = \delta_{-Q,P,0}(\omega) = \delta(\omega^{\parallel}) \quad (6)$$

$$\omega \mathbf{x}_c = -R\omega^{\perp}, \omega \in \mathbf{L}_{-Q,P,0}(\omega). \quad (7)$$

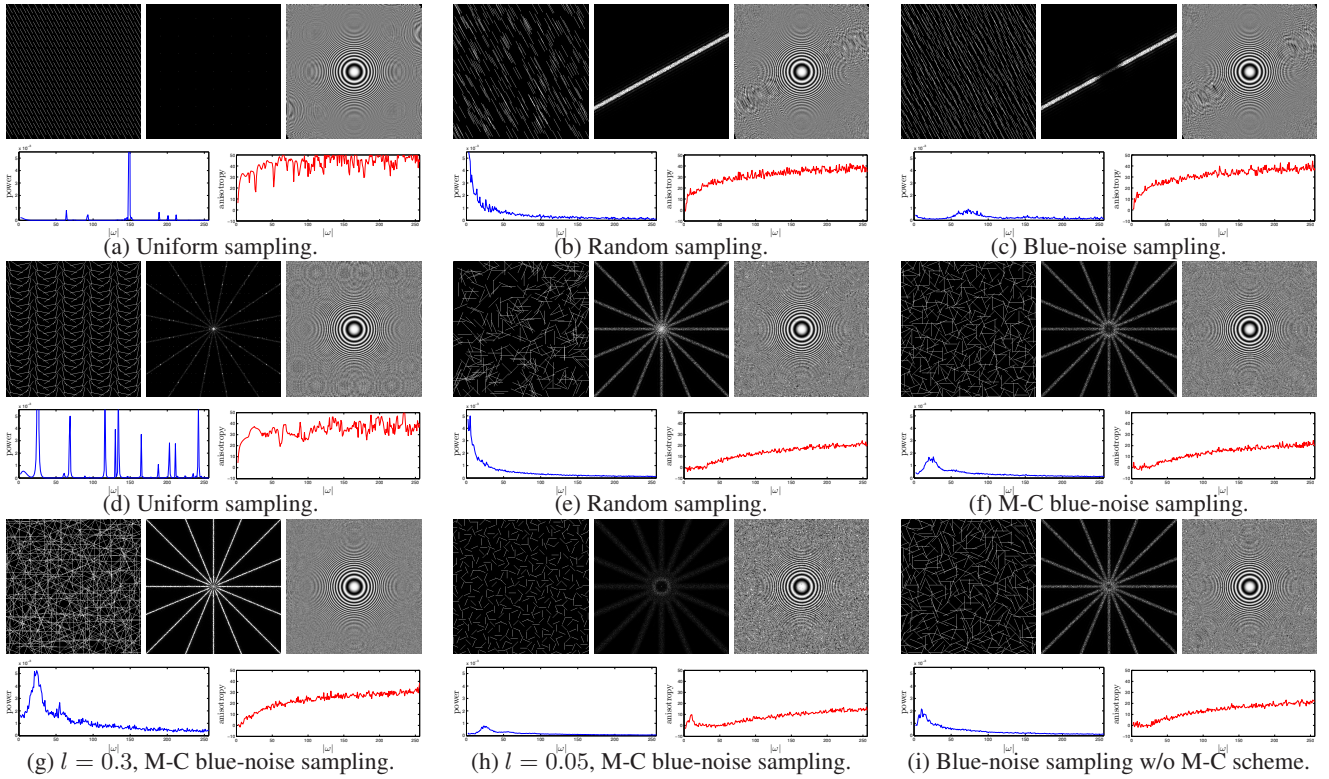


Figure 4: Different line segment sampling schemes with respect to \mathbf{x}_c . M-C means multi-class scheme. Top row samples one direction and the others sample eight evenly distributed directions. Within each group are the samples, Fourier spectrum, reconstruction image, radial mean of power spectrum and anisotropy. All cases contain about 10000 samples. r is 0.007 and 0.008 for blue-noise sampling with and without the multi-class scheme. The sample images show about 5% of the samples for better visualization. $l = 0.1$ unless otherwise noted.

A point sample is also a special case of a line segment sample when the length l goes to zero. With decreasing l , the weight $w(l, \omega^\parallel)$ is smoothed from a band toward being flat throughout the whole frequency domain, and the integral of the line impulse also converges to 0 accordingly. To match the integral of a point impulse (1), we scale the weight $w(l, \omega^\parallel)$ by $\frac{1}{l}$. The weight then converges to 1:

$$\lim_{l \rightarrow 0} \frac{1}{l} w(l, \omega^\parallel) = 1. \quad (8)$$

Therefore, the frequency content of a line segment sample lies on the transition from a point sample to a line sample. Short line segment samples act like point samples with minor anisotropy while long line segment samples introduce strong anisotropy similar to line samples. When comparing blue noise sampling with point samples (Fig. 2 (d)), line samples (Fig. 3 (f)) and line segment samples (Fig. 4 (f)), we can see that point sampling is aliasing-free but with more noise than line sampling, which exhibits angular aliasing. Line segment samples show effects that are in between them.

Multi-class blue-noise sampling among different directions.

Similar to line samples, the limited number of directions of line segment samples in the frequency domain also introduces angular aliasing, which can be similarly alleviated by using multiple groups of line segment samples with different directions as shown in Fig. 4 (f). The same multiple-direction scheme can be applied to uniform sampling (Fig. 4 (b)) and random sampling (Fig. 4 (d)).

An important difference between line sampling and line segment sampling is that the segment centers \mathbf{x}_c in groups with different directions should be multi-class blue-noise sampled [Wei 2010].

The multi-class scheme can keep the blue-noise properties for every group and their union. This is because the overlapping parts of line segment samples from different groups cover non-zero areas besides the origin in the frequency domain. The independently generated blue-noise samples from different groups will cause significant offsets toward white noise in the low frequency area around the origin (Fig. 4 (i)). The Poisson disk distance r of the union can be different from the value of individual groups. The reported r in this paper is for the union of all groups. Within a group, the Poisson disk distance is scaled to $r\sqrt{n_c}$, where n_c is the number of groups. The scaling factor of $\sqrt{n_c}$ can be changed to a different factor by the user [Wei 2010].

Line segment sampling scheme. The line segment samples are generated in a way similar to line samples. The first step is to sample directions. Then a set of line segment samples are generated for each direction to achieve blue-noise properties within the frequency bands. The difference is that we should perform Poisson disk sampling on the centers of line segment samples \mathbf{x}_c instead of on $-R$. Furthermore, the multi-class scheme is required to maintain blue-noise properties in the low frequency area around the origin.

As the frequency bands of line segment samples with different directions cover non-zero areas around the origin, the samples within the same group can be jittered in direction to reduce aliasing artifacts. An analysis and discussion of this will be presented next.

3.3 Jittering of Line Segment Samples

The non-zero width of the frequency band makes a crucial difference between line sampling and line segment sampling. The fre-

quency content of line segment samples with different directions can overlap in non-zero areas besides the origin. Based on this observation, we propose to jitter the directions of line segment samples within the same group. The frequency content Eq. (5) of line segment samples can be evaluated as the expectation of the jittering:

$$\begin{aligned} \mathcal{F}(\delta_{P,Q,R,\mathbf{x}_c,l}(\mathbf{x})) &= \mathcal{F}(\delta(\mathbf{x} - \mathbf{x}_c)) \tilde{w}(l, \boldsymbol{\omega}, \boldsymbol{\Omega}) \\ \tilde{w}(l, \boldsymbol{\omega}, \boldsymbol{\Omega}) &= l \int_{\boldsymbol{\Omega}} \text{sinc}(l \mathbf{d}^{\parallel} \boldsymbol{\omega}) \rho(\mathbf{d}^{\parallel}) d\mathbf{d}^{\parallel}, \end{aligned} \quad (9)$$

where $\mathbf{d}^{\parallel} \in \boldsymbol{\Omega}$ is the range of jittering and $\rho(\mathbf{d}^{\parallel})$ is the normalized *pdf* (probability density function) of \mathbf{d}^{\parallel} in jittering. As shown in Fig. 5 (a), the frequency bands will cover more directions after jittering, which can significantly reduce aliasing artifacts in reconstruction when combined with multiple direction sampling. However, a side effect brought by jittering is the degradation of blue-noise properties. So we analyze the corresponding changes of frequency content.

Power spectrum bounding for jittering. Jittering will change the power spectrum of the line segment samples. A large power spectrum magnitude may introduce more aliasing and noise in the reconstruction. However, we can control the range of jittering to bound the increase of the power spectrum magnitude.

Suppose we have N line segment samples with center positions at $\{\mathbf{x}_i\}$, and the same length l and direction \mathbf{d}^{\parallel} . Their power spectrum according to Eq. (5) is:

$$\begin{aligned} \mathcal{P} \left(\sum_{i=1}^N \mathcal{F}(\delta_{P,Q,R,\mathbf{x}_i,l}(\mathbf{x})) \right) \\ = \mathcal{P} \left(\sum_{i=1}^N \mathcal{F}(\delta(\mathbf{x}_i)) \right) \left(\sum_{i=1}^N w(l, \boldsymbol{\omega}^{\parallel})^2 \right), \end{aligned} \quad (10)$$

where \mathcal{P} is the power spectrum operator. After jittering, we have the samples of $\{\mathbf{L}_{P_i,Q_i,R_i}(\mathbf{x})\}$ with lengths of $\{l_i\}$. Their power spectrum is bounded according to the Cauchy-Schwarz inequality:

$$\begin{aligned} \mathcal{P} \left(\sum_{i=1}^N \mathcal{F}(\delta_{P_i,Q_i,R_i,\mathbf{x}_i,l_i}(\mathbf{x})) \right) \\ \leq \mathcal{P} \left(\sum_{i=1}^N \mathcal{F}(\delta(\mathbf{x}_i)) \right) \left(\sum_{i=1}^N w(l_i, \boldsymbol{\omega}_i^{\parallel})^2 \right). \end{aligned} \quad (11)$$

There is equality in this relationship iff $\forall i, \mathbf{d}_{P_i,Q_i}^{\parallel} = \mathbf{d}_{P,Q}^{\parallel}, l_i = l$, and the jittering is zero or a constant. With fixed positions \mathbf{x}_c , the term $\mathcal{F}(\delta(\mathbf{x}_i))$ is not changed by jittering. By subtracting Eq. (10) from Eq. (11), we find that the increase of the power spectrum magnitude from jittering is bounded by $w(l_i, \boldsymbol{\omega}_i^{\parallel})$.

The change of w from jittering is composed of two parts, ϵ_i^l and $\epsilon_i^{\mathbf{d}^{\parallel}}$, corresponding to the lengths and directions respectively:

$$\epsilon_i = \left| w(l_i, \boldsymbol{\omega}_i^{\parallel}) - w(l, \boldsymbol{\omega}^{\parallel}) \right| \leq \epsilon_i^l + \epsilon_i^{\mathbf{d}^{\parallel}}. \quad (12)$$

ϵ_i^l is the upper bound with respect to the change from l to l_i :

$$\epsilon_i^l = \left| w(l_i, \boldsymbol{\omega}_i^{\parallel}) - w(l, \boldsymbol{\omega}_i^{\parallel}) \right| \leq \epsilon_i^l, \quad (13)$$

where ϵ_i^l is the deviation of segment length $|l_i - l|$. The bound of the weight is linearly dependent on the change in length. $\epsilon_i^{\boldsymbol{\omega}^{\parallel}}$ is

related to the change from direction $\mathbf{d}_{P,Q}^{\parallel}$ to $\mathbf{d}_{P_i,Q_i}^{\parallel}$:

$$\begin{aligned} \epsilon_i^{\mathbf{d}^{\parallel}} &= \left| w(l, \boldsymbol{\omega}_i^{\parallel}) - w(l, \boldsymbol{\omega}^{\parallel}) \right| = \min \left(\pi l^2 |\boldsymbol{\omega}| \epsilon_i^{\mathbf{d}}, l + \frac{l}{\pi} \right) \\ \epsilon_i^{\mathbf{d}} &= \sqrt{1 - \left(\mathbf{d}_{P_i,Q_i}^{\parallel} \left(\mathbf{d}_{P,Q}^{\parallel} \right)^T \right)^2}. \end{aligned} \quad (14)$$

If $|\boldsymbol{\omega}|$ is small, $\epsilon_i^{\boldsymbol{\omega}^{\parallel}}$ is bounded by $\pi l^2 |\boldsymbol{\omega}| \epsilon_i^{\mathbf{d}}$, which is linearly dependent on the change from $\mathbf{d}_{P,Q}^{\parallel}$ to $\mathbf{d}_{P_i,Q_i}^{\parallel}$. This is in accord with intuition and experiments which show that greater jittering leads to greater changes. If $|\boldsymbol{\omega}|$ is large, $\epsilon_i^{\boldsymbol{\omega}^{\parallel}}$ is always less than $l + \frac{l}{\pi}$. In general, the upper bound of the power spectrum can be controlled by the range of jittering. There will not be any frequency impulses introduced by jittering.

Benefits of jittering. *Jittering line directions helps to reduce angular aliasing, but also reduces blue-noise properties.* Our analysis above (Eq. (14)) reveals the relationship between jittering and power spectrum bounding. The bounding proposed in Eq. (14) is conservative without any assumption on the *pdf* of $\rho(\mathbf{d}^{\parallel})$ in Eq. (9). Usually we perform jittering uniformly in the range with a constant value of $\rho(\mathbf{d}^{\parallel})$.

If we want to cover all directions in the frequency domain by jittering, much higher sampling rates are needed as shown in Fig. 5 (g), otherwise the frequency properties are close to white noise. The jittering should be used together with the multiple direction scheme. Without multiple direction sampling, jittering is equivalent to random direction sampling as shown in Fig. 5 (h), where the blue-noise properties are greatly weakened. Even high sampling rates can help very little, as shown in Fig. 5 (i).

We have also discussed power spectrum bounding according to the jittering of line segment lengths l . Note that sometimes we cannot generate samples with the same length. For example, if we use motion vectors as line segment samples to reconstruct rendered images of motion blur, the length of the vectors cannot be the same everywhere. Fortunately, the linear dependency between the weight and length as indicated in Eq. (13) tells us that smooth changes of line segment lengths can still provide good blue-noise properties locally.

Power spectrum fall-off in high frequency areas. It is observed in Fig. 4 and Fig. 5 that the power spectrum from non-uniform sampling always falls off in high frequency areas, while the anisotropy goes up at the same time. That is because the frequency band of the weight in Eq. (5) maintains a constant width throughout the domain. The frequency band must cross the origin and not diverge in high frequency areas, which makes it difficult for the accumulated frequency content of line segment samples to converge to a non-zero constant. That is a significant difference from point samples, which have a constant magnitude over the entire frequency domain as shown in Eq. (27).

Consequently, there is always a peak before the power spectrum falls off. The blue-noise properties of line segment sampling can be evaluated by the position and magnitude of the peak. Usually reducing anisotropy can degrade blue-noise properties. Jittering is an effective scheme because its influence on the power spectrum and anisotropy is different. From the comparisons in Fig. 4 (f) and Fig. 5 (b)(d), we find that jittering can significantly reduce anisotropy with only a minor loss of blue-noise properties.

Line segment sampling scheme with jittering. We first generate a set of parallel line segment samples for each group as de-

scribed in Section 3.2, from which we determine the range of jittering. Then the directions of all line segment samples are jittered with a constant probability in the range. The jittering range controls the balance between aliasing and noise under a fixed number of samples. To reconstruct signals with low frequencies, a small jittering range is often used for its low variance. Otherwise, a large jittering range is better because it is aliasing free. As we mainly focus on the frequency content of line segment samples in the spatial domain, we leave a thorough analysis in the angular domain for future work. In this paper, we uniformly sample several directions and choose the range of jittering manually. We can see the improvement brought by the jittering scheme, but a solid quantitative analysis still requires much more theoretical work.

3.4 Extensions to High Dimensional Spaces

Our analysis on line sampling (Section 3.1) and line segment sampling (Section 3.2 and Section 3.3) in the $2D$ space can be generalized to high dimensional spaces as mD affine subspace sampling and m -Ball sampling, respectively. Similar conclusions about blue-noise sampling can be determined accordingly.

Affine subspace sampling. In the nD space, an mD affine subspace $m \leq n$ can be defined as

$$\mathbf{L}_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x}) = \left\{ \mathbf{x} \mid |\mathbf{d}^\perp \mathbf{x} + \mathbf{R}| = 0 \right\}, \quad (15)$$

where \mathbf{d}^\perp is a matrix of dimensions $(n-m) \times n$, and \mathbf{R} is a vector with a dimension of $(n-m) \times 1$. It will degenerate to point sampling when $m=0$. The samples have non-zero area in the nD space when $m=n$. Without loss of generality, we assume \mathbf{d}^\perp is normalized: $\mathbf{d}^\perp \mathbf{d}^{\perp T} = \mathbf{I}$, where \mathbf{I} is the identity matrix. A sample in $\mathbf{L}_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x})$ is an mD impulse

$$\delta_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x}) = \begin{cases} \delta(\mathbf{d}^\perp \mathbf{x} + \mathbf{R}), & m < n \\ 1, & m = n, \end{cases} \quad (16)$$

whose Fourier transform is

$$\mathcal{F}(\delta_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x})) = e^{2\pi i((\mathbf{d}^\perp \boldsymbol{\omega})^T \mathbf{R})} \delta_{\mathbf{d}^\perp, \mathbf{0}}(\boldsymbol{\omega}), \quad (17)$$

where \mathbf{d}^\parallel is a matrix of dimensions $m \times n$ satisfying $\mathbf{d}^\parallel \mathbf{d}^{\parallel T} = \mathbf{I}$ and $\mathbf{d}^\perp \mathbf{d}^{\perp T} = \mathbf{0}$.

The form of Eq. (17) is consistent with Eq. (3). It is composed of the product of an exponential function and an impulse function. The line samples analyzed in Section 3.1 are special case of affine subspace sampling in high dimensional space. When $n=2$ and $m=1$, Eq. (17) is reduced to Eq. (3). The frequency features we introduced for line samples in Section 3.1 also hold true for this generalization based on the same analysis. As shown in Eq. (16), the frequency content of an mD affine subspace sample lies on the linear subspace impulse $\delta_{\mathbf{d}^\perp, \mathbf{0}}(\boldsymbol{\omega})$ perpendicular to the direction of the sample. The frequency content in the linear subspace is also an $(n-m)D$ Fourier basis function with parameter $\mathbf{d}^\perp \boldsymbol{\omega}$ and frequency \mathbf{R} . Consequently, the sampling scheme is also similar to line samples. Poisson disk sampled \mathbf{R} for parallel samples with the same direction \mathbf{d}^\parallel can produce blue-noise properties on the line subspace impulse $\delta_{\mathbf{d}^\perp, \mathbf{0}}(\boldsymbol{\omega})$.

m -Ball sampling. A generalized line segment in high dimensional space could be of various shapes. We propose to use an m -Ball because its orientation does not need to be considered.

If the sample is an m -Ball with center \mathbf{x}_c and radius b :

$$\delta_{\mathbf{d}^\perp, \mathbf{R}, \mathbf{x}_c, b}(\mathbf{x}) = \delta_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x}) H\left(b - \left| \mathbf{d}^\parallel (\mathbf{x} - \mathbf{x}_c) \right| \right) \quad (18)$$

then the corresponding frequency content is

$$\begin{aligned} \mathcal{F}(\delta_{\mathbf{d}^\perp, \mathbf{R}, \mathbf{x}_c, b}(\mathbf{x})) &= \mathcal{F}(\delta(\mathbf{x} - \mathbf{x}_c)) w(b, \boldsymbol{\omega}^\parallel) \\ w(b, \boldsymbol{\omega}^\parallel) &= b \left(b \left| \boldsymbol{\omega}^\parallel \right| \right)^{-\frac{n}{2}} \text{BesselJ}\left(\frac{n}{2}, 2\pi b \left| \boldsymbol{\omega}^\parallel \right| \right), \end{aligned} \quad (19)$$

where BesselJ is the Bessel function of the first kind.

As expected, Eq. (19) is the product of a point sample and a non-impulse weighting function, which is similar to the frequency content of a line segment sample as shown in Eq. (5). A line segment sample in Section 3.2 is indeed a l -Ball sample. Eq. (5) is a special case of Eq. (19) where $n=2$, $m=1$ and $l=2b$. We can also generalize the frequency analysis in Section 3.2 accordingly. From Eq. (19), the m -Ball sample is also equivalent to a weighted point sample in the nD space. The magnitude of the weight $w(b, \boldsymbol{\omega}^\parallel)$ is bounded by b . It decreases continuously along \mathbf{d}^\parallel with respect to the distance from the linear subspace $\mathbf{L}_{\mathbf{d}^\perp, \mathbf{0}}$ and is constant along \mathbf{d}^\perp . The Poisson disk sampled center \mathbf{x}_c will have blue-noise properties along \mathbf{d}^\perp . Multiple direction sampling can maintain blue-noise properties throughout the frequency domain and alleviate angular aliasing. The jittering scheme can also benefit the sampling scheme by providing a compromise between angular aliasing and noise. The relationships among point sampling, affine subspace sampling and m -Ball sampling still hold true, similar to what we discussed in Section 3.2 for the $2D$ space.

3.5 Extensions to General Non-point Samples

Our analysis for line segment samples can also be extended to other non-point samples of arbitrary shapes. Lines and line segments are two regular shapes widely use in graphics. We will show how the frequency content is determined by the shape of the samples.

Given a general shape \mathbf{S} in the nD space with center \mathbf{x}_c and lying in the mD affine subspace $\mathbf{L}_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x})$, it can be represented by the indicator function of $H_{\mathbf{d}^\parallel, \mathbf{x}_c, \mathbf{S}}(\mathbf{x})$:

$$H_{\mathbf{d}^\parallel, \mathbf{x}_c, \mathbf{S}}(\mathbf{x}) = \begin{cases} 1, & \mathbf{d}^\parallel (\mathbf{x} - \mathbf{x}_c) \in \mathbf{S} \\ 0, & \text{else.} \end{cases} \quad (20)$$

$\mathbf{d}^\parallel (\mathbf{x} - \mathbf{x}_c)$ is the position in the affine subspace $\mathbf{L}_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x})$ relative to \mathbf{x}_c . A sample of this shape is also defined as an mD impulse:

$$\begin{aligned} \delta_{\mathbf{d}^\perp, \mathbf{R}, \mathbf{x}_c, \mathbf{S}}(\mathbf{x}) &= \delta_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x}) H_{\mathbf{d}^\parallel, \mathbf{x}_c, \mathbf{S}}(\mathbf{x}) \\ &= \delta_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x}) (\delta_{\mathbf{d}^\parallel, -\mathbf{d}^\parallel \mathbf{x}_c}(\mathbf{x}) \otimes H_{\mathbf{d}^\parallel, \mathbf{0}, \mathbf{S}}(\mathbf{x})) \\ &= \delta(\mathbf{x} - \mathbf{x}_c) \otimes H_{\mathbf{d}^\parallel, \mathbf{0}, \mathbf{S}}(\mathbf{x}). \end{aligned} \quad (21)$$

The above derivation shows that the impulse of the shape can be obtained by convolving a point impulse with the shape \mathbf{S} . Its frequency content is thus their product:

$$\begin{aligned} \mathcal{F}(\delta_{\mathbf{d}^\perp, \mathbf{R}, \mathbf{x}_c, \mathbf{S}}(\mathbf{x})) &= \mathcal{F}(\delta(\mathbf{x} - \mathbf{x}_c)) w_{H, \mathbf{d}^\parallel, \mathbf{S}}, \\ w_{H, \mathbf{d}^\parallel, \mathbf{S}}(\boldsymbol{\omega}) &= \mathcal{F}(H_{\mathbf{d}^\parallel, \mathbf{0}, \mathbf{S}}(\mathbf{x})) = \int_{\mathbf{S}} e^{-2\pi i(\hat{\boldsymbol{\omega}}^T \hat{\mathbf{x}}_c)} d\hat{\mathbf{x}}, \end{aligned} \quad (22)$$

where $\hat{\boldsymbol{\omega}} = \mathbf{d}^\parallel \boldsymbol{\omega}$ and $\hat{\mathbf{x}}_c = \mathbf{d}^\parallel \mathbf{x}_c$. The frequency content of the shape sample is equivalent to the weighted frequency content of a point sample in the nD space. The weight is the integral of the

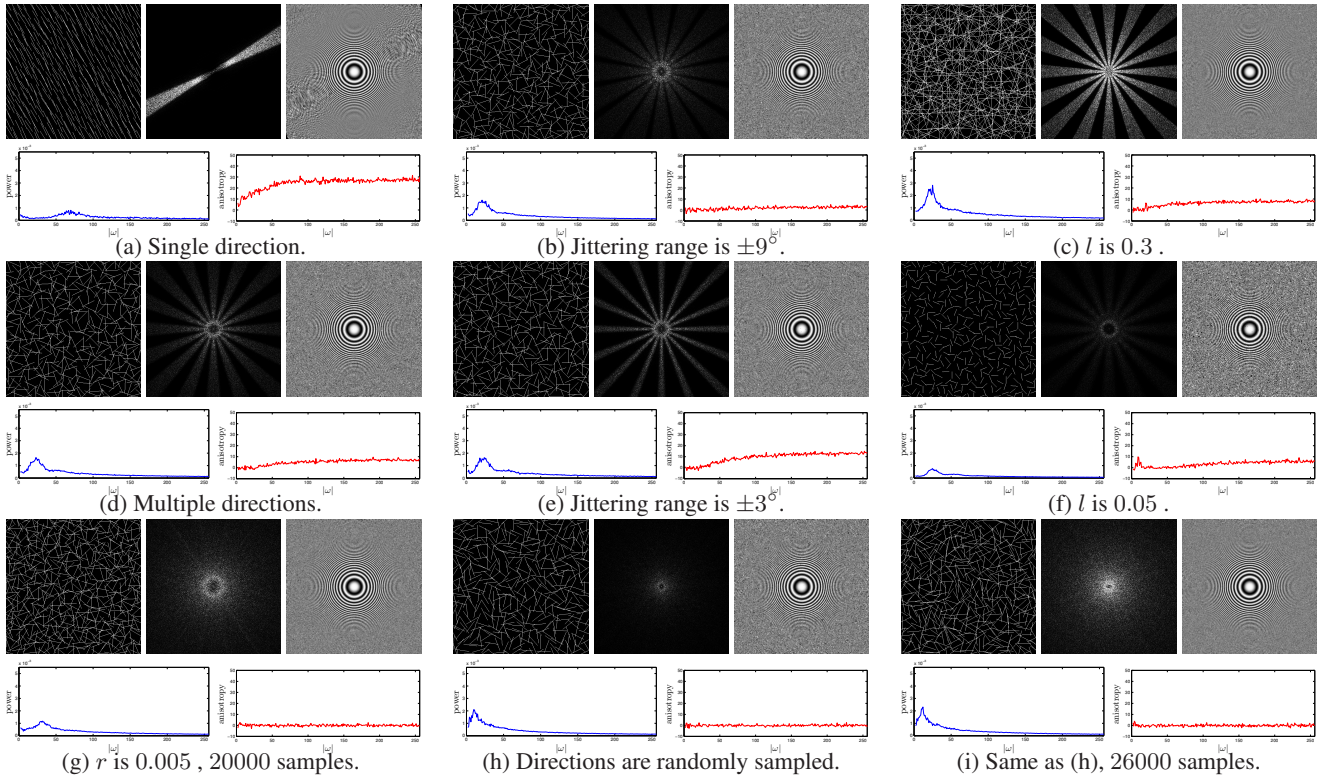


Figure 5: Comparisons on jittering for line segment sampling. Within each group are the samples, Fourier spectrum, reconstruction image, radial mean of power spectrum and anisotropy. Unless otherwise noted, the images have about 10000 samples. The r is 0.008 for blue-noise sampling. The sample images show about 5% of the samples for better visualization. \mathbf{x}_c is multi-class blue-noise sampled. The samples are evenly distributed into eight groups of different directions for the multiple direction sampling scheme. The length l is 0.1 and the jittering range is 12° .

frequency content in the affine subspace $\mathbf{L}_{\mathbf{d}^\perp, \mathbf{R}}(\mathbf{x})$. The integral is computed with the frequency parameter $\hat{\omega}$ within the domain of \mathbf{S} . \mathbf{x}_c determines the translation of the integral domain. Eq. (22) is reduced to Eq. (19) if the shape \mathbf{S} is an m -Ball.

Obviously, the weight $w_{H_{\mathbf{d}^\perp, \mathbf{S}}}(\omega)$ in Eq. (22) is dependent on the shape \mathbf{S} , which determines the frequency properties of the samples. The frequency content of various samples, including point samples (Eq. (27)), line samples (Eq. (3) and Eq. (17)) and line segment samples (Eq. (5) and Eq. (17)), only differs in this weight.

If the weight in Eq. (22) is bounded and very smooth, there will be a good chance to get blue-noise properties in the shape sampling, as we proposed in Section 3.1 to Section 3.4. It is natural that convex shapes are good candidates for non-point samples. Jittering is also possible by bounding the change of \mathbf{d}^\perp and the shape \mathbf{S} . The bounding in Section 3.3 is just a special case of line segment sampling in the $2D$ space.

If the shape $\mathbf{S}(\mathbf{x})$ is generated by the convolution of two shapes $\mathbf{S}(\mathbf{x}) = \mathbf{S}_1(\mathbf{x}) \otimes \mathbf{S}_2(\mathbf{x})$, its frequency content can easily be computed as the product between them:

$$\mathcal{F}(H_{\mathbf{d}^\perp, \mathbf{0}, \mathbf{S}}(\mathbf{x})) = \mathcal{F}(H_{\mathbf{d}^\perp, \mathbf{0}, \mathbf{S}_1}(\mathbf{x})) \mathcal{F}(H_{\mathbf{d}^\perp, \mathbf{0}, \mathbf{S}_2}(\mathbf{x})). \quad (23)$$

Eq. (23) tells us that the frequency content of a sample can be considered as a combination of multiple physical processes, just as in the temporal light field reconstruction that we will discuss in Section 4.4.

4 Applications

Our sampling algorithm is summarized here. We first uniformly sample a few directions. For each direction, we generate a group of samples with blue-noise properties by $1D$ and $2D$ dart throwing for line samples and line segment samples respectively. Then the directions of all samples are uniformly jittered if required. We apply this sampling algorithm to several problems. The reference images in Section 4.2, 4.3 and 4.4 are generated by path tracing with a substantial number of low discrepancy point samples.

4.1 Image Reconstruction

Although line segment sampling has never been applied to image reconstruction, we use this application to compare different sampling schemes due to its simplicity.

The comparisons for an image of detailed fabric are shown in Fig. 6. For both line sampling and line segment sampling, the blue-noise sampling schemes (Fig. 6 (c)(f)) produce the best results, as seen from the noise of random sampling (Fig. 6 (b)(e)) and aliasing of uniform sampling (Fig. 6 (a)(d)). Because only two directions are sampled, the angular aliasing is quite obvious for parallel samples in Fig. 6 (c)(f), which can be improved by jittering as shown in Fig. 6 (g). We increase the sampling rate for line segment samples to make the accumulated lengths of all samples approximately the same between line sampling and line segment sampling. By comparing Fig. 6 (c) and (f), we can see more and shorter samples alleviate aliasing. The reconstruction from blue-noise point sampling with a high sampling rate is aliasing-free but noisy. In general, the line segment sampling scheme makes it possible to achieve a good balance among aliasing, noise and sampling rate.

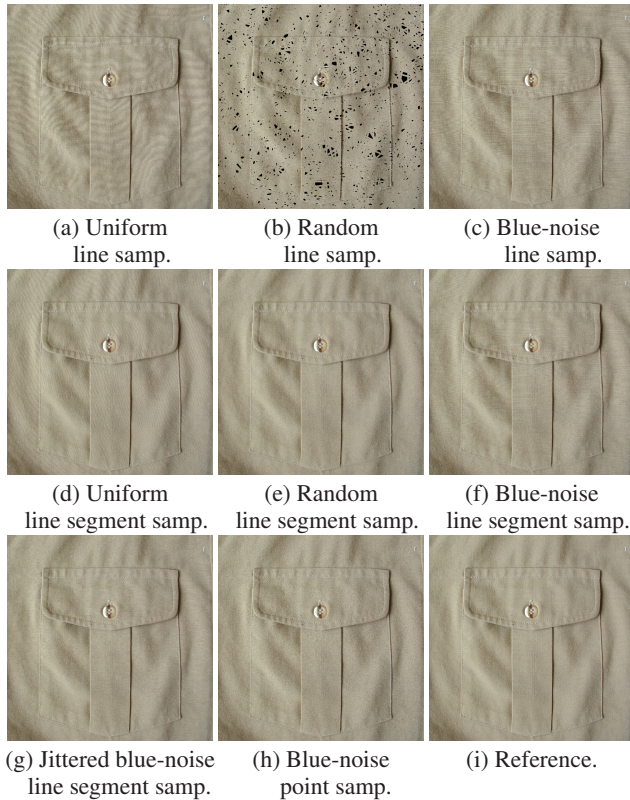


Figure 6: Comparisons of image reconstruction. The images of detailed fabric have a resolution of 1200×1200 . Two directions are sampled. Line sampling: $r = 0.005$ and 500 samples. Line segment sampling: $l = 0.1$, $r = 0.01$ and 5000 samples.

4.2 Motion Blur

Traditional motion blur rendering performs point sampling in the temporal-spatial domain, and shoots a ray from each sample for a visibility test. Recently, line segment sampling as shown in Fig. 7 has been employed in stochastic rasterization [Gribel et al. 2011] to achieve better quality as well as low computational cost. The basic idea is to divide the image into a set of tiles, each of which covers a small rectangular region. Usually a tile covers 32×32 pixels, so the rendered 512×512 image is divided into 16×16 tiles. Within each tile, line segments that cross the whole tile are sampled instead of the point samples employed in the traditional rendering pipeline. The visibility test and shading are computed by point sampling along the line segments.

In our implementation, the rendering framework is the same as that proposed by Gribel et al. [2011] except that we only modify the generation of line segments in the tiles. Gribel et al. [2011] sample two directions in a tile. For each direction, parallel line segments are sampled at uniform intervals. We use four directions in our experiments, and demonstrate that our blue-noise sampling scheme can improve rendering quality under the same sampling rate. In Fig. 9 (a), aliasing is caused by both the high-frequency textures and non-linear motions. Our blue-noise sampling scheme (Fig. 9 (b)) converts the aliasing into noise. We also allow a tradeoff between aliasing and noise by jittering according to the requirements of the application (Fig. 9 (c)).

4.3 Depth of Field

The line segment sampling for motion blur rendering [Gribel et al. 2011] can be directly extended to render depth of field as shown in




Figure 7: Motion blur rendering based on stochastic rasterization with line segment sampling in the image space [Gribel et al. 2011]. The triangle travels along time t . The line sample l on the image plane intersects the moving triangle, generating the area in the $t \times l$ space as shown on the right. The evaluation of visibility and shading within this area gives the rendering result of the motion blur.




Figure 8: Depth of field rendering with line segment sampling on the image. For each line segment on the image plane, we sample a parallel straight line on the camera lens to generate a plane. The plane intersects the scene geometry after the refraction on the lens. The intersection generates the area in $u \times l$ space. The evaluation including visibility and shading within this area produces the depth of field rendering result. This scheme is a straightforward extension of the motion blur rendering algorithm [Gribel et al. 2011].


Fig. 8. The difference is that each line segment sample on the image corresponds to a line segment sample on the camera lens. The two line segment samples on the image and lens should be parallel and form a plane that intersects the scene geometry.

The chess scene in Fig. 10 is mapped with high-frequency textures. Aliasing artifacts are generated in both in-focus and out-of-focus areas if we uniformly sample the line segments (see Fig. 10 (a)). Our blue-noise sampling (Fig. 10 (b)) reduces the stratified aliasing, but the effects are not smooth enough for the out-of-focus blur as expected. The jittering scheme helps smooth out the remaining aliasing and generates results much more consistent to the ground truth.

Besides the textures, the complex geometries also generate aliasing artifacts with the uniform sampling scheme (Fig. 11 (a)). The blue-noise sampling scheme (Fig. 11 (b)) removes almost all the aliasing artifacts, but still keeps the high-frequency features consistent to the ground truth as shown in Fig. 11 (d). The jittering scheme introduces some randomness to the result (Fig. 11 (c)).

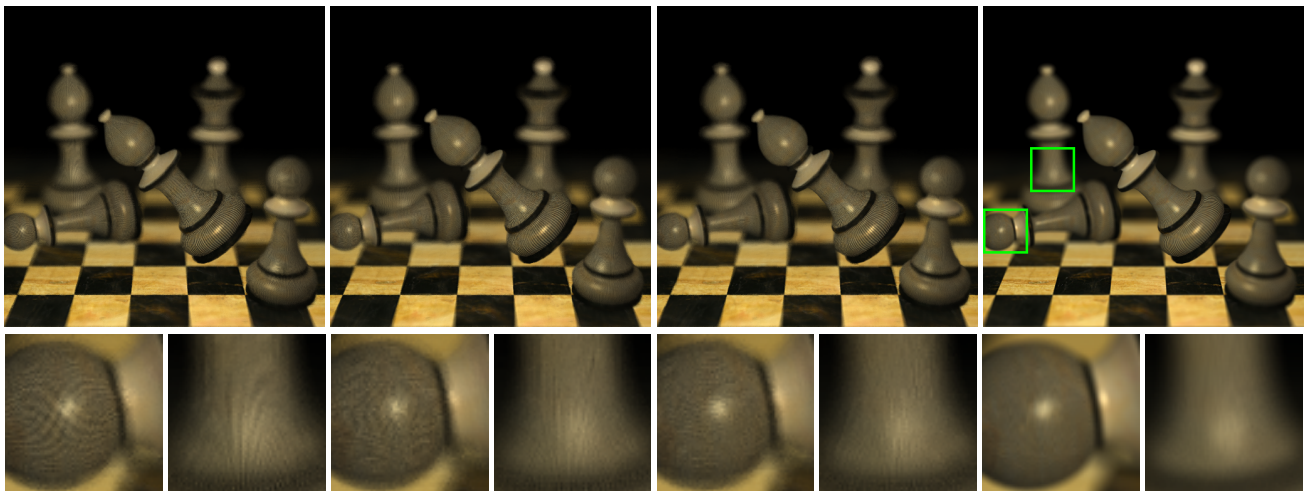
4.4 Temporal Light Field Reconstruction

The temporal light field reconstruction method proposed by Lehtinen et al. [2011] is efficient in rendering distribution effects such as depth of field and motion blur. The key idea is to perform point sampling in the temporal light field, and then reconstruct the image




(a) Uniform sampling. (b) Blue-noise sampling. (c) Blue-noise sampling w. jittering. (d) Reference.

Figure 9: Comparisons for motion blur rendering of polyhedra. The image is divided into square tiles of resolution 32. Within each tile, we sample four directions each with 32 line segment samples.



(a) Uniform sampling. (b) Blue-noise sampling. (c) Blue-noise sampling with jittering. (d) Reference.

Figure 10: Comparisons for depth of field rendering of a chess scene. The image is divided into square tiles of resolution 32. Within each tile, we sample eight directions each with 32 line segment samples.



(a) Uniform sampling. (b) Blue-noise sampling. (c) Blue-noise sampling with jittering. (d) Reference.

Figure 11: Comparisons for depth of field rendering of a bench. The image is divided into square tiles of resolution 32. Within each tile, we sample four directions each with 32 line segment samples.

from all samples with sheared filtering. Each sample is a $5D$ point:

$$\begin{cases} \mathbf{s} = (\mathbf{x}, \mathbf{u}, t)^T \\ \mathbf{u}^T \mathbf{u} \leq r_{len}^2 \text{ and } t \in (t_0, t_1) \end{cases} \quad (24)$$

where $\mathbf{x} = (x, y)^T$ is the position in the image, $\mathbf{u} = (u, v)^T$ is the position on the lens, t is the time stamp, r_{len} is the radius of the camera lens, and (t_0, t_1) is the range of the time period.

Even though each sample is a $5D$ point, the sampling space is actually not a full $5D$ space as $\mathbf{x}_u = \frac{d\mathbf{x}}{d\mathbf{u}}$ and $\mathbf{x}_t = \frac{d\mathbf{x}}{dt}$ are determined by \mathbf{s} . From sheared filtering, the shading value of a sample position \mathbf{x}' on the image will be reconstructed by gathering all samples which can potentially “move” to \mathbf{x}' . Such “movement” exists only if \mathbf{u}' and t' satisfy

$$\begin{cases} \mathbf{x}' = \mathbf{x} + \mathbf{x}_u (\mathbf{u}' - \mathbf{u}) + \mathbf{x}_t (t' - t) \\ \mathbf{u}'^T \mathbf{u}' \leq r_{len}^2 \text{ and } t' \in (t_0, t_1). \end{cases} \quad (25)$$

Reconstruction with sheared filtering performs reconstruction from point samples with anisotropic kernels. Another equivalent understanding of sheared reconstruction is to do reconstruction with isotropic kernels from non-point samples whose shape is determined by $\{\mathbf{x}_u, \mathbf{x}_t, r_{len}, (t_0, t_1)\}$, as shown in Fig. 12.

The quality of reconstruction is determined by the frequency content of those non-point samples. Fixing \mathbf{u} , $\{t' \in (t_0, t_1)\}$ generates a line segment from \mathbf{x}' of length l . On the other hand, fixing t and having $\{\mathbf{u}' | \mathbf{u}'^T \mathbf{u}' \leq r_{len}^2\}$ results in \mathbf{x}' covering the circle of confusion (COC) with radius b . Then the combination of all possible t' and \mathbf{u}' is a convolution between the line segment and the circle. We can get the frequency content of the non-point sample according to Eq. (5), Eq. (19) and Eq. (23):

$$\begin{aligned} \mathcal{F}(\delta_{P,Q,R,\mathbf{x}_c,l}(\mathbf{x})) &= \mathcal{F}(\delta(\mathbf{x} - \mathbf{x}_c)) w_{mb}(l, \omega^{\parallel}) w_{dof}(b, \omega^{\parallel}) \\ w_{mb}(l, \omega^{\parallel}) &= l \operatorname{sinc}(l\omega^{\parallel}) \\ w_{dof}(b, \omega^{\parallel}) &= \frac{1}{|\omega^{\parallel}|} \operatorname{BesselJ}(1, 2\pi b |\omega^{\parallel}|) \\ \mathbf{x}_c &= \mathbf{x} + \mathbf{x}_t ((t_0 + t_1)/2 - t) - \mathbf{x}_u \mathbf{u} \\ l &= |\mathbf{x}_t| (t_1 - t_0) \\ \omega^{\parallel} &= \frac{\mathbf{x}_t^T \boldsymbol{\omega}}{|\mathbf{x}_t|}. \end{aligned} \quad (26)$$

Similar to our conclusion about line segment sampling in Section 3.2, the frequency properties of the non-point samples with the same $\{\mathbf{x}_u, \mathbf{x}_t, r_{len}, (t_0, t_1)\}$ are determined by \mathbf{x}_c . Blue-noise point sampling of \mathbf{x}_c can produce high-quality reconstructions.

Based on our analysis, the quality of samples is determined by their position distribution in the image space. That conclusion can be validated by the comparisons in Fig. 1. In Fig. 1 (a), the samples are generated by low discrepancy sampling in the $5D$ space of $(\mathbf{x}, \mathbf{u}, t)$. The samples have frequency content close to blue-noise properties in the $5D$ space. But in the image space for reconstruction, the \mathbf{x}_c of those samples are distributed irregularly, and the frequency content does not hold blue-noise properties anymore. As a result, the reconstruction has much noise.

In Fig. 1 (b), we also sample in the $5D$ space, but we use a much higher sampling rate. After generating a large number of initial samples, we discard most of them according to their value of \mathbf{x}_c . By discarding samples, the \mathbf{x}_c of the remaining samples leads to blue-noise properties in the image space. After preprocessing, the




Figure 12: Temporal light field reconstruction from point samples is equivalent to image reconstruction from non-point samples. The original light path represents a sample with the given sample u on the lens and the sample t of the time stamp. Different samples of u' and t' produce the green light path. All of the possible u' and t' will generate the light paths covering the red area on the image plane. This red area is a non-point sample produced by the orange light path. This non-point sample is the convolution between the COC of the lens and the motion vector on the image.

samples used for reconstruction have much better frequency properties in the image space, so the variance in reconstruction is reduced.

The comparison between Fig. 1 (a) and (b) delivers an interesting insight. If the sampling and reconstruction are handled in different spaces, the quality of reconstruction is determined by the frequency properties of the samples in the reconstruction space instead of in the sampling space. That is consistent to the conclusion in Sec. 5 of [Mitchell 1991]. Based on that insight, we show another scheme in Fig. 1 (e) for even less noise in reconstruction. The samples are generated with low discrepancy sampling of \mathbf{x}_c in the image space, with \mathbf{u} and t fixed to $\mathbf{0}$ and $(t_0 + t_1)/2$ respectively without any randomness. But this scheme cannot converge to the ground truth with a higher sampling rate, because the geometries only visible with different values of \mathbf{u} and t cannot be rendered. Note that we mention Fig. 1 (e) only to discuss noise here; an analysis of physical correctness is out of the scope of this paper.

Fig. 1 (b) introduces a simple method to generate high quality samples to represent a light field while keeping the sampling rate unchanged. In other words, the traditional widely used low discrepancy samples in the high dimensional light field space are inefficient in light field representation and reconstruction.

As the computation cost of sample generation is determined by the sampling rate in $5D$ light field space, the preprocessing of sample generation and rejection is quite heavy in Fig. 1 (b). But the sampling scheme in Fig. 1 (b) could be a good strategy for applications whose sampling and reconstruction can be decoupled, such as refocusing and relighting, as also mentioned by Lehtinen et al. [2011]. We can build a small set of high quality samples in the preprocessing which can significantly improve the rendering quality without introducing extra computation or storage in the reconstruction phase as shown in the comparison between Fig. 1 (c) and (d).

5 Conclusion and Future Work

We have presented a frequency analysis of line segment sampling in the $2D$ space, and extended the conclusions to high dimensional spaces and samples of arbitrary non-point shapes. Based on this analysis, we proposed a line segment sampling scheme to best preserve the blue-noise properties of samples and minimize anisotropy. This sampling scheme outperforms other sampling schemes in several applications such as uniform sampling, random sampling or some straightforward extensions of blue-noise point sampling.

As previous work mainly focus on point samples, we regard our work as the first frequency analysis of non-point samples and hope that it can stimulate further work along this direction. Our work began with studying the frequency properties of a line segment sample. The anisotropic weighting is the key observation and all our conclusions are derived from it. The basic frequency analysis can help to understand the characteristics of line segment sampling, and provide the preliminaries for further exploration of more efficient sampling schemes.

Our current analysis is based on the assumptions that samples are of the same shape and similar sizes. How to use samples of different shapes or dramatically different sizes will be very useful in some applications. For example, we can use line segment samples for areas with less angular aliasing and employ point samples elsewhere. The sampling rate of line segment samples of the same direction can be adapted along that direction according to the lengths of samples. For example, in motion blur rendering with sheared filtering, the sampling rates along the motion direction should be much lower than in the perpendicular direction. We leave these for future work.

Acknowledgements

We would like to thank the reviewers for their valuable comments. We would like to thank Stephen Lin for paper proofreading. We would like to thank Li-Yi Wei for discussions. We would also like to thank Jiawen Chen for sharing us the code of temporal light field reconstruction [Lehtinen et al. 2011]. Kun Zhou is partially supported by NSFC (No. 61272305) and the 973 program of China (No. 2009CB320801). Guofu Xie and Wencheng Wang were partially supported by the Knowledge Innovation Program of the Chinese Academy of Sciences.

References

- AKENINE-MÖLLER, T., MUNKBERG, J., AND HASSELGREN, J. 2007. Stochastic rasterization using time-continuous triangles. In *GH*, 7–16.
- BALZER, M., SCHLÖMER, T., AND DEUSSEN, O. 2009. Capacity-constrained point distributions: a variant of lloyd’s method. *ACM Trans. Graph.* 28, 3 (July), 86:1–86:8.
- BARRINGER, R., GRIBEL, C. J., AND AKENINE-MÖLLER, T. 2012. High-quality curve rendering using line sampled visibility. *ACM Trans. Graph.* 31, 6 (Nov.), 162:1–162:10.
- BOWERS, J., WANG, R., WEI, L.-Y., AND MALETZ, D. 2010. Parallel poisson disk sampling with spectrum analysis on surfaces. *ACM Trans. Graph.* 29, 6 (Dec.), 166:1–166:10.
- BRIDSON, R. 2007. Fast poisson disk sampling in arbitrary dimensions. In *ACM SIGGRAPH 2007 sketches*, ACM, New York, NY, USA, SIGGRAPH ’07.
- CHESLACK-POSTAVA, E., WANG, R., AKERLUND, O., AND PELLACINI, F. 2008. Fast, realistic lighting and material design using nonlinear cut approximation. *ACM Trans. Graph.* 27, 5 (Dec.), 128:1–128:10.
- COHEN, M. F., SHADE, J., HILLER, S., AND DEUSSEN, O. 2003. Wang tiles for image and texture generation. *ACM Trans. Graph.* 22, 3 (July), 287–294.
- COOK, R. L. 1986. Stochastic sampling in computer graphics. *ACM Trans. Graph.* 5, 1 (Jan.), 51–72.
- DUNBAR, D., AND HUMPHREYS, G. 2006. A spatial data structure for fast poisson-disk sample generation. *ACM Trans. Graph.* 25, 3 (July), 503–508.
- DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. X. 2005. A frequency analysis of light transport. *ACM Trans. Graph.* 24, 3 (July), 1115–1126.
- EBEIDA, M. S., DAVIDSON, A. A., PATNEY, A., KNUPP, P. M., MITCHELL, S. A., AND OWENS, J. D. 2011. Efficient maximal poisson-disk sampling. *ACM Trans. Graph.* 30, 4 (July), 49:1–49:12.
- EBEIDA, M. S., MITCHELL, S. A., PATNEY, A., DAVIDSON, A. A., AND OWENS, J. D. 2012. A simple algorithm for maximal poisson-disk sampling in high dimensions. *Comp. Graph. Forum* 31, 2pt4 (May), 785–794.
- EGAN, K., TSENG, Y.-T., HOLZSCHUCH, N., DURAND, F., AND RAMAMOORTHY, R. 2009. Frequency analysis and sheared reconstruction for rendering motion blur. *ACM Trans. Graph.* 28, 3 (July), 93:1–93:13.
- EGAN, K., HECHT, F., DURAND, F., AND RAMAMOORTHY, R. 2011. Frequency analysis and sheared filtering for shadow light fields of complex occluders. *ACM Trans. Graph.* 30, 2 (Apr.), 9:1–9:13.
- FATTAL, R. 2011. Blue-noise point sampling using kernel density model. *ACM Trans. Graph.* 30, 4 (July), 48:1–48:12.
- GRIBEL, C. J., DOGGETT, M., AND AKENINE-MÖLLER, T. 2010. Analytical motion blur rasterization with compression. In *HPG*, 163–172.
- GRIBEL, C. J., BARRINGER, R., AND AKENINE-MÖLLER, T. 2011. High-quality spatio-temporal rendering using semi-analytical visibility. *ACM Trans. Graph.* 30, 4 (July), 54:1–54:12.
- HACHISUKA, T., JAROSZ, W., WEISTROFFER, R. P., DALE, K., HUMPHREYS, G., ZWICKER, M., AND JENSEN, H. W. 2008. Multidimensional adaptive sampling and reconstruction for ray tracing. *ACM Trans. Graph.* 27, 3 (Aug.), 33:1–33:10.
- HAVRAN, V., BITTNER, J., HERZOG, R., AND SEIDEL, H.-P. 2005. Ray maps for global illumination. In *EGSR*, 43–54.
- JAROSZ, W., ZWICKER, M., AND JENSEN, H. W. 2008. The beam radiance estimate for volumetric photon mapping. *Comp. Graph. Forum* 27, 2 (Apr.), 557–566.
- JAROSZ, W., NOWROUZEZAHRAI, D., SADEGHI, I., AND JENSEN, H. W. 2011. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Trans. Graph.* 30, 1 (Feb.), 5:1–5:19.
- JAROSZ, W., NOWROUZEZAHRAI, D., THOMAS, R., SLOAN, P.-P., AND ZWICKER, M. 2011. Progressive photon beams. *ACM Trans. Graph.* 30, 6 (Dec.).
- JENSEN, H. W., AND BUHLER, J. 2002. A rapid hierarchical rendering technique for translucent materials. *ACM Trans. Graph.* 21, 3 (July), 576–581.
- JONES, T. R., AND PERRY, R. N. 2000. Antialiasing with line samples. In *EGRW*, 197–206.
- KOPF, J., COHEN-OR, D., DEUSSEN, O., AND LISCHINSKI, D. 2006. Recursive wang tiles for real-time blue noise. *ACM Trans. Graph.* 25, 3 (July), 509–518.

LAGAE, A., AND DUTRÉ, P. 2005. A procedural object distribution function. *ACM Trans. Graph.* 24, 4 (Oct.), 1442–1461.

LAGAE, A., AND DUTRÉ, P. 2008. A comparison of methods for generating Poisson disk distributions. *Computer Graphics Forum* 27, 1 (March), 114–129.

LEHTINEN, J., AILA, T., CHEN, J., LAINE, S., AND DURAND, F. 2011. Temporal light field reconstruction for rendering distribution effects. *ACM Trans. Graph.* 30, 4 (July), 55:1–55:12.

LI, H., WEI, L.-Y., SANDER, P. V., AND FU, C.-W. 2010. Anisotropic blue noise sampling. *ACM Trans. Graph.* 29, 6 (Dec.), 167:1–167:12.

LLOYD, S. 1983. An optimization approach to relaxation labeling algorithms. *Image and Vision Computing* 1, 2 (May), 85–91.

MCCOOL, M., AND FIUME, E. 1992. Hierarchical poisson disk sampling distributions. In *GI*, 94–105.

MEHTA, S. U., WANG, B., AND RAMAMOORTHI, R. 2012. Axis-aligned filtering for interactive sampled soft shadows. *ACM Trans. Graph.* 31, 6 (Nov.), 163:1–163:10.

MITCHELL, D. P. 1987. Generating antialiased images at low sampling densities. *SIGGRAPH Comput. Graph.* 21, 4 (Aug.), 65–72.

MITCHELL, D. P. 1991. Spectrally optimal sampling for distribution ray tracing. *SIGGRAPH Comput. Graph.* 25, 4 (July), 157–164.

NEHAB, D., AND SHILANE, P. 2004. Stratified point sampling of 3D models. In *PBG*, 49–56.

NOVÁK, J., NOWROUZEZHAI, D., DACHSBACHER, C., AND JAROSZ, W. 2012. Progressive virtual beam lights. *Comp. Graph. Forum* 31, 4 (June), 1407–1413.

NOVÁK, J., NOWROUZEZHAI, D., DACHSBACHER, C., AND JAROSZ, W. 2012. Virtual ray lights for rendering scenes with participating media. *ACM Trans. Graph.* 31, 4 (July), 60:1–60:11.

OSTROMOUKHOV, V., DONOHUE, C., AND JODOIN, P.-M. 2004. Fast hierarchical importance sampling with blue noise properties. *ACM Trans. Graph.* 23, 3 (Aug.), 488–495.

OSTROMOUKHOV, V. 2007. Sampling with polyominoes. *ACM Trans. Graph.* 26, 3 (July).

ÖZTIRELI, A. C., AND GROSS, M. 2012. Analysis and synthesis of point distributions based on pair correlation. *ACM Trans. Graph.* 31, 6 (Nov.), 170:1–170:10.

ÖZTIRELI, A. C., ALEXA, M., AND GROSS, M. 2010. Spectral sampling of manifolds. *ACM Trans. Graph.* 29, 6 (Dec.), 168:1–168:8.

PANG, W.-M., QU, Y., WONG, T.-T., COHEN-OR, D., AND HENG, P.-A. 2008. Structure-aware halftoning. *ACM Trans. Graph.* 27, 3 (Aug.), 89:1–89:8.

RAMAMOORTHI, R., MAHAJAN, D., AND BELHUMEUR, P. 2007. A first-order analysis of lighting, shading, and shadows. *ACM Trans. Graph.* 26, 1 (Jan.).

SOLER, C., SUBR, K., DURAND, F., HOLZSCHUCH, N., AND SILLION, F. 2009. Fourier depth of field. *ACM Trans. Graph.* 28, 2 (May), 18:1–18:12.

SUN, X., ZHOU, K., LIN, S., AND GUO, B. 2010. Line space gathering for single scattering in large scenes. *ACM Trans. Graph.* 29, 4 (July), 54:1–54:8.

TZENG, S., PATNEY, A., DAVIDSON, A., EBEIDA, M. S., MITCHELL, S. A., AND OWENS, J. D. 2012. High-quality parallel depth-of-field using line samples. In *HPG*, 23–31.

WEI, L.-Y., AND WANG, R. 2011. Differential domain analysis for non-uniform sampling. *ACM Trans. Graph.* 30, 4 (July), 50:1–50:10.

WEI, L.-Y. 2008. Parallel poisson disk sampling. *ACM Trans. Graph.* 27, 3 (Aug.), 20:1–20:9.

WEI, L.-Y. 2010. Multi-class blue noise sampling. *ACM Trans. Graph.* 29 (July), 79:1–79:8.

WHITE, K. B., CLINE, D., AND EGBERT, P. K. 2007. Poisson disk point sets by hierarchical dart throwing. In *RT*, 129–132.

ZHOU, Y., HUANG, H., WEI, L.-Y., AND WANG, R. 2012. Point sampling with general noise spectrum. *ACM Trans. Graph.* 31, 4 (July), 76:1–76:11.

Appendix: Quantitative measurement of Poisson disk sampling

Blue-noise properties can be achieved by Poisson disk point sampling, which randomly generates the samples such that no two samples are placed closer than a given Poisson disk radius r . The power spectrum of Poisson disk samples is quantitatively measured based on the distribution of the displacements among the samples, which is referred to as differential domain analysis [Wei and Wang 2011].

The frequency content of a point sample at \mathbf{x}_c has a constant unit magnitude:

$$\mathcal{F}(\delta(\mathbf{x} - \mathbf{x}_c)) = e^{-2\pi i(\boldsymbol{\omega}^T \mathbf{x}_c)}. \quad (27)$$

Given a set of ideal Poisson disk samples, the distribution of the displacements $\mathbf{d}_{i,j} = \mathbf{x}_j - \mathbf{x}_i$ follows the distribution

$$\rho(\mathbf{d}) = \begin{cases} 0, & |\mathbf{d}| < r \\ 1, & \text{else.} \end{cases} \quad (28)$$

Without loss of generality, $r = 1$ is assumed. After normalization, this Poisson disk sampling provides a unit impulse on the original point, and the power spectrum elsewhere is

$$Pow(\boldsymbol{\omega}) = 1 - {}_0F_1 \left[\frac{n}{2} + 1, -(\pi|\boldsymbol{\omega}|)^2 \right] \quad (29)$$

where ${}_0F_1$ is a confluent hypergeometric function. The power spectrum $\mathcal{P}(\boldsymbol{\omega})$ is close to zero when $|\boldsymbol{\omega}| < \frac{1}{r}$, so the Nyquist limit ν_N of ideal Poisson disk sampling is approximately $\frac{1}{r}$.

The effects of Poisson disk sampling are shown in Fig. 2. The uniform sampling in Fig. 2 (a) generates aliasing. The random sampling in Fig. 2 (b) shows white-noise properties, leading to substantial noise in reconstruction. The blue-noise sampling in Fig. 2 (c) provides the best reconstruction result with low noise and an absence of aliasing.

The comparison in Fig. 2 also provides an explanation in the frequency domain. Uniform sampling generates impulses and shows much anisotropy in the spectrum. Random sampling maintains a similar magnitude throughout the frequency domain. The blue-noise attenuates the frequency content near the original point while having little anisotropy.