

Aliasing-Free Blue Noise Sampling

Daniel Heck Thomas Schlömer Oliver Deussen

University of Konstanz
August 2012

Abstract

In this paper we revisit the problem of blue noise sampling with a strong focus on the spectral behavior of the sampling patterns. We use the mathematical relationship between the radial power spectrum and the radial distribution function to synthesize two types of blue noise patterns: *ideal* blue noise patterns that have a power spectrum in form of a step function and produce almost no coherent aliasing, and *effective* blue noise patterns that have a high effective Nyquist frequency and produce a controlled amount of aliasing. We give a definition for this effective Nyquist frequency in stochastic sampling and propose an error metric that characterizes the amount and spectral distribution of aliasing. We show that our blue noise sets avoid most of the artifacts caused by oscillations in the power spectra of existing blue noise patterns. Finally, we present a new algorithm for constructing point sets with a given power spectrum.

1 Introduction

A fundamental problem in computer graphics is that high image frequencies cannot always be removed by prefiltering which necessarily leads to some form of aliasing. One solution is to use sampling patterns with a *blue noise* power spectrum, which reduces the visibility of this aliasing by mapping it to incoherent noise.

Even though the blue noise properties are defined in the frequency domain, most methods for constructing blue noise patterns work purely in the spatial domain, usually by placing geometric constraints on the sample positions. Most of the resulting sampling patterns work very well if the image being sampled is dominated by low frequencies. An issue that has been ignored so far, however, is that these sampling patterns can produce artifacts at higher frequencies (Figure 1). Ironically, these artifacts are in fact low-frequency aliasing—exactly the kind of image error blue noise sampling was originally designed to prevent. As can already be seen in Figure 1, this aliasing is caused by oscillations in the high-frequency region of the sampling patterns' power spectrum.




Figure 1: A bird's eye view of the standard zone plate reveals coherent aliasing for blue noise patterns that have peaks in their power spectrum. Ideal blue noise patterns, as defined in Section 3.1, prevent coherent aliasing by mapping all high frequencies to white noise.

Our goal is to construct sampling patterns that conform more closely to the original definition of blue noise: a wide zero-region and a flat high-frequency region. As a reference, we define *ideal blue noise* as a power spectrum that has the shape of a step function (see the last column of Figure 1). To optimize the blue noise properties of such a spectrum, we show how to derive the maximum frequency at which the step can occur. The point sets corresponding to this maximum frequency differ significantly from the well-known Poisson disk patterns prevalent in computer graphics.

To assess the blue noise qualities of a sampling pattern we introduce two useful measures for the shape of the spectrum: the effective Nyquist frequency ν_{eff} and the spectrum oscillation Ω . These two measures suggest two broad categories of blue noise patterns:

- *low aliasing blue noise* is characterized by low oscillation in the spectrum, which guarantees that aliasing is mapped to broadband noise. This class contains stochastic and jittered grid sampling as well as the proposed ideal blue noise. The challenge for this category is to maximize the effective Nyquist frequency without causing oscillations.
- *high effective Nyquist blue noise* is characterized by a high value of ν_{eff} , which guarantees a large range of low frequencies that can be sampled without aliasing. This comes at the cost of more oscillation at higher frequencies and therefore an increased risk of coherent aliasing. The class contains most blue noise patterns proposed in recent years. The challenge for this category is to keep the oscillation as low as possible.

For this second category, we design *effective blue noise* sampling patterns whose power spectra are flat except for a single peak. The resulting patterns achieve a good tradeoff between ν_{eff} and Ω : the effective Nyquist frequency is comparable to that of many high-quality blue noise patterns, but by controlling the shape of the spectrum, we can limit the oscillation and hence the amount of coherent aliasing that can be introduced during image-plane sampling.

The algorithm used to construct sampling patterns that match a given power spectrum is discussed in Section 4. It is conceptually similar to the recent method by Zhou et al. [ZHWW12] in that it exploits the connection between Fourier analysis and spatial statistics to translate the construction problem into the spatial domain. However, in contrast to Zhou et al., our method is based on a one-dimensional statistic called the radial distribution function, which also underlies our theoretical arguments in Section 3.

We evaluate the presented sampling patterns in detail in Section 5 and show that for image-plane sampling, they perform as well or better than commonly used sample points. We conclude with a summary and a discussion of open questions in Section 6.

2 Background and Related Work

Aliasing is often unavoidable in computer graphics, since many signals aren't band-limited and exact prefiltering is only rarely possible. Stochastic sampling

can reduce the visibility of this aliasing by replacing it with incoherent noise [DW85, Co086]. Better results can generally be obtained with *blue noise* sampling. The original motivation for blue noise sampling was mainly empirical and followed from psychovisual arguments regarding the visibility of aliasing. Yellot [Yel83] proposed that aliasing is least conspicuous if

1. the power spectrum of the sample points is noisy and without concentrated spikes, and
2. the spectrum is close to zero for low frequencies.

A more rigorous motivation for blue noise sampling can be obtained using arguments from signal processing [Mit91].

2.1 Constructing Blue Noise Point Sets

Constructing sampling patterns with blue noise properties is a nontrivial task that has received a lot of attention in the last 20 years [LD08]. The problem is usually approached by arranging points using certain geometric constraints such as large mutual distances between sample points, centroidal Voronoi tessellations, etc. Constructing blue noise patterns is still an active area of research; a few recent publications are [SGBW10, Fat11, SHD11, CYC*12].

Most construction algorithms work in the spatial domain, even though the “blue noise” property itself is defined in the Fourier domain. The only construction algorithms we are aware of that take frequency properties into direct account are the paper by Mitsa and Parker [MP92] and the recent work by Zhou et al. [ZHWW12]. Zhou et al. construct point sets matching a given power spectrum by performing a gradient descent optimization on an energy derived from the autocorrelation function. A similar approach based on simulated annealing was proposed in the context of solid state physics [RT97]. We present a third algorithm for constructing point sets with a given spectrum in Section 4.

2.2 Analysis of Sampling Patterns

Fourier analysis is the standard method for studying the properties of sample patterns [Uli93, LD08]. In Appendix A, we summarize our conventions for the Fourier transform and power spectrum and list the most important variables. We also normalize many of the quantities used in this paper to make their value independent of the number of sample points N .

As an alternative to Fourier analysis, it is also possible to characterize point sets using spatial statistics [IPSS08]. Spatial statistics are widely used in many domains that study the arrangement of point-like objects, such as solid-state physics, geology, and astronomy, but have seen little use in computer graphics. Two main statistics are particularly relevant in our case: the *autocorrelation function* and the *radial distribution function* (RDF).

The autocorrelation $A_f(\mathbf{d}) = (f(\mathbf{x}) \star \bar{f}(-\mathbf{x}))(\mathbf{d})$ measures the self-similarity of a signal f under translation by a vector \mathbf{d} . For a sampling pattern s , it is easy

to show that its power spectrum is related to its autocorrelation via a Fourier transform

$$P(\boldsymbol{\nu}) = N^{-1} \mathcal{F}[A_s(\boldsymbol{d})].$$

Up to normalization, $A_s(\boldsymbol{d})$ is therefore equivalent to the “differential distribution function” $p(\boldsymbol{d})$ proposed by Wei and Wang [WW11].

The RDF measures the distribution of point distances and is also known as the pair-correlation function [IPSS08]. It gives the probability of finding another point at a certain distance from a reference point. The bottom part of Figure 1 shows the RDFs for several prominent sampling patterns. Lau and Ulichney [LUA03] did use RDF diagrams to qualitatively illustrate the spatial distribution of points, but we aren’t aware of other applications in computer graphics. The RDF $g(r)$ is related to the *radial power spectrum* $P(\nu)$ by a Hankel transform

$$P(\nu) = 1 + n\mathcal{H}[g(r)] \quad (1)$$

(see Appendix B). We will use the RDF and its relation to the radial power spectrum extensively in the remainder of this paper.

RDF and autocorrelation both measure the distribution of inter-point distances and are closely related; the main difference is that the RDF is a function of the scalar distance whereas the autocorrelation is a function of the vector distance. We primarily focus on the RDF in this paper because it is easier to handle analytically. Its dependence on absolute distances is not a restriction, since sampling applications generally call for isotropic point sets.

3 Low-Oscillation Blue Noise

The example in Figure 1 demonstrates that low-frequency noise occurs in the sampled image if the power spectrum oscillates; we will explain this in detail in Section 5.2. In this section we construct two classes of blue noise patterns that prevent such artifacts by limiting the amount of oscillation.

3.1 Ideal Blue Noise

We begin with an idealization of blue noise which we call *ideal blue noise*. The power spectrum of ideal blue noise is zero in low frequencies and constant in high frequencies, i.e.,

$$P_{\text{ideal}}(\nu; \nu_0) = n \frac{\delta(\nu)}{2\pi\nu} + H(\nu - \nu_0).$$

Here we have included the DC peak and used the Heaviside step function H (cf. Figure 2a). It has been shown that it is possible to construct point distributions with such a power spectrum [ZHWW12] but what is the largest value of ν_0 for which $P_{\text{ideal}}(\nu; \nu_0)$ is realizable?

There are two necessary conditions a power spectrum $P(\nu)$ must fulfill to be realizable by a point distribution [CTS03, UST06]: since both $P(\nu)$ and the associated RDF $g(r)$ are non-negative by definition, we must have

$$g(r) \geq 0, \quad P(\nu) \geq 0. \quad (2)$$




Figure 2: Radial power spectra, corresponding RDFs, 2D power spectra and cutouts from the corresponding points sets for ideal blue noise and effective blue noise. (a) A radial power spectrum that is a perfect step is only possible up to a maximum frequency ν_{\max} . For frequencies $\nu_0 > \nu_{\max}$ the corresponding RDF becomes negative. (b) For a higher effective Nyquist frequency, we have to allow the power spectrum to be above 1 by introducing a (potentially smoothed) energy peak around a frequency $\nu_0 > \nu_{\max}$. To keep the corresponding RDF positive, we have to increase the height p_0 of the peak if we increase ν_0 .




Figure 3: The spectrum of the sampled signal \hat{I}_s is obtained by convolving the original spectrum with the Fourier transform of the sampling function. Aliasing occurs when the replicated spectra \hat{I} overlap the central spectrum on the right-hand side.

Because both functions are linked via a Hankel transform, these two conditions severely limit the range of realizable power spectra. Our experiments so far suggest that these *realizability conditions* are not only necessary but also sufficient: for every power spectrum/RDF pair that fulfills Eq. (2), we have been able to find corresponding point distributions.

For the ideal blue noise spectrum P_{ideal} the associated RDF is

$$g_{\text{ideal}}(r) = 1 - \frac{2\pi\nu_0^2}{n} \text{jinc}(2\nu_0 r).$$

This follows from Eq. (14) and the Hankel transforms in Table A.3. Solving $g_{\text{ideal}}(r) \geq 0$ for ν_0 , we obtain

$$\nu_0 \leq \nu_{\text{max}} = \sqrt{n/\pi}.$$

In other words, ν_{max} is the highest realizable position of the step: Figure 2(a) demonstrates that for $\nu_0 > \nu_{\text{max}}$, the RDF g_{ideal} becomes negative. The only way to move the position of the step further to the right is to increase the sample density n . The figure also shows one resulting point set for $\nu_0 = \nu_{\text{max}}$. Note that the point distribution differs significantly from usual Poisson disk patterns in that it contains many closely spaced point pairs. We will evaluate this sampling pattern in more detail in Section 5.

3.2 Quantifying Blue Noise

By comparing the power spectra in Figure 1, we see that ideal blue noise has a narrower zero region than classic blue noise patterns. Increasing the zero region beyond ν_{max} requires us to sacrifice some of the flatness of $P(\nu)$ in the high-frequency region. We will discuss a controlled way to do this in the next section. As a preparation, we introduce two numerical measures that characterize the shape of the power spectrum: the *effective Nyquist frequency* ν_{eff} measures the size of the zero region and indicates the range of frequencies that can be represented with almost no aliasing; the *oscillation* Ω measures the amount of oscillation in the power spectrum, and therefore the risk of coherent aliasing.

For uniform sampling with sample spacing d , the range of frequencies that can be reconstructed without aliasing is given by the Nyquist frequency $\nu_c = 1/2d$. This simple relationship between spectral properties and sample distances unfortunately breaks down for non-uniform sampling, but it is possible to define the equivalent of the Nyquist frequency directly in the frequency domain.

In general, the spectrum of a sampled image I_s is given by $\hat{I}_s = \hat{I} \star S$, where \hat{I} is the original image spectrum and S the Fourier transform of the sample set (cf. Appendix A.1). This relationship is visualized in Figure 3. This diagram suggests that the range of aliasing-free frequencies roughly equals half the radius of the zero-region in the power spectrum. To formalize this idea, we consider the average energy in the power spectrum up to a certain frequency ν

$$P_{\text{avg}}(\nu) = \frac{1}{\pi\nu^2} \int_{|\mathbf{v}'| < \nu} P(\mathbf{v}') d\mathbf{v}'$$

and define the *effective Nyquist frequency* ν_{eff} as the largest frequency so that $P_{\text{avg}}(2\nu_{\text{eff}})$ stays below a given threshold

$$\nu_{\text{eff}} = \max\{\nu : P_{\text{avg}}(2\nu) \leq E_\tau\}. \quad (3)$$

We found that a threshold of $E_\tau = 0.1$ gives useful results. As expected, the effective Nyquist frequency of ideal blue noise is $\nu_{\text{eff}} \approx \nu_0/2$.

To measure the amount of oscillation of the power spectrum, we use the standard deviation of $P(\mathbf{v})$ from the 1-level

$$\Omega = \left(\frac{1}{|R|} \int_R |P(\mathbf{v}) - 1|^2 d\mathbf{v} \right)^{1/2}.$$

Here, R is the integration domain and $|R|$ its area. We exclude the zero-region of the power spectrum by integrating over the ring $R = \{\mathbf{v} : 2\nu_{\text{eff}} \leq |\mathbf{v}| \leq 40\nu_{\text{hex}}\}$.

This definition is conceptually similar to the strict Nyquist frequency in the case of uniform sampling: Frequencies below ν_{eff} can be sampled and reconstructed with little error. Frequencies above ν_{eff} , on the other hand, are replaced by aliasing. In this case, the magnitude of Ω determines whether this noise is, on average, closer to white noise (Ω small) or colored noise (Ω large).

3.3 Effective Blue Noise

The ideal blue noise patterns from Section 3.1 prevent coherent aliasing by keeping the power spectrum flat, but this comes at a cost: The effective Nyquist frequency of these patterns cannot be higher than $\sqrt{n/4\pi}$, which is about 58% of the maximum Nyquist frequency of a hexagonal lattice. In this section, we discuss a class of blue noise patterns that offer a much higher Nyquist frequency (up to 86% of ν_{hex}) by introducing some oscillation into the power spectrum. By directly influencing the shape and the amount of oscillation, we obtain *effective blue noise* patterns that combine a high effective Nyquist frequency with comparatively low oscillation.

3.3.1 Increasing the Effective Nyquist Frequency

The main problem when constructing blue noise patterns with an effective Nyquist frequency higher than the ν_{\max} of ideal blue noise is that we cannot choose $P(\nu)$ arbitrarily due to the constraints imposed by the realizability conditions in Eq. (2). In general, finding power spectra that are realizable is a nontrivial problem because the space of functions that obey the realizability conditions is not easy to parametrize [GP06, UST06].

We briefly note two ways of constructing realizable power spectra with a higher effective Nyquist frequency that we ultimately discarded due to the lack of control over the shape of the spectrum. One option is to subject an ideal blue noise pattern to an optimization method such as Lloyd’s method [Llo82, BSD09, CYC*12]. Even though this gradually increases ν_{eff} , all of these methods operate in the spatial domain and give us no direct control over the height and shape of oscillations in the power spectrum. Another option is to exploit the duality between RDF and power spectrum: As we show in Appendix D, it is possible to generate point sets whose power spectrum matches the RDF of another (dual) point set. This approach gives access to a large set of realizable power spectra, but it also lacks the direct control over the spectrum that we seek.

3.3.2 Designing Effective Blue Noise Power Spectra

To obtain more direct control over the effective Nyquist frequency and the amount of oscillation, we study the following generalization of the ideal blue noise spectrum

$$P_{\text{eff}}(\nu; \nu_0; p_0) = n \frac{\delta(\nu)}{2\pi\nu} + G_\sigma \star (p_0 \delta(\nu - \nu_0) + H(\nu - \nu_0)).$$

Compared to P_{ideal} we added a Dirac peak of power p_0 at the step frequency ν_0 . Both this peak and the step function are optionally convolved with a Gaussian kernel G_σ with standard deviation σ . Figure 2(b) illustrates the shape of this power spectrum.

This family of blue noise spectra has three interesting properties:

1. Aside from the single peak at ν_0 , there are no oscillations.
2. The width and height of the peak can be controlled by adjusting the smoothing radius σ and the peak energy p_0 .
3. The ideal spectrum P_{ideal} is included as a special case.

Not all spectra in this family are realizable, however, and due to the realizability conditions, the parameters cannot be adjusted independently. The main challenge therefore is to find combinations of the three parameters ν_0 , p_0 , and σ that are realizable and yield good sampling patterns.

We explored this family of effective blue noise patterns empirically and searched for configurations where at the same time

- ν_0 is as high as possible,




Figure 4: We can explore the space of effective blue noise point sets by configuring the power spectrum from Figure 2(b). (a) Ideal blue noise with a step of constant power and $\nu_{\text{eff}} = 0.58$. (b) We can increase ν_{eff} by moving energy to higher frequencies—in this case a Dirac peak of height $P(\nu_0) = 16$ —but only at the cost of coherent aliasing (moiré patterns around ν_0). (c) Smoothing the peak to $P(\nu_0) \approx 4$ yields acceptable aliasing at the cost of a slightly reduced ν_{eff} . (d) The highest effective Nyquist frequency we can produce without severely contaminating frequencies below or above ν_0 . (e) The best compromise we found between a high ν_{eff} and coherent aliasing.

- coherent aliasing is kept at an acceptable level (by appropriate choice of p_0 and σ), and
- energy at frequencies below and above ν_0 are free of impurities, i.e., $P(\nu) \approx 0, \nu < \nu_0$ and $P(\nu) \approx 1, \nu > \nu_0$.

Figure 4 shows some of the results we generated. Ideal blue noise with $\nu_0 = \nu_{\max}$ is used as a starting point and is depicted to the left (a). We then increase ν_0 to $1.33\nu_{\max}$ (b, c) and $1.66\nu_{\max}$ (d, e). For each of those pairs, we show one result with a small σ and one with a large σ .

Generally, the more we increase the peak height, the further we are able to push ν_0 and thus the effective Nyquist frequency. If we take this too far, however, strong aliasing can show up in the sampled image. The zone plate renderings at the top of Figure 4 demonstrate that visible moiré patterns emerge once $P(\nu_0) > 4$. We can reduce these artifacts by increasing the amount of smoothing σ . This has two effects: It decreases the height of the peak and increases its width, which means that aliasing is scattered over a wider range of frequencies. The highest effective Nyquist frequency we could produce without severely contaminating frequencies below or above ν_0 is depicted in column (d).

The best compromise we found between a high effective Nyquist frequency and coherent aliasing is shown in (e); in the remainder of the paper we will refer to it as *effective blue noise*. This configuration yields an effective Nyquist frequency comparable to classic blue noise patterns but without sharing their high-frequency oscillations that may lead to additional aliasing (cf. Table 1).

4 Spectrum Matching Algorithm

Before we evaluate the performance of the new proposed sampling patterns in Section 5, we briefly describe the algorithm we used to synthesize sampling patterns from the power spectra we designed in the previous section. As mentioned in Section 2.1, similar algorithms have already been proposed by other researchers. In our tests, the annealing approach by Rintoul and Torquato [RT97] converged only slowly and didn't give good results for reasonably large point sets. The paper by Zhou et al. [ZHWW12] wasn't available when we started this work, but judging from the examples in that paper, the results seem to be very similar. We look forward to comparing both approaches in more depth in the future.

From an algorithmic point of view, our method is very similar to that of Zhou et al.: Both algorithms synthesize the point sets by reformulating the problem in the spatial domain, and both iteratively update the positions of all points by applying a force to each point

$$\mathbf{x}'_i = \mathbf{x}_i + h \times \mathbf{F}_i, \quad (4)$$

where h is a step size parameter and F is a force that depends on the current point positions. The main difference between both algorithms is how they calculate the forces in Eq. (4): Zhou et al. propose a force based on gradient descent, whereas ours is motivated geometrically.

We now describe our approach in more detail. First, the “target” power spectrum to be matched $P_t(v)$ is transformed into an equivalent RDF $g_t(r)$ using a numerical Hankel transform (Appendix B). Since both functions are one-dimensional, our approach can only synthesize isotropic point sets. The current point set is initialized with a random distribution of points. To evolve the point set towards the target distribution, we let all particles attract or repel each other using forces of the form

$$\mathbf{F}_i = \sum_{j \neq i} f(|\mathbf{x}_i - \mathbf{x}_j|) \frac{\mathbf{x}_i - \mathbf{x}_j}{|\mathbf{x}_i - \mathbf{x}_j|^2}. \quad (5)$$

The function $f(r)$ determines the degree of attraction or repulsion and is defined as follows:

$$f(r) = \underbrace{\int_0^r g(x) dx}_{G(r)} - \underbrace{\int_0^r g_t(x) dx}_{G_t(r)}. \quad (6)$$

It may seem surprising that f is defined in terms of the *integral* of the RDF difference. This can be explained as follows: Since the RDF measures the density of points at a certain distance, $G(r)$ measures the average point density in circular regions of radius r . Consider a fixed distance r_0 . If $G(r_0) > G_t(r_0)$, the current point distribution contains too many point pairs that are closer than r_0 , compared to the target distribution. In this case the point set as a whole has to spread out, so $f(r_0)$ should be repulsive to make room. Conversely, if $G(r_0) < G_t(r_0)$, too many pairs have a distance greater than r_0 , so the points have to move closer together; in this case, $f(r_0)$ should be attractive. In both cases, the choice $f(r) = G(r) - G_t(r)$ fulfills this condition.

In each full iteration, we first calculate all the forces \mathbf{F}_i and then update the positions according to Eq. (4). The main parameter during each iteration is the step size h , which is chosen adaptively depending on the largest force as in Zhou et al. [ZHWW12]. We use an additional temperature parameter T to reduce the step size whenever the optimization gets stuck

$$h = T \frac{F_{\max}}{\sqrt{N}}, \quad F_{\max} = \max_i \|\mathbf{F}_i\|.$$

The energy $E = \|g(r) - g_t(r)\|_2$ tracks the progress of the optimization. The optimization is considered stuck if E hasn’t reduced during the last 20 iterations; in this case, the temperature is reduced by a constant factor. The algorithm terminates once the temperature has fallen below 10^{-3} .

The number of bins we use for calculating the RDFs and the force f is set to $N_{\text{bins}} = N$. Significantly fewer bins lower the resolution of the force and can lead to inaccurate results, and significantly more bins lead to noisier RDF estimates. We consider the RDF in the interval $[0, 0.5)$; smaller intervals can be used to speed up some of the computations [WW11], but we have found that this reduces the quality of the results for RDFs that decay slowly.

Since a certain amount of noise is inevitable when estimating RDFs from finite point sets, we optionally smooth the histograms using a Gaussian kernel.

	<i>Method</i>	ν_{eff}	Ω	δ_{min}	δ_{avg}	Q_6	Note
Low Alias.	Random	0	0.05	0.01	0.47	0.35	F
	Jittered Grid [Coo86]	0.24	0.27	0.05	0.59	0.37	F
	Dart Throwing [LD08]	0.58	0.99	0.76	0.80	0.42	
	<i>Ideal Blue Noise</i>	0.58	0.03	0.09	0.64	0.36	F
Eff. Nyq.	<i>Effective Blue Noise</i>	0.86	0.75	0.55	0.80	0.40	
	CCCVT Centroids [BSD09]	0.89	1.22	0.75	0.88	0.53	O
	El. Halftoning [SGBW10]	0.89	1.29	0.74	0.88	0.52	O
	FPO [SHD11]	0.90	4.06	0.93	0.93	0.47	O
Hex.	CVT Centroids [Llo82]	0.98	2.82	0.80	0.94	0.85	R,O
	Rank-1 [DKD08]	1	11.1	0.99	0.99	0.997	R

Table 1: Comparison of several frequency and spatial statistics of sampling patterns. The last column marks methods that are (F)ree of coherent aliasing and methods that are either (R)egular or show strong (O)scillations in their power spectrum. The results are averages over ten sets containing 4096 points.

There is no simple rule for choosing the optimal width σ of the Gaussian, since this involves a tradeoff between reducing noise vs. keeping relevant information in the RDF. In our experiments, good values for σ were between 0 and $16/N_{\text{bins}}$. All the results in this paper have been generated with $\sigma = 8/N_{\text{bins}}$, and no parameters had to be adjusted manually.

5 Evaluation

We now evaluate our ideal and effective blue noise patterns and study some of their properties. First, we briefly compare their main characteristics to other sampling patterns. We then use two classes of synthetic test images to evaluate their performance during image-plane sampling: cosine waves and a variation of the standard zone plate. For cosine waves we include a theoretical error analysis and show how the power spectrum affects the error introduced by sampling. The zone plate is used to demonstrate how our sampling patterns behave over large frequency ranges.

5.1 Properties

Table 1 summarizes the most important characteristics of our blue noise sets and compares them with other well-known sampling patterns. We have sorted the sampling patterns into the two categories mentioned in the introduction: those with low aliasing, and those with a high effective Nyquist frequency. For reference, we have also included two point sets that are closer to the hexagonal lattice.

The effective Nyquist and oscillation measures from Section 3.2 are shown in the first two columns. We see that for traditional blue noise patterns, a high

ν_{eff} comes at the cost of a high oscillation as well. This interdependence was sometimes referred to as the noise-aliasing tradeoff [DW85, Gla95]. The point sets constructed in this paper demonstrate that it is possible to achieve high values of ν_{eff} with only little oscillation, so the tradeoff is not strict.

The smallest and average nearest-neighbor distance is measured by δ_{\min} and δ_{avg} [SHD11]. The value of δ_{\min} corresponds to the Poisson disk radius while δ_{avg} roughly measures how uniform the points are distributed. Our point sets have comparatively low values for both measures. Even though a large separation between samples can be beneficial, as exemplified by the success of Poisson disk patterns, it is not a prerequisite for efficient blue noise sampling.

Q_6 is the *bond-orientational order*, which measures how close to a hexagonal arrangement a point distribution is [KTT00]. Point sets with $Q_6 < 0.5$ have a high amount of “orientational disorder” or anisotropy. Both presented point sets are irregular in this sense.

5.2 Analysis of Sampling Artifacts

In this section we derive an estimate for the error introduced by blue noise sampling, or more specifically, for the spectrum of this error. This analysis will clarify how the shape of the power spectrum affects the amount and the spectral distribution of aliasing in the resulting image and why the low-oscillation patterns constructed in Section 3 perform better than other blue noise patterns. A similar analysis was already performed by Dippé and Wold [DW85], but they only derived the sampling error for constant functions.

Recall from Section 3.2 that the spectrum of a sampled image I_s is given by $\hat{I}_s(\mathbf{v}) = S \star \hat{I}(\mathbf{v})$, with S being the Fourier transform of the sampling pattern and \hat{I} the Fourier transform of the original image. For sampling patterns with a finite number of points, the function S has two components: a DC peak at the origin and a “noisy” remainder S' , so we can write

$$\hat{I}_s(\mathbf{v}) = [n\delta(\mathbf{v}) + S'(\mathbf{v})] \star \hat{I}(\mathbf{v}).$$

Likewise, we will denote by $P'(v)$ the power spectrum without the DC peak. The error introduced by sampling is the difference between the sampled and the original image

$$E(\mathbf{v}) = |\hat{I}_s(\mathbf{v}) - \hat{I}(\mathbf{v})|^2 = |S' \star \hat{I}(\mathbf{v})|^2. \quad (7)$$

Note that we are primarily interested in the low-frequency region of $E(\mathbf{v})$ up to the Nyquist frequency of the pixel grid, which describes the frequency distribution of aliasing. The high-frequency region of $E(\mathbf{v})$ is removed by the reconstruction process. For a constant image function, Eq. (7) gives us $E(\mathbf{v}) \propto |S'(\mathbf{v})|^2 = P'(\mathbf{v})$, which is the result by Dippé and Wold.

As an important and illustrative class of non-constant images, we now consider plane cosine waves of the form

$$I_f(\mathbf{x}) = \cos(2\pi\mathbf{f} \cdot \mathbf{x}), \quad \hat{I}_f(\mathbf{v}) = \frac{1}{2}(\delta(\mathbf{v} - \mathbf{f}) + \delta(\mathbf{v} + \mathbf{f})).$$




Figure 5: Illustration of sampling error for sinusoids of increasing frequencies. The plots in the rightmost column show the frequency distribution of the sampling error. Since we are only interested in the overall trend of this error, we plot $\sqrt{E_f}$ to make the plots easier to read. The frequency axis is relative to v_{px} , the Nyquist frequency of the pixel grid. Note how the shape of the error curve corresponds to noise in the image: flat curves show up as white noise, whereas peaks in the curve lead to different forms of colored noise. (Also see the enlarged versions in Appendix E.)

Plugging this into Eq. (7) gives the sampling error

$$E_f(\mathbf{v}) = \frac{1}{2} |S'(\mathbf{v} - \mathbf{f}) + S'(\mathbf{v} + \mathbf{f})|^2.$$

This error is bounded from above by the following term which depends only on the power spectrum:

$$E_f(\mathbf{v}) \leq \frac{P'(\mathbf{v} - \mathbf{f}) + P'(\mathbf{v} + \mathbf{f})}{2} + \sqrt{P'(\mathbf{v} - \mathbf{f})P'(\mathbf{v} + \mathbf{f})} \quad (8)$$

The proof is elementary and shown in Appendix C. Even though this is only an upper bound, we will see that it accurately predicts the frequency characteristics of the sampling error.

A few things can be read off directly from this formula. Obviously, the sampling error is obtained by combining shifted copies of the power spectrum. If $P'(v)$ has a peak around frequency f_0 , sampling a cosine with the same frequency will cause copies of this peak to overlap at the origin. The width of this peak then determines the appearance of this aliasing: A broad peak leads to broadband aliasing which appears as unstructured noise, whereas a narrow peak produces low-frequency aliasing that appears as structured noise.

This is illustrated in Figure 5 which compares the upper bound of $E_f(v)$ with actual sampled images. As can be seen, ideal blue noise and effective blue noise produce the cleanest images at low frequencies. Dart throwing performs especially badly since its power spectrum is non-zero for all frequencies, so $E_f(v)$ can never be zero. Ideal blue noise produces high-frequency noise quite early, but as we increase the signal frequency, this noise quickly takes the shape of white noise, as shown by the flat response in the last two rows. In all these cases, the visual appearance matches the predicted sampling error in the curves to the right.

Both effective blue noise and farthest-point optimization (FPO) can cause coherent aliasing for certain image frequencies, as demonstrated in the last two rows of Figure 5. The slow decay of a traditionally constructed sampling pattern like FPO means that every peak in the spectrum can potentially cause artifacts. For effective blue noise, on the other hand, only one range of frequencies is affected because it contains a single peak only.

5.3 Application to Image-Plane Sampling

We test the image-plane sampling qualities of our blue noise patterns using a variant of the common zone plate test function [Mit90]. This function is a radial sine wave of the form $z(r) = \frac{1}{2}[1 + \cos(\alpha r^2)]$. Aliasing is unavoidable when sampling z since the local frequency of z increases linearly with the distance r from the origin.

One problem with the standard definition of the zone plate is that it is hard to configure. Since we would like to configure the zone plate renderings with respect to particular frequencies (such as peaks in the power spectrum or the Nyquist frequency of the pixel grid ν_{px}), we used the following generalized




Figure 6: Here, we compare the image-plane sampling qualities of ideal and effective blue noise patterns with strong competitors within their respective categories. (*Left*) If the guarantee for a low amount of coherent aliasing is desired, ideal blue noise performs significantly better than traditional dart throwing by keeping valuable low frequency content much cleaner. (*Right*) A similar observation can be made for methods that aim for a high effective Nyquist frequency: the method by Schömer et al. [2011] offers one of the highest effective Nyquist frequencies, but low-frequency content appears not quiet as clean as with effective blue noise sampling. In addition, effective blue noise shows no coherent aliasing beyond the narrow range of frequencies around $2\nu_{\text{eff}}$. (Also see the enlarged versions in Appendix E.)

zone plate function

$$z'(r) = \frac{1}{2}(1 + \cos[v(\alpha r + \phi)])$$

The frequency ν , the parameter α , and the phase ϕ are defined as

$$\begin{aligned} \nu &:= \nu_{\text{cut}} \nu_{\text{px}} & \nu_{\text{cut}} &\geq 0, \\ \alpha &:= \min(r/r_{\text{cut}}, 1) & r_{\text{cut}} &> 0, \\ \phi &:= \max(r - r_{\text{cut}}, 0). \end{aligned}$$

Here, ν_{cut} denotes the maximum frequency that is produced by z' relative to ν_{px} , and r_{cut} the desired distance of this maximum frequency to the origin. For all zone renderings in this paper, we have tiled toroidal sets of 4096 points over the image-plane, and used a Lanczos filter with a support of width 4 for resampling.

Figure 1 already demonstrated that—in contrast to traditional blue noise patterns—ideal blue noise produces virtually no coherent aliasing, except for a small amount close to the step. This observation is confirmed by Figure 6, which directly compares ideal blue noise to dart throwing, which is the strongest competitor in the low-aliasing category. We focus on the low-frequency region by setting $\nu_{\text{cut}} = 5, r_{\text{cut}} = 1$. When using only one sample per pixel, the difference is negligible, but already at 4 spp, the improvement for low frequencies is evident.

The right-hand side of Figure 6 shows a similar comparison for methods with a high effective Nyquist frequency. The method by Schlömer et al. [SHD11] offers one of the highest effective Nyquist frequencies (cf. Table 1), but low-frequency image content appears not quite as clean as with effective blue noise sampling. Effective blue noise also shows no coherent aliasing beyond the narrow range of frequencies around $2\nu_{\text{eff}}$.

6 Conclusion

In this paper, we have revisited the problem of blue noise sampling with a stronger emphasis on the shape of the power spectrum than most other recent publications. Starting from the observation that oscillations in the power spectrum can lead to aliasing artifacts we have proposed two improvements to standard blue noise patterns:

- *ideal blue noise*, which has a step-like power spectrum, and
- *effective blue noise*, which has a single configurable peak but is otherwise flat.

Both blue noise types perform well in image-plane sampling tasks: low frequencies are represented at least as clean as with other sampling patterns, but high frequencies are less likely to lead to aliasing, because they are mapped to white noise more efficiently.

The main tradeoff is in the representation of middle frequencies, and in practical applications this should determine the choice of sampling pattern. There are good reasons to choose sampling patterns with higher oscillations, or even regular sampling patterns, for example if we know that the image being sampled contains little energy in frequency ranges that can cause aliasing, or if the artifacts that do occur are negligible compared to the increased image quality in medium frequencies. The advantage of sampling patterns like ideal blue noise is that they guarantee a low amount of coherent aliasing at the cost of a lower effective Nyquist frequency. Patterns like effective blue noise trade a better representation of low frequencies for a higher risk of coherent aliasing in middle frequencies. This is basically the classical noise-aliasing tradeoff, but we have shown that we gain considerable control over this tradeoff by the ability to directly shape the power spectrum.

We are confident that putting a stronger focus on the spectral properties of sampling patterns instead of their geometric properties is a fruitful avenue for further research. Some interesting open questions are:

- We have only studied a very limited subset of realizable blue-noise spectra, namely step functions and step functions with a single bulge at the transition. Differently shaped power spectra might perform even better for image-plane sampling.
- All currently known methods for constructing point sets from the power spectrum rely on iterative optimization and are therefore restricted to offline processes. Generating similar patterns on-line would require a completely different algorithmic approach.
- We introduced a new way to predict the frequency distribution of the aliasing error in Section 5.2. Can we derive more general predictions from this analysis or extend it to other interesting classes of signals?
- The sampling patterns we have proposed have the unique property that only a single range of image frequencies can produce coherent aliasing. Is it possible to exploit this by treating those frequencies specially during prefiltering or postfiltering?

Acknowledgments

We thank Victor Ostromoukhov for the stimulating discussions on blue noise sampling that ultimately led to this paper.

References

- [Bra99] Bracewell R. N.: *The Fourier Transform and Its Applications*, 3rd ed. McGraw-Hill, New York, 1999.
- [BSD09] Balzer M., Schlömer T., Deussen O.: Capacity-constrained point distributions: A variant of Lloyd's method. *ACM Trans. Graph.* 28, 3 (2009), 86:1–8.

- [Coo86] Cook R. L.: Stochastic sampling in computer graphics. *Computer Graphics (Proc. of SIGGRAPH 86)* 5, 1 (1986), 51–72.
- [CTS03] Crawford J., Torquato S., Stillinger F. H.: Aspects of correlation function realizability. *Journal of Chemical Physics* 119, 14 (Oct. 2003), 7065–7074.
- [CYC*12] Chen Z., Yuan Z., Choi Y.-K., Liu L., Wang W.: Variational blue noise sampling. *IEEE Transactions on Visualization and Computer Graphics* 18 (2012), 1784–1796.
- [DKD08] Dammertz H., Keller A., Dammertz S.: Simulation on rank-1 lattices. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 2008, pp. 205–216.
- [DW85] Dippé M. A. Z., Wold E. H.: Antialiasing through stochastic sampling. *Computer Graphics (Proc. of SIGGRAPH 85)* 19, 3 (1985), 69–78.
- [Fat11] Fattal R.: Blue-noise point sampling using kernel density model. *ACM Trans. Graph.* 30, 4 (Aug. 2011), 48:1–48:12.
- [Gla95] Glassner A. S.: *Principles of Digital Image Synthesis*, first ed., vol. 1. Morgan Kaufmann, San Francisco, CA, USA, Mar. 1995.
- [GP06] Giraud B. G., Peschanski R.: On positive functions with positive Fourier transforms. *Acta Physica Polonica B* 37 (Feb. 2006), 331.
- [IPSS08] Illian J., Penttinen A., Stoyan H., Stoyan D.: *Statistical Analysis and Modelling of Spatial Point Patterns*. John Wiley & Sons, 2008.
- [KTT00] Kansal A. R., Truskett T. M., Torquato S.: Nonequilibrium hard-disk packings with controlled orientational order. *Journal of Chemical Physics* 113, 12 (2000), 4844.
- [LD08] Lagae A., Dutré P.: A comparison of methods for generating Poisson disk distributions. *Computer Graphics Forum* 27, 1 (2008), 114–129.
- [Llo82] Lloyd S. P.: Least square quantization in PCM. *IEEE Transactions on Information Theory* 28, 2 (1982), 129–137.
- [LUA03] Lau D. L., Ulichney R., Arce G. R.: Blue- and green-noise halftoning models. *IEEE Signal Processing Magazine* 20, 4 (July 2003), 28–38.
- [Mit90] Mitchell D. P.: The anti-aliasing problem in ray tracing. *Advanced Topics in Ray Tracing, ACM SIGGRAPH 1990 Course Notes*, 1990.
- [Mit91] Mitchell D. P.: Spectrally optimal sampling for distribution ray tracing. *Computer Graphics (Proc. of SIGGRAPH 91)* 25, 4 (1991), 157–164.

- [MP92] Mitsa T., Parker K. J.: Digital halftoning technique using a blue-noise mask. *J. Opt. Soc. Am. A* 9, 11 (Nov. 1992), 1920–1929.
- [RT97] Rintoul M. D., Torquato S.: Reconstruction of the structure of dispersions. *Journal of Colloid and Interface Science* 186 (1997), 467–476.
- [SGBW10] Schmaltz C., Gwosdek P., Bruhn A., Weickert J.: Electrostatic halftoning. *Computer Graphics Forum* 29, 8 (Dec. 2010), 2313–2327.
- [SHD11] Schlömer T., Heck D., Deussen O.: Farthest-point optimized point sets with maximized minimum distance. In *High Performance Graphics 2011* (2011), Eurographics Association, pp. 135–154.
- [Uli93] Ulichney R. A.: *Digital Halftoning*. MIT Press, 1993.
- [UST06] Uche O. U., Stillinger F. H., Torquato S.: On the realizability of pair correlation functions. *Physica A* 360, 21 (2006), 21–36.
- [WW11] Wei L.-Y., Wang R.: Differential domain analysis for non-uniform sampling. *ACM Trans. Graph.* 30, 4 (July 2011), 50:1–50:10.
- [Yel83] Yellot Jr. J. I.: Spectral consequences of photoreceptor sampling in the rhesus retina. *Science* 221 (1983), 382–385.
- [ZHWW12] Zhou Y., Huang H., Wei L.-Y., Wang R.: Point sampling with general noise spectrum. *ACM Trans. Graph.* 31, 4 (2012).

A Conventions and Notation

A.1 Power Spectrum

Given a point set $P = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$, the power spectrum is defined as

$$P(\mathbf{v}) = N^{-1}|S(\mathbf{v})|^2,$$

with the sample function s and its Fourier transform S

$$s(\mathbf{x}) = \sum_{i=1}^N \delta(\mathbf{x} - \mathbf{x}_i), \quad S(\mathbf{v}) = \sum_{i=1}^N e^{-2\pi\mathbf{v} \cdot \mathbf{x}_i}. \quad (9)$$

We denote the radial average of the power spectrum by $P(\nu)$. For conciseness, we will typically use the term “power spectrum” even for this 1D profile. Figure 1 shows plots of the low-frequency region for several popular blue noise power spectra.

A.2 Normalization

We normalize many of the quantities used in this paper to make their value independent of the number of sample points N . Instead of the absolute number of points, we use the *number density* $n = N/V$, which measures the average

$$\begin{aligned}
af(r) + bg(r) &\leftrightarrow aF(u) + bG(u) \\
f(ar) &\leftrightarrow a^{-2}F(u/a) \\
\text{disk}(r) &\leftrightarrow 2\pi \text{jinc}(2u) \\
H(r-a) &\leftrightarrow \delta(u)/2\pi u - 2\pi a^2 \text{jinc}(2ua) \\
\delta(r)/2\pi r &\leftrightarrow 1
\end{aligned}$$

Table 2: Common Hankel transform pairs used in the text.

number of points per unit area; for points distributed in the unit torus, we simply have $n = N$. Distances are normalized by the largest minimum distance of a hexagonal lattice, $d_{\text{hex}} = (2/\sqrt{3}n)^{1/2}$, and frequencies are normalized by its Nyquist frequency, $v_{\text{hex}} = \sqrt{N}/2$.

A.3 Hankel Transform

The Fourier transform of a circularly symmetric function $f(r) = f(r, \theta)$ is also circularly symmetric and called the *Hankel transform* \mathcal{H} [Bra99]. This transform is obtained from the the standard Fourier transform

$$F(u, \phi) = \int_0^\infty r f(r) \int_0^{2\pi} e^{-2\pi i r u \cos(\theta - \phi)} d\theta dr.$$

by rewriting the inner integral using the definition of the Bessel function J_0 ,

$$F(u) = \mathcal{H}[f(r)] = 2\pi \int_0^\infty r f(r) J_0(2\pi r u) dr. \quad (10)$$

The inverse Hankel transform is identical to the forward transform,

$$f(r) = \mathcal{H}[F(u)] = 2\pi \int_0^\infty u F(u) J_0(2\pi r u) du.$$

The main properties of the Hankel transform follow from the properties of Fourier transforms in 2D; we list the most important ones in Table A.3. The unit disk $\text{disk}(r)$ and the Heaviside step function $H(r)$ are defined as

$$\text{disk}(r) = \begin{cases} 1 & r \leq 1 \\ 0 & r > 1 \end{cases}, \quad H(r) = \begin{cases} 0 & r \leq 0 \\ 1 & r > 0 \end{cases}.$$

B Radial Distribution Function

The radial distribution function $g(r)$ measures how the distances between pairs of points are distributed: It is proportional to the probability $\rho(r)$ of finding another point at a certain distance r from an arbitrary reference point,

$$g(r) = \rho(r)/n.$$

The normalization by n^{-1} ensures that $g(r) \rightarrow 1$ as $r \rightarrow \infty$. For random points with point density n , the probability of finding another point at distance r from

any particle must also be n ; in this case, the RDF is constant and the power spectrum is flat with a DC peak at the origin,

$$g_{\text{Pois}}(r) = 1, \quad P_{\text{Pois}}(u) = 1 + n \frac{\delta(u)}{2\pi u}.$$

For other point distributions, $g(r)$ oscillates around 1; this oscillation indicates how much the point set differs from a random arrangement. Important information can be read off directly from the RDF, such as the minimum distance between points and the distance at which the points become uncorrelated.

Throughout the paper we exploit the close relationship between the RDF and the power spectrum of a point set. The power spectrum $P(\mathbf{v})$ is defined as

$$P(\mathbf{v}) = \frac{1}{N} S(\mathbf{v}) S^*(\mathbf{v}) = \frac{1}{N} \sum_{k,j} e^{-i2\pi\mathbf{v}\cdot(\mathbf{x}_k - \mathbf{x}_j)}. \quad (11)$$

By setting $\mathbf{r}_{jk} = \mathbf{x}_k - \mathbf{x}_j$ and pulling out the $k = j$ terms, we can write

$$P(\mathbf{v}) - 1 = \frac{1}{N} \sum_{k \neq j} e^{-i2\pi\mathbf{v}\cdot\mathbf{r}_{jk}} = \int_{\mathbb{R}^2} e^{-i2\pi\mathbf{v}\cdot\mathbf{r}} \rho(\mathbf{r}) \, d\mathbf{r}. \quad (12)$$

The function $\rho(\mathbf{r})$ measures the density of points at distance \mathbf{r} . To recover Eq. (11), we can set $\rho(\mathbf{r}) = \sum_{k \neq j} \delta(\mathbf{r} - \mathbf{r}_{jk}) / N$, but we usually treat $\rho(\mathbf{r})$ as a continuous function.

For isotropic point sets, the density is circularly symmetric, so $\rho(\mathbf{r}) = \rho(r)$. In this case, Eq. (12) turns into a Hankel transform

$$P(v) - 1 = \mathcal{H}[\rho(r)] = n\mathcal{H}[g(r)]. \quad (13)$$

For numerical computations, a different notation is preferable. Since $g(r) \rightarrow 1$ as $r \rightarrow \infty$, the Hankel transform becomes easier to evaluate if we rewrite Eq. (13) as

$$\begin{aligned} P(v) &= 1 + n\mathcal{H}[g(r) - 1] + n \frac{\delta(v)}{2\pi v}, \\ g(r) &= 1 + n^{-1}\mathcal{H}\left[P(v) - 1 - n \frac{\delta(v)}{2\pi v}\right]. \end{aligned} \quad (14)$$

Another advantage of this representation is that the main components of a typical power spectrum become evident, namely the DC peak at $v = 0$ and the oscillation around 1, which is due to the Hankel transform of $g(r) - 1$.

C Sampling Error Bounds

This section proves the inequality from Eq. (8). To simplify the notation, we set $a \equiv S'(v - f)$ and $b \equiv S'(v + f)$; both values are complex numbers.

$$\begin{aligned} E_f(v) &= \frac{1}{2} |a + b|^2 \\ &= \frac{1}{2} \left[|a|^2 + |b|^2 + (a\bar{b} + \bar{a}b) \right] \\ &\leq \frac{1}{2} \left[|a|^2 + |b|^2 + 2|a||b| \right] \end{aligned}$$

The identity $|S'(v)|^2 = P'(v)$ now gives the stated inequality.




Figure 7: Comparison of dart throwing to its dual. The two point sets are dual in the sense that the radial power spectrum of the first point set has the shape of the RDF of the second point set and vice versa.

D Duality of RDF and Power Spectrum

An interesting way to investigate realizable power spectra is to consider the relation between RDF and power spectrum from Eq. (14). If we leave out the DC peak, power spectrum and RDF become *dual* in the following sense

$$P(\nu) = 1 + n\mathcal{H}[g(r) - 1], \quad g(r) = 1 + n^{-1}\mathcal{H}[P(\nu) - 1].$$

A direct consequence of this duality is that if $g(r)$ is a realizable RDF, then $P(\nu) := g(\nu/n)$ is a realizable power spectrum (this follows directly from the scaling properties of the Hankel transform). What is beautiful about this observation is that we immediately have access to large number of realizable RDFs and power spectra, namely those that correspond to classical blue noise patterns (Figure 1). For example, if we use the RDF of a dart throwing point set with $g_{\text{dart}}(r)$, we can construct a *dual* point set such that $P_{\text{dual-dart}}(\nu) = g_{\text{dart}}(\nu/n)$ (see Figure 7). Although this approach does not provide a way to control the amount of aliasing induced by the realized point set, it does give the guarantee that the power spectrum is realizable.

E Closeups

In the following, you find enlarged versions of the most important renderings from Figures 5 and 6. Note that it is difficult to faithfully reproduce these images due to the resampling process of a PDF viewer or printer.




Figure 8: Results for the low aliasing methods in Figure 5.




Figure 9: Results for the high effective Nyquist methods in Figure 5.




Figure 10: Results for the low aliasing methods in Figure 6.




Figure 11: Results for the high effective Nyquist methods in Figure 6.