Towards a Standardized Spectral Analysis of Point Sets with Applications in Graphics

Thomas Schlömer Oliver Deussen
University of Konstanz, Germany

May, 2010

Abstract

We investigate common pitfalls in the spectral analysis of point sets based on amplitude/power spectrum and radial statistics. We demonstrate the sensitivity of these measurements to the type of Fourier transform and formulate recommendations for crucial analysis and formatting parameters. Following these guidelines elevates comparability between different point generation methods with respect to their spectral characteristics.

1 Introduction

Spectral analysis is an important aspect in the evaluation of two-dimensional point sets in computer graphics domains such as sampling or non-photorealistic rendering. Fourier amplitude spectra, mean periodograms, and evaluating radial statistics all help to reveal correlations between points that may become problematic during their later application. As these evaluations are largely qualitative, however, special care has to be taken with respect to the exact analysis and formatting parameters, the utilized type of Fourier transform, or the applied mapping operators. Changing any of these variables may yield very different results which prevents the unbiased comparability of different point generation methods.

This report is an attempt to standardize these settings by formulating recommendations for the evaluation of amplitude spectrum/mean periodogram, radially averaged power spectrum, and anisotropy. It can be seen as a continuation of the established work by Ulichney [8] and Lagae and Dutré [3].

2 Spectral Analysis

Point sets generated via deterministic methods can be interpreted as finite-energy signals which possess a Fourier transform. Hence, they are characterized by the corresponding Fourier amplitude (magnitude) spectrum. This is not the case for point sets generated with non-deterministic methods due to the random fluctuations inherent to the specific method. Instead, these methods can be regarded as stationary stochastic processes where the main signal characteristic is complemented by noise. Such processes do not have finite energy but finite average power and hence are characterized by the corresponding power density spectrum [6].

2.1 Power Spectrum Estimation

The *power density spectrum* $P(f)$ of a stationary stochastic process is the Fourier transform of its autocorrelation function. Since the autocorrelation function is seldomly known, a spectral estimate of $P(f)$, $\hat{P}(f)$, must be obtained. Bartlett's method [6] of averaging K *periodograms* yields an unbiased and consistent estimator for the power spectrum, is simple to implement, and exhibits variance that decreases linearly with K.

For real signals, the periodogram is the square of the magnitude of the Fourier transform. Hence, for a given point set $\{x_0, \ldots, x_{n-1}\} \subset [0,1)^2$ it can be obtained by

$$\frac{1}{n}\left|\mathcal{F}\sum_{i=0}^{n-1}\delta(x-x_i)\right|^2,$$

where $\mathcal{F}$ denotes the Fourier transform and δ Dirac's delta function. The spectrum estimate $\hat{P}(f)$ is then captured by the *mean periodogram* which results from averaging K periodograms. In an implementation it will be discretized using a square resolution depending on the frequency domain size N.

2.2 Radial Statistics

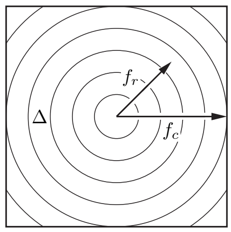

Ulichney [8] derives two useful one-dimensional statistics from a power spectrum estimate $\hat{P}(f)$ which help to reveal directional artifacts. The first is the *radially averaged power spectrum*

$$P_r(f_r) = \frac{1}{N_r(f_r)}\sum_{i=1}^{N_r(f_r)}\hat{P}(f),$$

which can be obtained by partitioning $\hat{P}(f)$ into concentric annuli of width Δ, and averaging the $N_r(f_r)$ frequency samples within each annulus of central radius f_r. Consequently, the annuli start to exceed the spectral estimate at a critical frequency f_c half the domain size.

The second statistic is the *anisotropy*

$$A_r(f_r) = \frac{s^2(f_r)}{P_r^2(f_r)},$$

where $s^2(f_r)$ denotes the variance of the frequency samples and is given by

$$s^2(f_r) = \frac{1}{N_r(f_r)-1}\sum_{i=1}^{N_r(f_r)}\left(\hat{P}(f) - P_r(f_r)\right)^2.$$

The anisotropy is a measure for the radial symmetry of the spectrum and is usually plotted in decibels.

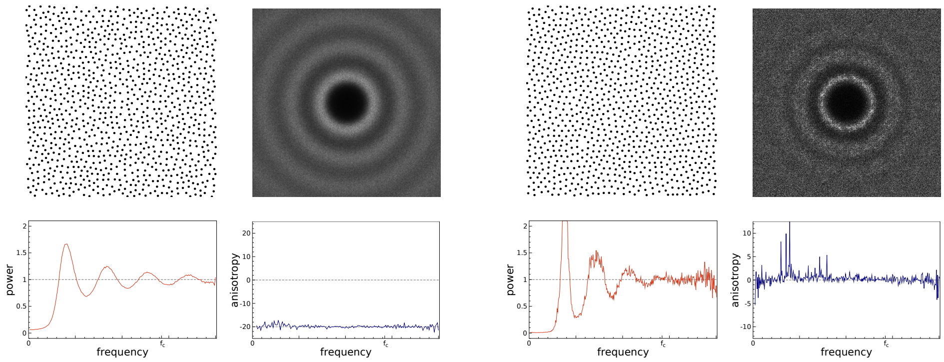

(a) Dart throwing

(b) G-hexominoes [4]

Figure 1: Exemplary spectral analysis of point sets with a blue noise property obtained via *(a)* a stochastic, and *(b)* a deterministic generation method. The averaging of multiple periodograms makes the statistics smoother in the case of the stochastic generation method.

As an example, Figure 1(a) shows mean periodogram (top right), radial power (bottom left) and anisotropy (bottom right) for a stochastic point generation method, in this case traditional dart throwing. The mean periodogram was obtained by averaging $K = 100$ periodograms which yields very smooth radial statistics. For comparison, Figure 1(b) shows Fourier amplitude spectrum, radial power and anisotropy for a deterministic method aiming at the same frequency response, in this case the tile-based method by Ostromoukhov [4]. As expected, the radial statistics are less smooth in this case. The anisotropy plot also reveals slight correlations in the source point set and shows peaks at several distinct frequencies.

3 Choosing the Parameters

The significance of the above measurements depends on a series of analysis and formatting parameters, among them the domain size of the Fourier transform N, the number of periodograms K, the annuli width Δ, and others. Moreover, some care has to be taken when applying the Fourier transform as sometimes the inherent discretization error of a DFT/FFT prevents an accurate analysis. In the following, we discuss each of the relevant analysis parameters and operators.

3.1 Domain Size

To capture a method's characteristic frequency response, its amplitude spectrum/mean periodogram should be of a resolution $N \times N$ that spans a frequency domain large enough to cover all relevant frequencies. For an arbitrary set of n points, this implies $N = 2\lfloor w \frac{1}{d_{\max}} \rfloor$ where $d_{\max} = (2/(\sqrt{3}n))^{1/2}$ is the maximal mutual minimum distance between points, derived from the packing density of circles in the plane [3]. The scalar $w \geq 1$ determines by how much the periodogram window exceeds the relevant frequency range in order to allow a proper view on the response. In our experience, a suitable choice is $w = 4$ which is also the value employed in this report.

3.2 Fourier Transform

The Fourier transform (FT) involved in the amplitude spectrum/periodogram computation is typically performed in its discrete variant as a FFT in order to exploit its computational efficiency [2]. It can be applied by translating the point set into a discrete periodic sample image where each sample is represented as an impulse of value 1, and where its location is rounded to the nearest pixel grid location. The square resolution $M \times M$ of this sample image is often chosen to be equal to the periodogram resolution $N \times N$ which, however, may be insufficient for an accurate spectral analysis.

Figure 2 demonstrates this by plotting mean periodogram and radial power based on a DFT in comparison to reference solutions obtained via a continuous FT. In the first row for each generation method, the resolution for the DFT equals the domain size of the CFT ($M = mN$ with $m = 1$), in the second row it has been increased by a factor of $m = 4$. If the resolution is insufficient, mean periodograms obtained via the DFT (left) significantly deviate from the reference solution as indicated by the difference images next to them, in particular in higher frequency regions. This observation is also supported by the radial power plots (right) where the results based on the DFT (red) are superimposed over a CFT ground truth (gray).

The discretization error of a DFT can become more severe for some deterministic point sets like, e.g., a regular grid or some Rank-1 lattices whose Fourier transform is just a set of peaks. In such a case, correct peaks may be complemented by "aliasing" peaks introduced by the discretization. A DFT also implies periodic signals which does not fit e.g. non-toroidal point sets.

In an analysis scenario accuracy is more important than performance which is why we generally recommend the use of a continuous Fourier transform.

3.3 Number of Periodograms

As already mentioned, Bartlett's method of averaging K periodograms yields an unbiased estimator for the power spectrum with variance that decreases linearly with K. Thus, a larger K yields smoother radial statistics than a smaller K (cf. Figure 3). According to our experience, $K = 10$ usually suffices for a proper spectral analysis which is in line with the current convention [3, 8].

Note that since the anisotropy is plotted in decibels according to the relation $x_{dB} = 10 \log_{10} x$, a value of $-K dB$ implies background noise. For this reason, anisotropy plots should contain a reference line at the appropriate noise level, e.g. at $0, -10, -20$ for $K = 1, 10, 100$ (cf. Figure 1). Anisotropy close to the noise level indicates good radial symmetry for the specific generation method.

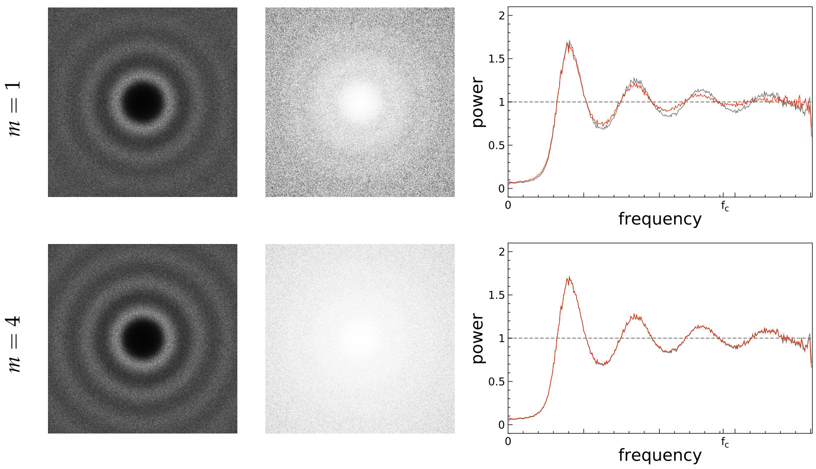

(a) Dart throwing

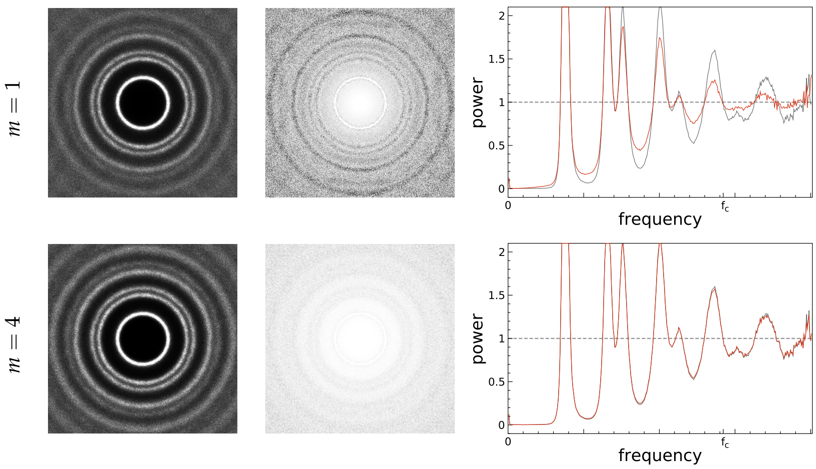

(b) Lloyd's method

Figure 2: Comparing analysis results based on a discrete FT and a continuous FT for *(a)* point sets obtained via dart throwing, and *(b)* random point sets optimized by Lloyd's method until full convergence.

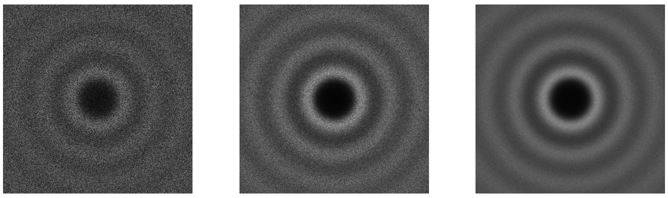

(a) $K = 1$ (b) $K = 10$ (c) $K = 100$

Figure 3: Mean periodograms obtained by averaging K periodograms for point sets generated via traditional dart throwing.

3.4 Annulus Width

The validity of the radial statistics—radially averaged power spectrum and anisotropy—obviously depends on the user-chosen annulus width Δ. In particular, Δ should not be too large as wide annuli effectively smooth the graphs and may hide subtle correlations in the analyzed point sets. Lagae and Dutré [3] suggest a width of approximately one frequency sample which is sufficiently small to reveal even subtle correlations. This width yields $\sqrt{2} \cdot N/2$ annuli where $N/2$ is the Fourier domain size as discussed above. Hence, at the critical frequency of $f_c = N/2$ the annuli start to exceed the spectral estimate and yield less reliable statistics. This should be marked appropriately (cf. Figure 1).

3.5 Formatting/Display

In order to ensure full comparability, we recommend some final touches on the analysis results.

Tone mapping Amplitude spectra and mean periodograms should conform to the same logarithmic tone mapping. The renderings in this report were generated using the mapping $x \mapsto \log_2(1 + \alpha x)$ with $\alpha = 0.25$. This works very well for a large variety of methods, both deterministic and stochastic.

Anisotropy scale When plotting the anisotropy, the axis minimum and maximum value should not exceed the background noise level by more than a factor of ≈ 1.25 for $K \geq 10$. Otherwise, the graph may get significantly compressed, suppressing potential artifacts.

DC peak The DC peak may be removed from all plots as it provides no insight into the spectral characteristics of a specific generation method.

4 Conclusion

We formulated recommendations for the spectral analysis of two-dimensional point sets with applications in graphics. Following these guidelines avoids common pitfalls and, hopefully, facilitates comparability. Spectral analysis remains an important tool among the variety of other valuable measurements such as (toroidal) minimum distance [3], sampling efficiency [5], or discrepancy [7]. To further support comparability, we provide an implementation of all recommendations via an open source project at [1].

Acknowledgements The authors would like to thank Daniel Heck for insightful comments and suggestions, in particular for pointing out the CFT/DFT discrepancy.

References

[1] psa—a tool for the analysis of two-dimensional point sets with applications in graphics, 2010. http://psa.googlecode.com.

[2] Matteo Frigo and Steven G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. Special issue on "Program Generation, Optimization, and Platform Adaptation".

[3] Ares Lagae and Philip Dutré. A comparison of methods for generating Poisson disk distributions. *Computer Graphics Forum*, 27(1):114–129, 2008.

[4] Victor Ostromoukhov. Sampling with polyominoes. In *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2007)*, volume 26, page 78. ACM, 2007.

[5] Daniel P. Petersen and David Middleton. Sampling and reconstruction of wave-number-limited functions in n-dimensional Euclidean spaces. *Information and Control*, 5(4):279 – 323, 1962.

[6] John G. Proakis and Dimitris G. Manolakis. *Digital Signal Processing*. Prentice-Hall, 1996.

[7] Peter Shirley. Discrepancy as a quality measure for sample distributions. In *Proceedings of Eurographics*, pages 183–194. Elsevier Science Publishers, 1991.

[8] Robert Ulichney. *Digital halftoning*. MIT Press, 1987.