



THESIS FOR DOCTOR OF PHILOSOPHY (PHD)

**Sampling and Variance Analysis for
Monte Carlo Integration in Spherical Domain**

AUTHOR: **Gurprit SINGH**

8th September, 2015

JURY MEMBERS

Prof. Jean-Michel DISCHLER

Dr. Nicolas HOLZSCHUCH

Prof. Michael KAZHDAN

Dr. Bruno LÉVY

Prof. Victor OSTROMOUKHOV (*directeur*)

rapporteur

rapporteur

examineur

examineur

examineur

©2015 Gurprit Singh / Gurpreet Singh Bagga
ALL RIGHTS RESERVED

To my parents:
Gurmail & Surinder

ACKNOWLEDGEMENTS

This dissertation would not have been possible without the continuous support of my advisor, Prof. Victor Ostromoukhov. Through out the period of my stay at LIRIS lab Victor showed great patience and responded to my countless questions. I am grateful to Victor for motivating me to gain an in-depth knowledge of the topics I wanted to work on. I will always remember his words: “*your project is your baby, you have to take care of it.*” Victor helped me not only with my work but also in many administrative tasks which could have taken for ever to complete due to my poor french (yes, I still don’t speak fluent french).

I would also like to thank David Coeurjolly for his cool and friendly nature through out our whiteboard discussions. Sylvain Lefebvre for introducing me to research before my PhD, his positive attitude and problem solving instincts are qualities I shall aspire to for the rest of my life. I would like to thank Nicolas Bonnel for bearing me in the lengthy mathematical discussions that greatly improved my understanding, Jean-Claude Iehl for many hours of fruitful discussions on topics related to rendering, Guillaume Bouchard for introducing me to wall climbing, rope walking and for sharing his enormous knowledge on rendering-related topics, Vincent Nivoliers for arranging all the fun-filling activities, dinners and picnics (whether it is in Nancy or Lyon), Jonathan Dupuy for always helping me in my writing and other members of the lab including Jean-Philippe Farrugia for coordinating several team meetings over all my PhD years. This thesis work would not have been fun without my co-PhD students including Gael, Florent, Adrien, Ricardo, Zohra, Blandine and Vinh. We shared the same office for more than two years. Special thanks to Brigitte, Isabelle for keeping the lab administration running and their always-smiling faces that helped ease my french conversations with them. I also thank Mathieu Desbrun for supporting me in various stages of my PhD, whether it is for publications or for scholarship applications, Misha Kazhdan for introducing me to the beauty of representation theory and Katherine Breeden for sharing her incredible knowledge in English and giving time from her busy schedule to help me out in paper writing during deadlines.

A big thanks to Delphine and her family for sharing their love and affection with me at all times. I sincerely thank my brothers, Gurjeet and Mandeep, for making sure that I concentrate on my thesis without worrying about my parents

and without getting homesick. Finally, I am greatly obliged to my parents Sh. Gurmail Singh and Shmt. Surinder Kaur, who always backed me in all important decisions, whether it is moving away from them to France or spending more than two years at home for the preps of IIT entrance exams with no job. I salute their patience as they had to spend three continuous years without my visit to their place during my PhD. This thesis work is a tribute to my parents for their infinite support, love and patience.

Contents

Contents	vii
List of Figures	ix
1 Introduction	5
1.1 Contributions	7
1.2 Overview	9
2 State of the Art	11
2.1 Sampling in Euclidean domain	13
2.2 Sampling in Spherical domain	37
2.3 Quality criteria for sample distributions	47
2.4 Discussion	50
3 Mathematical Background	53
3.1 Preliminaries	54
3.2 Vector Spaces	55
3.3 Group Theory	59
3.4 Representation Theory	64
3.5 Spherical Harmonics	67
3.6 Hemispherical analysis using SH	68
4 Spectral Analysis of Sampling Patterns	71
4.1 Spectral Analysis in Euclidean domain	74
4.2 Spectral Analysis in Spherical domain	85
4.3 Spectral Analysis in Hemispherical domain	93
4.4 Homogeneous sampling patterns	98
4.5 Discussion	101

5	Variance analysis of Monte Carlo Integration	103
5.1	Monte Carlo estimator	103
5.2	Background on SH	105
5.3	Variance in spectral form	105
5.4	Discussion	109
6	Theoretical Convergence analysis	113
6.1	Best and worst case variance	114
6.2	Convergence rate analysis	115
7	Experiments and Results	123
7.1	Implementation details	124
7.2	Case studies	126
7.3	Rendering Results	130
7.4	Discussions	133
8	Conclusions	139
8.1	Design principles	140
8.2	Quadrature rules on the sphere	141
8.3	Future directions	142
A	Simplified variance equation	143
A.1	Contribution of DC peak in Variance	143
A.2	Contribution of second term	144
B	Unbiased homogeneous sampling	145
C	Spectral Analysis of spherical cap	147
C.1	Spherical cap function in SH terms	147
C.2	Spherical cap power spectrum	151
	Bibliography	153

List of Figures

2.1	Left: Micro structures of cones Center: Monkey eye photoreceptor (cones) distribution Right: Optical transform of monkey eye. Image Courtesy, J. I. Yellot [108]	15
2.2	Illustration of how to compute the radially averaged power spectrum from the power spectrum of the signal in (a). Each circle represents one particular frequency. In (b) a radially averaged power spectrum of ideal <i>blue noise</i> is shown which is computed by taking the radial average over different values over the radial circle at each frequency. Images inspired from Ulichney [97]	17
3.1	Spherical harmonics basis functions (real part).	67
4.1	Zone plate test to show the aliasing (structural-artifacts) due to the Fourier peaks in the regular sampling patterns (b), which turns into noise on simple jittering (c)	72
4.2	Illustration of jittered (a) and regular (c) sampling patterns with their corresponding power spectra in (b) and (d), respectively. The high energy peaks of regular sampling patterns are surrounded by circular (orange) rings in (d) for illustration purposes. In a similar manner, the DC peak in (b) is surrounded by a circular (orange) ring.	73
4.3	Illustration of Top row: non-Toroidal (unwrapped) vs Bottom Row: Toroidal (wrapped) Poisson Disk samples with corresponding power spectra and the zoom-in of the low frequency region around the DC peak of the respective power spectra.	76
4.4	Illustration of multiple realizations of Poisson Disk samples (a) with corresponding Fourier power spectra with respect to each realization.	78

4.5	Illustration of sampling patterns in the Left column, with their corresponding power spectra in the Center column. In the Right column we show a heatmap of each power spectrum to emphasize their characteristics in the low frequency region around the DC peak, surrounded by a circle (orange) in (b), (e) and (h), located at the center of the power spectra images, in the Center and Right columns). Each peak in the regular pattern power spectrum (k) is also surrounded by circles (in orange) for illustration purposes.	80
4.6	Illustration of the radial mean power spectra of whitenoise (b) and jittered sampling (d) from the corresponding power spectra in (a) and (c), respectively. Radial mean power is computed by averaging the power spectra values along the circular rings, in blue, for a given frequency ρ	81
4.7	Illustration of sampling patterns in the Left column, with their corresponding power spectra in the Center column. In the Right column we show a heatmap of each power spectrum to emphasize their characteristics in the low frequency region around the DC peak which is located at the center of the power spectra images, in the Center and Right columns).	83
4.8	Illustration of the effects of high energy Fourier peaks in rendering a simple <i>zone plate</i> function.	84
4.9	Illustration of <i>angular</i> power spectrum of a white noise distribution in (a) with the same power spectrum shown in a <i>square-shape</i> (b).	86
4.10	Illustration of healpix regular (a) and spherical Fibonacci lattice (b) point sets with respective mean angular power spectra.	88
4.11	Illustration of <i>square-shaped</i> angular power spectrum of Healpix regular (left) with heatmaps of power (center) and square-root of power values (right).	89
4.12	Illustration of <i>square-shaped</i> angular power spectrum of spherical Fibonacci (left) with heatmaps of power (center) and square-root of power values (right).	90
4.13	Illustratin direct lighting evaluation with (b) or without (c) random perturbation of secondary rays. Image courtesy: Dr. Jean-Claude Iehl.	92
4.14	Comparison of the mean angular power spectra of various state of the art sampling patterns in the spherical (left) and hemispherical (right) domain.	94

4.15	Illustration of the high energy peaks in the low frequency region of <i>square-shaped</i> and <i>radial</i> angular power spectra of whitenoise, jittered and hemispherical fibonacci lattice point sets, respectively. .	95
4.16	Illustration of <i>square-shaped</i> angular power spectrum of hemispherical Fibonacci (left) with heatmaps of power (center) and square-root of power values (right).	97
4.17	Illustration of the distribution of a spherical harmonic coefficient ($l = 4, m = 3$) computed using 1024 samples over 8192 realizations in the complex plane.	100
5.1	(a) Illustration of the power spectra of a sampler and an integrand on the same axis. (b) With increase in number of samples, sampler's power spectra scales along the frequency axis. In (c) and (d), we show the corresponding product of sampler and integrand power spectra.	110
6.1	Illustration of the <i>mean angular</i> power spectra <i>vs</i> frequency ($\omega = l/\alpha\sqrt{N}$) of signals that can be considered as best (a) and worst (b) cases in numerical integration with respect to any sampling power spectrum.	115
6.2	Step	117
6.3	Quadratic (b=2)	119
7.1	Illustration of spherical sampling methods.	124
7.2	Poisson Disk	125
7.3	CCVT	125
7.4	Bounds on the mean angular power spectra of samplers and on the variance convergence rates of our test integrands. Log-Log plots are shown in the Center and Right columns. Left: mean angular power spectra of different sampling patterns (blue), bounded by a theoretical spectral profile (green). On the frequency axis, units corresponds to the frequency $\alpha\sqrt{N}$. Center: The variance curve of a spherical harmonic basis function with $l = 4, m = 0$, (blue), with bounds (green) computed using the bounds of the corresponding mean angular power spectrum. Right: The variance in integration of a spherical cap, using the same visualization.	128

7.5	Cornell box rendered with ambient occlusion. <i>Mean squared error</i> (MSE) values are computed w.r.t a reference image (shown in Fig. 7.7, for White noise (Left), Jittered sampling (Center) and Poisson Disk (Right) for a given number of shading rays used to sample directions on the visible hemisphere at each hitpoint in the scene.	131
7.6	Cornell box rendered with ambient occlusion. <i>Mean squared error</i> (MSE) values are computed w.r.t a reference image (shown in Fig. 7.7, for White noise (Left), CCVT (Center) and Regular Healpix (Right) sampling for a given number of shading rays used to sample directions on the visible hemisphere at each hitpoint in the scene.	132
7.7	Reference image	133
7.8	Cornell box rendered with ambient occlusion. Left: Spherical Fibonacci (SF) point set is used at each hitpoint. Centre: SF point set is used with uniform and random rotation at each hitpoint. Right: Healpix regular point set is used at each hit point with no rotation. <i>Mean squared error</i> (MSE) values are computed w.r.t a reference image (shown in Fig. 7.7	134
7.9	Comparison of the variance in MC integration for different integrand signals. Here we use experimental data from 7.4 in (a) and (b). Here we present comparisons in the (hemi-)spherical domain : (a) a spherical harmonic basis function (Y_4^0). Inset illustrates the gray scale of the absolute values of the function, (b) a spherical cap function ($\theta_0 = 60$), where the white shade in the inset represents non-zero constant value region (c) a Cornell box scene with all corresponding rendered images shown in Fig. 7.5 and 7.6.	135
7.10	Comparison between the power spectra at the low frequency region between (a) Poisson Disk and (b) jittered samplers. In (c) and (d) we show the zoom-in of the low frequency region ($\omega \in]0, 1[$) for the Poisson Disk and the jittered samplers respectively.	137
C.1	We illustrate the power profile (Power vs l) of a spherical cap function (in blue) for $\theta_0 = 60$ with the corresponding bound (in orange). We show three plots of the same power spectrum (a), with different zoom-in plots, (b) and (c), to better see the upper bound and to show that the power of the spherical cap is well bounded by the upper bound we derived in Sec. C.2.	151

LIST OF SYMBOLS

Symbol	Definition
\mathbb{D}, N	sampling domain (toroidal unit square \mathcal{T}^d , unit sphere \mathcal{S}^2 , unit hemisphere \mathcal{H}^2), with number of samples N
\mathcal{S}^d	Spherical domain of dimension d
$\mu(\mathbb{D})$	Lebesgue measure of domain \mathbb{D}
θ, ϕ	colatitude $\theta \in [0, \pi]$ and longitude $\phi \in [0, 2\pi]$, on the sphere or the hemisphere
l, m	degree and order of spherical harmonics, both values are integers with $l \geq 0$ and $m \in [-l, l]$.
$Y_l^m(\theta, \phi)$	Spherical harmonic function of degree l and order m whose inner product over the sphere is 4π
$P_l^m(\cos \theta)$	associated Legendre polynomial of degree l and order m .
$G(x)$	arbitrary integrable function on \mathbb{D} .
$\mathcal{P}_G(\cdot)$	power spectrum of G in the toroidal domain or angular power spectrum of G in the spherical domain.
$\check{\mathcal{P}}_G(\cdot)$	radial mean of the expected value of \mathcal{P}_G in the toroidal domain or expected value of \mathcal{P}_G in the spherical domain.
$\bar{g}, \ g\ $	conjugate and magnitude operators for a complex variable g , respectively.
$\langle \cdot, \cdot \rangle$	inner product operator
$\langle \cdot \rangle, \text{Var}(\cdot)$	expectation and variance operators
X	bold face letter represents a random variable

Abstract

This dissertation introduces a theoretical framework to study different sampling patterns in the spherical domain and their effects in the evaluation of global illumination integrals. Evaluating illumination (light transport) is one of the most essential aspect in image synthesis to achieve realism which involves solving multi-dimensional space integrals. Monte Carlo based numerical integration schemes are heavily employed to solve these high dimensional integrals. One of the most important aspect of any numerical integration method is sampling. The way samples are distributed on an integration domain can greatly affect the final result. For example, in images, the effects of various sampling patterns appears in the form of either structural artifacts or completely unstructured noise. In many cases, we may get completely false (biased) results due to the sampling pattern used in integration.

The distribution of sampling patterns can be characterized using their Fourier power spectra. It is also possible to use the Fourier power spectrum as input, to generate the corresponding sample distribution. This further allows spectral control over the sample distributions. Since this spectral control allows tailoring new sampling patterns directly from the input Fourier power spectrum, it can be used to improve error in integration. However, a direct relation between the error in Monte Carlo integration and the sampling power spectrum is missing. In this work, we propose a variance formulation, that establishes a direct link between the variance in Monte Carlo integration and the power spectra of both the sampling pattern and the integrand involved.

To derive our closed-form variance formulation, we use the notion of *homogeneous* sample distributions that allows expression of error in Monte Carlo integration, only in the form of variance. Based on our variance formulation, we develop an analysis tool that can be used to derive theoretical variance convergence rates of various state-of-the-art sampling patterns. Our analysis give insights to design principles that can be used to tailor new sampling patterns based on the integrand.

Chapter 1

Introduction

From elementary mathematics, we have learnt how to compute the integration of simple (*e.g.*, \sin , \cos) and complicated mathematical functions. To compute similar integrals using computers is not that straight forward. In computers, we cannot compute any continuous function directly. To perform this computation, it is required to discretize the function with points (samples) followed by performing weighted average of these sample values, that gives the final value of the function at given coordinates. In a mathematical sense, this process is termed as *numerical integration* or numerical quadrature.

How we sample the domain of integration is of paramount importance to numerical integration techniques. From computational physics, computational chemistry to computation finance and computer generated imagery, in all fields, numerical integration plays a key role. For low dimensional integrals, many efficient quadrature rules exist. However, for very high dimensional problems, it is not possible to use these efficient quadrature rules as these rules can have very

high time and space complexity. Monte Carlo and Quasi-Monte Carlo integration rules are the best in computing high-dimensional integrals.

Numerical integration techniques are used in various fields. In this dissertation, we focus our attention to the field of computer graphics, and more specifically to image synthesis. In images, a pixel is considered as the basic building block of computer generated imagery. Even though, to a common eye, an image can be seen as continuous, it is composed of discrete pixels. To obtain an image, each pixel needs to be further discretized (sampled) using a set of points (samples). A color corresponding to each point can be computed by using the state-of-the-art algorithms (*e.g.*, *ray tracing*), then a weighted average of these samples (reconstruction) is stored in a pixel as a color value.

Computation of a “real” life-like image requires more sophisticated algorithms. One of the many ways to add *realism* to an image is via computing proper global illumination in a scene. Evaluation of global illumination requires computing scattering directions of light at each point inside a virtual scene. If the point where we compute the scattering direction is a surface, then we can assume a hemisphere around this point and consider that the light scatters away from the surface from this hemisphere. In case, we have a scene with *fog* or *mist*, at each point where we compute the scattering direction, light is arriving from all the directions. Therefore, for a scene with *fog*, we assume a sphere around each point inside a scene and consider that the light scatters away from this point in all the directions.

1.1 Contributions

The central theme of this dissertation is to study the distribution of samples on the surface of a unit sphere and hemisphere. Even though, in global illumination algorithms spherical and hemispherical sampling have been well exploited, the frequency content of the spherical and hemispherical sampling patterns and its effect in the variance during integration, is not thoroughly studied. In this dissertation, first we perform a spectral analysis of various *state-of-the-art* sampling methods on the sphere and the hemisphere. We use the notion of *homogeneous* sample distributions, which has been extensively studied in physics and spatial statistics [40, 20, 21]. Homogeneous sample points have same statistical properties over the domain. We use the notion of homogeneous sampling to express error in Monte Carlo integration only in terms of variance. This complete work has been successfully published by Pilleboue, Singh and colleagues [79].

We analyze the effects of various spherical and hemispherical sampling methods on the variance for Monte Carlo integration of high dimensional global illumination integrals. We propose a theoretical framework that relates the frequency content (power spectrum) of the sampling pattern and the integrand involved with the variance in Monte Carlo integration. Thanks to this framework, we are able to derive the variance convergence rate of various samplers with the increase in number of samples. We propose a set of quintessential design principles that can be used to tailor new sampling patterns based on the integrand under study, resulting in improved numerical integration of the underlying integrand.

1.1.1 Publications

The theoretical framework that we discuss in this dissertation has been successfully published in SIGGRAPH 2015 as: “*Variance Analysis for Monte Carlo Integration*”

with authors: *Adrien Pilleboue*[†], ***Gurprit Singh***[†], *David Couerjolly*, *Michael Kazhdan*, *Victor Ostromoukhov*.

[†] *Adrien Pilleboue* and ***Gurprit Singh*** are joint first authors of this publication.

In the published work on Variance Analysis for Monte Carlo Integration, we build a theoretical framework in both the Euclidean and spherical domains that relates the power spectrum of the sampler and integrand involved, with the variance in Monte Carlo integration. Here, the Euclidean domain analysis is performed by Adrien Pilleboue [78] whereas a similar spherical domain analysis is performed by Gurprit Singh, the author of this dissertation, in this thesis work.

Previously, I also co-authored a research paper that was published in SIGGRAPH 2014 as: “*Fast Tile-based Adaptive Sampling with User-Specified Fourier Spectra*”.

with authors: *Florent Wachtel*, *Adrien Pilleboue*, *David Couerjolly*, *Katherine Breeden*, ***Gurprit Singh***, *Gaël Cathelin*, *Fernando De Goes*, *Mathieu Desbrun*, *Victor Ostromoukhov*.

Although, *Florent Wachtel* has been the main investigator of this very large project which lasted three years, many other people substantially contributed in conceptual clarification of the key issues, implementation, tests, improvements, presentation of the work. It has been a team work, in which I fully participated. That project

can be seen as a preliminary case study, which stressed the importance of the frequency content of the sampler, and ultimately led to the present analysis work. This publication is not included as a contribution in this dissertation. However, we give details of this publication in Section 2.1.7 while discussing the state-of-the-art work in Chapter 2.

1.2 Overview

To give a brief overview of this dissertation, there are eight chapters in this dissertation including this one. We have already introduced our intuitive ideas on which our theoretical framework has been developed in this dissertation. We give a not-too-long survey (Chapter 2) of various sampling methods proposed in the literature. Our literature survey shows that there is not enough work done to study the frequency content of various stochastic samplers on the sphere and especially on the hemisphere. Therefore, we have devoted a complete chapter (Chapter 4) to study the spectral properties of various sampling patterns (stochastic and deterministic) on the sphere and hemisphere.

In Chapter 4, we also introduce the notion of *homogeneous* sample distributions on the sphere. Later on, we use this notion of homogeneity in Chapter 5, to derive our closed-form variance expression that relates the variance in Monte Carlo integration with the power spectra of the sampling pattern and the integrand involved. In our derivation, we use the *representation theory* to formally express the effects of rotation on the spherical harmonic coefficients. We give a brief

background on spherical harmonics and the representation theory in Chapter 3. We also provide a background on group theory and vector spaces which are prerequisites for the representation theory.

In Chapter 6, we use our mathematical model to derive variance convergence rates of various sampling patterns for a given class of functions. We also provide tools to bound the sampling power spectra and lay down some design principles for future samplers. At the end, we show our experiments and results in Chapter 7 and conclude our thesis in Chapter 8.

Chapter 2

State of the Art

In this chapter, we do a survey on various *state-of-the-art* sampling methods developed for image synthesis. Regular sampling is the simplest to generate. However, it is not recommended to use regular sampling methods for numerical integration. We study the problems caused by regular sampling patterns and discuss various solutions proposed by researchers to avoid the artifacts caused by regular samples in Monte Carlo numerical integration schemes. Arguably the best approach to algorithmically generate point distributions is to find a good balance between density control and spatial irregularity. There are many Monte Carlo (MC) and Quasi-Monte Carlo (QMC) based [71, 55, 49] sampling patterns proposed in the literature to find this balance between density control and spatial irregularity.

We classify various sampling methods in two classes: deterministic and non-deterministic methods. Samplers which generate a same sequence or point sets at each initialization falls into the category of deterministic samplers, whereas, sam-

plers that are completely stochastic, *i.e.*, which generate new set of samples with assigned characteristics at each initialization, falls into the non-deterministic class. We further divide these classes into subclasses, based on the approaches used to generate the samples. We subdivide non-deterministic samplers into relaxation (optimization) and dart throwing based approaches. Most of the non-deterministic samplers are relaxation and dart throwing based, others can be considered as an extension of these two. To generate purely deterministic samplers, Number theory (radical inverse function) based approaches are heavily employed. Most of the radical inverse based samplers are mainly categorized as *low discrepancy* samplers and they are the core of the QMC integration schemes.

There are some tile-based approaches that can be used to generate samples that mimic the non-deterministic behaviour. Tile-based methods were originally introduced to overcome the time complexity of relaxation and dart throwing based algorithms. In computer graphics community two tile-based approaches are well known: First approach uses a set of precomputed tiles, with each tile composed of multiple samples, and later use these tiles, in a sophisticated way, to pave the sampling domain. Second approach employed tiles with *one sample per tile* and uses some relaxation-based schemes, with look-up tables, to improve the overall quality of samples.

Recently, researchers have noticed the importance of having a *spectral control* on sampling methods. The idea behind having a sampler with spectral control is to be able to generate samples from the pre-assigned Fourier characteristics. Consequently, the samples can spatially arrange themselves according to the

characteristics assigned in the spectral domain [110]. Our contribution in this dissertation is inspired from the work on spectral control. In Chapter 5, we derive a closed-form expression that relates the variance in MC integration with the spectral characteristics of the sampling pattern and the integrand involved. In the following sections, we start our discussion with the sampling methods in the Euclidean domain (Section 2.1). We introduce *blue noise* characteristics and study their importance in image synthesis. Then, we discuss various sampling methods that are employed to achieve blue noise characteristics, which includes: relaxation based, dart throwing based and/or tiling based methods. In Section 2.2, we discuss sampling on the sphere. The problem of distributing a large number of points uniformly over the surface of a sphere has been intensively studied by many mathematical researchers, biologists, chemists and physicists. We focus our attention on various spherical sampling algorithms proposed in computer graphics community to improve the image quality. At the end, in Section 2.3, we briefly overview various quality criteria proposed in the literature to estimate the quality of various sampling patterns.

2.1 Sampling in Euclidean domain

To improve the quality of image synthesis, various sampling methods have been developed. Since the early ages of electronic imagery [111], the error due to *regular sampling* patterns, called aliasing, is one of the major issues in image synthesis. Spatial aliasing error is often manifested in the form of *staircasing*

along the edges in an image. In the moving objects, regular patterns in temporal sampling can cause *Strobing* effects, or the jerky motion of an object in an image. Defects from sampling can arise in other aspects of the displayed image, such as color. By increasing the sampling rate higher and higher frequencies can be represented. However, any regular sampling rate has an associated frequency limit above which aliasing still occurs. Crow [18] proposed prefiltering as a recognized cure to aliasing. Dippé and Wold [25] proposed *non-uniform* sampling followed by filtering to avoid aliasing effects in images. Usage of *irregular* or *stochastic* sampling patterns convert high frequency defects (like aliasing or false patterns) into featureless noise. The type of randomness used in the sampling process controls the spectral characteristics of the noise. Dippé and Wold analyzed in particular, Poisson and jittered sampling using purely randomized samples, which can be generalized to a *minimum distance Poisson process* where the randomly distributed samples are separated by a minimum distance criteria.

The motivation to introduce stochastic sampling method comes from a careful study of the human visual system. Studies have shown that the human visual system is more sensitive to noise in the intermediate frequencies [84]. Since the human visual system is less sensitive of high frequency noise, reducing low and mid-frequency noise while increasing high frequency noise can produce perceptually better pictures. Yellot [108] studied the distribution of cones in an extrafoveal region of the eye of a rhesus monkey, which has a photoreceptor distribution similar to that in the human eye. Yellot took the optical Fourier transform of this distribution, which is called a *Poisson Disk distribution*. An illustration is shown




Figure 2.1: **Left:** Micro structures of cones **Center:** Monkey eye photoreceptor (cones) distribution **Right:** Optical transform of monkey eye. Image Courtesy, J. I. Yellot [108]

in Fig. 2.1.

Motivated from the work by Yellot, Cook [16] performed a careful Fourier analysis of white noise, Poisson Disk and jittered sampling patterns. He advocated the use of sampling patterns that lacks low frequency content. Although, Cook in his theoretical analysis mentioned that the Fourier transform of Poisson Disk has no low frequency content, Dippé and Wold have correctly shown an offset in the low frequency zone of the spectrum of Poisson Disk using the *flat field response noise spectrum*.

Dippé and Wold, and Cook also suggested adaptive and importance based stochastic sampling techniques to improve the quality of distributed ray tracing. In the following year, Mitchell [68] proposed an algorithm for fast generation of adaptive nonuniform sampling patterns. Mitchell used the knowledge of how the human eye perceives noise as a function of contrast and color to guide the

adaptive sampling algorithm. Finally, to generate the digital picture, Mitchell proposed a nonuniform reconstruction from samples and advocates resampling at the display pixel rate to solve the problem of noise filtering efficiently.

2.1.1 Blue Noise

Even though Dippe and Wold, Cook and Mitchell clearly indicated the importance of no low frequency content in the sampling patterns' Fourier spectra, an optimal sampling Fourier spectrum was still missing. Ulichney [96], in 1987, was the first to provide qualitative characterization of a good sampling pattern, which he called *Blue Noise*. Before going further, we briefly discuss the naming convention for different kinds of noise.

Color of noise The nature of the spectral distribution property of various types of noise is mainly described by a color name. The most common example is *white noise*, so named because its power spectrum is flat across all frequencies with equal energy distributed in all frequency bands, much like the visible frequencies in the white light. *Pink noise* has the power spectrum such that the power spectral density is inversely proportional to the frequency. Green noise, as its name suggests, consists of primarily mid-range frequencies. Ulichney also mentioned the curious case of *brown noise*, named for the spectrum associated with the Brownian motion [36]. He introduced *blue noise*, the high frequency complement of the pink noise. Pink noise occurs very frequently in nature and thus is used for physical simulation and biological distribution.




Figure 2.2: Illustration of how to compute the radially averaged power spectrum from the power spectrum of the signal in (a). Each circle represents one particular frequency. In (b) a radially averaged power spectrum of ideal *blue noise* is shown which is computed by taking the radial average over different values over the radial circle at each frequency. Images inspired from Ulichney [97]

Characteristics of Blue noise Ulichney performed a careful study of rectangular and hexagonal regular grid patterns to improve the quality of digital halftones. He estimated the power spectrum of various grid patterns using the Bartlett's method [7] of averaging periodograms. He mentioned that a desirable attribute of a well-produced halftone of a gray level is *radial symmetry*. Therefore, he investigated a radially averaged power spectra of various patterns (Fig. 2.2(a)). Ulichney advocated three important features for an ideal radial power spectrum; First, its peak should be at the principal frequency. Second, the principal frequency marks a sharp transition level below which little or no energy exists. And finally, the uncorrelated high-frequency fluctuations are characterized by high-frequency white noise. In Fig. 2.2(b), taken from Ulichney's article, all these three regions are marked. The radial power spectrum with above features is called the *blue noise*

power spectrum. In the following subsections, we discuss various methods used to generate sampling patterns with blue noise characteristics.

2.1.2 Poisson Disk distributions

Yellott [108] have shown that a rhesus monkey has a Poisson Disk distribution. On the basis of his study, he proposed that the least conspicuous form of aliasing would be produced if the spectrum of the sampling pattern had two properties; First, the spectrum should be noisy and lack any concentrated spikes of energy. Secondly, the spectrum should have a deficiency of low-frequency energy. This causes aliasing noise to be concentrated in higher, less conspicuous frequencies. Since the *blue noise* spectrum proposed by Ulichney [96, 97] is in accordance with Yellot's findings, many algorithms have been proposed for efficient generation of Poisson Disk distributions. Most of the algorithms proposed are *dart throwing* and *relaxation* based and the other techniques are more or less derived from or based on these two main approaches.

As Ripley [82] mentioned, several point processes could be referred to as "Poisson-disk", but by strict definition, a true Poisson-disk process is realized by generating complete patterns with Poisson statistics until one is found that meets the minimum-distance constraint. Cook [16] proposed the first *dart throwing* algorithm for generating Poisson Disk distributed point sets. Random samples are continually tested and only those that satisfy the minimum distance constraint relative to samples already in the distribution are accepted. The main source of inefficiency of the dart throwing method is a rejection sampling mechanism

where a large number of samples is attempted but only a small percentage of them is inserted into the distribution. This algorithm cannot guarantee that a maximal distribution can be generated. As the allowable area for new insertions gradually shrinks, the probability that attempted samples will fall inside this area becomes progressively smaller. This also means that the algorithm does not have a guaranteed termination.

Mitchell [68] discovered that Poisson disk distribution can be generated cheaply on the fly by an algorithm inspired by the Floyd-Steinberg halftoning algorithm [32]. Later, Mitchell [69] proposed an algorithm which is known as the *best candidate algorithm*. The best candidate algorithm works by trying $m \times k$ samples when placing the k -th new sample, where m is a supplied parameter. From all $m \times k$ samples attempted, the one that is farther away from all previous $k - 1$ samples is chosen. The algorithm does its best to place samples well away from each other but it does not enforce any particular distribution radius r . There is the probability, however small, that a sequence of unfavourable sampling outcomes will make the best candidate sample be arbitrarily close to some other previous sample.

McCool and Fiume [66] describe a more practical variant of dart throwing, where the dart radius is gradually decreased as more samples are placed. They used Lloyd's algorithm [60, 61]—for the first time in the computer graphics community—for the generation of sampling point sets. Their algorithm slowly decreases the radius of the distribution at each iteration until the final desired radius is reached or a desired number of samples is generated. For each intermediate

radius, it makes a finite number of attempts to place new samples, proportional to the iteration number, before proceeding to the next smaller radius. What makes the decreasing radius algorithm an approximate Poisson-disk sampling method is that it uses radii that are larger than r for most of the iterations.

For many years, dart-throwing was the only available method for accurate Poisson-disk sampling. Its inefficiency led to the development of approximate Poisson Disk sampling algorithms. The situation changed with the development of efficient dart-throwing methods. These new methods take advantage of a spatial data structure to guide in the placement of samples. The data structure encodes the regions of space where the insertion of samples is allowed. This avoids to a great extent the expensive procedure of having to blindly test new samples by trial and error. Every time a sample is inserted in the distribution, the spatial data structure is updated to remove the portion of space occupied by the new sample. The spatially guided methods, used in Computer Graphics, were developed specifically for two-dimensional sample distributions and do not extend well to higher dimensions.

Jones [47] proposed the first spatially guided method for Poisson disk sample generation. The method uses a Voronoi tessellation as the spatial data structure with the samples at the centroid of the Voronoi cells. The Voronoi cell of a sample is randomly selected and a new sample is inserted in the available area of the cell that falls outside a circle of radius $2r$ with the original sample at the centre. A weighted binary tree helps in the selection of samples, with the Voronoi cells as the leaves and with the available areas of the Voronoi cells as the weights. This

ensures that sample placement is done with a uniform probability distribution – the tree is randomly traversed top to bottom, with the area weights giving the probability of selecting the left or right child of each tree node. The placement of a new sample requires the computation of the intersection between the Voronoi cell (a polygon) and the circle of radius $2r$, which can be reduced on case basis. A rejection sampling method cannot be avoided but the probability of a new sample being accepted is much larger than the probability of it being rejected. Although Voronoi tessellations can be extended to three dimensions, placing a new sample in the available area of a three-dimensional Voronoi cell requires the computation of the intersection between the cell (a polytope) and a sphere of radius $2r$. This is a more complex procedure than its two-dimensional equivalent and cannot be reduced to a small number of simple cases.

Dunbar and Humphrey [27] proposed two variants; an *approximate* $O(N)$ and an *accurate* $O(N \log N)$ method, that results from collapsing their *scaloped sector* data structure into a single arc of a circle with radius $2r$. With this transformation, every new sample is always placed at a distance of exactly $2r$ from some other previous sample. The “scaloped sector” spatial data structure, proposed by the authors, is bounded by two arcs of circles of different radii and centred at distinct points. The available area around each sample can be represented as the disjoint union of several scaloped sectors. Similar to the method by Jones [47], a weighted binary tree is used to select a scaloped sector for the placement of a new sample, resulting in a spatial uniform probability distribution. A rejection sampling strategy is avoided as sampling inside a scaloped sector is always guar-

anted to generate a valid Poisson-disk sample. It is not known how the scalloped sector data structure can be extended to three dimensions.

Yet another spatially guided method in two dimensions was proposed by White and colleagues [103]. The authors used a quadtree to signal the allowable sample insertion space. An auxiliary uniform grid stores neighbouring information about samples and is used to check for minimum distance conflicts for every new sample. The cells in the grid have lateral size $2r$ and all the possibly conflicting samples of a newly inserted sample are found by looking in the grid cell where the new sample falls plus the eight surrounding cells. This method can easily be generalised to higher dimensions but it does not scale well due to the need for a uniform grid. The memory size of the grid is $O(r^{-d})$ for d dimensions and this can become intractable for small r .

Bridson [11] proposed a Poisson disk generation algorithm in arbitrary dimensions, similar to White and colleagues [103]. Bridson used a simple d -dimensional grid to evaluate samples at an optimal distance from neighboring samples based on local rejection. An active list of samples is kept. At each iteration, a sample from the active list is randomly chosen and several *dart throwing* attempts try to insert a new sample inside a hypersphere of radius $4r$ centred on the chosen sample. The new sample is added to the grid and to the active list while the previously chosen sample is removed from the list if *dart throwing* did not succeed after some number k of attempts. The method does not distribute samples uniformly because every new sample is always placed inside a hyperspherical neighbourhood of some previous sample. The time complexity of Bridson's algorithm is of the order

$O(N)$.

Wei [100] proposed a parallel sampling method that can run on a GPU. The method uses a multi-resolution strategy where uniform subdivisions of the domain with increasing resolution are considered one at a time. The cells in each resolution level are then arranged into distinct cell groups in such a way that the insertion of new samples inside each group cell can proceed independently from the insertion of samples in the other cells of the same group. This allows sample insertion to be parallelised for each of the groups of any resolution level. Sample insertion is done by making k dart-throwing attempts inside every group cell. Although the sampling inside each group is random, the sequence of groups visited for every resolution level follows a pre-determined order. This violates the uniform sampling condition because samples inside a group cannot be placed until all previous groups at the same resolution level have been sampled. For a more comprehensive summary of Poisson sampling methods developed till the year 2008, we refer the interested readers to the survey by Lagae and Dutre [52].

To overcome shortcomings of all previously proposed Poisson Disk algorithms, Gamito and Maddock [35] presented an algorithm in d -dimensions where samples are generated inside the canonical domain $\mathbb{D} = [0, 1]^d$, such that the samples can be subsequently modified by the application of any Euclidean transform without changing the Poisson Disk nature of the distribution. Their algorithm generates Poisson Disk distributions with the samples being uniformly distributed in \mathbb{D} and with the distance between every pair of samples being equal to or greater than a specified distance $2r$, where r is the distribution radius. The algorithm also

generates maximal distributions, in the sense that no new samples can be further inserted in \mathbb{D} without violating the minimum distance constraint relative to other samples. Exceptions occur when there is a point $x \in \mathbb{D}$ that is at an almost equal distance of $2r$ to three or more samples. Depending on the maximum subdivision level of the tree, the algorithm may fail to place an additional valid sample at x . In the worst case, if x is at a distance of exactly $2r$ to three or more samples, the algorithm will fail to place a sample at x irrespective of the maximum subdivision level. This algorithm is easily extendable on surfaces (*e.g.*, sphere).

Ebeida and colleagues [29] addressed the problem of generating a uniform Poisson disk sampling and proposed an algorithm that is both maximal and unbiased over bounded non-convex domains. The authors proposed the first provably correct algorithm with time and space dependent only on the number of points produced. The algorithm works in two phases, both based on classical *dart throwing*. The first phase uses a background grid of square cells to rapidly create an unbiased, near-maximal covering of the domain. The second phase completes the maximal covering by calculating the connected components of the remaining uncovered voids, and by using their geometry to efficiently place unbiased samples that cover them. This second phase converges quickly, overcoming a common difficulty in dart throwing methods. The deterministic memory requirements for this algorithm is $O(N)$ and the expected running time is $O(N \log N)$, where N is the output size, the number of points in the final point set.

2.1.3 Relaxation based methods

Researchers have also proposed many *relaxation* based methods for the generation of blue noise sample distributions. The early algorithms proposed were inspired from the traditional artistic technique of stippling, which involves placing small dots of ink onto paper such that their density give the impression of tone. The artist tightly controls the relative placement of the stipples on the paper to produce even tones and avoid artifacts, leading to long creation times for the drawings.

Deussen and colleagues [23] proposed an approach for stippling that computes an initial dot distribution by a specialized halftoning technique, given a reference image. Otherwise the user creates an initial distribution manually. After the initial point distribution is obtained, the points are moved by an iteration scheme known as Lloyd's method [60, 61]. In the first step, Deussen and colleagues compute the Voronoi diagram [72] of the points. This assigns a Voronoi region to each point. The boundaries of these regions are (possibly open) polygons. In the second step, boundaries of the region to be stippled, are intersected with the Voronoi regions. And finally, each point is moved to the center of the Voronoi region. This dot set is modified automatically or semi automatically to generate a final distribution similar to a stipple drawing. Deussen and colleagues used points of different size and shape to resemble manual drawing.

Secord [87] proposed another iterative method which is a direct descendant of the one described by Deussen and colleagues [23]. Central to Secord's approach is the use of centroidal Voronoi diagrams to produce good distributions of points. A centroidal Voronoi diagram has the interesting property that each generating

point lies exactly on the centroid of its Voronoi region. Secord also incorporate the idea of a density function which weights the centroid calculation. Regions with higher densities pack generating samples (points) closer than regions with lower density values.

Deussen et al. and Secord uses iterative method by Lloyd [60, 61] which is a powerful and flexible technique that is commonly used to enhance the spectral properties of existing distributions of points or similar entities. However, the results from Lloyd's method are satisfactory only to a limited extent. First, if the method is not stopped at a suitable iteration step, the resulting point distributions will develop regularity artifacts. A reliable universal termination criterion to prevent this behavior is unknown. Second, the adaptation to given heterogeneous density functions is suboptimal, requiring additional application-dependent optimizations to improve the results.

Balzer and colleagues [6] present a variant of Lloyd's method, termed capacity constrained Voronoi tessellation (CCVT), which reliably converges towards distributions that exhibit no regularity artifacts and precisely adapt to given density functions. Like Lloyd's method it can be used to optimize arbitrary input point sets to increase their spectral properties while avoiding its drawbacks. They apply the so called capacity constraint that enforces each point in a distribution to have the same capacity. Intuitively, the capacity can be understood as the area of the point's corresponding Voronoi region weighted with the given density function. By demanding that each point's capacity is the same, Balzer et al. ensure that each point obtains equal importance in the resulting distribution. This is a direct

approach to generating uniform distributions, whereas Lloyd's method achieves such distributions only indirectly by relocating the sites into the corresponding centroids. Based on this capacity constraint, an iterative optimization process is performed following the algorithm proposed by Balzer and Heck [5]. The computational complexity per iteration of the CCVT algorithm is $O(N^2 + Nm \log \frac{m}{N})$, where N gives the number of samples (sites) and m is the capacity of each site. Capacity of a site is a number of points assigned to each site.

CCVT has superior quality than Lloyd relaxation, but its time complexity is very high. It could run orders of magnitude slower with the increase in the demand of samples. To address this performance issue, Li and colleagues [58] present a fast CCVT algorithm. Li and colleagues proposed many algorithmic innovations that removes several performance bottlenecks of the original CCVT algorithm. They proposed median site swap to improve the swapping of points within sites in the original CCVT. They also mentioned that by performing a coherent initialization for the input distribution to CCVT algorithm the time complexity can be noticeably improved.

DeGoes and colleagues [22] proposed a novel approach to improve the CCVT algorithm by treating the concept of capacity constrained Voronoi tessellation as an optimal transport problem. This insight leads to a continuous formulation that enables to enforce the capacity constraints exactly, unlike Balzer et al [6]. de Goes and colleagues exploit the variational nature of this formulation to design an efficient optimization technique of point distributions via constrained minimization in the space of power diagrams. This approach leads to a high-quality blue noise

point sets with improved spectral and spatial properties with a computational complexity of $O(N \log N)$.

Chen and colleagues [12] present another approach to address the capacity-constrained Voronoi tessellation (CCVT) algorithm in the continuous setting, by casting the blue noise sampling generation as a variational problem. This variational framework allows generation of high-quality blue noise characteristics while precisely adapting to given density functions. Based on an accurate evaluation of the gradient of an energy function, an efficient optimization is developed which delivers significantly faster performance than the previous optimization-based methods. This framework can be easily extended to generating blue noise point samples on manifold surfaces and for multi-class sampling. The optimization formulation also allows to naturally deal with dynamic domains, such as deformable surfaces, and to yield blue noise samplings with temporal coherence.

Schlömer and colleagues [85] describes an optimization procedure to obtain blue noise samples. Their algorithm iteratively enlarges the minimum distance between points and thereby improves the blue noise characteristics of the point set. The basic algorithm is as follows: each step takes a single point from a set of points X and attempts to move it to a new position that is as far away from the remaining points as possible, i.e., the farthest point. Therefore, they call this optimization procedure, *farthest point optimization* (FPO). One full iteration consists of moving each point in X once. This iteration scheme converges, and each full iteration increases the average minimum distance. In general, the farthest point of a set of points is the center of the largest circle that can be placed in

the domain under consideration without covering any of the points. This largest empty circle can be computed efficiently using the Delaunay triangulation which corresponds to the largest circumcircle of the triangles. The algorithm moves each point in such a way that the point spacing increases monotonically until convergence. The computational complexity of FPO is $O(N \log N)$.

Schmaltz and colleagues [86] introduced a new approach inspired by the physical principals of *electrostatics*. They used the idea of repelling forces between equally charged particles. These forces create a homogeneous distribution in flat regions, while attracting forces from the image brightness values ensure a high approximation quality. Their model is transparent and uses only two parameters: One steers the granularity of their halftoning approach, and the other its regularity. They evaluate both a discrete and continuous version of their algorithm. Their algorithm is global in the sense that it can be used for image dithering, stippling, screening and sampling.

Fattal [30] propose a new approach for generating stochastic blue-noise point distributions that formulates the problem as sampling a statistical mechanics interacting particle model. In this model a radially-symmetric kernel function is placed around every point to produce an approximate density function. The difference between this approximation and the given target point density function assigns an energy value to the points configuration. Rather than minimizing this energy, the algorithm use it to define a Boltzmann-Gibbs statistical model that introduces randomness. Thus, this framework unifies randomness with the requirement for uniform point spacing that achieves the enhanced blue noise

spectral properties.

2.1.4 Anisotropic Blue Noise

Previously mentioned samples generate isotropic blue noise samples. There are many applications including stippling, visualization, surface texturing, and object distribution that requires anisotropic blue noise sample distribution. Feng and colleagues [31] present a practical approach to generate stochastic anisotropic samples with Poisson-disk characteristic over a two-dimensional domain. In contrast to isotropic samples, Feng and colleagues developed anisotropic samples as non-overlapping ellipses whose size and density match a given anisotropic metric. To generate these samples with the desired properties a set of non-overlapping ellipses is constructed whose distribution closely matches the underlying metric. This set of samples is used as input for a generalized anisotropic *Lloyd relaxation* [60] to distribute noise samples more evenly. Instead of computing the Voronoi tessellation explicitly, they introduce a discrete approach which combines the Voronoi cell and centroid computation in one step. The algorithm supports automatic packing of the elliptical samples, resulting in textures similar to those generated by anisotropic reaction-diffusion methods. For quality measurement of uniformly distributed samples the authors use Fourier analysis tools. The resulting samples have nice blue noise property where low frequencies in the power spectrum are reduced to a minimum.

Li and colleagues [59] also proposed an algorithm to generate anisotropic blue noise sample. They also extend the Fourier spectrum analysis, which is for

uniform samples, to evaluate quality of anisotropic and adaptive sampling methods. To generate anisotropic sampling, Li and colleagues extend dart throwing and relaxation—the two classical methods for isotropic blue noise sampling—for the anisotropic setting, while ensuring both high-quality results and efficient computation. On the verification side, Li and colleagues introduce approaches based on warping and sphere sampling that extend Fourier spectrum analysis for adaptive and/or anisotropic samples.

2.1.5 Multi Class Blue Noise

Sampling methods discussed in previous sections so far has been mainly focused on blue noise sampling with a single class of samples. This could be insufficient for common natural as well as man-made phenomena requiring multiple classes of samples, such as object placement, imaging sensors, and color stippling patterns.

Wei [101] extends blue noise sampling to multiple classes where each individual class as well as their unions exhibit blue noise characteristics. Wei propose two flavors of algorithms to generate such multi-class blue noise samples, one extended from traditional Poisson *hard* disk sampling for explicit control of sample spacing, and another based on his *soft* disk sampling for explicit control of sample count. A *hard* disk, centered on each sample, can neither deform nor intersect another, while a *soft* disk can intersect another, but subject to an energy penalty which, when minimized, produces uniform distribution. Both hard and soft disk algorithms support uniform and adaptive sampling, and are applicable to both discrete and continuous sample space in arbitrary dimensions.

Oztireli and Gross [75] introduce another set of methods for analysis and synthesis of general multi-class point distributions based on the statistical measure *pair correlation function* (PCF). To explore the nature of this measure, they introduce an analysis based on the interpretation of a measure as a mean in a high dimensional vector space that they call the *pair correlation space* (PCS). The vector for a given point in the PCS simply measures the distribution of its distance to all other points. This analysis propose an irregularity measure and allows to explain distributions and existing synthesis algorithms in a unified way. It also show that the PCF provides a compact representation for the characteristics. Following this analysis, Oztireli and Gross propose two general synthesis algorithms. The first one is a generalization of dart-throwing for arbitrary PCFs and the other is a gradient descent based fitting of the PCFs. The output of the first algorithm is used as the input for the second to facilitate convergence. These algorithms can generate point distributions with desired characteristics extracted from example distributions or it can be synthesized. The example point sets and generated point sets can be of different dimensions and sizes, contain multiple classes, and reside on non-Euclidean domains. The proposed algorithms have the computational complexity of order $O(N)$, where N is the number of output points.

2.1.6 Spectral control

Most construction algorithms work in the spatial domain, even though the blue noise property itself is defined in the Fourier domain. Parker and colleagues [76] proposed an algorithm for manipulating the power spectra of blue noise halftone

patterns. A blue noise spectrum defined by a step function is used as input to influence the sample distributions.

Zhou and colleagues [110] construct point sets matching a Fourier power spectrum function by performing a gradient descent optimization on an energy derived from the autocorrelation function. Such a Fourier power spectrum function can be either obtained from a known sampling method, or completely constructed by the user. The key idea is to convert the Fourier spectrum function into a *differential distribution function* that describes the samples' local spatial statistics. The algorithm then use a gradient descent solver to iteratively compute a sample set that matches the target differential distribution function. It can be easily modified to achieve adaptive sampling, and is extendable to GPU.

Heck and colleagues [41] also studied directly the Fourier properties of sampling patterns to perform analysis of blue noise sampling. They started their analysis from the observation that oscillations in the power spectrum of a sampling pattern can cause aliasing artifacts in the resulting images. They synthesize two new types of blue noise patterns: *step* blue noise with a power spectrum in the form of a *step function* and *single-peak* blue noise with a wide zero-region and no oscillations except for a single peak. They study the mathematical relationship of the radial power spectrum to a spatial statistic known as the *radial distribution function* to determine which power spectra can actually be realized and to construct the corresponding point sets. The method discussed in Section 2.1.5 by Oztireli and Gross [75] also offers spectral control on top of multi-class blue noise distributions.

2.1.7 Tile-based Blue Noise Sampling

Researchers also proposed many tile-based methods to overcome the computational complexity of *dart-throwing* and/or *relaxation* based approaches in generating blue noise sampling patterns. The first tile-based Poisson Disk sampling methods used Wang tiles and was proposed by Hiller and colleagues [44]. This Wang tile-based approach was later extended by Cohen and colleagues [15] for image and texture generation. Wang tiles have colours assigned to their edges in specific ways. A Wang tile can only be placed next to another if they share the same colour along the common edge. This allows non-periodic tilings of the plane to be created. The generation of Poisson-disk samples inside each tile must respect the minimum distance constraint across the edges of the tile relative to all other tiles that share the same edge colour. The authors achieve this by using several steps of Voronoi relaxation [61].

In the initial Wang tile methods, the tiling had to be computed in advance inside some finite region of space. Lagae and Dutré [53] introduced procedural tiling rules that allow a Wang tile to be assigned on the fly in a consistent way to any arbitrary point in space. This leads to the creation of infinite non-periodic tilings of Poisson-disk samples. Later, Lagae and Dutré [54] introduced procedural tiling rules for *corner tiles*. Corner tiles have colors associated to their corners instead of their edges. They can enforce the minimum distance constraint across tiles that share a common corner. The same authors also extended corner tiles to three-dimensions, creating corner cubes [51].

Methods for non-uniform Poisson-disk sampling have been proposed based

on tile distributions. Kopf and colleagues [50] apply subdivision rules to Wang tiles in order to create sample distributions with varying density across space. Their technique utilizes a set of carefully constructed progressive and recursive blue noise Wang tiles. The use of Wang tiles enables the generation of infinite non-periodic tilings. The progressive point sets inside each tile are able to produce spatially varying point densities. Recursion allows to adaptively subdivide tiles only where high density is required, and makes it possible to zoom into point sets by an arbitrary amount, while maintaining a constant apparent density.

Ostromoukhov and colleagues [74] use Penrose tiling and perform hierarchical subdivision to create a sufficiently large number of sample points. Each Penrose tile has a single sample inside, which is subject to a Voronoi relaxation together with the samples from other tiles to reduce sampling artifacts. These points are numbered with *structural indices* using the Fibonacci number system, and these numbers are used to threshold the samples against the local value of the importance density. Pre-computed correction vectors, obtained using relaxation, are used to improve the spectral characteristics of the sampling pattern. The technique is fast ($O(N)$), deterministic and can be used for importance sampling.

Later, Ostromoukhov [73] present another fast hierarchical importance sampling method with blue-noise properties. This approach is based on self-similar tiling of the plane or the surface of a sphere with rectifiable polyominoes. Sampling points are associated with polyominoes, one point per polyomino. Each polyomino is recursively subdivided until the desired local density of samples is reached. A numerical code generated during the subdivision process is used for

thresholding to accept or reject the sample. The exact position of the sampling point within the polyomino is determined according to a structural index, which indicates the polyomino's local neighborhood. The variety of structural indices and associated sampling point positions are computed during the off-line optimization process, and tabulated. Consequently, the sampling itself is extremely fast. The method allows both deterministic and pseudo-non-deterministic sampling.

Wachtel and colleagues [99] introduce a first tile-based method for adaptive two-dimensional sampling with user-specified spectral properties. They selected hexagons as the building blocks of their tiling due to their additional symmetry axis, better intrinsic spectral properties, and only one adjacency relationship between neighboring blocks [39]. Wachtel and colleagues build a hierarchical and deterministic subdivision process based on trihexes (*i.e.*, connected agglomerates of three hexagons) as prototiles to break the symmetries of the regular hexagonal lattice—much like polyominoes were unions of squares introduced to create more complex tiling structures on a square lattice. By simply using barycenters of these trihexes as sample points, the Fourier spectrum of the underlying sampling pattern appears to be lack of spurious high energy Fourier peaks. To remove the remaining frequency peaks in the resulting Fourier spectrum they devised a hierarchical, deterministic, and area-preserving *border shuffling* procedure that mimics a stochastic process. By combining these features, they create non-overlapping tiles whose barycenters form a high quality blue noise distribution. A lookup table of sample points, computed offline using any existing procedure that optimizes point sets to shape their Fourier spectrum, is then used to populate the tiles.

The resultant method gives a linear-time, adaptive, and high-quality sampling of arbitrary density functions that conforms to the desired spectral distribution.

Tile based methods can be used, for example, to overcome the defaults of Poisson disk sampling which enforces its radial power spectrum to have an offset in the low frequency zone. By using the tile-based approach, user can enforce a proper subdivision of the sampling domain \mathbb{D} which can be later sampled with Poisson disk. The resulting Poisson disk sampling results in a radial power spectrum with no low frequency content. Recently, Yan and colleagues [107] published a complete survey on blue noise sampling techniques. We refer interested readers to Yan and colleagues survey for sampling methods that might not have been covered in this chapter.

2.2 Sampling in Spherical domain

In image synthesis, spherical and hemispherical domain plays a crucial role in the evaluation of global illumination integrals. Many algorithms are proposed to perform sampling on the sphere and hemisphere. Researchers have studied various 3D surface sampling techniques that can be used directly for sampling on a sphere. However, since the 2D spherical domain is very simple, these surface sampling algorithms can be considered burdensome in their general settings for sphere sampling. Therefore, researchers proposed many simple algorithms, to sample the sphere directly or indirectly using a mapping function. Many deterministic sampling methods are also proposed for sphere sampling in the Quasi-Monte

Carlo numerical integration literature. In this section, we look at these approaches one by one.

2.2.1 Sampling 3D surfaces

In the geometry remeshing literature, many algorithms are proposed that can be used directly to generate samples on a unit sphere. To give an example, for surface sampling, Alliez and colleagues [2] generate seed points by error diffusion on a triangle mesh to generate an isotropic surface remeshing. This algorithm can be used to sample a sphere. However, to keep our focus on only sphere sampling, we do not cover the literature work on surface geometry remeshing in this dissertation and refer interested readers—for an early work on remeshing—to a complete survey by Alliez and colleagues [1].

Algorithms based on dart throwing for surface sampling are also proposed in the literature. Bowers and colleagues [8] proposed a parallel dart throwing algorithm for efficient Poisson disk sampling on arbitrary manifold surfaces. The authors extend the parallel grid cell sampling approach by Wei [100] from Euclidean to geodesic distance metric. Since the geodesic distance cannot be smaller than the Euclidean distance, their algorithm can draw samples directly on surfaces without any parameterization. To account for geodesic distance, they propose a fast approximation that is easy to compute and accurate for close-by sample points.

Cline and colleagues [14] present dart throwing algorithms to generate maximal Poisson disk point sets directly on 3D surfaces. They optimize dart throwing

by efficiently excluding areas of the domain that are already covered by existing darts. In the case of triangle meshes, the algorithm shows dramatic speed improvement over comparable sampling methods. They demonstrated the ability to directly handle several surface types, including triangle meshes, spheres, Bezier patches, subdivision surfaces and implicits. Their algorithm can also be used to handle non-uniform densities, ellipsoid placement, and point sets spaced according to geodesic distance. Optimized dart throwing is quite fast in the uniform case, but it can run much more slowly when producing non-uniform distributions because of the acceleration grid.

Peyrot and colleagues [77] proposed an algorithm to perform direct Poisson disk sampling for surface meshes. Their algorithm is focused on sampling triangular meshes, while satisfying good blue noise properties, they also preserve features. The algorithm starts by initializing a list of vertices available for the sampling (the so-called *candidate vertices*) with all the vertices of a densely sampled mesh M_{sub} . After each valid dart throwing giving a valid sample s , the candidate vertices located into the sphere relative to s (depending on its radius R) are removed from the set of available vertices. If a dart throwing fails for a sample s , this latter is removed from available vertices to avoid a useless second test at this location. To efficiently evaluate the region of the mesh relative to a given sphere associated to a sample s , their Dijkstra-based algorithm is able to stop its region growing process under given conditions.

Turk [95] proposed a relaxation based algorithm which uses a variant of Lloyd relaxation defined for polygon meshes. The method starts by placing points

randomly on the surface. The points then repel each-other, eventually reaching a uniform distribution. However, it is difficult to choose the proper number of iterations to achieve the right balance between uniformity and randomness, as too few iterations might produce insufficient uniformity whereas too many iterations might produce a highly regular result.

Xu and colleagues [106] proposed a blue noise sampling technique for surfaces that is motivated from the CCVT algorithm. The core idea behind the method is to compute a Capacity-Constrained Delaunay Triangulation (CCDT), namely, given a simple polygon P in the plane, and the desired number of points N , compute a Delaunay triangulation of the interior of P with N Steiner points, whose triangles have areas which are as uniform as possible. This is computed iteratively by alternating update of the point geometry and triangulation connectivity. The vertex set of the CCDT is shown to have good blue noise characteristics, comparable in quality to those of state-of-the-art methods, achieved at a fraction of the runtime. The CCDT method may be applied also to an arbitrary density function to produce non-uniform point distributions.

Xu and colleagues [105] proposed a more generalized approach to generate blue noise samples on surfaces. It is based on the concept of Capacity-Constrained Surface Triangulation (CCST), which approximates the underlying continuous surface as a well-formed triangle mesh with uniform triangle areas. The algorithm takes a triangle mesh and the number of sample points as input, and iteratively alternates between optimization of the geometry (positions) of the points and optimization of their topology (connectivity) until convergence. Since the method

is relaxation-based, it allows precise control over the number of sample points. The algorithm alternates between a geometry optimization phase, which minimizes the variance of triangle capacities, and an edge-flipping based retriangulation phase, until convergence. When applied to a uniform distribution, the result has been shown to possess superior blue noise characteristics.

As relaxation-based sampling methods, both CCST and CCDT could be regarded as the dual versions of CCVT method. CCVT operates on Voronoi diagrams while CCST/CCDT generates the dual Delaunay triangulations. All of these three methods start with some random initializations and converge to some local minima of area variance that possess spatial uniformity of points after several alternations between geometry and topology phases. These local minima also avoid regularity artifacts that are caused by global minima, namely all the triangles/cells that are equilateral, thus achieving blue noise characteristics. The iterative algorithm of CCVT could be categorized into Lloyd's relaxation method, typically requiring about 100 alternations between geometry and topology phases. In contrast, CCST/CCDT method aims at directly minimizing variance of Delaunay triangle areas, leading the algorithm converges in less than 10 alternations between geometry and topology phases, thus is much faster than CCVT method. The CCST algorithm can also be regarded as the generalization of the CCDT method for planar surfaces to curved surfaces.

There are some tile-based methods proposed in the literature for surface sampling. As discussed in Section 2.1.7, Ostromoukhov [73] proposed a polyomino's based tiling method that can be easily extended on the surface of a sphere. Li and

colleagues [57] generalize the concept of Wang tiles to surfaces. Once computed, the tiling can be used to quickly distribute points evenly over the surface, but how this method can be used for non-uniform densities was not demonstrated. To overcome distortion in the tiling, Lloyd's method is applied to the initial point distribution made by the tiling.

All the methods discussed in this section are for 3D surfaces which can be used to sample a sphere.

2.2.2 Sampling on the sphere

Researchers have also proposed many algorithms to directly sample the surface of a sphere. Rakhmanov and colleagues [80] investigated the energy of arrangements of N points on the surface of the sphere in \mathbb{R}^3 , by interacting through a power law potential. They proposed an area-regular sphere partitioning scheme for the purpose of obtaining bounds for the extremal (equilibrium) energy for such points. They devise a scheme to distribute points over a nearly regular spherical hexagonal net, which they called the *generalized spiral points*.

Saff and Kuijlaars [83] studied the fundamental nature of distributing many points on the sphere. They noticed that the general pattern for optimal configuration of the points on the sphere is the same. For large number of points, they observed experimentally that on partitioning the sphere into Dirichlet (Voronoi) cells, all points but exactly 12 of the Voronoi cells for an optimal configuration are *hexagonal*. That is, with large enough N , points appear to arrange themselves according to a hexagonal pattern that is slightly perturbed in order to fit on the

sphere.

Cui and Freeden [19] proposed a concept of generalized discrepancy, which involves pseudo differential operators to give a criterion of equidistributed point sets on the sphere. Cui and Freeden extend the famous Koksma-Hlawka inequality of Numerical Analysis from the unit hypercube to the sphere, and evaluate the uniformity properties of some deterministic sequences on the sphere. Recently, Choirat and Seri [13] derive the properties of this new class of statistics, called the *generalized discrepancies* introduced by Cui and Freeden, for testing equidistribution on the sphere and to investigate their computational aspects.

Yershova and LaValle [109] studied deterministic sampling method on the sphere. The authors proposed a general framework for performing deterministic uniform sampling over spheres and the three- dimension rotation group $SO(3)$. They developed and implemented a particular sequence which extends the layered Sukharev grid sequence [93] designed for the unit cube. They tested the performance of the sequence in PRM-like motion planning algorithms, which demonstrated that this sequence is a useful alternative to a random sampling. This is in addition to the advantages that this sequence has over random sampling, such as deterministic resolution completeness guarantees and the regular lattice structure. For a complete survey of recent developments in the context of spherical designs and minimal energy point configurations on sphere, we encourage readers to see the compilation by Brauchart and Grabner [10].

Another very promising sphere surface sampling was proposed by Gorski and colleagues [37], called HEALPix. The inherent data structure contains an equal

area quadrangulation of the sphere. It was originally developed to address the data processing and analysis needs of the present generation of cosmic microwave background (CMB) experiments. HEALPix with an associated library is used for large area surveys in the form of discretized spherical maps. HEALPix possesses the following three essential properties; First, the sphere is hierarchically tessellated into curvilinear quadrilaterals with the lowest resolution partition comprised of 12 base “pixels”. Resolution of the tessellation increases by division of each “pixel” into four new ones. Second, areas of all “pixels” at a given resolution are identical and, Thirdly, all “pixels” are distributed on lines of constant latitude. This property is essential for all harmonic analysis applications involving spherical harmonics. But the spherical harmonic pixelization on the sphere presented in HEALPix is inexact, and do not lead to a sampling theorem on the sphere.

2.2.3 Sampling theorem on the sphere

Sampling theorem on the sphere state that all of the information contained in a band-limited signal may be represented by a finite set of samples in the spatial domain. On the sphere, unlike Euclidean space, the number of samples required in the harmonic and spatial domains differ, with different sampling theorems on the sphere requiring a different number of samples in the spatial domain.

The fundamental property of any sampling theorem is the number of samples required to represent a band-limited signal. To represent exactly a signal on the sphere band-limited at degree L , all sampling theorems on the sphere require $O(L^2)$ samples. There are several sampling theorems proposed on the sphere.

For an equiangular sampling of the sphere, the Driscoll and Healy (DH) sampling theorem [26] has become the standard. It requires approximately $4L^2$ samples on the sphere. The complexity of DH sampling theorem and the corresponding algorithms to compute fast spherical harmonic transforms is $O(L^3)$. It also require a precomputation or otherwise restricted use of Wigner recursions.

In Gauss-Legendre sampling theorem, sample positions are given by the roots of Legendre functions. It requires approximately $2L^2$ samples on the sphere. A simple separation of variables gives algorithm with the complexity $O(L^3)$. Similar to DH sampling theorem, Gauss-Legendre sampling theorem also require a precomputation or otherwise restricted use of Wigner recursions.

McEwen and Wiaux [67] recently developed a sampling theorem and the corresponding fast algorithms by factoring of rotations and then associating the sphere with the torus through a periodic extension. MW sampling theorem requires less than half the number of samples of other equiangular sampling theorems on the sphere and an asymptotically identical, but smaller, number of samples than the Gauss-Legendre sampling theorem. The complexity of their algorithms scale as $O(L^3)$, however, the continual use of fast Fourier transforms reduces the constant prefactor associated with the asymptotic scaling considerably, resulting in algorithms that are fast.

An in-depth study of the above mentioned sampling theorems is out of the scope of this dissertation. However, we have given brief description in this subsection for completeness of our survey.

2.2.4 Sphere sampling for illumination integrals

As mentioned earlier, in the evaluation of global illumination integrals hemispherical and spherical sampling plays a crucial role in deciding the scatter direction of incoming and/or outgoing rays. Shirley [90] proposed a low distortion map between disk and square that can be directly used to map samples from a square domain on to the hemisphere. Arvo [3, 4] proposed a spherical triangle sampling technique that reduces the noise level in illumination evaluation upto a greater extent.

Shao and Badler [89] proposed a novel spherical sampling technique based on the Archimedes' Theorem. Archimedes' theorem states that; Globally, the area of a sphere equals the area of every right circular cylinder circumscribed about the sphere excluding the bases. Locally, the axial projection of any measurable region on a sphere on the right circular cylinder circumscribed about the sphere preserves area. The proposed algorithm based on these axioms is simple and efficient for generating uniformly distributed samples on the unit sphere. The implementation is straightforward and may be easily extended to include stratified sampling for variance reduction.

Dimov and colleagues [24] address the problem for generation of uniformly distributed random samples over hemisphere and sphere and study the parallel sampling scheme for hemisphere and sphere. First they apply the symmetry property for partitioning of hemisphere and sphere. The domain of solid angle subtended by a hemisphere is divided into a number of equal sub-domains. Each sub-domain represents solid angle subtended by orthogonal spherical triangle with

fixed vertices and computable parameters. Dimov and colleagues introduce two new algorithms for sampling of orthogonal spherical triangles. The first sampling algorithm generates a sample by mapping of the unit square onto orthogonal spherical triangle. The second algorithm directly compute the unit radius vector of a sampling point inside to the orthogonal spherical triangle.

Ureña and colleagues [98] proposed an area preserving spherical rectangle parametrization to reduce variance in scattering computation from planar rectangular emitters. Marques and colleagues [65] introduced spherical Fibonacci lattices to improve QMC sampling in the hemispherical domain. Marques and colleagues also studied the worst-case error for spherical Fibonacci point sets for illumination integrals.

2.3 Quality criteria for sample distributions

As discussed before, Fourier analysis has been extensively used to study the frequency content of the sampling patterns. Fourier spectrum obtained by averaging the periodograms [7] and the corresponding radially averaged spectrum can be considered as a good measure for quality for uniform sampling patterns. Li and colleagues [59] introduce adaptive and anisotropic sampling analysis based on warping and sphere sampling to extend the Fourier spectrum analysis from uniform to non-uniform sample distributions.

Bowers and colleagues [8] proposed an algorithm for analyzing the spectral distribution quality of surface samples. In planar sampling, the distribution quality

is typically measured in terms of the radial means and anisotropy of the Fourier power spectrum [52]. Bowers and colleagues extend these concepts to arbitrary manifold surfaces by employing spectral mesh basis functions, derived as the eigenfunctions of the discrete mesh Laplacian operator [56]. These functions define a Fourier-style basis set that exists on the mesh surface, and hence can be used to evaluate the power spectrum of samples distributed over the mesh. Because the spectral mesh basis obeys surface geodesic distance and does not require any mesh parametrization, it provides a convenient way to study and compare the quality of different surface sampling algorithms.

Wei and Wang [102] proposed a more robust method for analyzing non-uniform sample distributions. The key insight of their algorithm is that standard Fourier analysis, which depends on samples' spatial locations, can be reformulated into an equivalent form that depends only on the distribution of their location differentials. They call this *differential domain analysis*. The main benefit of this reformulation is that it bridges the fundamental connection between the samples' spatial statistics and their spectral properties. In addition, it allows to generalize their method with different computation kernels and differential measurements. Using this analysis, they can quantitatively measure the spatial and spectral properties of various non-uniform sample distributions, including adaptive, anisotropic, and non- Euclidean domains.

Shirley [91] introduced the notion of discrepancy to the computer graphics community to compute the quality of sampling patterns. Even though, Fourier analysis tools can be used to study the behaviour of various uniformly distributed

sampling distributions but this approach only provides a quantitative information in the form of two-dimensional frequency spectrum. Computation of discrepancy allows assigning a single quality number, the *discrepancy*, to the point set. This allows to order point sets according to their *discrepancy* and compare which one is the *best*.

The quality criteria derived from the discrepancy of sample positions is related to the Koksma- Hlawka inequality. The Koksma-Hlawka inequality is a tight error bound on the approximation of an integral by the sample average of integrand values:

$$\left| \frac{1}{N} \sum_{j=1}^N F(x_j) - \int_0^1 F(x) dx \right| \leq \mathcal{D}(x_j) \mathcal{V}(F). \quad (2.1)$$

In this inequality $\mathcal{D}(x_j)$ is the discrepancy of the points $0 \leq x_j \leq 1$ and $\mathcal{V}(F)$ is the total variation of the function F ,

$$\mathcal{D}(x_j) = \sup_{0 \leq t \leq 1} \left\{ \left| \frac{1}{N} \chi_{[0,t]}(x_j) - t \right| \right\}, \quad (2.2)$$

$$\mathcal{V}(F) = \sup_{0=y_0 < y_1 < \dots < y_n=1} \left\{ \sum_{k=1}^n |F(y_k) - F(y_{k-1})| \right\}, \quad (2.3)$$

where $\chi_{[a,b]}$ is a characteristic function which non-zero only in the range $[a, b]$ and 1 otherwise. Hickernell [43] gives a detailed overview of this inequality and mentions that although the Koksma-Hlawka inequality originally derived for a particular integration domain \mathbb{D} , and a particular space of integrands, F , the name may be applied to similar inequalities that have been derived for other \mathbb{D} and F . In these inequalities the integration error is bounded by a product of two terms, the discrepancy of the sample points, and the variation of the integrand.

In other words, the Koksma-Hlawka inequality splits the error into the part due to the quality of the sample points and the part due to the roughness of the integrand. When F is a reproducing kernel Hilbert space, the Koksma-Hlawka inequality is straightforward to derive, and there is a simple formula for the discrepancy. This discrepancy also has several other useful interpretations, including, (i) how the proportion of sample points in a box deviates from the volume of the box, (ii) the average-case integration error, and (iii) a goodness-of-fit statistic. Integration lattices and digital sequences are two popular families of low discrepancy sample points. These sets typically give better convergence rates for the discrepancy than a simple random sample. The Koksma-Hlawka inequality plays a key role in the development of quasi-Monte Carlo methods. It has also influenced the study of experimental design and led to the creation of *uniform designs*.

2.4 Discussion

In this chapter, we discussed various sampling methods to generate blue noise samples. We also studied various sampling methods that can offer good spectral control. The benefit of having a good spectral control over samples is that we can design a Fourier profile according to the integrand and can generate samples with respect to that profile. This can be very beneficial to reduce the integration error in Monte Carlo integration of the integrand involved. However, there is a missing connection between the error in Monte Carlo integration and the Fourier spectrum of sampling patterns.

Recent work has conducted a comprehensive analysis of error in Monte Carlo integration. In particular, Durand [28] investigated error in integration and relates error to the spectral properties of the associated sampling patterns. Ramamoorthi and colleagues [81] focused specifically on visibility to assess error due to sampling patterns in soft shadow rendering. Subr and Kautz [92] proposed a mathematical formulation that relates the variance in Monte Carlo integration directly to the variance of sampling Fourier coefficients taken over multiple realizations. However, this is still not sufficient as the variance of sampling Fourier coefficients can not be characterized according to the requirements.

In this dissertation, we establish a direct connection between the error in integration and the power spectrum of a sampling pattern in the spherical domain. We use the notion of *homogeneous* sample distributions in the spherical domain, which has been extensively studied in physics and spatial statistics [40, 20, 21]. Homogeneous samples have same properties over the whole domain. Thanks to the homogeneous sampling, the error in MC integration can be expressed in the form of variance only. In Chapter 5, we develop a direct relation between the variance in MC integration and the power spectra of both the sampling pattern and the integrand involved in the spherical domain. This work has been successfully published by Pilleboue, Singh and colleagues [79].

Chapter 3

Mathematical Background

In this chapter, we discuss some basic mathematical tools that we use to lay down the platform of our work on variance analysis for Monte Carlo integration (Chapter 5). Our mathematical model proposed in Chapter 5 uses the core ideas from representation theory. Therefore, in this chapter, we introduce the background on representation theory in a step-by-step manner while providing all prerequisites that are needed to get a grasp on this theory. We keep our discussion to the point by directly introducing definitions of various entities without going in depth details. We encourage interested readers to refer to an introductory text by Herstein [42] and Hoffmann and Kunze [46] to brush up their linear and abstract algebra, respectively.

We first introduce briefly a light background on vector spaces and group theory which is essential to understand the representation theory. We later discuss the results which are used to derive our variance formulation in a compact manner. Later, we discuss spherical harmonics which has been heavily employed in this

work to understand the frequency content of various sampling patterns on the sphere and the hemisphere. We assume that the reader is familiar with the elementary algebra of real and complex numbers and with the basic ideas from set theory and topology.

3.1 Preliminaries

Before we dig into our mathematical introduction we give some preliminary definitions and introduce the notion of *mapping*.

3.1.1 Mapping

Let A and B be sets; a *mapping* or *function* f from A to B is a *rule* that associate each element in A to a maximum one element in B . That is, a function never has one element from A associated to more than one element in B . We denote that f is a mapping from A to B by $f : A \rightarrow B$.

Definition 3.1.1. *The mapping $f : A \rightarrow B$ is one-to-one (written as 1-1) or injective if each element in A is associated to only one element of B .*

Definition 3.1.2. *The mapping $f : A \rightarrow B$ is many-to-one or onto or surjective if and only if, for a given element of B there exists an element in A .*

Definition 3.1.3. *The mapping $f : A \rightarrow B$ is said to be a bijection or 1-1 correspondence if f is both 1-1 and onto.*

Definition 3.1.4. If $f : A \rightarrow B$ and $g : B \rightarrow C$, then the **composition** (or [product), denoted by $f \circ g$, is the mapping $f \circ g : A \rightarrow C$ defined by $(f \circ g)(a) = f(g(a))$ for every $a \in A$.

In abstract algebra, there exist structure-preserving mappings between algebraic structures. We describe various algebraic mathematical properties that holds for all structures like rings, groups, algebras.

Definition 3.1.5. A **homomorphism** is a structure preserving map between algebraic structures.

Homomorphism exists for different structures like rings, groups, algebras. Later in this chapter, we define homomorphism for structures, like groups and inner product spaces which is more suited to our context.

Definition 3.1.6. A **isomorphism** between two algebraic structures is a homomorphism whose inverse is also a homomorphism.

An isomorphism thus defines an equivalence between two structures that enables to perform all operations on the first structure also with the second one, and then relate the result back to the first.

3.2 Vector Spaces

In this section, we give a brief background on vector space. We start with the definition of a *field*;

Definition 3.2.1. A set A , together with some operations on the objects in that set which behave like ordinary addition, subtraction, multiplication, and division of numbers in the sense that they obey the rules of elementary algebra, constitutes a *field*.

To give an example, the complex space \mathbb{C} is a *field*.

Definition 3.2.2. A *vector space* (or linear space) consists of the following:

- (a) a field Υ of scalars;
- (b) a set V of objects, called vectors;
- (c) a rule (or operation), called *vector addition*, which associates with each pair of vectors \vec{v}_1, \vec{v}_2 in V a vector $\vec{v}_1 + \vec{v}_2$ in V , called the *sum of \vec{v}_1 and \vec{v}_2* , in such a way that:
 - (i) *addition is commutative*, $\vec{v}_1 + \vec{v}_2 = \vec{v}_2 + \vec{v}_1$
 - (ii) *addition is associative*, $\vec{v}_1 + (\vec{v}_2 + \vec{v}_3) = (\vec{v}_1 + \vec{v}_2) + \vec{v}_3$
 - (iii) *there is a unique vector $\mathbf{0}$ in V , called the zero vector, such that $\vec{v} + \mathbf{0} = \vec{v}$ for all \vec{v} in V .*
 - (iv) *for each vector \vec{v} in V there is a unique vector $-\vec{v}$ in V such that $\vec{v} + (-\vec{v}) = \mathbf{0}$;*
- (d) a rule (or operation), called *scalar multiplication*, which associates with each scalar c in Υ and vector \vec{v} in V a vector $c\vec{v}$ in V , called the *product of c and \vec{v}* , in such a way that:
 - (i) $1\vec{v} = \vec{v}$ for every \vec{v} in V ;

$$(ii) (\mathbf{c}_1\mathbf{c}_2)\vec{v} = \mathbf{c}_1(\mathbf{c}_2\vec{v});$$

$$(iii) \mathbf{c}(\vec{v}_1 + \vec{v}_2) = \mathbf{c}\vec{v}_1 + \mathbf{c}\vec{v}_2;$$

$$(iv) (\mathbf{c}_1 + \mathbf{c}_2)\vec{v} = \mathbf{c}_1\vec{v} + \mathbf{c}_2\vec{v};$$

It is important to observe that a vector space is a composite object consisting of a *field*, a set of ‘vectors’, and two operations with certain special properties.

Definition 3.2.3. Let V be a vector space over the field Υ . A **subspace** of V is a subset W which is itself a vector space over Υ with the operations of vector addition and scalar multiplication on V .

Definition 3.2.4. The **direct sum** $V \oplus W$ of two linear space V and W consists of order pairs (\vec{v}, \vec{w}) for $\vec{v} \in V$ and $\vec{w} \in W$, with addition and scalar multiplication defined componentwise.

Definition 3.2.5. Let V_1 and V_2 be vector spaces over the field Υ . A **linear transformation from V_1 into V_2** is a function T from V_1 into V_2 such that:

$$T(\mathbf{c}\vec{v}_1 + \vec{v}_2) = \mathbf{c}(T\vec{v}_1) + T\vec{v}_2$$

for all \vec{v}_1 and \vec{v}_2 in V and all scalars \mathbf{c} in Υ .

Linear Transformations preserves the operations of addition and multiplications.

Definition 3.2.6. If V is a vector space over the field Υ , a linear transformation f from V into the scalar field Υ is called a **linear functional** on V .

If we start from scratch, this means that f is a function from V into Υ such that:

$$f(\mathbf{c}\vec{v}_1 + \vec{v}_2) = \mathbf{c}f(\vec{v}_1) + f(\vec{v}_2)$$

for all vectors \vec{v}_1 and \vec{v}_2 in V and all scalars \mathbf{c} in Υ . The concept of linear functional is important in the study of finite-dimensional spaces because it helps to organize and clarify the discussion of subspaces linear equations, and coordinates. Interested readers are encouraged to read more concise text from Hoffmann and Kunze [46].

Example 3.2.1. Here is an important example of a linear functional. Let n be a positive integer and Υ a field. If M is an $n \times n$ matrix with entries in Υ , the **trace** (Tr) of M is the scalar:

$$Tr(M) = M_{11} + M_{22} + \cdots + M_{nn}.$$

Definition 3.2.7. Let Υ be the field of real numbers or the field of complex numbers, and V a vector space over Υ . An **inner product** on V is a function which assigns to each ordered pair of vectors \vec{v}_1, \vec{v}_2 in V a scalar $\langle \vec{v}_1, \vec{v}_2 \rangle$ in Υ in such a way that for all $\vec{v}_1, \vec{v}_2, \vec{v}_3$ in V and all scalars \mathbf{c} :

1. $\langle \vec{v}_1 + \vec{v}_2, \vec{v}_3 \rangle = \langle \vec{v}_1, \vec{v}_3 \rangle + \langle \vec{v}_2, \vec{v}_3 \rangle$
2. $\langle \mathbf{c}\vec{v}_1, \vec{v}_2 \rangle = \mathbf{c} \langle \vec{v}_1, \vec{v}_2 \rangle$
3. $\langle \vec{v}_2, \vec{v}_1 \rangle = \overline{\langle \vec{v}_1, \vec{v}_2 \rangle}$, the bar denoting complex conjugation;
4. $\langle \vec{v}, \vec{v} \rangle > 0$ if $\vec{v} \neq 0$.

Definition 3.2.8. An **inner product space** is a real or complex vector space, together with a specified inner product on that space.

3.2.1 Unitary Operators

In this section, we consider the concept of an isomorphism between two inner product spaces. If V_1 and V_2 are vector spaces, an isomorphism of V_1 onto V_2 is a one-one linear transformation from V_1 onto V_2 , *i.e.*, a 1-1 correspondence between the elements of V_1 and those of V_2 , which ‘preserves’ the vector space operations. Now an inner product space consists of a vector space and a specified inner product on that space. Thus, when V_1 and V_2 are inner product spaces, we shall require an isomorphism from V_1 onto V_2 not only to preserve the linear operations, but also to preserve inner products. An isomorphism of an inner product space onto itself is called a **unitary operator** on that space.

Definition 3.2.9. *Let V_1 and V_2 be inner product spaces over the same field, and let T be a linear transformation from V_1 into V_2 . We say that T **preserves inner products** if $\langle T\vec{v}_1, T\vec{v}_2 \rangle = \langle \vec{v}_1, \vec{v}_2 \rangle$ for all \vec{v}_1 and \vec{v}_2 in V_1 . In this case, T is called the **unitary transformation**. An **isomorphism** V_1 onto V_2 is a vector space isomorphism T of V_1 onto V_2 which also preserves inner products.*

3.3 Group Theory

We first discuss what a *group* really composed of and we also see the axioms which makes a set, a *group*. Group theory has a vast back ground, and in this section we focus on only the properties and groups we need to present our background on *representation theory* in Section 3.4.

3.3.1 Groups

To understand the notion of Groups, we start with a collection of objects A and endow this collection with an algebraic structure by assuming that we can combine, in different ways, these elements to obtain, once more, elements of this set A . These ways of combining elements of A we call *operations* on A . Now, we can try to regulate the nature of A by imposing some constraints or rules on how these operations behave on A . These rules are usually called the *axioms* defining the particular structure on A . *Groups* are one of the basic axiomatic algebraic systems that allows us to work with a collection without ending up in a nonsensical contradiction computing within the framework. That is, *Groups* are designed with such *axioms* that would always be well-behaved and would end up with a consistent expression. *Fields* and *Rings* are also some examples of such basic axiomatic algebraic systems. In this section, we focus only on *Groups*.

Definition 3.3.1. *A nonempty set H is said to be a **group** if in H there is defined an operation $*$ such that:*

- (a) Closed: $h_1, h_2 \in H$ implies that $h_1 * h_2 \in H$. This property is described by saying that H is closed under $*$.
- (b) Associativity: Given $h_1, h_2, h_3 \in H$ then $h_1 * (h_2 * h_3) = (h_1 * h_2) * h_3$. That is, associative law holds in H .
- (c) There exists a special element $e \in H$ such that $h * e = e * h = h$ for all $h \in H$.

(d) For every $h_1 \in H$ there exists an element $h_2 \in H$ such that $h_1 * h_2 = h_2 * h_1 = e$. We write this element h_2 as h_1^{-1} and call it the inverse of h_1 in H .

These are the four defining postulates of a group which are called *group axioms*.

Definition 3.3.2. A group H is said to be a **finite group** if it has a finite number of elements. The number of elements in H is called the order of H and is denoted by $|H|$.

Definition 3.3.3. A group H is said to be **abelian** if $h_1 * h_2 = h_2 * h_1$ for all $h_1, h_2 \in H$.

Definition 3.3.4. A relation \equiv on a set A is called an **equivalence relation** if, for all $a_1, a_2, a_3 \in A$, it satisfies:

(a) $a_1 \equiv a_1$ (reflexivity)

(b) $a_1 \equiv a_2$ implies that $a_2 \equiv a_1$ (symmetry)

(c) $a_1 \equiv a_2, a_2 \equiv a_3$ implies that $a_1 \equiv a_3$ (transitivity).

Remark Of course, equality $=$, is an equivalence relation, so the general notion of equivalence relation is a generalization of that of equality. In a sense, an equivalence relation measures equality with regard to some attribute.

Definition 3.3.5. If \equiv is an equivalence relation on set A , then $[a]$, the **class** of a , is defined by $[a] = \{b \in A \mid b \equiv a\}$.

Definition 3.3.6. Let H, H' be two groups; then the mapping $\varphi : H \rightarrow H'$ is a **homomorphism** if $\varphi(h_1 * h_2) = \varphi(h_1) * \varphi(h_2)$ for all $h_1, h_2 \in H$.

Definition 3.3.7. The homomorphism $\varphi : H \rightarrow H'$ is called **monomorphism** if φ is 1-1. A monomorphism that is onto is called an **isomorphism**. An isomorphism from H to H itself is called an **automorphism**.

Definition 3.3.8. If V is a vector space over the field Υ , the **general linear group** of V , written $GL(V)$ is the group of all automorphisms of V , i.e. the set of all bijective linear transformations $V \rightarrow V$, together with functional composition as group operation.

3.3.2 Lie Groups

Groups provide the modern language to characterize symmetries: the invariance of a system under reversible changes. For example, permutations preserve the number and labels of elements while they allow to describe all different orderings and how to transition between them. Lie groups, which have the additional structure of a smooth manifold, enable to describe continuous symmetries, where the group elements are indexed using a continuous variable. A Lie group is a group with a smooth structure provided by an underlying manifold.

Definition 3.3.9. A Lie group G is a manifold that has a group structure so that multiplication and inverse operations in the group structure are smooth as an operation on a manifold.

Lie groups form a very nice class of topological groups, and the compact Lie groups have a particularly well-developed theory. Basic examples of compact Lie groups—that we study in our work—are:

1. circle group T and the torus groups T^n ,
2. the special orthogonal group $SO(n)$

Definition 3.3.10. *The **circle group**, denoted by T , is the multiplicative group of all complex numbers with absolute value 1, i.e., the unit circle in the complex plane.*

Definition 3.3.11. *The **torus group**, denoted by T^n is an n -dimensional torus or an n -dimensional compact manifold.*

Example 3.3.1. *The 1-torus is just the circle: $T^1 = \mathcal{S}^1$.*

Definition 3.3.12. *The **orthogonal group** of dimension n , is the group of distance-preserving transformations of a Euclidean space of dimension n that preserve a fixed point, where the group operation is given by composing transformations. Equivalently, it is a group of $n \times n$ orthogonal matrices.*

Definition 3.3.13. *The **special orthogonal group** $SO(n)$ is an important subgroup of the orthogonal group with all orthogonal matrices having determinant equal to 1 or -1. The $SO(n)$ is also called a rotation group and has been studied extensively in lower dimensions as $SO(2)$ and $SO(3)$ in two and three dimensions.*

3.4 Representation Theory

Group theory has been factored into two parts; First, there is the study of the structure of abstract groups. Second, is the companion question: given a group H , how can we describe all the ways in which H may be embedded in (or mapped to) a linear group $GL(V)$. This is the subject matter of representation theory. In this section, we briefly review some aspects of the Representation theory.

In this work, we use the representation theory language to express the changes in the Fourier coefficients of a signal when it is transformed by the elements of a group (*e.g.*, rotation). We review some basic concepts from this theory, before deriving the lemmas that lead to a closed-form expression that we would use to derive the variance expression for Monte Carlo integration in Chapter 5.

Definition 3.4.1. *Given a complex inner-product space $(V, \langle \cdot, \cdot \rangle)$ and compact group H , we say that the map $\rho : H \rightarrow U(V)$ (with $U(V)$ the group of unitary transformations on V) is a representation if:*

$$\rho(h_0 \cdot h) = \rho(h_0) \circ \rho(h), \quad \forall h_0, h \in H. \quad (3.1)$$

Notation Given a representation $\rho : H \rightarrow U(V)$, $h \in H$, and $v \in V$, we will write:

$$h(v) \equiv \rho(h)(v). \quad (3.2)$$

Definition 3.4.2. *Given a representation $\rho : H \rightarrow U(V)$ and a subspace $W \subset V$, we say that W is a sub-representation if $h(w) \in W$ for all $w \in W$ and all $h \in H$.*

Definition 3.4.3. Given a representation $\rho : H \rightarrow U(V)$ we say that V is an irreducible representation if the only sub-representations are the trivial ones, $W = \{0\}$ and $W = V$.

Lemma 1 (Schur's Lemma). If V and W are irreducible representations of H and $\rho : V \rightarrow W$ is a H -module homomorphism, then

1. Either ρ is an isomorphism, or $\rho = 0$.
2. If $V = W$, then $\rho = \lambda \cdot I$ for some $\lambda \in \mathbb{C}$, where I is the identity.

Now, we give the closed-form expressions for the average:

Proposition 2. Given an irreducible representation $\rho : h \rightarrow U(V)$, for any $x, y, v, w \in V$, we have:

$$\int_H \langle x, h(y) \rangle \overline{\langle v, h(w) \rangle} dh = \frac{\mu(H)}{\dim(V)} \langle x, v \rangle \overline{\langle y, w \rangle}. \quad (3.3)$$

where \langle , \rangle represents an inner product operator.

Proof Fixing $y, w \in V$, let $B_{y,w} : V \times V \rightarrow \mathbb{C}$ be the map:

$$B_{y,w}(x, v) = \int_{h \in H} \langle x, h(y) \rangle \overline{\langle v, h(w) \rangle} dh. \quad (3.4)$$

It is not hard to show that this map is linear in the first argument, conjugate-linear in the second, and H -equivariant. (That is, for any $h_0 \in H$ we have $B_{y,w}(h_0(x), h_0(y)) = B_{y,w}(x, y)$). Thus, by Schur's Lemma [88, 33], $B_{y,w}$ is a scalar multiple of the inner-product on V :

$$B_{y,w}(x, v) = \lambda_{y,w} \langle x, v \rangle. \quad (3.5)$$

Noting that this satisfies $B_{y,w}(x, v) = \overline{B_{x,v}(y, w)}$, it follows that:

$$B_{y,w}(x, v) = \lambda \langle x, y \rangle \overline{\langle v, w \rangle}, \quad (3.6)$$

for some constant $\lambda \in \mathbb{C}$ that is independent of v and w . Finally, letting $\{v_1, \dots, v_n\}$ be an orthonormal basis for V , we can express the integral of the square norm of the trace of $\rho(h)$ as:

$$\begin{aligned} \int_H \|\text{Tr}(\rho(h))\|^2 dh &= \int_H \left\| \sum_{i=1}^n \langle h(v_i), v_i \rangle \right\|^2 dh \\ &= \sum_{i,j=1}^n B_{v_j, v_i}(v_j, v_i) \\ &= \text{dim}(V) \cdot \lambda. \end{aligned}$$

Since the trace is the character of the representation, it follows by the orthogonality of characters [88, 33] that $\int_H \|\text{Tr}(\rho(h))\|^2 dh = \mu(H)$, which gives:

$$\lambda = \frac{\mu(H)}{\text{dim}(V)}. \quad (3.7)$$

Thus, we get:

$$\int_H \langle x, h(y) \rangle \overline{\langle v, h(w) \rangle} dh = \frac{\mu(H)}{\text{dim}(V)} \langle x, v \rangle \overline{\langle y, w \rangle}. \quad (3.8)$$

□

Proposition 3. *Leveraging Schur's Lemma in a similar manner, it follows that if $\rho_1 : h \rightarrow U(V_1)$ and $\rho_2 : h \rightarrow U(V_2)$ are two irreducible representations that are not isomorphic, then for any $v_1, w_1 \in V_1$ and $v_2, w_2 \in V_2$:*

$$\int_H \langle v_1, h(w_1) \rangle \overline{\langle v_2, h(w_2) \rangle} dh = 0. \quad (3.9)$$

We use Propositions 2 and 3 in Chapter 5 to derive the variance closed-form expression in the spherical domain.




Figure 3.1: Spherical harmonics basis functions (real part).

3.5 Spherical Harmonics

On the sphere, the most commonly used tool for spectral analysis is the spherical harmonics (SH) [38, 104], which is the Fourier analog on the sphere, and is given by:

$$Y_l^m(\theta, \phi) := \sqrt{(2 - \delta_{0m}) \frac{(2l+1)(l-m)!}{\mu(\mathcal{S}^2)(l+m)!}} P_l^m(\cos \theta) \exp(im\phi). \quad (3.10)$$

Here, δ_{ij} is the Kronecker delta function, $\mu(\mathcal{S}^2) = 4\pi$, is the Lebesgue measure of a unit sphere, $Y_l^m(\theta, \phi)$ is the spherical harmonic basis function of degree l and order m and $P_l^m(x)$ denotes the *associated Legendre Polynomials*, for $x \in [-1, 1]$. SH (Y_l^m) are orthonormal basis functions, such that, any integrable function G on \mathcal{S}^2 can be decomposed into SH components. as:

$$G(x) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \mathcal{S}_G(l, m) Y_l^m(x), \quad (3.11)$$

where $\mathcal{S}_G(l, m)$ are the (complex) spectral coefficients of $G(x)$. It can be easily shown that:

$$\int_{\mathcal{S}^2} \|G(x)\|^2 d\omega = \sum_{l=0}^{\infty} \sum_{m=-l}^l \|\mathcal{S}_G(l, m)\|^2. \quad (3.12)$$

which is the Parseval's theorem on the sphere. Analogously, the inner product between any two arbitrary functions, $G(x)$ and $F(x)$ defined over a unit sphere, is related to their corresponding spectral coefficients by:

$$\int_{\mathcal{S}^2} G(x) \bar{F}(x) dx = \sum_{l=0}^{\infty} \sum_{m=-l}^l \mathcal{S}_G(l, m) \cdot \overline{\mathcal{S}_F(l, m)}, \quad (3.13)$$

where $\mathcal{S}_G(l, m) = \langle G, Y_l^m \rangle$ is the (l, m) -th spherical harmonic coefficients of $G(x)$. The angular mean power spectrum at a frequency band l is defined as the average energy distributed over different m for a given l , as follows:

$$\check{\mathcal{P}}_G(l) := \frac{1}{2l+1} \sum_{m=-l}^l \|\mathcal{S}_G(l, m)\|^2. \quad (3.14)$$

The *mean* angular power spectrum $\check{\mathcal{P}}_G(l)$ is invariant under a rotation of the coordinate system ([48, 64]), as they contain a sum over all orders m . A spherical harmonic power spectrum can be defined with or without the averaging factor. In our formulation, we prefer to work with the average power per degree, or power spectral density, $\check{\mathcal{P}}_G(l)$, as this ensures that the spectral coefficients of a spherical Dirac delta function are constant and independent of degree l ([45]).

3.6 Hemispherical analysis using SH

In the hemispherical domain, we can apply the mathematical procedure developed on the sphere, using spherical harmonics. This is possible due to the fact that the sphere is a Riemannian double-cover of *projective 2-space* and all local geometric calculations in the spherical domain are applicable on hemispherical functions if

we associate the hemisphere to the *projective 2-space*. Here we use the fact that, any function on *projective 2-space* can be extended to an even function on the sphere.

Projective 2-space (\mathbb{P}^2) is defined to be the set of points on the sphere, modded out by the relation that for all $x \in \mathbb{P}^2$, x and $-x$ belongs to the same equivalence class. Then given a function on the projective 2-space, $G \in L^2(\mathbb{P}^2)$, it can be turned into a function on the sphere by setting the value at the points x and $-x$ to be the value of G on the equivalence class $\{x, -x\}$. In particular, this means that the space of functions on *projective 2-space* is the same as the space of *even* functions on the sphere. Thus, rotation can be defined on $L^2(\mathbb{P}^2)$ by treating the function as an even function on the sphere, rotating the even function (which remains even after rotation) and then considering the corresponding function on *projective 2-space*. In that case, all the results derived in the spherical domain follow, including the *homogeneous property*. The space of even functions on the sphere is precisely the space of functions spanned by spherical harmonic basis functions of even degree, Y_m^l with l even. The only assumption here is that we restrict our analysis to functions on the hemisphere with $G(x) = G(-x)$ for all points on the equator.

Chapter 4

Spectral Analysis of Sampling Patterns

Computer Graphics, with pixels being spatially discrete, inherently deals with sampling. Monte Carlo integration schemes are normally employed to evaluate the color of each pixel. The error in Monte Carlo integration is directly associated to the sampling patterns used to sample the integrand involved. In image synthesis, regular sampling patterns can cause structural artifacts, whereas, purely random samples can produce very high noisy appearance in the images. An illustration is shown in Fig. 4.1 for euclidean-space samplers. In Fig. 4.1, we render a simple function: $\sin(x^2 + y^2)$, which is also termed as the *zone-plate*. Reference image (Fig. 4.1(a)) is computed using nine samples per pixel followed by reconstruction filtering by Mitchell-Netravalli filter [70]. For Fig. 4.1(b) and 4.1(c) we use one sample per pixel with box filtering. Regular sampling (Fig. 4.1(b)) shows number of structural artifacts, whereas jittered sampling (Fig. 4.1(c)) turns these artifacts into *noise*. This can be explained from the power spectra of regular and jittered sampling.




Figure 4.1: Zone plate test to show the aliasing (structural-artifacts) due to the Fourier peaks in the regular sampling patterns (b), which turns into noise on simple jittering (c)

The power spectrum of regular sampling pattern contains high energy peaks (Fig. 4.2(b)), whereas in case of jittered sampling pattern the power spectrum (Fig. 4.2(a)) has no visible high energy peaks. This implies that sampling power spectra with high energy Fourier peaks can cause structural aliasing. This observation emphasizes a careful study of the spectral characteristics of various sampling patterns. Various tools have been designed to understand the inherent properties of different sampling patterns. For example, statistics such as Ripley's K and L statistics [82] are commonly used to model point distributions, while tools such as the variogram and autocorrelation are used for assessing distributions. In computer graphics, methods such as frequency domain analysis [25, 16, 69] are extensively used to better understand the characteristics of various sampling patterns.

In this chapter, we first look (Section 4.1) at how spectral (frequency) analysis is performed in the Euclidean domain. Later, we perform similar analysis in the spherical domain (Section 4.2) using spherical harmonics. There exist another

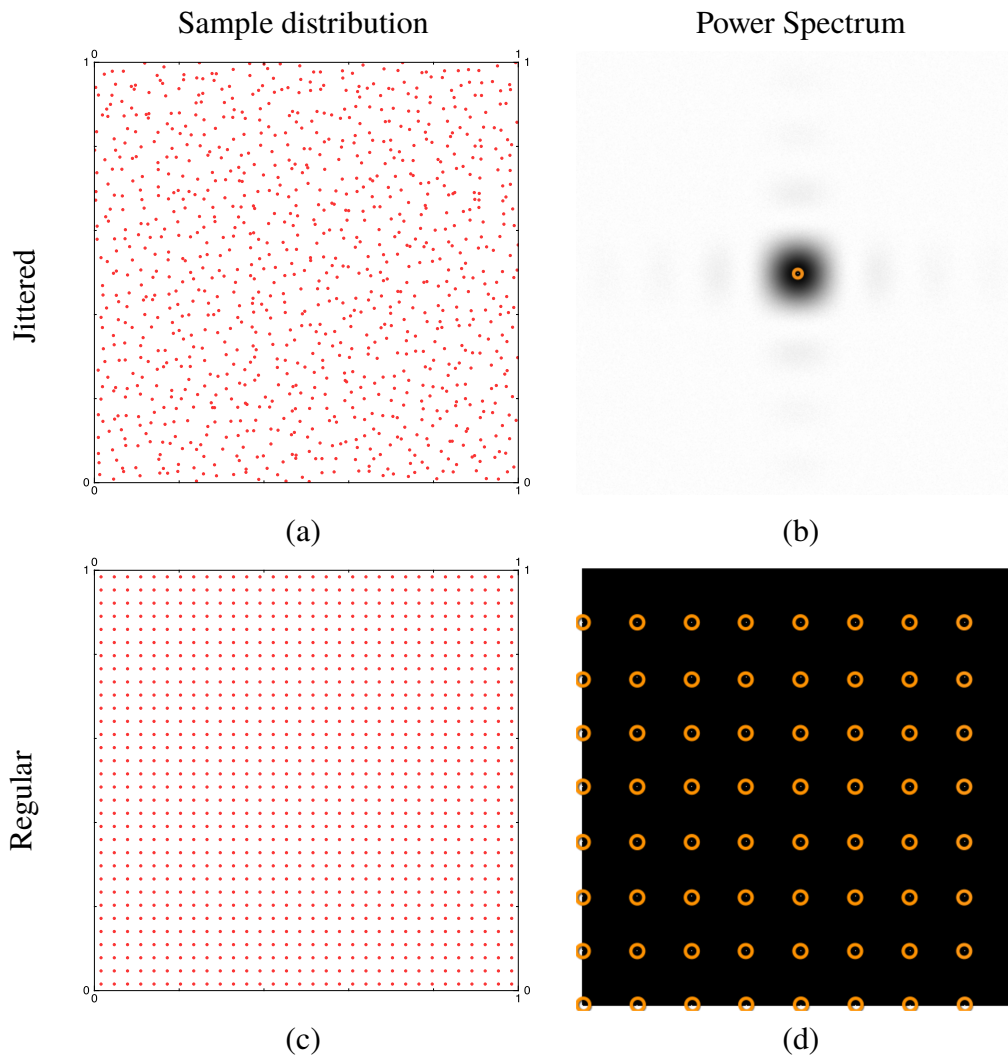


Figure 4.2: Illustration of jittered (a) and regular (c) sampling patterns with their corresponding power spectra in (b) and (d), respectively. The high energy peaks of regular sampling patterns are surrounded by circular (orange) rings in (d) for illustration purposes. In a similar manner, the DC peak in (b) is surrounded by a circular (orange) ring.

tool developed by Wei and colleagues [102] that can be used to perform analysis of sampling patterns in the *differential domain*. The advantage of differential domain analysis over frequency analysis, is that, the differential domain analysis can be performed for uniform and adaptive sampling methods in both the Euclidean as well as non-euclidean domains. However, as we focus only on uniform (non-adaptive) sampling patterns in the Euclidean and spherical domains, we restrict ourselves to the Fourier and spherical harmonics tools, respectively, to analyze the characteristics of various sampling patterns.

Based on our spectral analysis of sampling patterns over the sphere and the hemisphere, we introduce the notion of *homogeneous* sample distributions in spherical domain. We show that over multiple realizations (point sets) of a sampling pattern, any non-homogeneous sampling pattern can be made homogeneous by performing uniform and random rotations on each of its realization. We also show that by performing uniform and random rotations on a point set, the energy in each frequency band, l , of an angular power spectrum, gets equally distributed over each of its components. Consequently, the structural artifacts turns into structural noise (Fig. 4.13). A theoretical justification of this operation is given in Section 3.4 using representation theory.

4.1 Spectral Analysis in Euclidean domain

Fourier tools are extensively used to study the frequency content of sampling patterns in the Euclidean domain. To perform frequency analysis of samples, we

first compute the Fourier transform of a sampling pattern S . Then we compute a $2D$ power spectrum of S which represents the amplitude squared values of the Fourier transform coefficients of a sampling distribution, normalized by the number of samples N . An example is shown in Fig. 4.3. In mathematical notation, power spectrum can be written as:

$$\left\| \frac{1}{N} \sum_{i=0}^{N-1} \mathcal{F}_S(\delta(x - x_i)) \right\|^2, \quad (4.1)$$

where, \mathcal{F}_S denotes the Fourier transform on S .

Sample generation In the Euclidean domain, we generate samples on a square domain $[0, 1] \times [0, 1]$ to study their spectral characteristics. Here, we study only four state-of-the-art sampling patterns namely: White noise, jittered, Poisson Disk (or dart throwing) and regular. White noise is generated using a purely random generators. The resulting distribution can have many voids and clusters. Jittered samples are generated by first subdividing the domain into strata, followed by generating one sample in each of the stratum. Poisson Disk sampling is generated by following the dart throwing algorithm. The resulting distribution can be seen as close to *blue noise* distribution. And finally, regular sampling can be obtained by simply placing samples at regular intervals along the axes of the square domain.

Wrapped vs unwrapped distances While generating sampling patterns with the underlying characteristics, it is quintessential to make sure that the characteristics of the sampling pattern are also obeyed across the boundaries. For example, to generate Poisson Disk samples a minimum distance criteria is enforced between

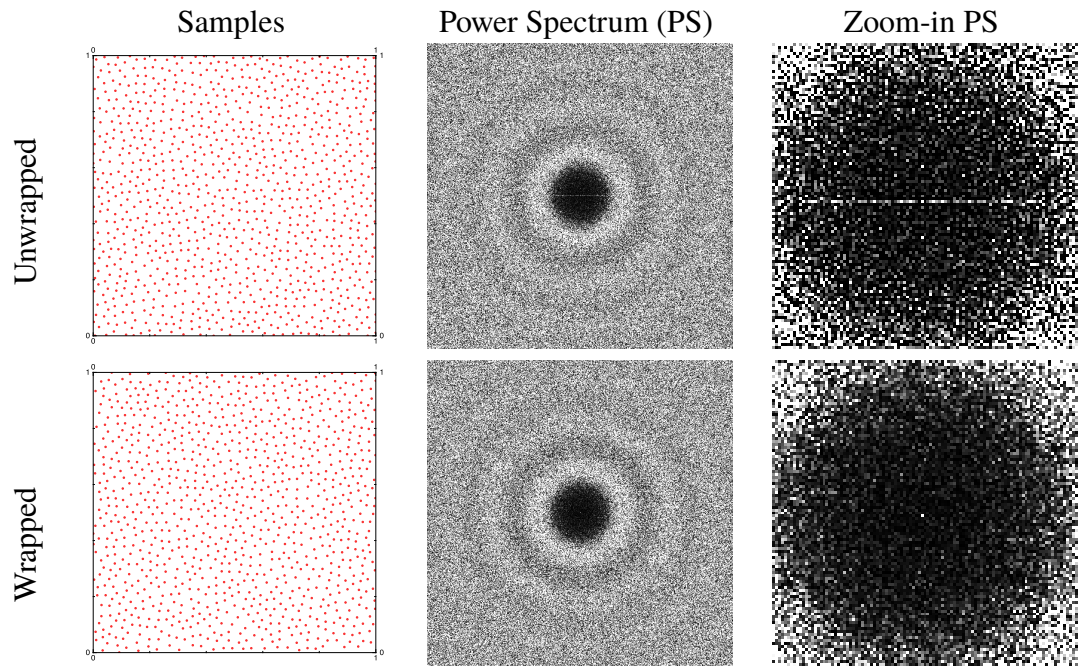


Figure 4.3: Illustration of **Top row:** non-Toroidal (unwrapped) vs **Bottom Row:** Toroidal (wrapped) Poisson Disk samples with corresponding power spectra and the zoom-in of the low frequency region around the DC peak of the respective power spectra.

each pair of neighboring samples. However, if this minimum distance criteria is not respected across the boundaries of the domain, the corresponding sampling power spectrum can get some harmful high energy peaks in the low frequency region around the DC peak. We can easily get rid of these high energy peaks by making sure that the Poisson disk samples obey the minimum distance criteria across the boundaries by using wrapped around distances. In literature, these samples corresponds to the toroidal domain (\mathcal{T}^2) sampling patterns.

As illustrated in Fig. 4.3, if we do not enforce minimum distance criteria while generating Poisson disk distribution, we obtain Poisson Disk distribution with

samples near the boundaries of the domain. As shown in Fig. 4.3(a), if we look with an hawk-eye, there are samples close to the left and right borders of the domain. These samples appear to be far off from each other, as they are on the opposite sides of the domain. However, if we fold this square sampling domain to make a torus shaped surface, we observe that these border samples are too close to each other and therefore, does not obey the minimum distance criteria any more. Similar observation can be made for Top and Bottom borders as well as the samples at the four corners of this square domain. The downside of having samples too close to the boundaries on the opposite sides is that, if we try to use this sample set as a *tile* and start tiling a plane with these *tiles*, the samples at the boundaries of two tiles would be too too close to each other and the Poisson Disk property of minimum distance between the samples would not be valid anymore. As a result, not all the samples we get on a tiling plane using these tiles, would be respecting the minimum distance criteria.

To overcome this issue, samples are generated with a wrapped around distance, *i.e.* in the toroidal domain. The idea is to make sure that samples at the boundaries of the domain are placed such that the minimum distance criteria is also preserved across the boundaries of the domain. As shown in Fig. 4.3(c), by placing samples in \mathcal{T}^2 , samples on the Right border are far from the boundary with respect to the samples at the exact opposite Left border of the domain. Similarly, samples at the four corners are placed such that no sample in a corner has an overlapping sample in its assigned radius across the corners (boundaries).

The result of placing samples in the toroidal and non-toroidal domains are




Figure 4.4: Illustration of multiple realizations of Poisson Disk samples (a) with corresponding Fourier power spectra with respect to each realization.

clearly visible in the Fourier domain. As we can see in Fig. 4.3(b) and 4.3(d), the effects of placing samples in a non-toroidal manner shows up in the power spectrum as a set of large magnitude peaks (Fig. 4.3(b)) visible in the low frequency zone (black region around the DC peak) as a horizontal line. To see the DC peak, zoom-in towards the center of the square power spectra images (only possible on the electronic version). However, when the samples are placed in the toroidal manner, we see no such peaks in the low frequency region of the Fourier spectra (Fig. 4.3(d)) of the corresponding sampling pattern. For illustration purposes, we show 1024 Poisson Disk samples in Fig. 4.3(a) and 4.3(b) but the corresponding power spectrum is generated using 4096 samples in both Fig. 4.3(b) and 4.3(d).

Mean Power Spectrum The power spectra shown in Fig. 4.3(b) and 4.3(d) appears pretty noisy. To obtain a neat power spectrum of a sampling pattern we

generate multiple sample sets (realizations), followed by taking average of the amplitude squared values of the Fourier transform of each realization, as shown in Fig. 4.4. In Fig. 4.5, we show square power spectrum of various sampling patterns with the corresponding heatmap to emphasize values in each of the low-, mid- and high-frequency regions. The zero frequency is the DC peak which is located at the center (of the image) of all the power spectra. The power spectrum of the Poisson Disk sampling pattern (Fig. 4.5(h)) appears to have zero value around the DC peak in the low frequency region. However, the corresponding heatmap of Poisson Disk (shown in blue, Fig. 4.5(i)) reveals that the power value is not zero or tending towards zero value, near the DC peak in the low frequency region. Consequently, we can observe a non-zero values with almost same power over the whole low frequency region around the DC peak. For jittered sampling the low frequency region around the DC peak appears to be zero for frequencies ($\omega \rightarrow 0$) going to zero (Fig. 4.5(f)). For regular sampling patterns, we see a set of high energy peaks in the power spectrum. The peaks are not visible on a printed paper however, a zoom-in look (electronic version) would show the peaks.

We have already illustrated (Fig. 4.1) the effects of these high energy peaks in the regular sampling pattern power spectrum which appears as artifacts where we render a zone plate function: $\sin(x^2 + y^2)$. Since the jittered sampler power spectrum distributes the Fourier energy of samples all over the spectrum, the resulting rendered image using jittered samples shows noise, which is less objectionable to human eye.




Figure 4.5: Illustration of sampling patterns in the **Left** column, with their corresponding power spectra in the **Center** column. In the **Right** column we show a heatmap of each power spectrum to emphasize their characteristics in the low frequency region around the DC peak, surrounded by a circle (orange) in (b), (e) and (h), located at the center of the power spectra images, in the **Center** and **Right** columns). Each peak in the regular pattern power spectrum (k) is also surrounded by circles (in orange) for illustration purposes.




Figure 4.6: Illustration of the radial mean power spectra of whitenoise (b) and jittered sampling (d) from the corresponding power spectra in (a) and (c), respectively. Radial mean power is computed by averaging the power spectra values along the circular rings, in blue, for a given frequency ρ .

Radial Mean Power Spectrum We can also study the characteristics of a sampling pattern from its radial mean of the square power spectrum. To compute the radial mean of the power spectrum, we first consider the DC peak as the center of the *circle* with radius ρ frequency, such that all coefficients surrounding the DC peak at frequency ρ are averaged to compute the radial mean at frequency ρ .

Intuitively, the procedure can be understood as taking average of all the Fourier coefficients that are at a same distance from the DC peak on the square power

spectrum. An illustration is shown in Fig. 4.6. For whitenoise, as the square power spectrum is flat, the energy is spread all over the frequencies. Consequently, the corresponding radial mean of the White noise is also flat (Fig. 4.6(b)). For jittered sampling, we observe a region with power values close to zero near DC peak (at the center of the image) and as we go towards mid- to high-frequencies the power values increases (Fig. 4.6(d)). After a characteristic frequency, the radial mean spectrum of a jittered sampling pattern appears to be flat, like whitenoise.

4.1.1 Low Discrepancy Samplers

We also study the frequency content of some of the Quasi-Monte Carlo based sampling patterns for which *discrepancy* parameter is used as the quality criteria (please refer to Section 2.3 for more details). The idea behind low discrepancy samplers is to generate samples such that the proportion of points in the sequence falling into an arbitrary set B is close to the measure of B . That is, each sample shall have approximately the same sub-area assigned to it in the sampling domain \mathbb{D} . Even though, this condition do impose uniformity, it cannot avoid regular structures in the sampling pattern. As a result a lot of high frequency peaks can be seen in the corresponding sampling power spectra. To better understand this behaviour, we study some well known low discrepancy samplers, namely: Halton, Hammerslay and Fibonacci lattice point sets.

In Fig. 4.7, we show the Halton, Hammerslay and Fibonacci lattice point sets. In Halton's point set, worm-like structures are very visible although, no regular structure seems to be present. As a result, the corresponding Fourier power




Figure 4.7: Illustration of sampling patterns in the **Left** column, with their corresponding power spectra in the **Center** column. In the **Right** column we show a heatmap of each power spectrum to emphasize their characteristics in the low frequency region around the DC peak which is located at the center of the power spectra images, in the **Center** and **Right** columns).




Figure 4.8: Illustration of the effects of high energy Fourier peaks in rendering a simple *zone plate* function.

spectrum energy of Halton point set is more spread out. However, some noisy regions are captured by the heatmap (Fig. 4.7(c)) of the Halton power spectrum. In case of Hammerslay (Fig. 4.7(d)), the spatial distribution of the point set is looking uniform. However, there are some unavoidable close points in the distribution. The overall structure when analyzed with Fourier spectrum has shown a number of high energy Fourier peaks, making an interesting pattern (Fig. 4.7(e,f)). We also study the Fibonacci lattice point set (Fig. 4.7(g)), which seems very similar to the rotated version of a regular sampling pattern. The corresponding power spectrum has a sea of high energy peaks (Fig. 4.7(h,i)).

The effects of these high energy peaks is shown in Fig. 4.8 where we render a zone plate function: $\sin(x^2 + y^2)$. Fibonacci lattice Fig. 4.8(c) shows structural artifacts, with less noise. These structural artifacts can be explained from its sampling power spectrum Fig. 4.7(h) which is full of peaks, whereas less noise in Fig. 4.8(c) is due to the regular structure of Fibonacci lattice point set. Halton appears more noisy and shows less artifacts Fig. 4.8(a) than Hammerslay and

Fibonacci lattice. This is because the power spectrum of Halton sequence distributes most of its Fourier energy all over the frequencies. The reference zone plate function is shown in Fig. 4.1(a). A detailed analysis of various samplers in the Euclidean space is performed by Pilleboue [78] in his thesis work. In this dissertation, we keep our focus in the (hemi-)spherical domain.

4.2 Spectral Analysis in Spherical domain

In the spherical domain, we perform a similar frequency analysis of different sampling patterns using *spherical harmonics* [38]. We follow our discussion from Section 3.5, where we introduce spherical harmonics. As we have shown in Fig. 3.1, the frequencies in spherical harmonic basis functions (Y_m^l) are distributed in an angular manner, with $|m| \leq l, \forall l \in \mathbb{Z} \cup 0$. Therefore, it is not possible to directly generate a *square* power spectrum similar to the euclidean case. However, for illustration purposes, later in this section, we show a *square-shaped* angular power spectra of some of the sampling patterns. To generate the *square-shaped* angular power spectra, we first compute the $\|\mathcal{S}_G(l, m)\|^2$ values for a given range of l with $|m| \leq l$, where $\mathcal{S}_G(l, m)$ is the (l, m) -th spherical harmonic coefficient of a function $G(x)$. This results in an angular power spectrum as shown in Fig. 4.9, with the DC peak ($l = 0, m = 0$) at the center of the image, frequency bands l increases down the vertical axis and the corresponding m components are spreading along the horizontal axis.

We then generate the *square-shaped* angular power spectrum (Fig. 4.9(b))




Figure 4.9: Illustration of *angular* power spectrum of a white noise distribution in (a) with the same power spectrum shown in a *square-shape* (b).

by repetitively using the angular power spectrum (Fig. 4.9(a)) in the missing quadrants. Note that, this *square-shaped* angular power spectrum is generated just for illustration purposes. The corresponding *mean* angular power spectrum can be obtained by simply averaging the $\|\mathcal{S}_G(l, m)\|^2$ values for all m in a given frequency band l , as follows:

$$\check{\mathcal{P}}_G(l) := \frac{1}{2l+1} \sum_{m=-l}^l \|\mathcal{S}_G(l, m)\|^2 . \quad (4.2)$$

The resulting $\check{\mathcal{P}}_G(l)$ is called the *mean* angular power spectrum which can be considered equivalent to the *radial* power spectrum we studied in the Euclidean case (Section 4.1).

4.2.1 Deterministic sampling on the sphere

All the sampling patterns we are studying on the sphere or hemisphere, in this dissertation, are computed directly on the sphere. No warping or mapping of any kind is used to compute these samples on the sphere. We give more details on the implementation of our spherical samplers in Chapter 7. To generate regular and jittered sampling patterns on the sphere we use the *Healpix* data structure [37]. To analyze a QMC (deterministic) sampler on the sphere and hemisphere, we choose (hemi-)spherical Fibonacci lattice point sets which has been recently introduced to the computer graphics community by Marques and colleagues [65]. This is the first QMC sampler that can be directly generated on the sphere or hemisphere.

As shown in Fig. 4.10, both Healpix regular and spherical Fibonacci sample distributions look very similar on the sphere. Therefore, we compared the spectral properties of spherical Fibonacci point sets with the Healpix regular point sets. Their *mean* angular power spectrum is showing a zero-value low -frequency region followed by high energy peaks. The corresponding *square-shaped* angular power spectrum is shown in Fig. 4.11(a,b,c) for Healpix regular, and in Fig. 4.12(a,b,c) for spherical fibonacci lattice point sets with their respective heatmap to visualize the high energy region.

Random rotations We perform a simple test on both Healpix regular and spherical Fibonacci lattice point sets. We perform uniform and random rotation on each of the point sets and average the respective (*square-shaped*) angular power spectrum over multiple rotations. We observed that with increase in the number




Figure 4.10: Illustration of healpix regular (a) and spherical Fibonacci lattice (b) point sets with respective mean angular power spectra.

of rotations, the high energy peaks in both the Healpix regular and spherical Fibonacci lattice, are getting distributed over the frequency band l . Each frequency band l of the spherical harmonics has $2l + 1$ rotational components, which are indexed with m . Therefore, in each frequency band l , the energy gets distributed over all $(2l + 1)$ rotational components. Consequently, the corresponding power of all the $(2l + 1)$ components in a frequency band l , becomes equal.

We experimentally observe this behaviour in Fig. 4.11 and 4.12 for Healpix regular and spherical Fibonacci lattice respectively. We perform a number of rotations of a point set and average it over all the rotations. The effects are shown in the *square-shaped* angular power spectra, over multiple rotations. The corresponding heatmaps shows that with the increase in the number of rotations, the high Energy peaks in the angular power spectrum starts getting uniformly distributed over different m , for a given l . However, the *mean* angular power spectrum do not change. Note that, the energy within the frequency band l remains the same, it only gets equally distributed within the $(2l + 1)$ components




Figure 4.11: Illustration of *square-shaped* angular power spectrum of Healpix regular (**left**) with heatmaps of power (**center**) and square-root of power values (**right**).




Figure 4.12: Illustration of *square-shaped* angular power spectrum of spherical Fibonacci (**left**) with heatmaps of power (**center**) and square-root of power values (**right**).

corresponding to that l .

To illustrate this effect on a real scene, we use a simple lighting setup, as shown in Fig. 4.13, where we compute direct lighting scene by first (Fig. 4.13(b)) using same regular hemispherical sampling pattern for secondary rays all over the scene. The resultant image is full of structural artifacts due to high correlation between different pixels. A simple perturbation of samples at each hitpoint, of these regular hemispherical patterns turns these artifacts into noise (Fig. 4.13(c)).

In Chapter 5, we introduce a theoretical framework that exploits the consequences of uniformly and randomly rotating the samples, for the Monte Carlo integration of spherical functions. We use the *representation theory* to formulate this operation of random rotation as a $SO(3)$ group operator (more details are given in Section 3.4).

4.2.2 Non-deterministic sampling on the sphere

We study four different irregular samplers on the sphere namely: white noise, jittered, Poisson Disk and CCVT [6]. All samplers are generated directly on the sphere with no warping at all. The implementation details of these samplers is given in Chapter 7. All the *mean* angular power spectra are plotted along *Power* vs. ω , where $\omega = \alpha\sqrt{N}$. Here, N is the number of samples and α is used to quantify the range of energy-free frequency with respect to the mean frequency.

In Fig. 4.14, we show the *mean* angular power spectrum of these sampling patterns on the sphere. The corresponding sample distributions are shown as insets. As we can see, the sampling patterns behaves similar to the Euclidean domain.



(a) Reference Image



(b) No Rotation



(c) With Rotations

Figure 4.13: Illustrating direct lighting evaluation with (b) or without (c) random perturbation of secondary rays. Image courtesy: Dr. Jean-Claude Iehl.

The computation of spherical harmonics is very time consuming. Therefore, we generate the radial angular power spectrum of spherical samples upto a restricted frequency. However, for hemispherical samples, in the next section, we are able to go to a reasonably high frequencies to compute the corresponding radial angular power spectrum. Note that, since the spherical domain is *closed*, it is naturally “toroidal” in nature.

4.3 Spectral Analysis in Hemispherical domain

To perform spectral analysis of hemispherical sampling patterns, we again used spherical harmonics. However, spherical harmonics cannot be directly employed on a hemispherical domain. This is because the spherical harmonics are no longer orthogonal on the hemisphere. Since the hemispherical domain is not *closed*, if we perform spectral analysis of hemispherical sampling patterns using spherical harmonics, we obtain a set of high energy peaks within the low frequency region. This is illustrated in Fig. 4.15 for white noise, jittered and hemispherical Fibonacci lattice point sets with the corresponding spatial distributions shown as an inset. We compute their *square-shaped* and *mean* angular power spectra using spherical harmonics. The high energy peaks can be seen in the low-frequency region near the $l = 0$ frequency.

The best way to compute spectral content of hemispherical sampling patterns using spherical harmonics is by first mapping the hemispherical samples to a projective 2-space such that the sampling pattern acts as an *even* function on




Figure 4.14: Comparison of the mean angular power spectra of various state of the art sampling patterns in the spherical (**left**) and hemispherical (**right**) domain.




Figure 4.15: Illustration of the high energy peaks in the low frequency region of *square-shaped* and *radial* angular power spectra of whitenoise, jittered and hemispherical fibonacci lattice point sets, respectively.

the sphere, more details are given in Section 3.6. This allows computation of angular power spectra of various hemispherical samplers directly using spherical harmonics.

Hemispherical sample generation We generate white noise and Fibonacci lattice points directly on the hemisphere whereas for jittered samples, we first generate jittered sampling pattern on a unit sphere and then consider only the upper hemisphere of that unit sphere. A similar procedure is also adopted for CCVT, Healpix regular and Poisson Disk sampling patterns. All hemispherical samplers are shown as insets in Fig. 4.14, with their corresponding *mean* angular power spectra.

Random rotations We perform a similar operation of uniform and random rotations on the hemispherical sampling patterns. To illustrate the effects of random rotation on the characteristic spectral properties of a sampling pattern, we use hemispherical Fibonacci lattice point set. As shown in Fig. 4.16, we observe a very similar behaviour as in the spherical domain (Fig. 4.12), on the hemisphere. Hemispherical Fibonacci point sets have a number of high energy peak in the spectrum. However, just after one rotation, the energy peaks within different $(2l + 1)$ rotational components of a frequency band l starts distributing their energy among all $(2l + 1)$ components of l . After a number of rotations, all the $(2l + 1)$ components appears to have equal energy in a given frequency band l .




Figure 4.16: Illustration of *square-shaped* angular power spectrum of hemispherical Fibonacci (**left**) with heatmaps of power (**center**) and square-root of power values (**right**).

4.4 Homogeneous sampling patterns

In both, the spherical and hemispherical domains, we observed that after performing a finite number of rotations, the spectral power of the spherical harmonic coefficients gets distributed within the $(2l + 1)$ components of a frequency band l . In global illumination evaluation, Cranley-Patterson [17] rotation is used to add perturbations by simply adding a random shift vector to each point modulo 1, over the Euclidean domain, which converts structural artifacts into noise. We instead, perform direct rotation of hemispherical samples on the (hemi-)sphere. The operation of rotation results in, what we call, *homogeneous* sample distributions.

First, we define the idea of *homogeneous* sample distributions in a general setting. A *homogeneous* sample distribution has statistical properties invariant to translation over the domain \mathbb{D} . The notion of homogeneous distribution of points is commonly used in physics [40] for compound gases. In the computer graphics community, this notion is similar to the “widesense stationary process” described by Dippé and Wold [25] which is discussed only for *Poisson sampling*. However, in our mathematical formulation (Chapter 5), we extend the domain of application of this *widesense stationary* notion by making any kind of sampling pattern translation invariant, that is, homogeneous. In MC integration, the consequence is that a homogeneous sampling pattern produces no bias, so that error only arises from the variance. For more details, please refer to Appendix B. Given this property, we can restrict our error analysis to variance only.

Many state-of-the-art sampling methods such as White noise, Poisson Disk,

and all optimization-based methods starting from a white noise distribution (*e.g.*, the methods of Schlömer et al. [85] and de Goes et al. [22]) are homogeneous. Surprisingly, jittered sampling and Latin hypercube sampling are non-homogeneous. This can be explained by the fact that both methods rely on a subdivision of the sampling domain that is fixed over all realizations. However, any sampling method can be transformed into a homogeneous sampling by uniformly and randomly translating each realization of the sampling pattern. Analogously, sampling patterns on the unit sphere (S^2) can be made homogeneous by uniformly and randomly rotating the distribution.

Spectral properties of homogeneous sampling As we observed in the case of spherical and hemispherical samples, uniform and random rotations of a point set can result in equi-distribution of energy within $(2l + 1)$ components of a frequency band l . We also observe that the the distribution of a spherical harmonic coefficient ($\mathcal{S}_S(l, m)$) of a sampling pattern, taken over multiple rotations of a sampler, results in a radial distribution of this coefficient values in the complex plane. An illustration is shown, in Fig. 4.17, for whitenoise, jittered, Healpix regular and spherical Fibonacci lattice point sets. This implies some interesting characteristics of the coefficients $\mathcal{S}_S(l, m)$ of the sampling pattern: First, the expected values of $\mathcal{S}_S(l, m)$, *i.e.* $\langle \mathcal{S}_S(l, m) \rangle$, is zero. Second, the phase of coefficients is uniformly distributed over $[0, 2\pi]$.

Since, the sample distributions are made homogeneous after performing uniform and random rotations of point sets, the above properties are associated to



Figure 4.17: Illustration of the distribution of a spherical harmonic coefficient ($l = 4, m = 3$) computed using 1024 samples over 8192 realizations in the complex plane.

the homogeneous sample distributions. In the Euclidean space, Subr and colleagues [92] have shown that an unbiased estimator can be obtained if the mean of the Fourier spectrum of the sampling pattern is zero everywhere, except at frequency zero (DC). By definition, homogeneous sampling patterns exhibit this property and therefore can be considered unbiased. As a result, homogeneous patterns manifest error only in terms of noise (variance).

4.5 Discussion

To summarize the contribution of this chapter we outline the conclusions drawn from our spectral analysis of various sampling patterns. We focus our spectral analysis mainly in the spherical and hemispherical domain. To our knowledge, this is the first time such a careful spectral analysis of various sampling patterns have been performed in the spherical and hemispherical domain. In image synthesis, spherical and hemispherical samples play a crucial role in the evaluation of global illumination integrals. In practice, for global illumination problems, samples are first generated on a square domain and then a low distortion mapping [90] is performed to project these square samples on to the sphere or a hemisphere. For our analysis, we generate samples directly on the sphere. To obtain samples on a hemisphere we consider an upper hemisphere of spherical samples.

We perform random rotations on the spherical and hemispherical samples that distributes the high energy peaks over the respective frequency bands l . This property can explain why the structural artifacts turn into noisy signals if we start

randomly rotating *regular* or *QMC* samplers while evaluating global illumination integrals.

We also lay down the basic platform to perform variance analysis for Monte Carlo integration. We develop the notion of homogeneous sampling patterns on the sphere which express error in integration only in terms of variance. Homogeneous sample distributions are the direct consequence of uniform and random rotations of spherical and hemispherical samples.

Chapter 5

Variance analysis of Monte Carlo Integration

In this chapter, we develop our mathematical framework that relates the variance in Monte Carlo (MC) integration directly with the power spectra of both the sampling pattern and the integrand involved. We use the notion of *homogeneous* sampling to propagate the error in MC integration into variance. This framework has been developed on the spherical domain which is easily extendable to the hemispherical domain.

5.1 Monte Carlo estimator

Monte Carlo integration is a numerical integration method to estimate the integral I of a function F . The MC estimator I_N averages over N stochastic samples taken over the integrand in a given sampling domain \mathbb{D} . We consider samples

$\{\mathbf{s}_1, \dots, \mathbf{s}_N\}$ as a set of equally weighted random variables. The resultant MC estimator \mathbf{I}_N of an integrand F is defined as:

$$\mathbf{I}_N := \frac{\mu(\mathbb{D})}{N} \sum_{k=1}^N F(\mathbf{s}_k). \quad (5.1)$$

A sampling pattern can be written as a random variable \mathbf{S} made up of N Dirac functions located at sample positions:

$$\mathbf{S}(x) = \sum_{k=1}^N \delta(x - \mathbf{s}_k), \quad (5.2)$$

which can be represented as an inner product [28] given by:

$$\mathbf{I}_N = \frac{\mu(\mathbb{D})}{N} \int_{\mathbb{D}} \mathbf{S}(x) F(x) dx. \quad (5.3)$$

Depending on the sampling strategy, different weights (w_i) can be assigned to each sample. On the spherical domain ($\mathbb{D} = \mathcal{S}^2$), MC estimator can be formulated as:

$$\mathbf{I}_N := \frac{\mu(\mathcal{S}^2)}{N} \int_{\mathcal{S}^2} S(x) F(x) dx, \quad (5.4)$$

where $\mu(\mathcal{S}^2) = 4\pi$, denotes the Lebesgue measure of a unit sphere with $x = (\theta, \phi)$ for $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$, and $dx = \sin \theta d\theta d\phi$. To relate variance of \mathbf{I}_N with the frequency content of F and \mathbf{S} , we first obtain the spectral form of the estimator \mathbf{I}_N using spherical harmonics. Before going further, we quickly discuss a brief background on spherical harmonics and the corresponding properties which we need to derive the variance relation.

5.2 Background on SH

In Section 3.5, we introduce spherical harmonics, which are the Fourier analog on the sphere. Since the Fourier transform cannot be directly applied on the sphere, spherical harmonics were designed to study the spectral properties of spherical signals directly on the sphere. For any integrable function, $G(x)$, it can be easily shown that:

$$\int_{S^2} \|G(x)\|^2 d\omega = \sum_{l=0}^{\infty} \sum_{m=-l}^l \|\mathcal{S}_G(l, m)\|^2. \quad (5.5)$$

where $\mathcal{S}_G(l, m)$ are the (complex) spectral coefficients of $G(x)$, which is the Parseval's theorem on the sphere. Analogously, the inner product between any two arbitrary functions, $G(x)$ and $F(x)$ defined over a unit sphere, is related to their corresponding spectral coefficients by:

$$\int_{S^2} G(x) \bar{F}(x) dx = \sum_{l=0}^{\infty} \sum_{m=-l}^l \mathcal{S}_G(l, m) \cdot \overline{\mathcal{S}_F(l, m)}, \quad (5.6)$$

where $\mathcal{S}_G(l, m) = \langle G, Y_l^m \rangle$ is the (l, m) -th spherical harmonic coefficients of $G(x)$. The angular mean power spectrum at a frequency band l is defined as the average energy distributed over different m for a given l , as follows:

$$\check{\mathcal{P}}_G(l) := \frac{1}{2l+1} \sum_{m=-l}^l \|\mathcal{S}_G(l, m)\|^2. \quad (5.7)$$

5.3 Variance in spectral form

With this background, we are now ready to represent the Monte Carlo estimator from Eq. (5.4) in its spectral form. By using Eq. (5.6) in Eq. (5.4) we obtain the

estimator in terms of the spherical harmonic coefficients:

$$\mathbf{I}_N = \frac{\mu(\mathcal{S}^2)}{N} \sum_{l=0}^{\infty} \sum_{m=-l}^l \mathcal{S}_{\mathbf{S}}(l, m) \cdot \overline{\mathcal{S}_F(l, m)}. \quad (5.8)$$

where $\mathcal{S}_{\mathbf{S}}(l, m)$ and $\mathcal{S}_F(l, m)$ are the complex harmonic coefficients of \mathbf{S} and F . Following the definition of Variance for a complex variable \mathbf{X} , $\text{Var}(\mathbf{X}) = \langle \|\mathbf{X}\|^2 \rangle - \|\langle \mathbf{X} \rangle\|^2$, and using Eq. (5.8), we obtain variance of the Monte Carlo estimator \mathbf{I}_N as:

$$\begin{aligned} \text{Var}(\mathbf{I}_N) &= \langle \|\mathbf{I}_N\|^2 \rangle - \|\langle \mathbf{I}_N \rangle\|^2 \\ &= \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l=0}^{\infty} \sum_{m=-l}^l \left\langle \|\mathcal{S}_{\mathbf{S}}(l, m) \cdot \overline{\mathcal{S}_F(l, m)}\|^2 \right\rangle - \left\| \left\langle \mathcal{S}_{\mathbf{S}}(l, m) \cdot \overline{\mathcal{S}_F(l, m)} \right\rangle \right\|^2. \end{aligned} \quad (5.9)$$

Here, we can easily show that the harmonic coefficients with $(l = 0, m = 0)$ do not contribute to the variance in Eq. (5.9). We provide a step-by-step derivation of the same in Appendix A. As a result, we can rewrite the Eq. (5.9) as follows:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l=1}^{\infty} \sum_{m=-l}^l \left\langle \|\mathcal{S}_{\mathbf{S}}(l, m) \cdot \overline{\mathcal{S}_F(l, m)}\|^2 \right\rangle - \left\| \left\langle \mathcal{S}_{\mathbf{S}}(l, m) \cdot \overline{\mathcal{S}_F(l, m)} \right\rangle \right\|^2, \quad (5.10)$$

where the degree l starts from 1 rather than zero. From hereafter, we consider only homogeneous sampling patterns in our mathematical derivation. This makes the estimator \mathbf{I}_N unbiased as a result of which the subtracting factor $\left\| \left\langle \mathcal{S}_{\mathbf{S}}(l, m) \cdot \overline{\mathcal{S}_F(l, m)} \right\rangle \right\|^2$ in Eq. (5.10) becomes zero at all non-zero frequencies ($l \neq 0$). We show a simple proof of this statement in Appendix A. Finally, the non-zero part $\left\| \left\langle \mathcal{S}_{\mathbf{S}}(0, 0) \cdot \overline{\mathcal{S}_F(0, 0)} \right\rangle \right\|^2$ gets cancelled out for DC frequency $l = 0$.

This results in further simplification of our variance relation given in Eq. (5.10):

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l=1}^{\infty} \sum_{m=-l}^l \left\langle \|\mathcal{S}_{\mathbf{S}}(l, m) \cdot \overline{\mathcal{S}_F(l, m)}\|^2 \right\rangle. \quad (5.11)$$

We rewrite above Eq. (5.11) by expanding the norm square term, which gives:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l, l'=1}^{\infty} \langle \mathcal{S}_{\mathbf{S}, F}(l, l') \rangle \quad (5.12)$$

In Eq. (5.12), $\mathcal{S}_{\mathbf{S}, F}(l, l')$ is given by:

$$\mathcal{S}_{\mathbf{S}, F}(l, l') = \sum_{\substack{m=-l \\ m'=-l'}}^{l, l'} \mathcal{S}_{\mathbf{S}}(l, m) \cdot \overline{\mathcal{S}_F(l, m)} \cdot \overline{\mathcal{S}_{\mathbf{S}}(l', m')} \cdot \mathcal{S}_F(l', m'). \quad (5.13)$$

Since the homogeneous sampling patterns, mentioned in Section ??, have statistical properties that are invariant to rotation on the sphere, it is equivalent to study the variance generated by the rotated version of each realization, with the average taken over all rotations. Formally, this is equivalent to taking average over the group of rotations ($SO(3)$), with $\tau(\mathbf{S})$ denotes the rotation of \mathbf{S} by an element $\tau \in \mathcal{S}^{SO(3)}$. Then, averaging equation Eq. (5.13) over all translations of \mathbf{S} , we get:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)^2}{N^2} \frac{1}{\mu(\mathcal{S}^{SO(3)})} \sum_{l, l'=0}^{\infty} \left\langle \int_{SO(3)} \mathcal{S}_{\tau(\mathbf{S}), F}(l, l') d\tau \right\rangle, \quad (5.14)$$

To simplify the integral part in Eq. (5.14) we employ the *representation theory* which has been briefly explained in Section 3.4. By using the fact that the span of the l -th frequency spherical harmonics, $V^l = \text{Span}\{Y_l^{-l}, \dots, Y_l^l\}$, is a $(2l + 1)$ -dimensional irreducible representation for the $SO(3)$ group and applying Propositions 2 and 3 from Chapter 3, we get:

$$\int_{SO(3)} \langle \tau(\mathbf{G}), Y_l^{m'} \rangle \cdot \langle Y_l^m, \tau(\mathbf{G}) \rangle d\tau = \delta_{ll'} \delta_{mm'} \mu(\mathcal{S}^{SO(3)}) \check{\mathcal{P}}_{\mathbf{G}}(l), \quad (5.15)$$

for any function $G(x) \in \mathcal{S}^2$. Here, $\check{\mathcal{P}}_G(l)$ is the angular mean power spectrum of the l -th spherical frequency of F as defined in Eq. (5.5). By plugging the angular mean power spectrum (5.5) in Eq. (5.15) and combining it with Eq. (5.14), we get:

$$\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l,l'=0}^{\infty} \left\langle \check{\mathcal{P}}_S(l) \left(\sum_{m=-l}^l \mathcal{S}_F(l, m) \cdot \overline{\mathcal{S}_F(l, m)} \right) \right\rangle \quad (5.16)$$

$$= \frac{\mu(\mathcal{S}^2)^2}{N^2} \sum_{l=0}^{\infty} \left\langle \check{\mathcal{P}}_S(l) \right\rangle \left(\sum_{m=-l}^l \|\mathcal{S}_F(l, m)\|^2 \right). \quad (5.17)$$

Here we normalize the angular mean power spectrum $\check{\mathcal{P}}_S(l)$ by $N/\mu(\mathcal{S}^2)$ to make sure that the angular mean power spectrum of any sampling pattern converges to value 1. This normalization ensures comparison within different sampling patterns with different N . At last, using the angular mean power spectrum definition from Eq. (5.5) with the normalization factor discussed in Eq. (5.16), we obtain:

$$\boxed{\text{Var}(\mathbf{I}_N) = \frac{\mu(\mathcal{S}^2)}{N} \sum_{l=0}^{\infty} (2l+1) \mathcal{P}_S(l) \check{\mathcal{P}}_F(l)} \quad (5.18)$$

which gives the final expression for the variance in terms of the angular mean power spectra of both the sampling pattern \mathbf{S} and the integrand F . Here, $\mathcal{P}_S(l) = \left\langle \check{\mathcal{P}}_S(l) \right\rangle$ which represents the averaged angular mean power spectrum—over multiple realizations (point sets)—of a sampling pattern.

Variance in spectral form on \mathcal{H}^2 We can use the same mathematical model derived in Eq. (5.18) in the hemispherical domain. As discussed in Section 3.6, we can map all hemispherical signals as even functions on the sphere. After that,

we can perform hemispherical signal analysis directly on the sphere while using only even l frequency bands of the spherical harmonics.

5.4 Discussion

In this chapter, we have derived a direct relation (Eq. (5.18)) between the variance in Monte Carlo integration with the power spectra of both the sampling pattern \mathbf{S} and integrand F involved. By looking at Eq. (5.18), we observe that variance is a sum of the product of the *mean angular* power of \mathbf{S} and F over all frequencies l . Consequently, to have a low variance in integration, at least one of $\check{\mathcal{P}}_{\mathbf{S}}(l)$ or $\check{\mathcal{P}}_F(l)$ should have minimal possible energy in the low frequency region of its corresponding power spectra.

To give an illustration, we plot both the $\check{\mathcal{P}}_{\mathbf{S}}(\cdot)$ and $\check{\mathcal{P}}_F(\cdot)$ on the same axis for a given number of samples N , as shown in Fig. 5.1(a). We assume the shape of the integrand's power spectrum, $\check{\mathcal{P}}_F(\cdot)$, (in orange) such that the energy of the spectrum is mostly concentrated in the low frequency zone and as we move towards mid to high frequencies the energy diminishes. This is a valid assumption since from a practical point of view, almost all display devices have a low pass filter that is applied to a signal (F) to avoid unusual artifacts, which means that we can display a signal only upto a particular bandwidth. That implies that most of the information, of the signals, which can be displayed on display devices is concentrated towards low- to mid-frequency regions.

Since the variance in integration (Eq. (5.18)) depends on the product of the

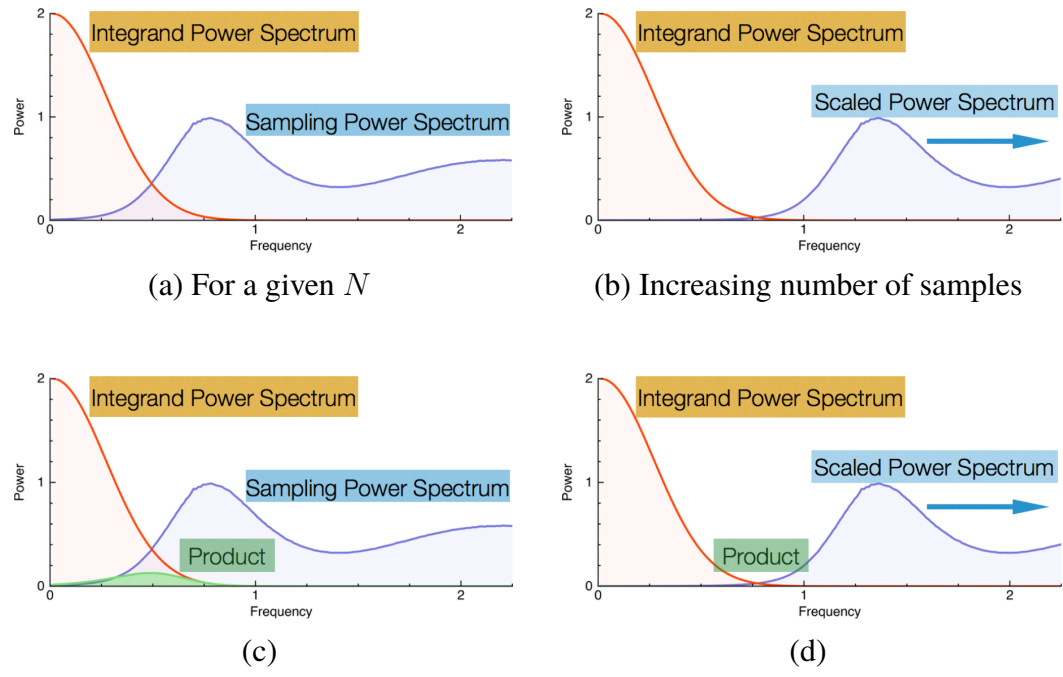


Figure 5.1: (a) Illustration of the power spectra of a sampler and an integrand on the same axis. (b) With increase in number of samples, sampler's power spectra scales along the frequency axis. In (c) and (d), we show the corresponding product of sampler and integrand power spectra.

power spectra of both the sampler and the integrand involved, the variance can be reduced by choosing a sampling pattern that has minimal possible energy in the low- to mid-frequency regions.

Now an obvious question that one can ask is, why the variance goes down as we increase the number of samples? This can be explained from the fact that as we increase the number of samples, the power spectrum of a sampler (in blue) scales along the frequency, as shown in Fig. 5.1(b). As a result, the low frequency part of a sampling power spectrum, where the energy is negligible, becomes more and

more dominating in the domain (low to mid frequency region) of the integrand. Consequently, the corresponding product of the sampler and integrand power spectra, in Fig. 5.1(c), diminishes with the increase in number of samples, as shown in Fig. 5.1(d).

With this illustration in Fig. 5.1, we have shown that as we increase the number of samples, the variance convergence rate—*i.e.*, the rate at which the variance diminishes with the increase in the number of samples—of a sampler depends on the low frequency region of the sampling power spectra. In Chapter 6, we focus our attention to analyze the low-frequency region of various sampling patterns. We classify the low-frequency region of various sampling patterns, and obtain bounds on the corresponding *angular mean* power spectra in the low-frequency region using simple shapes (profiles). These profiles are later used to derive the corresponding variance convergence rates in terms of the number of samples involved.

Chapter 6

Theoretical Convergence analysis

Using the framework derived in the spherical (Chapter 5) domain, we have shown that the variance in MC integration is related to the power spectra of the integrand and the sampling pattern. This implies that if we know the sampling and integrand power spectra then we can predict the variance and the corresponding variance convergence rate, as a function of the number of samples N , in MC integration. In practice however, the power spectrum of an integrand is unknown. Therefore, we restrict our analysis to a particular class of functions. We follow the work by Brandolini and colleagues [9], and restrict our analysis to integrable functions of the form $F(x)\chi_{\Omega}(x)$ with $F(x)$ smooth and Ω a bounded domain with smooth boundary (where, $\chi_{\Omega}(x)$ is the characteristic function of Ω). We consider a *best-case* function and a *worst-case* function, both from this class of functions to derive the best- and worst-case variance convergence rate, as the number of samples N increases.

Conversely, analytical models for the sampling power spectra are often avail-

able [34] in the Euclidean space. We show that these power spectra can be used in our framework to predict the variance of an integrand. We bound the variance convergence rate of these samplers—with known power spectra—by simply imposing bounds on the sampling power spectra in MC integration. For sampling power spectra with no analytical formulations, we provide tools for bounding the power spectra in Sec. 6.2, which have been developed a priori in the Euclidean domain by Pilleboue [78]. In this chapter, we extend these analysis tools on the sphere.

6.1 Best and worst case variance

First, we derive theoretical convergence rates of the variance for the best- and the worst-case functions from a given class of functions. We define our *best-case* integrand directly in the spectral domain with an angular power spectrum $\check{\mathcal{P}}_F(l)$, which is a constant a_F for $(l < l_0)$ and zero elsewhere. An illustration is shown in Fig. 6.1(a). The corresponding variance can be obtained using Eq. (5.18):

$$\text{Var}(\mathbf{I}_N) \leq \frac{a_F \mu(\mathcal{S}^2)^2}{N} \sum_{l=1}^{l_0} \mathcal{P}_{\mathbf{s}}(l), \quad (6.1)$$

where a_F is a non-zero positive constant, that can be derived by plugging in the best-case angular power spectrum.

From the class of functions provided by Brandolini and colleagues [9], we define our *worst-case* integrand to be a sum of zonal harmonics, $\mathcal{S}_F(l, m) = 0 \forall m \neq 0$, with spectral amplitude square zonal coefficients $\|\mathcal{S}_F(l, 0)\|^2$ exhibiting a decay rate of order $O(l^{-2})$ for $l > l_0$. The corresponding worst-case power




Figure 6.1: Illustration of the *mean angular* power spectra *vs* frequency ($\omega = l/\alpha\sqrt{N}$) of signals that can be considered as best (a) and worst (b) cases in numerical integration with respect to any sampling power spectrum.

spectrum \mathcal{P}_S of \mathbf{S} has a decay rate of order $O(l^{-3})$, as illustrated in Fig. 6.1(b). By substituting this worst-case decay rate in Eq. (5.18) to obtain the corresponding worst-case variance as follows:

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^2)^2 a'_F}{N} \sum_{l=1}^{\infty} \frac{\check{\mathcal{P}}_S(l)}{l^2}, \quad (6.2)$$

where a'_F is a non-zero positive constant that can be derived by plugging in the worst-case angular power spectrum. For our experiments, we consider a *spherical cap* function as our worst-case integrand as it falls into the given class of functions. The power spectrum of a spherical cap function exhibits a decay rate of order $O(l^{-3})$, whose derivation is given in our Appendix C.

6.2 Convergence rate analysis

To study the variance convergence rate analysis we have restricted our variance model Eq. (5.18) to a given class of functions and define the corresponding best- and worst-case (integrand) power spectra in this class. We then obtain the best-

case (Eq. (6.1) and worst-case (6.2)) variance, as a function of the sampling power spectrum $\check{\mathcal{P}}_{\mathbf{S}}(\cdot)$.

In this section, we study the effects of different shapes of $\check{\mathcal{P}}_{\mathbf{S}}(\cdot)$ on the variance in MC integration for both the best- and worst-case (integrands). We use simple shape profiles (quadratic, polynomial, step) for $\check{\mathcal{P}}_{\mathbf{S}}(\cdot)$ and derive the variance convergence rate associated to each of these power spectra profiles. Then, we classify the existing state-of-the-art sampling power spectra—with respect to the shape of the low frequency zone of these power spectra—in terms of these profiles. We also obtain corresponding upper and lower bounds on the sampling power spectra in terms of these profiles.

As mentioned by Pilleboue, Singh and colleagues [79], to analyze the variance convergence rate of a sampling pattern we first identify, which part of our variance formulation (Eq. (6.1) and (6.2)) depends on N . Since the distribution of distances in a sampling pattern gets affected by the sampling density, the *mean neighborhood distance* between samples can be approximated as the d -dimensional root of the representative area of each sample. In the Fourier domain, the corresponding frequency of this *mean distance* is $\sqrt[d]{N}$, which corresponds to the first peak in the power spectrum of a regular grid pattern, and also, approximatively, to the first bump of the typical blue-noise spectrum. As N increases, the low frequency region of the sampling power spectrum becomes more and more significant with respect to the integrand power spectrum. This explains the noticeable variance reduction observed for sampling patterns exhibiting no low-frequency content. Mathematically, this means that the radial mean of a power spectrum is dependent

on a factor $\sqrt[d]{N}$. In the following subsection, we use this factor ($\sqrt[d]{N}$), for the spherical domain ($d = 2$), to derive variance convergence rates.

6.2.1 Convergence analysis in Spherical domain

To start our analysis, the mean power spectrum produced by the sampler must be known, or at least estimated. Therefore, we cover only the classical stochastic samplers, from the simplest one (White noise) to the most sophisticated ones, *e.g.*, CCVT [6], for which the radial power spectra can be estimated. We classify different sampling patterns with respect to the shape of their corresponding radial power spectra in the low-frequency zone. We study three shapes of the power spectra, given as follows:

Constant power spectra Here, $\check{\mathcal{P}}_S(l) := \gamma$, where γ is a constant value. As the radial power spectrum $\check{\mathcal{P}}_S(l)$ is independent of N , by simply plugging $\check{\mathcal{P}}_S(l) = \gamma$ in Eq. (6.1) and Eq. (6.2), we can deduce the best- and worst-case variance convergence rates, which is given of order $O(N^{-1})$.

Step power spectra The step profile corresponds to a constant power spectrum with zero low-frequency part (Fig. 6.2), *i.e.*, from $l = 0$ to $\alpha\sqrt{N}$, for a given $\alpha \in \mathbb{R}^+/\{0\}$, where α is used to quantify the range of energy-free frequency with respect to the mean frequency.




Figure 6.2: Step

The resultant step profile is given by:

$$\check{\mathcal{P}}_{\mathbf{S}}(l) := \begin{cases} 0, & \text{if } l < \alpha\sqrt{N} \\ \gamma, & \text{otherwise.} \end{cases} \quad (6.3)$$

For $N > l^2/\alpha^2$, the variance of MC integration of the best-case F is zero. We normalize the l term with $\alpha\sqrt{N}$, where α is a positive constant and N represents the number of samples. This allows us to bring all power spectra with different N , to the same scale for all N . The expression for the variance, $\text{Var}(\mathbf{I}_N)$, due to step power spectrum in the worst-case can be approximated by plugging Eq. (6.3) in Eq. (6.2):

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^2)^2 a'_F}{N} \sum_{l=\lceil \alpha\sqrt{N} \rceil}^{\infty} \frac{1}{l^2}. \quad (6.4)$$

where a'_F is a hidden constant from $O(l^{-2})$. The summation term in Eq. (6.4) can be symbolically solved to deduce the worst-case variance convergence rate of sampling patterns with a step power spectra, which is of order $O(N^{-1.5})$.

Quadratic power spectra A more general radial power spectrum can be constructed by considering a polynomial ramp of degree b from the frequency zero to $\alpha\sqrt[b]{N}$. After the frequency $\alpha\sqrt[b]{N}$, the radial power spectrum is a constant function of γ , which is similar to the *Constant* and *Step* profiles. More formally, the polynomial ramp can be written as:

$$\check{\mathcal{P}}_{\mathbf{S}}(l) := \begin{cases} \frac{\gamma l^b}{\alpha^b \sqrt[b]{N}^b} & \text{if } l < \alpha\sqrt[b]{N} \\ \gamma & \text{otherwise.} \end{cases} \quad (6.5)$$

When we plug Eq. (6.5) directly in the best Eq. (6.1) or worst-case Eq. (6.2) variance formulation, we were not able to obtain a general variance expression for the best-case and the worst-case. This is because, the general polynomial profile Eq. (6.5) expression cannot be symbolically solved for an unknown degree b .

Therefore, we study the polynomial profile defined in Eq. (6.5), for a given polynomial degree, in our case for ($b = 2$), which gives a quadratic profile (Fig. 6.3). Using this quadratic profile, the power spectrum in the best-case variance Eq. (6.1) can be obtained as follows:




Figure 6.3: Quadratic ($b=2$)

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^2)^2}{N} \sum_{l=1}^{l_0} \frac{\gamma a_F l^2}{(\alpha\sqrt{N})^2} \quad (6.6)$$

$$= O(N^{-2}) . \quad (6.7)$$

Using the quadratic profile power spectrum for the worst-case variance Eq. (6.2), we obtain:

$$\text{Var}(\mathbf{I}_N) \leq \frac{\mu(\mathcal{S}^2)^2}{N} \left(\sum_{l=1}^{\lfloor \alpha\sqrt{N} \rfloor} \frac{\gamma a_F (l/(\alpha\sqrt{N}))^2}{l^2} + \sum_{l=\lceil \alpha\sqrt{N} \rceil}^{\infty} \frac{\gamma a'_F}{l^2} \right), \quad (6.8)$$

which can be symbolically solved to obtain a convergence rate of order $O(N^{-1.5})$. Now, we have the variance convergence rates of various simple radial power spectral profiles. The inspiration behind studying these simple profiles (constant, step, quadratic) is the following: First, since we do not know the exact analytical expression of the radial power spectra of various sampling patterns, we can approximate these radial power spectra using these simple profiles. Second, if we

can bound the sampler's radial power spectrum using any one of these profiles, it would give us the variance convergence rate of that sampler, which is derived using the profile used to bound the radial power spectra.

6.2.2 Discussion

We have proposed a classification of radial power spectra of various sampling patterns with respect to the shape of their radial power spectra in the low-frequency zone. We can use this classification as a theoretical tool to bound arbitrary radial power spectra using these simple profiles (e.g., linear, quadratic, etc.). Thus, the corresponding variance convergence rate of the sampling patterns can be deduced from the bounds of their respective profiles. For some sampling patterns, the corresponding analytical formulation is known. For example:

White noise The radial power spectrum of white noise has a constant profile with $\gamma = 1$ exhibiting a variance convergence rate of order $O(N^{-1})$, which is well known in literature.

Poisson Disk For Poisson Disk distributions, Torquato and colleagues [94] observed that for a sufficiently large N , the minima of the normalized mean power spectrum can be asymptotically seen near the DC peak, with a value of 0.05. Torquato and colleagues made this observation while studying Poisson Disk samples in the Euclidean domain. We observed that a similar offset 0.055 is obtained for the *mean* angular power spectrum of Poisson Disk samples in the

spherical domain. Following our theoretical analysis, the power spectrum of the Poisson disk distribution is lower bounded by a constant profile with $\gamma = 1/20$ for both Euclidean and spherical domains. Consequently, the variance convergence rate of the Poisson Disk sampling patterns is similar to the convergence rate of the white noise sampler, which can also be bounded by the constant profile. The variance convergence rates of white noise and Poisson Disk samples is of order $O(N^{-1})$.

In the case of jittered, regular and CCVT samplers, we have seen an improved variance convergence rates, compared to Poisson Disk. One of the similarity between jittered, regular and CCVT samplers is that all of them have mean angular power spectra tending to zero as the frequency parameter (ω) goes to zero. On top of that, in case, of regular and CCVT sampling there is a low-frequency region with no energy. All these power spectra are shown in Fig. 4.14 with regular sampling power spectra shown in Fig. 4.10. This implies that the ideal mean angular power spectrum of a sampler must have zero-valued power in the low-frequency zone. Poisson Disk sampling patterns is always categorized as *blue noise* due to the resemblance of its power spectrum shape with *blue noise*. However, it falls behind even jittered sampling patterns due to the offset in the low-frequency zone of its power spectra. Therefore, it would be desirable to have sampling patterns with mean angular power spectra tending towards zero as the frequencies tends to zero with no energy in the low-frequency zone.

In Chapter 7, we perform experiments on other sampling patterns and bound their radial power spectra with the profiles discussed in this chapter. An illustration

of the variance convergence rates of various samplers is shown in Fig. 7.4 with other results.

Chapter 7

Experiments and Results

The theoretical framework proposed in this dissertation is based on the observations we made while performing spectral analysis (Chapter 4) of various sampling patterns. Based on this theoretical framework proposed in Chapter 5 we develop some analytical tools in Chapter 6 to derive theoretical variance convergence rates of various sampling patterns. In this chapter, we use our framework to perform experiments on the state-of-the-art sampling patterns on the sphere.

We first, give implementation details on how various sampling patterns are obtained on the sphere. We then use these spherical samplers to study the bounds on the variance convergence rates of these samplers. We also evaluate some ambient occlusion integrals to show that the observations made while deriving variance convergence rates can also be generalized when the underlying integrand is not known.




Figure 7.1: Illustration of spherical sampling methods.

7.1 Implementation details

We first look at the implementation details of our spherical samplers. We present several experiments by comparing existing sampling methods: white noise, jittered sampling, Poisson Disk and capacity constraint methods. The underlying data structure used to perform sampling on the sphere for some samplers is the *Healpix* data structure [37] which is an equal area quadrangulation of the sphere.

White noise or purely random samples are generated on the sphere by simply using the *drand48* C programming language built-in function from the *stdlib.h* header. To implement regular sampling on a sphere, we perform sampling directly on the sphere using the *Healpix* data structure [37] which is an equal area quadrangulation of the sphere. To implement jittered sampling, we use the *Healpix* quads as strata and randomly place samples in each stratum. An illustration of regular and jittered sampling with the corresponding *healpix* stratas are shown in Fig. 7.1.

Poisson Disk Sampling For Poisson Disk sampling, there exist many algorithms in the literature that mimic the dart throwing approach on the surface of the sphere [14, 77, 35]. In our implementation, we first generate a dense set of samples on the sphere, as in the work of Li and colleagues [59], and then reject samples that are too close, as in the work of Bridson [11]. We achieve a distribution close to the boundary sampling described in the work of Dunbar and Humphreys [27].




Figure 7.2: Poisson Disk

CCVT We also implement a spherical version of the CCVT algorithm by Balzer and colleagues [6] using the geodesic distance on the sphere. An illustration is shown on the right side. All sphere samplers are generated on a multi-processor machine, with an Intel Core i7 CPU 980@3.3g GHz processor.




Figure 7.3: CCVT

To plot the radial power spectra of various sampling patterns (shown in Fig. 7.4(a,d,g,j)), we compute the average power spectrum over 1000 realizations of a point set, as suggested by Schlömer and colleagues [85]. We plot power spectra for $N = 4096$ in the spherical domain. Later, we use these power spectra to estimate lower and upper bounds on the variance of our test integrands in the spherical (Fig. 7.4) domain.

7.2 Case studies

In Chapter 6, we propose an analysis tool that can be used to bound the radial power spectra of various sampling patterns using simple profiles (constant, quadratic, etc.). We also propose a best-case and worst-case variance convergence rates for the respective profiles. In this section, we experimentally test our tools and bound various state-of-the-art sampling (radial) power spectra on the sphere. As mentioned in Chapter 6, we consider a class of integrable (smooth) functions with bounded domain and smooth boundaries [9]. From this class of functions we choose a spherical harmonic Y_l^m ($l = 4, m = 0$) basis function for best-case study since it is smooth and bandwidth-limited. For our worst-case study, we choose a spherical cap function, with spherical cap size $\theta_0 = 60$ degrees, in the spherical domain.

To perform variance analysis we compute the variance of our test (best- and worst-case) integrands using different sampling patterns. The variance is computed over 1000 trials for cases where the samplers (white noise and jittered) are not prohibitively slow. For spherical CCVT and Poisson Disk, we keep the number of trials between 200 to 1000, depending on the speed of the sampler. As the spectral profile of these test integrands is known, our variance prediction model can be used to estimate the bounds on the variance in integration of each function by simply using the bounds derived from the corresponding sampling power spectra.

First, we look at our test integrands which we chose for our best-case and

worst-case variance convergence analysis:

Spherical harmonic basis function (Y_l^m) In the spherical domain, we chose a SH basis function Y_l^m with ($l = 4, m = 0$), as a best-case which is a band limited function. Y_0^4 is a smooth function with compact support and can be considered as a *best-case* for the class of functions discussed in the paper. The power spectral profile of any SH basis function, Y_m^l , is given by:

$$\mathcal{P}_F(l) = \frac{\mu(\mathcal{S}^2)}{(2l + 1)}, \quad (7.1)$$

since $\|Y_l^m\|^2 = \mu(\mathcal{S}^2)$.

Spherical cap function Spherical cap function is our worst-case test integrand for a given class of functions. We compute the angular power spectrum, in Appendix C, of spherical cap function with size $\theta_0 = 60$ which is given by:

$$\mathcal{P}_F(l) = 4\pi^2 \left(\frac{|P_{l-1}^0(\cos(\theta_0)) - P_{l+1}^0(\cos \theta_0)|^2}{(2l + 1)^2} \right), \quad (7.2)$$

where $P_l^0(\cos(\theta_0))$ is the Legendre Polynomial function. We numerically compute the above expression for the power spectrum of spherical cap and use it in our variance formulation (Eq. (5.18)), to obtain the variance with respect to a given sampler.

To obtain the bounds on the variance of these test integrands, we first bound the mean angular power spectra ($\mathcal{P}_S(l)$) of various sampling patterns using the simple profiles discussed in Section 6.2. To start with, we first compute the bounds on the white noise mean angular power spectrum ($\mathcal{P}_S(l)$). Since, $\mathcal{P}_S(l) = \gamma$

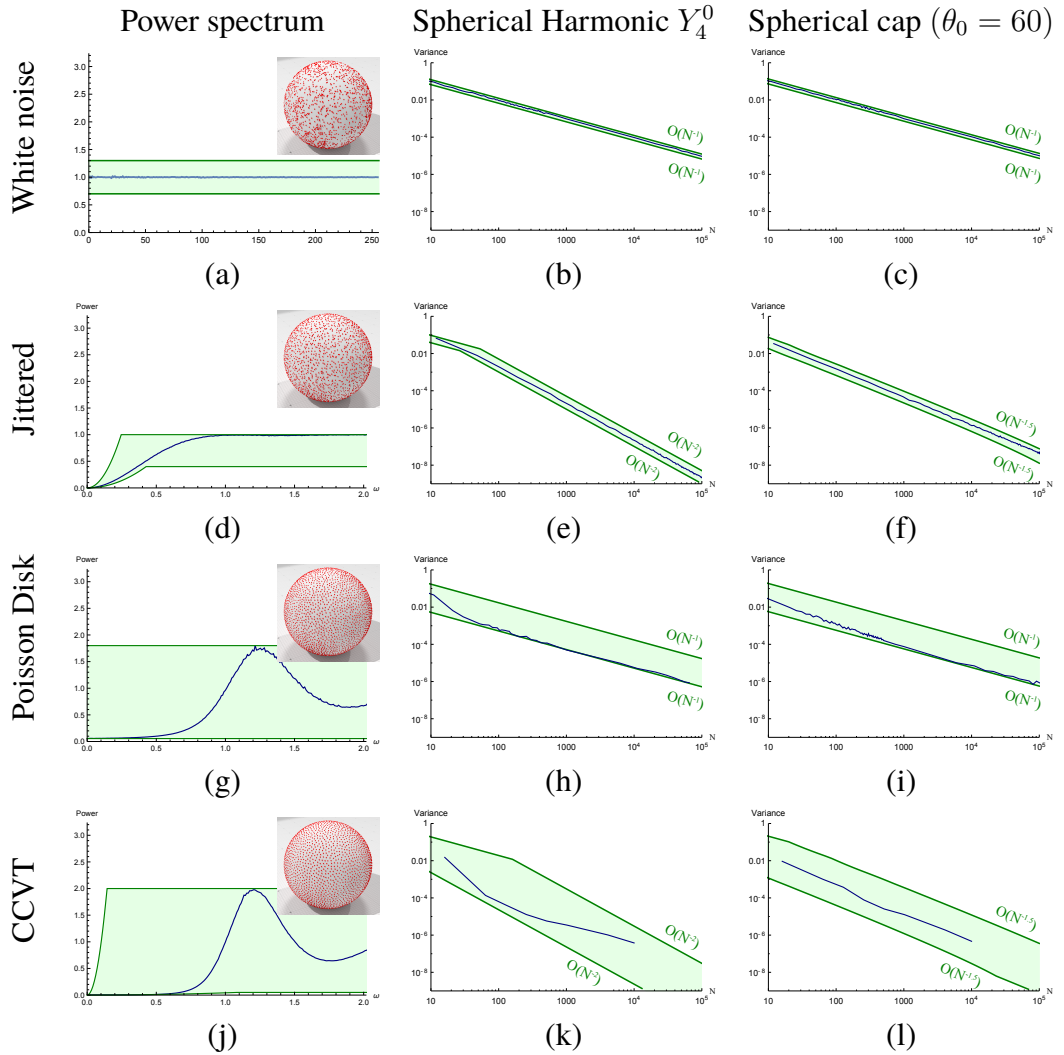


Figure 7.4: Bounds on the mean angular power spectra of samplers and on the variance convergence rates of our test integrands. Log-Log plots are shown in the **Center** and **Right** columns. **Left:** mean angular power spectra of different sampling patterns (blue), bounded by a theoretical spectral profile (green). On the frequency axis, units corresponds to the frequency $\alpha\sqrt{N}$. **Center:** The variance curve of a spherical harmonic basis function with $l = 4, m = 0$, (blue), with bounds (green) computed using the bounds of the corresponding mean angular power spectrum. **Right:** The variance in integration of a spherical cap, using the same visualization.

for whitenoise, it can be upper and lower bounded by any constant profile. For white noise, we normalize the mean angular power spectrum such that $\gamma = 1$. Consequently, the parameters (α, γ) of the upper bound are $\alpha_u = 1.0, \gamma_u = 1.3$, and for the lower bound are $\alpha_l = 1.0, \gamma_l = 0.7$. Note that, for $\gamma = 1$ both the upper and lower variance bounds overlaps with the variance convergence rate of a given integrand.

By substituting $\mathcal{P}_S(l) = 1$ and the mean angular power spectrum expression of our Y_4^0 from Eq. (7.1) in the variance formulation Eq. (5.18) we can directly predict the variance of spherical basis function using white noise. Similarly, by using the mean angular power spectrum of spherical cap test integrand from Eq. (C.12) with $\mathcal{P}_S(l) = 1$ in the variance formulation we can directly predict the variance of spherical cap function. The corresponding bounds on Y_4^0 and spherical cap functions can be obtained by simply using $\mathcal{P}_S(l) = \gamma_u = 1.3$ and $\mathcal{P}_S(l) = \gamma_l = 0.7$ in the variance formulation. All bounds in the case of white noise, are shown in Fig. 7.4(a,b,c).

In a similar way, constant profiles are used to bound the mean angular power spectrum of Poisson Disk sampling pattern. The parameters used to bound the power spectrum are: $\alpha_u = \sqrt{2.75}, \gamma_u = 1.8, \alpha_l = \sqrt{2.75}, \gamma_l = 0.055$. The bounds on our test integrands, spherical cap and Y_0^4 are similarly as explained in the case of white noise. For jittered sampling and CCVT sampler we use quadratic profile to bound their mean angular power spectra. The parameters used in the case of jittered sampler are: $\alpha_u = 0.2, \gamma_u = 1.0, \alpha_l = \sqrt{0.6}, \gamma_l = 0.4$. For CCVT the quadratic ($b = 2$) profile parameters used to bound the mean angular power

spectrum are given by: (j) $\alpha_u = \sqrt{0.05}$, $\gamma_u = 2.0$, $\alpha_l = \sqrt{2.85}$, $\gamma_l = 0.05$.

We illustrate the variance computed via our framework in Fig. 7.4. Our experimental results are consistent with our variance predictive model (Chapter 6), and the variance is correctly bounded in the spherical domain.

As shown in Fig. 7.4, we observe that tighter bounds on the sampling power spectra can result in tighter bounds on the variance. For example, the lower bound on the jittered sampling power spectrum is tighter than the upper bound which results in a tighter lower bound on the associated variance. For Poisson Disk sampling, the spectral profile is bounded by a constant profile and consequently, the convergence rate of the variance is similar to white noise. Also, as the lower bound on the Poisson Disk sampling is tighter than the upper bound, the corresponding variance achieves a tight lower bound. In Fig. 7.4 for the sphere, all the power spectral profiles are plotted as $(\mathcal{P}_F \text{ vs } \omega)$ with $\omega = l/(\alpha\sqrt{N})$. Note that, we only bound the variance convergence rate. This is why the variance values are normalized (with the variance value for $N = 1$) for all of our test integrands.

7.3 Rendering Results

We also implement ambient occlusion in a Cornell box scene to study the behavior of various state-of-the-art sampling patterns (Fig. 7.5 and 7.6). As most state-of-the-art samplers are not directly extendable to the hemisphere, we sample a complete sphere and consider samples only from the visible part. We designed the Cornell box scene to enhance the noise level in the whole box.




Figure 7.5: Cornell box rendered with ambient occlusion. *Mean squared error* (MSE) values are computed w.r.t a reference image (shown in Fig. 7.7, for White noise (**Left**), Jittered sampling (**Center**) and Poisson Disk (**Right**) for a given number of shading rays used to sample directions on the visible hemisphere at each hitpoint in the scene.




Figure 7.6: Cornell box rendered with ambient occlusion. *Mean squared error* (MSE) values are computed w.r.t a reference image (shown in Fig. 7.7, for White noise (**Left**), CCVT (**Center**) and Regular Healpix (**Right**) sampling for a given number of shading rays used to sample directions on the visible hemisphere at each hitpoint in the scene.

For rendering, we chose a regular sampling pattern on the image plane to keep the aliasing coherent throughout all samplers. This is why our reference image, shown as an inset in Fig. 7.9(c), shows some structural artifacts. The reference Cornell box image is computed by shooting one




Figure 7.7: Reference image

ray from the center of each pixel followed by shooting 32k secondary rays from each primary hit point. The secondary ray directions were sampled via a jittered sampling pattern. A similar process was adopted to generate the Cornell box images using other samplers. This is done to compare the variance coming through the hemispherical sampling patterns from each primary ray hit point.

In Fig. 7.8, we also compare the regular patterns of Spherical Fibonacci (SF) and Healpix point sets on a cornell box scene. Regular patterns on the sphere can be obtained using different ways. In this dissertation, we use SF lattice and the Healpix data structure to obtain regular sample points on the sphere. This comparison (Fig. 7.8) shows that the structural artifacts due to SF lattice pointset are way more visible than the Healpix point set. Although, after rotation the structural artifacts turns into noise.

7.4 Discussions

To provide a side by side comparison of the variance convergence rate of different samplers, we plot the convergence rate curves in Fig. 7.9. As predicted by our




Figure 7.8: Cornell box rendered with ambient occlusion. **Left:** Spherical Fibonacci (SF) point set is used at each hitpoint. **Centre:** SF point set is used with uniform and random rotation at each hitpoint. **Right:** Healpix regular point set is used at each hit point with no rotation. *Mean squared error* (MSE) values are computed w.r.t a reference image (shown in Fig. 7.7)

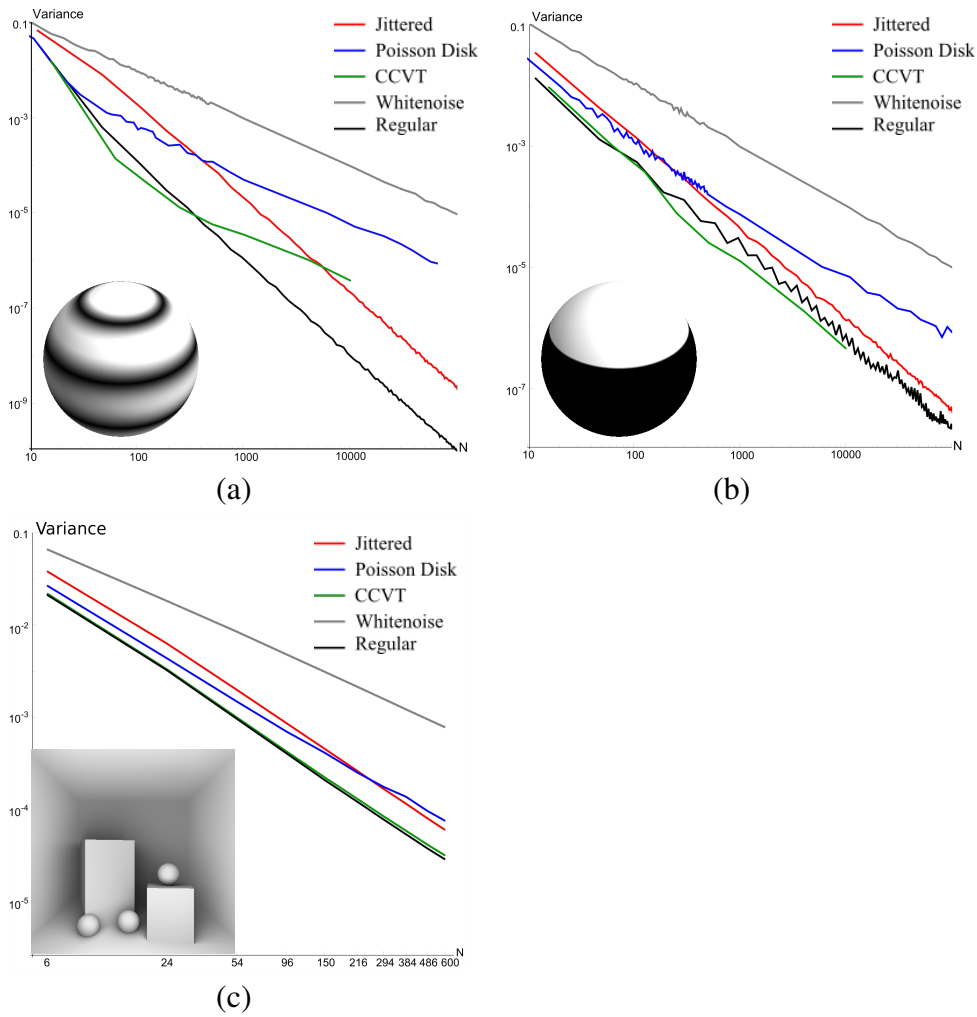


Figure 7.9: Comparison of the variance in MC integration for different integrand signals. Here we use experimental data from 7.4 in (a) and (b). Here we present comparisons in the (hemi-)spherical domain : (a) a spherical harmonic basis function (Y_4^0). Inset illustrates the gray scale of the absolute values of the function, (b) a spherical cap function ($\theta_0 = 60$), where the white shade in the inset represents non-zero constant value region (c) a Cornell box scene with all corresponding rendered images shown in Fig. 7.5 and 7.6.

model, in the Euclidean domain, we can see that the asymptotic behavior of CCVT [6], Healpix regular and jittered sampling is better than both the Poisson Disk and the white noise sampling methods. The jittered sampling convergence rate overtakes the Poisson Disk sampling convergence rate just after 300 samples (as shown in Fig. 7.9). Similar observations can also be made in the variance convergence rates of a Cornell box scene. Even though jittered sampling pattern shows higher variance compared to Poisson Disk samples when the number of samples are low, after 294 samples variance due to Poisson Disk becomes higher than jittered sampling pattern.

One of the interesting observations that comes into limelight is regarding the Poisson Disk (dart throwing based) sampling pattern, which is always categorized as *blue noise* in the literature. In this work, we have shown—both mathematically and experimentally—that the convergence rate of Poisson Disk sampling patterns is no better than white noise. This is due to an offset in the low-frequency region ($\omega \in]0, 0.5[$) of the mean angular power spectrum of Poisson Disk, shown in Fig. 7.10(c). However, Poisson Disk samples obtained using tile based approaches can help get rid of this offset in the low-frequency region, resulting in an improved variance convergence rate. We also observe that jittered sampling patterns are as good as *blue noise* samples (*e.g.*, CCVT) when it comes to the variance convergence rate, though the variance is higher. This is due to the reason that the angular mean power spectrum of jittered samples approaches to zero as the frequency (ω) tends to zero (Fig. 7.10(d)).

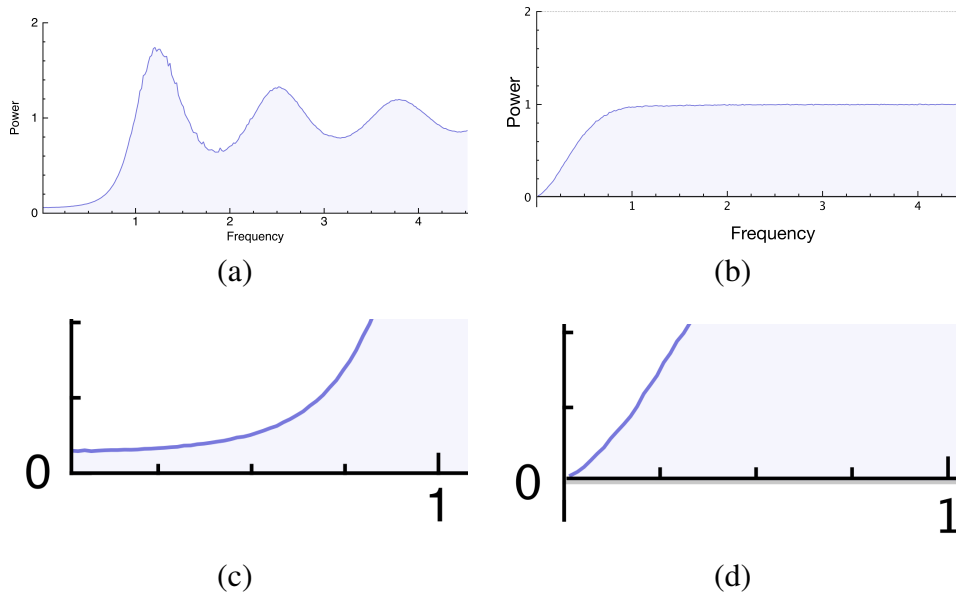


Figure 7.10: Comparison between the power spectra at the low frequency region between (a) Poisson Disk and (b) jittered samplers. In (c) and (d) we show the zoom-in of the low frequency region ($\omega \in]0, 1[$) for the Poisson Disk and the jittered samplers respectively.

Summary: We have shown that the variance in MC integration can be upper and lower bounded if we have some knowledge of the samplers' mean angular power spectra. We test our theory on number of state-of-the-art samplers with different integrands. We are also able to predict the variance in MC integration if we know exactly the power spectrum of the integrand involved.

Among various spherical samplers discussed in this dissertation (homogenized) jittered sampler is recommended for rendering purposes (or other numerical integrations) over Poisson Disk samplers, even when the number of samples required are not too high. This is because the convergence rate of Poisson Disk samples is the same as the purely random samples (white noise). Homogenized

regular samples seems to give good results in the cornell box scenes (Fig. 7.5 and Fig. 7.6) due to the fairly low frequency content of the scene. For a scene with highly complex material properties regular patterns can give aliasing artifacts (false signals) due to the presence of peaks in its power spectra.

Chapter 8

Conclusions

In this work, we develop a framework to analyze the effects of various spherical sampling patterns on the variance of MC integration. We also propose a mathematical background to study hemispherical signals and samplers using spherical harmonics by simply mapping hemispherical signals in the projective space. In the evaluation of global illumination integrals, we use the notion of *homogeneous* sample distributions on the sphere that allows manifestation of error only in terms of variance during MC integration. We show that the variance is directly related to the sampling and integrand power spectra. We also show that for a given class of functions, we obtain the worst-case variance convergence rates.

One of the limitations of our framework is that it is designed for only stochastic samplers and assumes some knowledge of the power spectrum of the sampling pattern under study. Therefore, sampling patterns with varying power spectra (e.g. Sobol), that is, which depend on the number of samples, cannot be handled by our framework. We would like to explore this issue in future. Another promising

direction for future research is to develop a similar framework for deterministic sampling methods like, Sobol or other low-discrepancy sequences that can be generated directly on the sphere.

Additionally, our framework can be used to study the effects of various sampling patterns with known and/or unknown closed-form power spectral formulation. We show that by using simple shapes (quadratic, constant, etc.) for the mean angular power spectra, existing state-of-the-art sampling power spectra can be easily bounded. This also results in bounding the associated variance for MC integration. We use our framework to theoretically derive best- and worst-case variance convergence rates—for a given class of functions—of various state-of-the-art sampling methods.

8.1 Design principles

In our theoretical framework (Sec. 6.2), we propose some principles to design new sampling patterns. We also advocate the importance of the low frequency zone of the corresponding sampling power spectra. We emphasize the fact that an ideal sampling power spectrum must attain zero value as the frequency parameter tends to zero. This property ensures that the distribution is sufficiently uniform. In addition, the shape of the power spectrum near the zero frequency zone is also very important. Moreover, the flatter the shape of the power spectral profile in this low frequency zone, the faster the variance converges in MC integration for favorable integrands. Ideally, the power spectrum must contain a range of low

frequencies that are perfectly zero.

Our experimental results in the Euclidean and spherical domains show that no existing methods fulfill these design principles. Sampling methods based on packing and hard local conditions like Poisson Disk and [85], do not have a power spectrum converging to zero near zero frequency. This can be explained by the lack of uniformity in the distribution which can be corrected using proper subdivision of the domain through tile-based methods. For example, the tiled version of Poisson Disk sampling in Wachtel et al. [99] has power spectrum converging to zero towards zero frequency. Similarly, all methods based on the capacity constraint [6] seem to guarantee a power spectrum that converges to zero. However, these methods still have residual energy spread along low frequencies which can be improved.

8.2 Quadrature rules on the sphere

In our framework, we perform Monte Carlo integration—instead of using optimal adaptive quadrature rules—on the spherical and hemispherical domains despite their low dimensional nature. We do this for a couple of reasons: First, our present framework is designed only for non-adaptive sampling patterns. Adaptive or importance-based sampling do not support the *homogeneity* of the sampling patterns, and designing a variance analysis framework in such a case needs further research; Second, as most production renderers are Monte Carlo based, our framework provides an analysis tool that can be employed in the production pipeline

for choosing a sampling pattern that reduces the overall variance in integration of (hemi-)spherical signals. New sampling patterns can also be designed using our framework.

8.3 Future directions

Generating a sampling pattern with absolutely no energy in the low frequency range seems to be a very challenging problem. We expect to use the theoretical results and tools developed in this paper to help design such sampling patterns. Another important avenue of research is to develop sampling methods with spectral control directly in the spherical domain. Finally, taking into account more sophisticated sampling strategies, such as adaptive (or importance) sampling and filtering, could be an interesting avenue for future work.

Appendix A

Simplified variance equation

A.1 Contribution of DC peak in Variance

As shown in Eq. (5.9), we have the general closed-form expression of variance given by:

$$\begin{aligned} \text{Var}(\mathbf{I}_N) &= \langle \|\mathbf{I}_N\|^2 \rangle - \|\langle \mathbf{I}_N \rangle\|^2 \\ &= \frac{\mu(\mathcal{S}^2)^2}{N^2} \left(\sum_{l=0}^{\infty} \sum_{m=-l}^l \langle \|\mathcal{S}_S(l, m) \cdot \overline{\mathcal{S}_F(l, m)}\|^2 \rangle - \|\langle \mathcal{S}_S(l, m) \cdot \overline{\mathcal{S}_F(l, m)} \rangle\|^2 \right). \end{aligned} \tag{A.1}$$

In this appendix, we show that the DC peak contribution is zero. Let us introduce a temporary notation $\text{Var}(\mathbf{I}_N, l, m)$ that represents the contribution of the (l, m) -th component in the variance $\text{Var}(\mathbf{I}_N)$ given in Eq. (A.1). We compute the

contribution of the DC peak ($l = 0, m = 0$) component in the variance as follows:

$$\text{Var}(\mathbf{I}_N, 0, 0) = \left\langle \|\mathcal{S}_S(0, 0) \cdot \overline{\mathcal{S}_F(0, 0)}\|^2 \right\rangle - \left\| \left\langle \mathcal{S}_S(0, 0) \cdot \overline{\mathcal{S}_F(0, 0)} \right\rangle \right\|^2 \quad (\text{A.2})$$

$$= \langle \mathcal{S}_S(0, 0)^2 \cdot \mathcal{S}_F(0, 0)^2 \rangle - \langle \mathcal{S}_S(0, 0) \cdot \mathcal{S}_F(0, 0) \rangle^2 \quad (\text{A.3})$$

$$= \langle \mathcal{S}_S(0, 0)^2 \rangle \cdot \mathcal{S}_F(0, 0)^2 - \langle \mathcal{S}_S(0, 0) \rangle^2 \cdot \mathcal{S}_F(0, 0)^2 \quad (\text{A.4})$$

$$= (\langle \mathcal{S}_S(0, 0)^2 \rangle - \langle \mathcal{S}_S(0, 0) \rangle^2) \cdot \mathcal{S}_F(0, 0)^2 \quad (\text{A.5})$$

Here, $(\langle \mathcal{S}_S(0, 0)^2 \rangle - \langle \mathcal{S}_S(0, 0) \rangle^2)$ can be seen as the variance of the spherical harmonic basis function Y_l^m with $l = 0, m = 0$. Since, Y_0^0 is a constant function, the corresponding variance would be zero resulting in the $\text{Var}(\mathbf{I}_N, 0, 0) = 0$. This makes the $\text{Var}(\mathbf{I}_N, 0, 0) = 0$, showing that DC peak component does not contribute to the variance of MC integration in the spherical domain.

A.2 Contribution of second term

In $\left\| \left\langle \mathcal{S}_S(l, m) \cdot \overline{\mathcal{S}_F(l, m)} \right\rangle \right\|^2$, as only \mathbf{S} is a random variable, we can simplify this term as follows:

$$\left\| \left\langle \mathcal{S}_S(l, m) \cdot \overline{\mathcal{S}_F(l, m)} \right\rangle \right\|^2 = \left\| \langle \mathcal{S}_S(l, m) \rangle \cdot \overline{\mathcal{S}_F(l, m)} \right\|^2.$$

For homogeneous sampling patterns $\langle \mathcal{S}_S(l, m) \rangle = 0$ for all $l \neq 0, m \neq 0$.

Consequently, for DC peak ($l = 0, m = 0$), we get:

$$\left\| \langle \mathcal{S}_S(l, m) \rangle \cdot \overline{\mathcal{S}_F(l, m)} \right\|^2 = 0$$

for all $l \neq 0, m \neq 0$.

Appendix B

Unbiased homogeneous sampling

In this appendix, we mathematically illustrate that the homogeneous sampling patterns are unbiased in nature. For this, we show that the expected value of the Monte Carlo estimator is equal to the integration of F . The expected value of the MC estimator can be written as:

$$\langle \mathbf{I}_N \rangle = \frac{\mu(\mathbb{D})}{N} \left\langle \int_{\mathbb{D}} \mathbf{S}(x) F(x) dx \right\rangle = \frac{\mu(\mathbb{D})}{N} \int_{\mathbb{D}} \langle \mathbf{S}(x) \rangle F(x) dx. \quad (\text{B.1})$$

To make \mathbf{S} homogeneous, we perform averaging over the group of motions, H (translations for the torus and rotations for the sphere) over the whole domain, this gives us:

$$\langle \mathbf{I}_N \rangle = \frac{\mu(\mathbb{D})}{N} \int_{\mathbb{D}} \frac{1}{\mu(H)} \left\langle \int_H [\tau(\mathbf{S})](x) d\tau \right\rangle F(x) dx. \quad (\text{B.2})$$

To simplify the inner integral, we use the fact that H acts transitively on \mathbb{D} , as a result, the average of $[\tau(\mathbf{S})](x)$ over all $\tau \in H$ is the average of $\mathbf{S}(x)$ over all $x \in \mathbb{D}$:

$$\frac{1}{\mu(H)} \int_H [\tau(\mathbf{S})](x) d\tau = \frac{1}{\mu(\mathbb{D})} \int_{\mathbb{D}} \mathbf{S}(x) dx. \quad (\text{B.3})$$

Plugging Eq. (B.3) back into Eq. (B.2) (with a change of variable), we get:

$$\langle \mathbf{I}_N \rangle = \frac{\mu(\mathbb{D})}{N} \int_{\mathbb{D}} \left\langle \frac{1}{\mu(\mathbb{D})} \int_{\mathbb{D}} \mathbf{S}(y) dy \right\rangle F(x) dx \quad (\text{B.4})$$

$$= \frac{1}{N} \int_{\mathbb{D}} \langle \mathbf{S}(y) \rangle dy \int_{\mathbb{D}} F(x) dx. \quad (\text{B.5})$$

In particular, when $\mathbf{S}(y) = \sum \delta(y - s_k)$ for $k = 1 : N$, we get:

$$\langle \mathbf{I}_N \rangle = \int_{\mathbb{D}} F(x) dx. \quad (\text{B.6})$$

Appendix C

Spectral Analysis of spherical cap

C.1 Spherical cap function in SH terms

In Chapter 6, we mentioned that a spherical cap function can be considered as our worst-case integrand for a given class of functions [9]. We use spherical cap function to perform our experiments in Chapter 7. In this section, we compute an analytical expression for the power spectrum of a spherical cap function. For simplicity sake, we define spherical cap function, $F_{\theta_0}(\theta)$, as a circularly symmetric spherical cap centered at the north pole as a function of the colatitude $\theta \in [0, \pi]$:

$$F_{\theta_0}(\theta) = \begin{cases} 1 & \text{if } \theta \leq \theta_0 \\ 0 & \text{otherwise} \end{cases}, \quad (\text{C.1})$$

where, $\epsilon < \theta_0 < \pi - \epsilon$, is a constant that controls the size of the spherical cap, for $\epsilon > 0$. Since F_{θ_0} is circularly symmetric, we can consider only the zonal ($m = 0$) harmonic basis functions. Hereafter, we drop the subscript θ_0 from F_{θ_0} to denote spherical cap function and express it as F . To obtain spectral coefficients

($\mathcal{S}_F(l, 0)$) of the spherical cap function, we simply take its inner product—denoted by $\langle \cdot, \cdot \rangle$ —with the zonal components of the spherical harmonics:

$$\mathcal{S}_F(l, 0) = \langle F_{\theta_0}(\theta), Y_l(\theta, \phi) \rangle \quad (\text{C.2})$$

$$= \int_{\theta=0}^{\pi} \int_{\phi=0}^{2\pi} F_{\theta_0}(\theta) Y_l(\theta, \phi) \sin \theta d\phi d\theta \quad (\text{C.3})$$

$$= \int_{\theta=0}^{\theta_0} \int_{\phi=0}^{2\pi} Y_l(\theta, \phi) \sin \theta d\phi d\theta \quad (\text{C.4})$$

$$= \sqrt{2l+1} \int_{\theta=0}^{\theta_0} \int_{\phi=0}^{2\pi} P_l^0(\theta, \phi) \sin \theta d\phi d\theta \quad (\text{C.5})$$

$$= \sqrt{2l+1} \int_{\theta=0}^{\theta_0} \int_{\phi=0}^{2\pi} P_l^0(\cos \theta) \sin \theta d\phi d\theta \quad (\text{C.6})$$

$$= \sqrt{2l+1} \int_{\theta=0}^{\theta_0} \left(\int_{\phi=0}^{2\pi} d\phi \right) P_l^0(\cos \theta) \sin \theta d\theta \quad (\text{C.7})$$

$$= \sqrt{2l+1} \int_{\theta=0}^{\theta_0} 2\pi P_l^0(\cos \theta) \sin \theta d\theta \quad (\text{C.8})$$

$$= 2\pi \sqrt{2l+1} \int_{\theta=0}^{\theta_0} P_l^0(\cos \theta) \sin \theta d\theta. \quad (\text{C.9})$$

Legendre polynomials satisfy:

$$(2l+1)P_l^0(x) = \frac{d}{dx} [P_{l+1}^0(x) - P_{l-1}^0(x)]. \quad (\text{C.10})$$

Substituting Eq. (C.10) in Eq. (C.9), we get:

$$\mathcal{S}_F(l, 0) = 2\pi \left(\frac{P_{l-1}^0(\cos(\theta_0)) - P_{l+1}^0(\cos \theta_0)}{\sqrt{2l+1}} \right), \quad (\text{C.11})$$

which is the expression for the spherical cap harmonic coefficients $\mathcal{S}_F(l, 0)$ in terms of l and θ_0 . As we have shown in Chapter 5, the variance in numerical integration for MC integration is directly related to the $\|\mathcal{S}_F(l, 0)\|^2$ of the integrand.

Here, our integrand is a spherical cap function and the corresponding $\|\mathcal{S}_F(l, 0)\|^2$ is given by:

$$\|\mathcal{S}_F(l, 0)\|^2 = 4\pi^2 \left(\frac{|P_{l-1}^0(\cos(\theta_0)) - P_{l+1}^0(\cos \theta_0)|^2}{2l + 1} \right), \quad (\text{C.12})$$

We use Eq. (C.12) in the *angular* power spectrum definition to numerically compute the bounds on the associated variance of the spherical cap for various state-of-the-art sampling patterns discussed in this dissertation. We would like to set an upper bound on $\|\mathcal{S}_F(l, 0)\|^2$ given in Eq. (C.12). This expression contains the shifted copies (with respect to l) of the associated Legendre Polynomial $P_l(x) := P_l^0(x)$. We can first upper bound the term $|P_{l-1}(\cos \theta_0) - P_{l+1}(\cos \theta_0)|$. Several sharp estimations can be found in the literature for the Legendre polynomial $P_l(x)$. A classical result for $x \in]-1, 1[$ and $l \in \mathbb{N}$, is the improved version of *Bernstein's inequality*, [62, 63], given by:

$$|P_l(x)| < \sqrt{\frac{2}{\pi(l + 1/2)}} \frac{1}{(1 - x^2)^{\frac{1}{4}}}. \quad (\text{C.13})$$

For $x = \cos \theta_0$, this inequality becomes:

$$|P_l(\cos \theta_0)| < \sqrt{\frac{2}{\pi(l + 1/2)}} \frac{1}{\sqrt{\sin \theta_0}} = \frac{2}{\sqrt{(2l + 1)\pi \sin \theta_0}}. \quad (\text{C.14})$$

We treat $l = 0$ as a special case. For $l > 0$, using the above inequality in Eq. (C.14) we can derive:

$$|P_{l-1}(\cos \theta_0)| < \frac{2}{\sqrt{(2l - 1)\pi \sin \theta_0}}, \quad (\text{C.15})$$

and,

$$|P_{l+1}(\cos \theta_0)| < \frac{2}{\sqrt{(2l + 3)\pi \sin \theta_0}} < \frac{2}{\sqrt{(2l - 1)\pi \sin \theta_0}}, \quad (\text{C.16})$$

without any loss of generality. From triangular inequality, we get:

$$|P_{l-1}(\cos \theta_0) - P_{l+1}(\cos \theta_0)| < |P_{l-1}(\cos \theta_0)| + |P_{l+1}(\cos \theta_0)|. \quad (\text{C.17})$$

Using Eq. (C.15) and Eq. (C.16) in the above expression, gives us:

$$|P_{l-1}(\cos \theta_0) - P_{l+1}(\cos \theta_0)| < \frac{4}{\sqrt{(2l-1)\pi \sin \theta_0}}. \quad (\text{C.18})$$

For $l > 0$, both left and right hand side of Eq. (C.18) are positive. Therefore, squaring followed by a division with $2l + 1 \geq 1$ on both sides, would not affect the inequality:

$$\frac{|P_{l-1}(\cos \theta_0) - P_{l+1}(\cos \theta_0)|^2}{2l + 1} < \frac{16}{(2l + 1)(2l - 1)\pi \sin \theta_0}. \quad (\text{C.19})$$

Plugging this in Eq. (C.12), gives us the upper bound on the power spectrum of the spherical cap function for any θ_0 and $l > 0$, which can be written as:

$$\boxed{\|\mathcal{S}_F(l, 0)\|^2 < \frac{64\pi}{\sin \theta_0} \left(\frac{1}{(2l + 1)(2l - 1)} \right)}. \quad (\text{C.20})$$

For $l = 0$: Associated Legendre polynomials have a property according to which $P_{-1}(x) = P_0(x)$, that gives:

$$|P_{l-1}(\cos \theta_0)| = |P_{-1}(\cos \theta_0)| = |P_0(\cos \theta_0)| < \frac{2}{\sqrt{\pi \sin \theta_0}}, \quad (\text{C.21})$$

$$|P_{l+1}(\cos \theta_0)| = |P_1(\cos \theta_0)| < \frac{2}{\sqrt{3\pi \sin \theta_0}} < \frac{2}{\sqrt{\pi \sin \theta_0}}, \quad (\text{C.22})$$

$$|P_0(\cos \theta_0) - P_1(\cos \theta_0)| < \frac{4}{\sqrt{\pi \sin \theta_0}}, \quad (\text{C.23})$$




Figure C.1: We illustrate the power profile (Power vs l) of a spherical cap function (in blue) for $\theta_0 = 60$ with the corresponding bound (in orange). We show three plots of the same power spectrum (a), with different zoom-in plots, (b) and (c), to better see the upper bound and to show that the power of the spherical cap is well bounded by the upper bound we derived in Sec. C.2.

$$|P_0(\cos \theta_0) - P_1(\cos \theta_0)|^2 < \frac{16}{\pi \sin \theta_0}, \quad (\text{C.24})$$

using the above inequality in Eq. (C.12), for $l = 0$, gives us the upper bound on:

$$\boxed{\|\mathcal{S}_F(0, 0)\|^2 < \frac{64\pi}{\sin \theta_0}}. \quad (\text{C.25})$$

The inequality in Eq. (C.20) does not guarantee a (very) tight bound on the power of the spherical cap coefficients. However, for our analysis we found that this bound can still be used to conjecture the rate of convergence for the worst-case error in integration.

C.2 Spherical cap power spectrum

To study the worst-case, we are looking for a function with squared-norm zonal spectral coefficients, $\|\mathcal{S}_F(l, 0)\|^2$, exhibiting a decay rate of order $O(l^{-2})$, as derived by Brandolini in [9]. We have shown in the previous subsection C.1, that a spherical cap function has this behaviour (Eq. (C.20)). The corresponding *angular*

power spectrum decay rate can be obtained using Eq. (C.20) in the definition of the *angular power spectrum*. The resulting decay rate of the angular power spectrum for a spherical cap is given by:

$$\boxed{\mathcal{P}_F(l) < \frac{C_0}{(2l+1)^2(2l-1)}}, \quad (\text{C.26})$$

where $C_0 = 64\pi/(\sin \theta_0)$. In Fig. C.1, we show an illustration of this upper bound on a spherical cap function of size $\theta_0 = 60$. We can derive a corresponding worst-case variance convergence rate using Eq. (C.26) given by:

$$\text{Var}(\mathbf{I}_N) < \frac{\mu(\mathcal{S}^2)^2 C_0}{N} \sum_{l=0}^{\infty} \frac{\check{\mathcal{P}}_{\mathbf{S}}(l)}{(2l+1)(2l-1)}. \quad (\text{C.27})$$

This variance bound can be used to study different shapes of the sampling power spectra ($\check{\mathcal{P}}_{\mathbf{S}}(l)$) to derive their convergence rate with respect to the number of samples N .

Bibliography

- [1] P. Alliez, G. Ucelli, C. Gotsman, and M. Attene. Recent advances in remeshing of surfaces. In *Shape Analysis and Structuring, Mathematics and Visualization*. Springer, 2008.
- [2] P. Alliez, E.C. Verdière, O. Devillers, and M. Isenburg. Isotropic surface remeshing. In *Proceedings of the Shape Modeling International 2003, SMI '03*, pages 49–, Washington, DC, USA, 2003.
- [3] J. Arvo. Stratified sampling of spherical triangles. In *Proc. SIGGRAPH '95*, pages 437–438. ACM, 1995.
- [4] J. Arvo. Stratified sampling of 2-manifolds. *SIGGRAPH 2001 Course Notes*, 29(2), 2001.
- [5] M. Balzer and D. Heck. Capacity-constrained Voronoi diagrams in finite spaces. In *Proceedings of the 5th Annual International Symposium on Voronoi Diagrams in Science and Engineering*, volume 2, pages 44–56, 2008.
- [6] M. Balzer, T. Schlömer, and O. Deussen. Capacity-constrained point distributions: A variant of Lloyd’s method. *ACM Trans. on Graphics*, 28(3):86:1–8, 2009.
- [7] M.S. Bartlett. *An Introduction to Stochastic Processes, with Special Reference to Methods and Applications*. Cambridge University Press, 1978.

- [8] J. Bowers, R. Wang, L-Y. Wei, and D. Maletz. Parallel Poisson disk sampling with spectrum analysis on surfaces. In *Proc. SIGGRAPH Asia '10*, pages 166:1–166:10. ACM, 2010.
- [9] L. Brandolini, L. Colzani, and A. Torlaschi. Mean square decay of Fourier transforms in Euclidean and non Euclidean spaces. *Tohoku Math. J. (2)*, 53(3):467–478, 2001.
- [10] J.S. Brauchart and P.J. Grabner. Distributing many points on spheres: minimal energy and designs. *ArXiv e-prints*, July 2014.
- [11] R. Bridson. Fast Poisson disk sampling in arbitrary dimensions. In *Proc. SIGGRAPH '07 Sketches*. ACM, 2007.
- [12] Z. Chen, Z. Yuan, Y.K. Choi, L. Liu, and W. Wang. Variational blue noise sampling. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1784–1796, October 2012.
- [13] C. Choirat and R. Seri. Computational aspects of Cui-Freeden statistics for equidistribution on the sphere. *Mathematics of Computation*, 82(284):2137–2156, 2013.
- [14] D. Cline, S. Jeschke, K. White, A. Razdan, and P. Wonka. Dart throwing on surfaces. In *Proc. EGSR '09*, pages 1217–1226. Eurographics Association, 2009.
- [15] M. F. Cohen, J. Shade, S. Hiller, and O. Deussen. Wang tiles for image and texture generation. In *ACM SIGGRAPH 2003 Papers*, pages 287–294, New York, USA, 2003.
- [16] R.L. Cook. Stochastic sampling in computer graphics. *ACM Trans. Graph.*, 5(1):51–72, 1986.

- [17] R. Cranley and T.N.L. Patterson. Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis*, 13(6):904–914, 1976.
- [18] F.C. Crow. The aliasing problem in computer-generated shaded images. *Commun. ACM*, 20(11):799–805, 1977.
- [19] J. Cui and W. Freeden. Equidistribution on the sphere. *SIAM Scientific Computing*, 18(2):595–609, 1997.
- [20] D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes: volume I: elementary Theory and Methods*, volume 1. Springer-Verlag New York, 2003.
- [21] D. J. Daley and D. Vere-Jones. *An introduction to the theory of point processes: volume II: general theory and structure*, volume 2. Springer-Verlag New York, 2007.
- [22] F. de Goes, K. Breeden, V. Ostromoukhov, and M. Desbrun. Blue noise through optimal transport. *Proc. SIGGRAPH Asia '12*, 31:171:1–171:10, 2012.
- [23] O. Deussen, Stefan Hiller, Cornelius van Overveld, and Thomas Strothotte. Floating points: A method for computing stipple drawings. *Computer Graphics Forum*, 19:40–51, 2000.
- [24] I. T. Dimov, A. A. Penzov, and S. S. Stoilova. Parallel Monte Carlo sampling scheme for sphere and hemisphere. In *Proceedings of the 6th International Conference on Numerical Methods and Applications*, NMA'06, pages 148–155, Berlin, Heidelberg, 2007.
- [25] M.A.Z. Dippé and E.H. Wold. Antialiasing through stochastic sampling. In *Proc. SIGGRAPH '85*, pages 69–78. ACM, 1985.

- [26] J.R. Driscoll and D.M. Healy. Computing Fourier Transforms and Convolutions on the 2-Sphere. *Advances in Applied Mathematics*, 15(2):202 – 250, 1994.
- [27] D. Dunbar and G. Humphreys. A Spatial Data Structure for Fast Poisson-disk Sample Generation. In *ACM SIGGRAPH 2006*, pages 503–508, New York, USA, 2006.
- [28] F. Durand. A frequency analysis of Monte-Carlo and other numerical integration schemes. *MIT CSAIL Technical report TR-2011-052*, 2011.
- [29] M.S. Ebeida, A.A. Davidson, A. Patney, P.M. Knupp, S.A. Mitchell, and J.D. Owens. Efficient Maximal Poisson-disk Sampling. In *ACM SIGGRAPH 2011*, pages 49:1–49:12, New York, USA, 2011.
- [30] R. Fattal. Blue-noise point sampling using kernel density model. *ACM SIGGRAPH 2011 papers*, 28(3):1–10, 2011.
- [31] L. Feng, I. Hotz, B. Hamann, and K.I. Joy. Anisotropic noise samples. *Visualization and Computer Graphics, IEEE Transactions on*, 14(2):342–354, March 2008.
- [32] R. Floyd and L. Steinberg. An adaptive algorithm for spatial gray scale. pages 36–37, 1975.
- [33] W. Fulton and J. Harris. *Representation Theory: A First Course*. Springer-Verlag, New York, 1991.
- [34] A. Gabrielli and S. Torquato. Voronoi and void statistics for superhomogeneous point processes. *Physical Review E*, 70(4):041105, 2004.
- [35] M.N. Gamito and S.C. Maddock. Accurate multi-Dimensional Poisson-disk sampling. *ACM Trans. on Graphics*, 29(1):8, 2009.
- [36] Martin Gardner. White and brown music, fractal curves and one-over-f fluctuation. pages 16–32, 1978.

- [37] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal*, 622(2):759, 2005.
- [38] H. Groemer. *Geometric Applications of Fourier Series and Spherical Harmonics*. Cambridge University Press, 1996. Cambridge Books Online.
- [39] B. Grünbaum and G.C. Shephard. *Tilings and Patterns*. W. H. Freeman & Co., New York, NY, USA, 1986.
- [40] J-P. Hansen and I.R. McDonald. *Theory of simple liquids*. Elsevier, 1990.
- [41] D. Heck, T. Schlömer, and O. Deussen. Blue noise sampling with controlled aliasing. *ACM Trans. on Graphics*, 32(3):25:1–25:12, 2013.
- [42] I.N. Herstein. *Abstract Algebra*. Wiley, 1996.
- [43] F.J. Hickernell. *Koksma-Hlawka Inequality*. John Wiley & Sons, Ltd, 2014.
- [44] S. Hiller, O. Deussen, and A. Keller. Tiled blue noise samples. In *Proceedings of the Vision Modeling and Visualization Conference 2001*, pages 265–272. Aka GmbH, 2001.
- [45] R.G. Hipkin. The statistics of pink noise on a sphere: applications to mantle density anomalies. *Geophysical Journal International*, 144(2):259–270, 2001.
- [46] K. Hoffman and R.A. Kunze. *Linear Algebra*. Prentice-Hall mathematics series. Prentice-Hall, 1971.
- [47] T.R. Jones. Efficient generation of Poisson-disk sampling patterns. *Journal of graphics, gpu, and game tools*, 11(2):27–36, 2006.

- [48] W.M. Kaula. Theory of statistical analysis of data distributed over a sphere. *Reviews of Geophysics*, 5(1):83–107, 1967.
- [49] A. Keller, S. Premoze, and M. Raab. Advanced (quasi) Monte Carlo methods for image synthesis. In *SIGGRAPH '12 Courses*, pages 21:1–21:46, New York, USA, 2012.
- [50] Johannes Kopf, Daniel Cohen-Or, Oliver Deussen, and Dani Lischinski. Recursive wang tiles for real-time blue noise. In *ACM SIGGRAPH 2006 Papers*, pages 509–518, New York, NY, USA, 2006.
- [51] A. Lagae and P. Dutré. Poisson Sphere Distributions. In L. Kobbelt, T. Kuhlen, T. Aach, and R. Westermann, editors, *Vision, Modeling, and Visualization 2006*, pages 373–379, Berlin, November 2006.
- [52] A. Lagae and P. Dutré. A Comparison of Methods for Generating Poisson Disk Distributions. *Computer Graphics Forum*, 27(1):114–129, 2008.
- [53] Ares Lagae and Philip Dutré. A Procedural Object Distribution Function. *ACM Trans. Graph.*, 24(4):1442–1461, October 2005.
- [54] Ares Lagae and Philip Dutré. An Alternative for Wang Tiles: Colored Edges Versus Colored Corners. *ACM Trans. Graph.*, 25(4):1442–1459, October 2006.
- [55] C. Lemieux. *Monte Carlo and Quasi Monte Carlo Sampling*. Springer, 2009.
- [56] B. Levy. Laplace-Beltrami Eigenfunctions towards an Algorithm that "Understands" Geometry. In *Proceedings of the IEEE International Conference on Shape Modeling and Applications 2006, SMI '06*, pages 13–, Washington, DC, USA, 2006.
- [57] H. Li, K.Y. Lo, M.K. Leung, and C.W. Fu. Dual Poisson-Disk Tiling: An Efficient Method for Distributing Features on Arbitrary Surfaces. *IEEE*

Transactions on Visualization and Computer Graphics, 14(5):982–998, September 2008.

- [58] H. Li, D. Nehab, L-Y. Wei, P.V. Sander, and C.W. Fu. Fast capacity constrained Voronoi tessellation. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D '10*, pages 13:1–13:1, New York, NY, USA, 2010.
- [59] H. Li, L.Y. Wei, P.V. Sander, and C.W. Fu. Anisotropic blue noise sampling. In *ACM SIGGRAPH Asia 2010 Papers*, pages 167:1–167:12, New York, NY, USA, 2010.
- [60] S.A. Lloyd. Least squares quantization in PCM. *Information Theory, IEEE Transactions on*, 28(2):129–137, 1982.
- [61] S.A. Lloyd. An optimization approach to relaxation labelling algorithms. *Image and Vision Computing*, 1(2):85–91, 1983.
- [62] G. Lohöfer. Inequalities for the associated Legendre functions. *J. Approx. Theory*, 95(2):178–193, 1998.
- [63] L. Lorch. Alternative proof of a sharpened form of Bernstein’s inequality for Legendre polynomials. *Applicable Analysis*, 14(3):237–240, 1983.
- [64] F.J. Lowes. Spatial power spectrum of the main geomagnetic field, and extrapolation to the core. *Geophysical Journal International*, 36(3):717–730, 1974.
- [65] R. Marques, C. Bouville, M. Ribardière, L. P. Santos, and K. Bouatouch. Spherical Fibonacci point sets for illumination integrals. *Computer Graphics Forum*, 32(8):134–143, 2013.
- [66] Michael McCool and Eugene Fiume. Hierarchical poisson disk sampling distributions. In *Proceedings of the Conference on Graphics Interface '92*, pages 94–105, San Francisco, CA, USA, 1992.

- [67] J.D. McEwen and Y. Wiaux. A novel sampling theorem on the sphere. *Signal Processing, IEEE Trans. on*, 59(12):5876–5887, 2011.
- [68] D.P. Mitchell. Generating antialiased images at low sampling densities. In *Proc. SIGGRAPH '87*, pages 65–72, 1987.
- [69] D.P. Mitchell. Spectrally optimal sampling for distributed ray tracing. In *Proc. SIGGRAPH '91*, volume 25, pages 157–164, 1991.
- [70] D.P. Mitchell and A.N. Netravali. Reconstruction filters in computer-graphics. *ACM SIGGRAPH '88*, pages 221–228, New York, NY, USA, 1988.
- [71] H. Niederreiter. *Random Number Generation and Quasi-Monte-Carlo Methods*. SIAM, 1992.
- [72] A. Okabe, B. Boots, and K. Sugihara. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, Inc., New York, NY, USA, 1992.
- [73] V. Ostromoukhov. Sampling with polyominoes. In *ACM SIGGRAPH 2007 Papers*, New York, NY, USA, 2007.
- [74] V. Ostromoukhov, C. Donohue, and P.M. Jodoin. Fast hierarchical importance sampling with blue noise properties. In *ACM Trans. on Graphics*, volume 23, pages 488–495, 2004.
- [75] A.C. Öztireli and M. Gross. Analysis and synthesis of point distributions based on pair correlation. *ACM Trans.on Graphics*, 31(6):170:1–170:10, 2012.
- [76] K. Parker, T. Mitsa, and R. Ulichney. A new algorithm for manipulating the power spectrum of halftone patterns. pages 471–475, 1991.

- [77] J.L. Peyrot, F. Payan, and M. Antonini. Feature-preserving direct blue noise sampling for surface meshes. In *Eurographics 2013*, page 4 pages, May 2013.
- [78] A. Pilleboue. *Analyse spatiale et spectrale des motifs d'échantillonnage pour l'intégration Monte Carlo*. PhD thesis, Université Claude Bernard Lyon 1, 10 2015.
- [79] A. Pilleboue, G. Singh, D. Coeurjolly, M. Kazhdan, and V. Ostromoukhov. Variance analysis for Monte Carlo integration. *ACM Trans. Graph. (Proc. SIGGRAPH)*, 34(4), 2015.
- [80] E.A. Rakhmanov, E.B. Saff, and Y.M. Zhou. Minimal discrete energy on the sphere. *Mathematical Research Letters*, page 647?662, 1994.
- [81] R. Ramamoorthi, J. Anderson, M. Meyer, and D. Nowrouzezahrai. A theory of Monte Carlo visibility sampling. *ACM Trans. on Graphics*, 31(5):121:1–121:16, 2012.
- [82] B.D. Ripley. Modelling spatial patterns. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 172–212, 1977.
- [83] E.B. Saff and A.B.J. Kuijlaars. Distributing many points on a sphere. *The Mathematical Intelligencer*, 19(1):5–11, 1997.
- [84] D. Sakrison. On the role of the observer and a distortion measure in image transmission. *Communications, IEEE Transactions on*, 25(11):1251–1267, Nov 1977.
- [85] T. Schlömer, D. Heck, and O. Deussen. Farthest-point optimized point sets with maximized minimum distance. In *Proc. Symp. High Performance Graphics '11*, pages 135–142. ACM, 2011.
- [86] C. Schmaltz, P. Gwosdek, A. Bruhn, and J. Weickert. Electrostatic halftoning. *Computer Graphics Forum*, 29(8):2313–2327, 2010.

- [87] A. Secord. Weighted Voronoi Stippling. In *Proceedings of the 2nd ACM International Symposium on Non-photorealistic Animation and Rendering*, NPAR '02, pages 37–43, New York, NY, USA, 2002.
- [88] J-P. Serre. *Linear representations of finite groups*. Springer-Verlag, New York, 1977.
- [89] M.Z. Shao and N. Badler. Spherical sampling by Archimedes' Theorem. *Technical Report, University of Pennsylvania*, 1996.
- [90] P. Shirley and K. Chiu. A low distortion map between disk and square. *J. Graph. Tools*, 2(3):45–52, December 1997.
- [91] P.S. Shirley. Discrepancy as a quality measure for sample distributions. In *Proc. Eurographics '91*, pages 183–194, 1991.
- [92] K. Subr and J. Kautz. Fourier analysis of stochastic sampling strategies for assessing bias and variance in integration. *ACM Trans. on Graphics*, 32(4):128:1–128:12, 2013.
- [93] A.G. Sukharev. Optimal strategies of the search for an extremum. *{USSR} Computational Mathematics and Mathematical Physics*, 11(4):119 – 137, 1971.
- [94] S. Torquato, O.U. Uche, and F.H. Stillinger. Random sequential addition of hard spheres in high euclidean dimensions. *Physical Review E*, 74(6):061308, 2006.
- [95] G. Turk. Re-tiling polygonal surfaces. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques*, ACM SIG-GRAPH '92, pages 55–64, New York, NY, USA, 1992.
- [96] R. Ulichney. *Digital Halftoning*. MIT Press, 1987.
- [97] R.A. Ulichney. Dithering with blue noise. *Proceedings of the IEEE*, 76(1):56–79, Jan 1988.

- [98] C. Ureña, M. Fajardo, and A. King. An area-preserving parametrization for spherical rectangles. *Computer Graphics Forum*, 32(4):59–66, 2013.
- [99] F. Wachtel, A. Pilleboue, D. Coeurjolly, K. Breeden, G. Singh, G. Cathelin, F. de Goes, M. Desbrun, and V. Ostromoukhov. Fast Tile-based Adaptive Sampling with User-specified Fourier Spectra. *ACM Trans. Graph.*, 33(4):56:1–56:11, July 2014.
- [100] L-Y. Wei. Parallel Poisson Disk sampling. In *ACM SIGGRAPH 2008 Papers*, pages 20:1–20:9, New York, NY, USA, 2008.
- [101] L-Y. Wei. Multi-class blue noise sampling. In *ACM SIGGRAPH 2010 Papers*, pages 79:1–79:8, New York, NY, USA, 2010.
- [102] L-Y. Wei and R. Wang. Differential domain analysis for non-uniform sampling. In *ACM SIGGRAPH 2011 Papers*, pages 50:1–50:10, New York, NY, USA, 2011.
- [103] K.B. White, D. Cline, and P.K. Egbert. Poisson disk point sets by hierarchical dart throwing. In *Interactive Ray Tracing, 2007. RT '07. IEEE Symposium on*, pages 129–132, Sept 2007.
- [104] M.A. Wiecek and F.J. Simons. Localized spectral analysis on the sphere. *Geophysical Journal International*, 162(3):655–675, 2005.
- [105] Y. Xu, R. Hu, C. Gotsman, and L. Liu. Blue noise sampling of surfaces. *Computers & Graphics*, 36(4):232–240, 2012.
- [106] Y. Xu, L. Liu, C. Gotsman, and S.J. Gortler. Capacity-constrained Delaunay triangulation for point distributions. *Computers & Graphics*, 35(3):510–516, 2011. Shape Modeling International (SMI) Conference 2011.
- [107] D.M. Yan, J.W. Guo, B. Wang, X.-P. Zhang, and P. Wonka. A survey of blue-noise sampling and its applications. *Journal of Computer Science and Technology*, 30(3):439–452, 2015.

- [108] J.I. Yellott. Spectral consequences of photoreceptor sampling in the rhesus retina. *Science*, 221(4608):382–385, 1983.
- [109] A. Yershova and S.M. LaValle. Deterministic sampling methods for spheres and $SO(3)$. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, volume 4, pages 3974–3980 Vol.4, April 2004.
- [110] Y. Zhou, H. Huang, L-Y. Wei, and R. Wang. Point sampling with general noise spectrum. *ACM Trans. on Graphics*, 31(4):76:1–76:11, July 2012.
- [111] V.K. Zworykin and G.A. Morton. *Television: The Electronics of Image Transmission in Color and Monochrome*. John Wiley & Sons, 1954.