

Editors:

Timothy J. Barth
Michael Griebel
David E. Keyes
Risto M. Nieminen
Dirk Roose
Tamar Schlick

Markus Holtz

Sparse Grid Quadrature in High Dimensions with Applications in Finance and Insurance

 Springer

Markus Holtz
Universität Bonn
Institut für Numerische
Simulation
Wegelerstr. 5
53115 Bonn
Germany
holtz@ins.uni-bonn.de

ISSN 1439-7358
ISBN 978-3-642-16003-5 e-ISBN 978-3-642-16004-2
DOI 10.1007/978-3-642-16004-2
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2010936374

Mathematics Subject Classification (2010): 65-02

© Springer-Verlag Berlin Heidelberg 2011

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Cover design: deblik, Berlin

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

This text deals with the numerical analysis and efficient numerical treatment of high-dimensional integrals using sparse grids and related dimension-wise integration techniques and discusses applications to finance and insurance. The focus on sparse grid quadrature methods on the one hand is thereby complemented by larger parts of this text which are devoted to dimension-wise decompositions of high-dimensional functions, such as the ANOVA or the anchored-ANOVA, on the other hand. The main intention is to cover these two areas of research, which have attracted independently of each other considerable attention in recent years, in a self-contained, easily accessible and unified way. In particular the interplay between the convergence behaviour of sparse grid methods with effective dimensions (which are a measure for the decay of the dimension-wise decompositions) and with coordinate transformations (which aim to improve the decay of the dimension-wise decompositions) is studied. The text moreover aims to investigate potential benefits but also limitations of these techniques with respect to applications from mathematical finance and insurance and to give some recommendations based on the theoretical and numerical results presented in the manuscript.

This manuscript mainly originated during my time from July 2004 to January 2009 at the Institute for Numerical Simulation at the University of Bonn where I worked in the area of sparse grid methods and high-dimensional integration problems. I had the chance to participate in several research projects, in which we investigated, partly in close collaboration with financial institutions, the use of these methods for applications from financial engineering and insurance. In addition I was involved in the years 2004–2008 in the teaching of the laboratory "Computational Finance" that was offered for master students with special focus on the computational aspects of option pricing problems from mathematical finance. This manuscript summarizes material and results which have been developed during these different research projects and teaching activities in this period of time.

I thank all people who contributed to this work. First of all, I would like to thank Michael Griebel for his continuous support over the years and then Thomas Gerstner for a very pleasant collaboration and many contributions to this manuscript.

I furthermore thank Ralf Goschnick and Marcus Haep from the Zurich Group Germany for their support with the development of asset-liability management models for life insurance policies, Ralf Korn from the University of Kaiserslautern for the introduction to the interesting problem to value performance-dependent options as well as Angela Kunoth from the University of Paderborn, Jan Baldeaux, Frances Kuo and Ian Sloan from the University of New South Wales and Christoph Schwab from the ETH Zurich for valuable comments and discussions. I would also like to thank all my colleagues Jutta Adelsberger, Jürgen Braun, Roberto Croce, Niclas Doll, Martin Engels, Christian Feuersänger, Jan Hamaekers, Frederik Heber, Alexander Hullmann, Lukas Jäger, Margit Klitz, Christina Kürten, Bram Metsch, Daniel Oeltz, Jens Oettershagen, Tim Osadnik, Eva-Maria Rau, Marc Schweitzer and Ralf Wildenhues from the Institute for Numerical Simulation of the University of Bonn who all contributed in some way to this manuscript. Finally, I appreciate the financial support from the program "Mathematics for innovations in industry and service" funded by the German Federal Ministry of Education and Research (BMBF) and from the "Sonderforschungsbereich 611" funded by the Deutsche Forschungsgemeinschaft (DFG).

Freiburg, July 2010

Markus Holtz

Contents

1	Introduction	1
2	Dimension-wise Decompositions	11
2.1	Classical ANOVA Decomposition	13
2.1.1	Effective Dimensions	16
2.1.2	Error Bounds	19
2.2	Anchored-ANOVA Decomposition	21
2.2.1	Effective Dimensions	22
2.2.2	Error Bounds	26
3	Dimension-wise Quadrature	29
3.1	Classical Multivariate Quadrature Methods	29
3.1.1	Monte Carlo	31
3.1.2	Quasi-Monte Carlo	34
3.1.3	Product Methods	38
3.2	Dimension-wise Quadrature Methods	43
3.2.1	Truncation and Discretization	43
3.2.2	Error and Costs	45
3.2.3	A priori Construction	47
3.2.4	Dimension-adaptive Construction	49
4	Sparse Grid Quadrature	51
4.1	Sparse Grid Methods	51
4.1.1	Classical Construction	52
4.1.2	Delayed Basis Sequences	55
4.1.3	Generalised Sparse Grids	59
4.1.4	Dimension-adaptive Sparse Grids	61
4.2	Optimal Sparse Grids in Weighted Spaces	63
4.2.1	Cost-Benefit Ratio	64
4.2.2	Cost Analysis	67
4.2.3	Error Analysis	69

4.2.4	Analysis of Error versus Cost	71
4.3	Relation to Dimension-wise Quadrature	74
5	Dimension Reduction and Smoothing	77
5.1	Dimension Reduction	77
5.1.1	Random Walk, Brownian Bridge, PCA	78
5.1.2	Linear transformations	82
5.2	Domain Decomposition	89
5.2.1	Root Finding	89
5.2.2	Hyperplane Arrangements	90
5.2.3	Conditional Sampling	98
6	Validation and Applications	101
6.1	Interest Rates Derivatives	104
6.1.1	Zero Coupon Bonds	104
6.1.2	Collateralized Mortgage Obligations	109
6.2	Path-dependent Options	114
6.2.1	Asian options	115
6.2.2	Barrier options	120
6.3	Performance-dependent Options	125
6.3.1	Framework and Pricing Formulas	125
6.3.2	Numerical Results	131
6.4	Asset-Liability Management in Life Insurance	136
6.4.1	Model and Integral Representation	136
6.4.2	Numerical Results	145
6.5	Summary and Discussion	149
7	Summary and Conclusions	153
A	Discrepancy and Tractability	157
A.1	Reproducing Kernel Hilbert Spaces	157
A.2	Notions of Discrepancy	161
A.3	Tractability Results	167
	References	172
Index	179

Chapter 1

Introduction

Financial institutions have to understand the risks that their financial instruments create as precisely as possible. To this end, mathematical models are developed which are usually based on tools from stochastic calculus. Most of the models are too complex to be analytically tractable and are hence analysed with the help of computer simulations which rely on efficient algorithms from scientific computing.

An important example is the quantitative, i.e. model-based, pricing of financial derivatives. Derivatives are financial instruments whose values are derived from the value of one or several underlying assets such as stocks, interest rates or commodities. A fundamental result from mathematical finance is that, under certain model assumptions, the prices of derivatives can be represented as expected values which in turn correspond to high-dimensional integrals

$$I_d = \int_{\mathbb{R}^d} g(\mathbf{z}) \varphi(\mathbf{z}) d\mathbf{z} \quad (1.1)$$

over the d -dimensional Euclidean space with the Gaussian weight function φ and $\mathbf{z} := (z_1, \dots, z_d)$. Alternatively, after a suitable transformation, high-dimensional integrals

$$I_d = \int_{[0,1]^d} f(\mathbf{x}) d\mathbf{x} \quad (1.2)$$

over the unit cube can be obtained. The dimension d depends on the number of sources of uncertainty respected by the model assumptions and is the larger the more random variables are involved. This way, high-dimensional integrals in hundreds of variables appear in many applications from finance. Since the integrals can in most cases not be calculated analytically, they have to be computed numerically up to a prescribed accuracy ε . Today, in 2010, more than 10% of the most powerful commercially available computer systems worldwide is owned by financial institutions and used for such purposes [106]. This share has more than doubled compared to the share in 2006 and there are by now more supercomputers working for finance than for, e.g., weather and climate research, defense or geophysics.

For an efficient computation of high-dimensional integrals, one of the key prerequisites is that the *curse of dimension* [7] can be avoided at least to some extent. The curse of dimension states that the cost to compute an approximation with a prescribed accuracy ε depends exponentially on the dimension d of the problem. This is one of the main obstacles for a conventional numerical treatment of high dimensional problems. Classical product quadrature methods for the computation of multivariate integrals [25] achieve with n evaluations of the integrand an accuracy of

$$\varepsilon(n) = O(n^{-r/d})$$

for functions with bounded derivatives up to order r . For fixed r , their convergence rates r/d thus deteriorate with increasing dimension and are already in moderate dimensions so small that high accuracies can no longer be obtained in practise. On the positive side, the case $r = d$ indicates that the problem of a high dimension can sometimes be compensated by, e.g., a high degree of smoothness. Also other aspects such as the concentration of measure phenomenon¹ or the superposition theorem of Kolmogorov² show that there is some chance to treat high-dimensional problems despite the curse of dimension. Furthermore, it is known from numerical complexity theory [154] that some algorithm classes can break the curse of dimension for certain function classes.

Randomised methods, so called *Monte Carlo methods*, are the most well-known examples of such classes of algorithms. Here, the integrand is approximated by the average of n function values at random points. Monte Carlo methods were first introduced to derivative pricing by Boyle [10] and are today the workhorses in the financial industry in particular for complex problems which depend on many variables. For square integrable functions f , the expected mean square error³ of the Monte Carlo method with n samples is

$$\varepsilon(n) = O(n^{-1/2}). \tag{1.3}$$

The convergence rate is thus independent of the dimension d , but quite low and a high accuracy is only achievable with a tremendously high number n of function evaluations. This slow convergence of the Monte Carlo method is one of the main reasons for the enormous need of the financial industry for computer resources.

Under more restrictive assumptions on the smoothness of the integrand it can be shown that faster rates of convergence can be attained by deterministic integration methods such as quasi-Monte Carlo methods [55, 97, 114] and sparse grid methods [16, 146, 168]. *Quasi-Monte Carlo methods* are number theoretic algorithms which approximate the integral by the average of n function values at deterministic,

¹ The concentration of measure phenomenon [96] says that every Lipschitz function on a sufficiently high dimensional domain is well approximated by a *constant* function.

² The theorem of Kolmogorov [89] shows that every continuous function of several variables can be represented by the superposition of continuous functions that depend on only *one* variable.

³ The error of a Monte Carlo estimate with n samples is approximately normally distributed with mean zero and standard deviation $\sigma(f)/\sqrt{n}$. Here, the term $n^{-1/2}$ describes the convergence rate and $\sigma^2(f)$ can be considered as the constant of the Monte Carlo method.

uniformly distributed points. For integrands f of bounded variation, their error can be shown to be

$$\varepsilon(n) = O(n^{-1}(\log n)^d), \quad (1.4)$$

see [114]. They thus converge with a rate of almost one, almost independent of the dimension and almost half an order faster than the Monte Carlo method. *Sparse grid methods* are deterministic methods based on polynomial exactness, which are constructed using certain combinations of tensor products of one-dimensional quadrature rules.⁴ In their simplest form [146], they achieve

$$\varepsilon(n) = O(n^{-r}(\log n)^{(d-1)(r+1)}) \quad (1.5)$$

for all integrands which have bounded mixed partial derivatives of order r . Their convergence rate is almost independent of the dimension and increases with higher smoothness r of the integrand. For analytic functions even spectral convergence is observed.

A difficulty in higher dimensions is that quasi-Monte Carlo and sparse grid methods still depend on the dimension d through the logarithmic terms in (1.4) and (1.5). Furthermore, the implicit constants in (1.4) and (1.5) depend on d and often grow exponentially with d . Moreover, it is known from numerical complexity theory that many classes of integration problems are intractable [154] with respect to these deterministic methods meaning that even the best quasi-Monte Carlo or the best sparse grid algorithm can not completely avoid the curse of dimension. For a large dimension d and a small or moderate number n of sample points, the asymptotic advantage of the deterministic numerical methods over the Monte Carlo method might thus not pay off.

Nevertheless, integrals from practise are often in different or smaller problem classes and thus may be tractable. Paskov and Traub [129] indeed observed in 1995 that quasi-Monte Carlo methods converge nearly independent of the dimension and faster than Monte Carlo for a 360-dimensional integration problem which was given to them by the investment bank Goldman Sachs. This empirical observation indicated that the long computing times required by the Monte Carlo method may be avoided by deterministic integration methods even in high dimensions. It initiated intensive research to generalise the result of Paskov and Traub to wider classes of problems and to explain the success of the quasi-Monte Carlo method despite the high dimension. In the following years also sparse grid methods were successfully applied to this particular integration problem [45, 46, 118, 130]. These methods were also clearly superior to Monte Carlo, showing a similar efficiency as quasi-Monte Carlo.

One explanation for the success of the deterministic quadrature methods, which is by now widely accepted, is based on the *analysis of variance* (ANOVA) representation of the integrand [17]. There, the function f on \mathbf{R}^d is decomposed into a sum

⁴ We refer with sparse grid methods to the generalised sparse grid approach [45, 66, 132, 164]. While the classical sparse grid method [146] is uniquely determined by the choice of the underlying univariate quadrature rule, generalised sparse grid methods leave in addition the choice of an underlying index set open.

of 2^d terms, where each term describes the relative importance of a subset of variables with respect to the total variance of f . It turned out that for most integrands from finance the importance of each variable is naturally weighted by certain hidden weights. With increasing dimension, the lower-order terms of the ANOVA decomposition continue to play a significant role, whereas the higher-order terms tend to be negligible [17, 160]. The integrands are of so-called low *effective dimension in the superposition sense* meaning that they can be well approximated by a sum of low-dimensional functions. Moreover, often coordinate transformations (usually interpreted as path generating methods), such as the Brownian bridge [112], can be used to exploit the underlying special structure of the problems from finance and to enforce the importance of the leading dimensions in this way. The corresponding integrands are then of so-called low *effective dimension in the truncation sense* meaning that only few of the variables have a significant impact on the output of the function.

The relation between the effective dimension and the performance of quasi-Monte Carlo methods has been investigated intensively in recent years. While the effective dimension has no impact on Monte Carlo methods, quasi-Monte Carlo methods profit from low effective dimensions in the superposition sense, since low dimensional projections of their points are especially well-distributed. They also profit from low effective dimensions in the truncation sense, since their points are usually better uniformly distributed in smaller dimensions than in higher ones.

Classical sparse grid methods can not utilize low effective dimensions. However, *dimension-adaptive sparse grid methods*, as recently introduced by Gerstner and Griebel [46], take advantage of low effective dimensions in a very general and automatic way by a dimension-adaptive grid refinement.

To describe classes of functions of low effective dimension Sloan and Woźniakowski introduced in [144] *weighted Sobolev spaces* of functions with bounded mixed regularity and proved that there exist quasi-Monte Carlo methods which can avoid the curse of dimension in such spaces, provided the function space weights decay sufficiently fast with growing dimension. Their results that integration is tractable in certain weighted spaces with respect to quasi-Monte Carlo were generalised, also to sparse grid methods [119], and complemented by constructive approaches, e.g. the CBC-construction of lattice rules [145], in a series of papers and are still in the focus of active research.

Today many different settings are known in which integration is tractable with respect to quasi-Monte Carlo and sparse grid methods [119]. However, these settings usually do not apply to applications, since most of them do not satisfy the smoothness assumptions on bounded mixed regularity. This issue was recently addressed in [60, 102] with the argument that the lower order terms in the ANOVA decomposition are in certain cases smoother than the original function. Since the higher order terms are small because of low effective dimension, this may explain the fast convergence of the deterministic methods despite the low regularity of the application problems.

This surveys most of the important results which have been obtained in recent years in the area of research that is concerned with the numerical analysis and with the efficient numerical treatment of high-dimensional integration problems from finance and insurance. It also describes the motivation and the focus of this manuscript which relates several different concepts of this active and emerging area of research. We thereby depart from most of the existing literature in the following two ways:

- We mostly address the arising integrals directly on \mathbf{R}^d and avoid transformations⁵ to the unit cube.
- We base the numerical analysis and the numerical methods mainly on the *anchored-ANOVA decomposition*⁶ instead of the (classical) ANOVA.

These approaches lead to valuable insights into the interplay between coordinate transformations, effective dimensions and the error of sparse grid methods. They also lead to the developments of dimension-wise quadrature methods, which are based on the anchored-ANOVA decomposition. This class of methods is designed to exploit low effective dimension and it includes the class of sparse grid methods as a special case. We study sparse grid methods in detail and discuss several improvements of these methods with respect to their efficiency and with respect to integrands from finance. We derive error bounds for sparse grid methods in weighted spaces and explain why sparse grid methods can profit from low effective dimension and from smoothness of the integrands much more efficiently than other approaches. For illustration we provide many numerical experiments which demonstrate for various applications from finance and insurance that the approaches presented in this manuscript can be faster and more accurate than Monte Carlo and quasi-Monte Carlo methods even for integrands with hundreds of dimensions.

We next describe the contents of this manuscript in more detail.

- Based on the anchored-ANOVA decomposition we define the notions of effective dimension in the classical and in the anchored case, and derive error bounds which relate these dimensions to approximation and integration errors. We determine the effective dimensions in the anchored and in the classical case for several applications from finance with hundreds of dimensions and indicate by theoretical arguments and by numerical experiments that the performance of sparse grid methods can be better explained with the help of the effective dimension in the anchored case than with the classical one.
- We furthermore present the general class of quadrature methods for the computation of high-dimensional integrals, which we refer to as *dimension-wise quadrature*.

⁵ Transformation to the unit cube introduce singularities which deteriorate the efficiency of methods that take advantage of higher smoothness, such as sparse grids.

⁶ The anchored-ANOVA decomposition expresses a multivariate function as superposition of its values on lines, faces, hyperplanes, etc., which intersect a certain anchor point and are parallel to the coordinate axes [135]. Only a finite number of function values is required for its calculation. The computation of the classical ANOVA decomposition is significantly more expensive, since here 2^d many high-dimensional integrals have to be computed.

ture methods. These quadrature methods are developed in two steps: First, the anchored-ANOVA decomposition is truncated either, a priori, based on function space weights or, a posteriori, in a dimension-adaptive fashion where important terms of the decomposition are automatically detected. This truncation introduces a modeling error which is controlled by the effective dimension in the anchored case. Then, the remaining terms are integrated using appropriate low-dimensional quadrature rules which may be different from term to term and may refine the approximation in a locally-adaptive way. This introduces a discretization error which only depends on the maximum order of the kept terms in the decomposition, but not on the nominal dimension d . We present numerical results using the CUHRE algorithm [8] for the integration of the low-order anchored-ANOVA terms and quasi-Monte Carlo methods for the higher-order ones. This way, we obtain mixed CUHRE/QMC methods which are to our knowledge the first numerical quadrature methods that can profit from low effective dimension by dimension-adaptivity and can at the same time deal with low regularity by local adaptivity. A correct balancing of modeling and discretization error is more difficult with these methods than with sparse grid methods. However, numerical experiments with an Asian option as a test function from finance with discontinuous first derivatives demonstrate that this disadvantage can in some cases more than be compensated by the benefits of the local adaptivity.

- We show that the dimension-wise quadrature methods includes the class of *sparse grid methods* as a special case if we use particular tensor product methods for the integration of the subterms. We explain that sparse grid methods can be interpreted as a refinement of the anchored-ANOVA decomposition by first expanding each term of the decomposition into an infinite basis and then truncating this expansion appropriately. This allows one to intertwine the truncation of the anchored-ANOVA series and the subsequent discretization and allows one to balance modeling and discretization error in an optimal way in the sense of [16] through the choice of the underlying index set. Such optimal index sets can be found in a dimension-adaptive fashion as in [46] or by using a priori information on the function space weights similar to [164]. We determine optimal index sets for integrands from weighted tensor products of Sobolev spaces and discuss cost and error bounds, which take into account the function space weights and recover known results in case of equal weights.
- Moreover, we study two special variants of the sparse grid method which can treat the integral (1.1) directly on \mathbf{R}^d avoiding the singular transformation to the unit cube. For moderate high-dimensional integrals with Gaussian weight and equally important dimensions, we define sparse grid methods based on the *delayed Genz-Keister sequence* using the recent approach from Petras [130]. For integrals with Gaussian weight, which have a high nominal, but a low effective dimension, we study dimension-adaptive sparse grid methods based on the *slowly increasing Gauss-Hermite sequence* combining ideas from [46, 118, 130]. We apply the latter method to several applications from finance and observe that this method can be superior to Monte Carlo, quasi-Monte Carlo and other sparse grid

methods [46, 116, 130, 146] by several orders of magnitude even in hundreds of dimensions.

- To further improve the performance of dimension-adaptive sparse grid methods for problems from finance, we study the impact of different path generating methods such as the Brownian bridge construction. In particular, we consider the *linear transformation method* from Imai and Tan [78]. Here, the main idea is that the integral (1.1) is invariant with respect to orthonormal transformations, such as rotations, and hence equals to

$$I_d = \int_{\mathbf{R}^d} g(\mathbf{Q}\mathbf{z})\varphi(\mathbf{z})d\mathbf{z}$$

for all orthogonal matrices $\mathbf{Q} \in \mathbf{R}^{d \times d}$. The linear transformation method aims to identify the matrix \mathbf{Q} which minimizes the effective dimension of the integrand for certain function classes and can in this way maximize the performance of dimension-adaptive sparse grid methods. We provide numerical experiments with several application problems from finance which illustrate the efficiency of this approach.

- We also address the difficulty that integrands from finance often have kinks or even jumps and do therefore not satisfy the high smoothness requirement of sparse grid methods. To overcome this obstacle we here investigate the approach first to identify all kinks and jumps and then to decompose the integration domain \mathbf{R}^d into subdomains Ω_i in which the integrand g is smooth. We thus shift the integration of one discontinuous function to the computation of

$$I_d = \sum_i \int_{\Omega_i} g(\mathbf{Q}\mathbf{z})\varphi(\mathbf{z})d\mathbf{z},$$

i.e., to the integration of several smooth functions. This way, we can regain the fast convergence of sparse grid methods in some cases with costs that depend on the number of terms in the sum and on the complexity of the decomposition. We show that this approach can be superior to standard methods for the pricing problems of barrier options and performance-dependent options.⁷ In the first case, we study the decomposition of the integration domain with the help of *conditional sampling* [55]. In the second case, we summarize the main results of [51, 54]. Here, the decomposition is performed using tools from *computational geometry* for the enumeration and decomposition of hyperplane arrangements in high dimensions.

- One of the most complex applications from finance and insurance is the simulation of *stochastic asset-liability management (ALM) models in life insurance*. In such models the development of the capital markets, the behaviour of the policyholders and the decisions of the company's management have simultaneously be taken into account as well as guarantees and option-like features of the insur-

⁷ Barrier options are financial derivatives which become worthless if the underlying asset crosses a specified barrier. Performance-dependent options are financial derivatives whose payoff depends on the performance of one asset in comparison to a set of benchmark assets.

ance products, see, e.g., [49]. New regulations, stronger competitions and more volatile capital markets have increased the demand for such simulations in recent years. The numerical simulation of such models is usually performed by Monte Carlo methods, which, however, often lead to unsatisfactorily long computing times of several days even on a supercomputer. Here, we follow [47] where it is shown that quasi-Monte Carlo and sparse grid methods can successfully be applied to these problems. To this end, we explain how the ALM simulation problem can be rewritten as a multivariate integration problem and then be solved by deterministic methods in combination with adaptivity and dimension reduction techniques. We provide various numerical experiments with a general ALM model framework, which incorporates the most important features of ALM simulations in life insurance such as the surrender of contracts, a reserve-dependent surplus declaration, a dynamic asset allocation and a two-factor stochastic capital market. The results demonstrate that in particular the quasi-Monte Carlo methods often converge faster, less erratic and produce more accurate results than Monte Carlo simulation even for small sample sizes n and complex models with many variables. We determine the effective dimensions and the important variables of different ALM models and indicate that ALM model problems are of very low effective dimension, or can be transformed to be of low effective dimension by coordinate transformations. This way, we also provide a theoretical explanation for the success of the deterministic quadrature methods.

The remainder of this monograph is organized as follows. In Chapter 2, we introduce the classical ANOVA and the anchored-ANOVA decomposition of a multivariate function f . Based on these decompositions, we then define the notions of effective dimensions in the classical and the anchored case and discuss error bounds for approximation and integration.

In Chapter 3, we start with a short survey of classical numerical methods for the computation of high-dimensional integrals. Then, we define the class of dimension-wise quadrature methods. These methods proceed dimension-wise and are constructed by truncation of the anchored-ANOVA decomposition and by integration of the remaining terms using one or several of the classical numerical methods. We derive cost and error bounds for the methods and discuss a priori and a posteriori approaches to exploit low effective dimension.

We specify some components of the dimension-wise quadrature methods in Chapter 4, which leads us to the class of sparse grid methods. We then consider these methods in more detail. We first define two special variants of sparse grid methods based on delayed Genz-Keister and on slowly increasing Gauss-Hermite sequences. Then, we discuss optimal index sets of sparse grid constructions for integrands from weighted Sobolev spaces.

The scope of Chapter 5 are approaches which can be used to improve the performance of sparse grid methods by dimension reduction and by the smoothing of the integrands. Here, we consider different path generating methods to reduce the dimension and discuss domain decompositions and conditional sampling to regain smoothness.

In Chapter 6, we finally present several applications from finance which can efficiently be treated by sparse grid methods. Using the pricing problems of different interest rate derivatives we study the effects of coordinate transformations and compare the performance of different sparse grid methods. Then, we consider path-dependent options, which lead to integrands with kinks or jumps. To overcome this obstacle we describe the approach to apply local adaptivity in the low-order anchored-ANOVA terms using the CUHRE algorithm and consider smoothing by conditional sampling. Moreover, we discuss the efficient pricing of performance-dependent options using domain decompositions to regain smoothness. Finally, we consider the simulation of stochastic asset-liability management models in life insurance using deterministic integration methods.

We conclude with a summary of the presented results and some remarks on areas of future research in Chapter 7.

In Appendix A, we formally define reproducing kernel Hilbert spaces and the notions of tractability and discrepancy and summarize some related results as complementary information.

Chapter 2

Dimension-wise Decompositions

In this chapter, we introduce the classical ANOVA and the anchored-ANOVA decomposition of a multivariate function f . Based on these decompositions, we then define different notions of effective dimensions of f and derive error bounds for approximation and integration.

We start with the introduction of general dimension-wise decompositions. To this end, let $\Omega \subseteq \mathbb{R}$ be a set and let

$$d\mu(\mathbf{x}) = \prod_{j=1}^d d\mu_j(x_j) \tag{2.1}$$

denote a d -dimensional product measure defined on Borel subsets of Ω^d . Here, $\mathbf{x} = (x_1, \dots, x_d)$ and $\mu_j, j = 1, \dots, d$, are probability measures on Borel subsets of Ω . Let $V^{(d)}$ denote the Hilbert space of all functions $f : \Omega^d \rightarrow \mathbb{R}$ with the inner product

$$(f, g) := \int_{\Omega^d} f(\mathbf{x})g(\mathbf{x}) d\mu(\mathbf{x}).$$

For a given set $\mathbf{u} \subseteq \mathcal{D}$, where $\mathcal{D} := \{1, \dots, d\}$ denotes the set of coordinate indices, the measure μ induces projections $P_{\mathbf{u}} : V^{(d)} \rightarrow V^{(|\mathbf{u}|)}$ by

$$P_{\mathbf{u}}f(\mathbf{x}_{\mathbf{u}}) := \int_{\Omega^{d-|\mathbf{u}|}} f(\mathbf{x})d\mu_{\mathcal{D}\setminus\mathbf{u}}(\mathbf{x}). \tag{2.2}$$

Thereby, $\mathbf{x}_{\mathbf{u}}$ denotes the $|\mathbf{u}|$ -dimensional vector containing those components of \mathbf{x} whose indices belong to the set \mathbf{u} and $d\mu_{\mathcal{D}\setminus\mathbf{u}}(\mathbf{x}) := \prod_{j \notin \mathbf{u}} d\mu_j(x_j)$. In this notation the case $\mathbf{u} = \emptyset$ is included for which it holds

$$P_{\emptyset}f(\mathbf{x}_{\emptyset}) := \int_{\Omega^d} f(\mathbf{x})d\mu(\mathbf{x}) =: If.$$

The projections define a decomposition of $f \in V^{(d)}$ into a finite sum according to

$$f(x_1, \dots, x_d) = f_0 + \sum_{i=1}^d f_i(x_i) + \sum_{\substack{i,j=1 \\ i < j}}^d f_{i,j}(x_i, x_j) + \dots + f_{1,\dots,d}(x_1, \dots, x_d)$$

which is often written in the more compact notation

$$f(\mathbf{x}) = \sum_{\mathbf{u} \subseteq \mathcal{D}} f_{\mathbf{u}}(\mathbf{x}_{\mathbf{u}}). \quad (2.3)$$

The 2^d terms $f_{\mathbf{u}}$ describe the dependence of the function f on the dimensions $j \in \mathbf{u}$ with respect to the measure μ . They are recursively defined by

$$f_{\mathbf{u}}(\mathbf{x}_{\mathbf{u}}) := P_{\mathbf{u}} f(\mathbf{x}_{\mathbf{u}}) - \sum_{\mathbf{v} \subset \mathbf{u}} f_{\mathbf{v}}(\mathbf{x}_{\mathbf{v}}) \quad (2.4)$$

and can also be given explicitly by

$$f_{\mathbf{u}}(\mathbf{x}_{\mathbf{u}}) = \sum_{\mathbf{v} \subseteq \mathbf{u}} (-1)^{|\mathbf{u}|-|\mathbf{v}|} P_{\mathbf{v}} f(\mathbf{x}_{\mathbf{v}}), \quad (2.5)$$

see [92]. The resulting decomposition (2.3) is unique for a fixed measure μ and orthogonal in the sense that

$$(f_{\mathbf{u}}, f_{\mathbf{v}}) = 0 \quad (2.6)$$

for $\mathbf{u} \neq \mathbf{v}$, see, e.g., [57, 135].

Example 2.1. For the case $d = 2$ the dimension-wise decomposition of f is given by

$$f(x_1, x_2) = f_0 + f_1(x_1) + f_2(x_2) + f_{1,2}(x_1, x_2)$$

with the constant term

$$f_0 = \int_{\Omega} \int_{\Omega} f(x_1, x_2) d\mu_1(x_1) d\mu_2(x_2),$$

the one-dimensional functions

$$f_1(x_1) = \int_{\Omega} f(x_1, x_2) d\mu_2(x_2) - f_0$$

$$f_2(x_2) = \int_{\Omega} f(x_1, x_2) d\mu_1(x_1) - f_0$$

and the highest-order term

$$\begin{aligned} f_{1,2}(x_1, x_2) &= f(x_1, x_2) - f_1(x_1) - f_2(x_2) - f_0 \\ &= f(x_1, x_2) - \int_{\Omega} f(x_1, x_2) d\mu_2(x_2) - \int_{\Omega} f(x_1, x_2) d\mu_1(x_1) + f_0. \end{aligned}$$

In the following we will specify the measure $d\mu$ in (2.1) using the Lebesgue and the Dirac measure which will lead us to the ANOVA and to the anchored-ANOVA

decomposition. Other choices are also possible, see e.g. [74], but not further investigated here.

2.1 Classical ANOVA Decomposition

For $\Omega = [0, 1]$ and the example of the Lebesgue measure $d\mu(\mathbf{x}) = d\mathbf{x}$ in (2.1), the space $V^{(d)}$ is the space of square integrable functions and the projections are given by

$$P_{\mathbf{u}}f(\mathbf{x}_{\mathbf{u}}) = \int_{[0,1]^{d-|\mathbf{u}|}} f(\mathbf{x}) d\mathbf{x}_{\mathcal{D}\setminus\mathbf{u}}.$$

The decomposition (2.3) then corresponds to the well-known analysis of variance (ANOVA) decomposition which is used in statistics to identify important variables and important interactions between variables in high-dimensional models. It goes back to [73] and has been studied in many different contexts and applications, e.g., [34, 66, 81, 155]. Recently, it has extensively been used for the analysis of quasi-Monte Carlo methods, see, e.g., [17, 97, 98, 102, 148] and the references cited therein.

In the ANOVA the orthogonality (2.6) implies that the variance

$$\sigma^2(f) := \int_{\Omega^d} (f(\mathbf{x}) - If)^2 d\mu(\mathbf{x})$$

of the function f can be written as

$$\sigma^2(f) = \sum_{\substack{\mathbf{u} \subseteq \mathcal{D} \\ \mathbf{u} \neq \emptyset}} \sigma^2(f_{\mathbf{u}}), \quad (2.7)$$

where $\sigma^2(f_{\mathbf{u}})$ denotes the variance of the term $f_{\mathbf{u}}$.¹ The values $\sigma^2(f_{\mathbf{u}})/\sigma^2(f)$, called global sensitivity indices in [147, 148], can then be used to measure the relative importance of the term $f_{\mathbf{u}}$ with respect to the function f .

Example 2.2. For the class of polynomials in two variables of the form

$$f(x_1, x_2) := a + bx_1 + cx_2 + dx_1x_2$$

with parameters a, b, c and $d \in \mathbb{R}$ one easily calculates that the terms of the ANOVA decomposition are given by

¹ Note that $If_{\mathbf{u}} := \int_{\Omega^d} f_{\mathbf{u}}(\mathbf{x}_{\mathbf{u}}) d\mu(\mathbf{x}) = 0$ for $\mathbf{u} \neq \emptyset$.

$$\begin{aligned}
f_0 &= a + \frac{b}{2} + \frac{c}{2} + \frac{d}{4} \\
f_1(x_1) &= (b + \frac{d}{2})(x_1 - \frac{1}{2}) \\
f_2(x_2) &= (c + \frac{d}{2})(x_2 - \frac{1}{2}) \\
f_{1,2}(x_1, x_2) &= \frac{d}{4}(2x_1 - 1)(2x_2 - 1).
\end{aligned}$$

For the specific case $a = 0$, $b = 12$, $c = 6$ and $d = -6$, i.e. for the polynomial

$$f(x_1, x_2) := 12x_1 + 6x_2 - 6x_1x_2$$

one obtains

$$\begin{aligned}
f_0 &= 15/2 \\
f_1(x_1) &= 9x_1 - 9/2 \\
f_2(x_2) &= 3x_2 - 3/2 \\
f_{1,2}(x_1, x_2) &= -6x_1x_2 + 3x_1 + 3x_2 - 3/2.
\end{aligned}$$

The variance of f is given by $\sigma^2(f) = 31/4$ and we see that $\sigma^2(f_1) = 27/4$, $\sigma^2(f_2) = 3/4$ and $\sigma^2(f_{1,2}) = 1/4$. Hence, the one-dimensional terms f_1 and f_2 explain about 87% and 10% of $\sigma^2(f)$, respectively. The highest-order term $f_{1,2}$ contributes the remaining 3% of the total variance.

Example 2.3. For the slightly modified polynomial

$$f(x_1, x_2) := 12x_1^2 + 6x_2^2 - 6x_1x_2$$

we obtain

$$\begin{aligned}
f_0 &= 9/2 \\
f_1(x_1) &= 12x_1^2 - 3x_1 - 5/2 \\
f_2(x_2) &= 6x_2^2 - 3x_2 - 1/2 \\
f_{1,2}(x_1, x_2) &= -6x_1x_2 + 3x_1 + 3x_2 - 3/2.
\end{aligned}$$

It holds $\sigma^2(f) = 35/4$, $\sigma^2(f_1) = 151/20$, $\sigma^2(f_2) = 19/20$ and $\sigma^2(f_{1,2}) = 1/4$. The first-order terms f_1 and f_2 explain about 86% and 11% of the variance of f and the second-order term $f_{1,2}$ about 3% of the variance.

Example 2.4. For given univariate functions $g_j \in L^2([0, 1])$, $j = 1, \dots, d$, let

$$I g_j := \int_{[0,1]} g_j(x) dx \quad \text{and} \quad \sigma^2(g_j) := \int_{[0,1]} (g_j(x) - I g_j)^2 dx.$$

For the classes of purely additive or multiplicative functions

$$f^+(\mathbf{x}) := \sum_{j=1}^d g_j(x_j) \quad \text{and} \quad f^*(\mathbf{x}) := \prod_{j=1}^d g_j(x_j)$$