

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Data-driven Simulations of Distributed Systems /

Permalink

<https://escholarship.org/uc/item/4qt7290r>

Author

Barajas-Solano, David Alonso

Publication Date

2013

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Data-driven Simulations of Distributed Systems

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Engineering Science (Applied Mechanics)

by

David Alonso Barajas-Solano

Committee in charge:

Professor Daniel M. Tartakovsky, Chair
Professor Yuri Bazilevs
Professor Juan Carlos del Álamo
Professor Michael Holst
Professor Sutanu Sarkar

2013

Copyright
David Alonso Barajas-Solano, 2013
All rights reserved.

The dissertation of David Alonso Barajas-Solano is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2013

DEDICATION

To Katie,
and to Kaitlyn, forever now.

EPIGRAPH

...repetition that involves good luck and inspiration is always a daring venture.

—Søren Kierkegaard

TABLE OF CONTENTS

Signature Page	iii
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Figures	viii
List of Tables	x
Acknowledgements	xi
Vita	xiii
Abstract of the Dissertation	xiv
Chapter 1	Introduction 1
	1.1 Epistemic uncertainty in distributed systems 1
	1.2 Uncertainty propagation 2
	1.3 Parameter identification 4
Chapter 2	Green functions for flow in heterogeneous composite media . . 5
	2.1 Introduction 5
	2.2 Regularized formulation of Green’s function problem . . 7
	2.2.1 Extensions of the regularization formulation . . . 10
	2.3 Numerical examples 11
	2.4 Conclusions 14
Chapter 3	Stochastic collocation methods for nonlinear systems 16
	3.1 Introduction 16
	3.2 Problem statement 19
	3.2.1 Probabilistic model parameterization 20
	3.3 Stochastic collocation 21
	3.3.1 Selection of quadrature rules 21
	3.3.2 Quadrature rules for multinormal weights 23
	3.3.3 Quadrature rules for other UQ methods 24
	3.4 Numerical experiments 24
	3.4.1 Elliptic problems 24
	3.4.2 Parabolic problems 28
	3.5 Conclusions 32

Chapter 4	Multi-level Monte Carlo for nonlinear systems	34
	4.1 Introduction	34
	4.2 Multi-level Monte Carlo	35
	4.3 Implementation of MLMC	39
	4.4 Numerical experiments	41
	4.5 Conclusions and further work	46
Chapter 5	Efficient Reconstruction of Contaminant Release History . . .	48
	5.1 Introduction	48
	5.2 Problem formulation	50
	5.2.1 Likelihood function	51
	5.3 Monte Carlo sampling	53
	5.3.1 Hybrid Monte Carlo (HMC)	53
	5.3.2 Generalized Hybrid Monte Carlo (gHMC)	55
	5.4 Numerical applications	56
	5.4.1 Continuous in space, discrete in time measurements	56
	5.4.2 Discrete in space, continuous in time measurements	59
	5.5 Application of gHMC to linear transport	62
	5.6 Conclusions and further work	71
Chapter 6	Inverse Modeling via Linear Functional Minimization	72
	6.1 Introduction	72
	6.2 Formulation of optimization problem	74
	6.3 Linearized functional minimization	77
	6.4 Forward solution	80
	6.5 Computational cost	81
	6.6 Application	82
	6.7 Conclusions and future work	88
Chapter 7	Conclusions	91
Appendix A	Computation of sensitivities for the hybrid Monte Carlo method	94
	A.1 Adjoint sensitivity analysis	94
	A.2 Ad-hoc sensitivity analysis	95
Bibliography	98

LIST OF FIGURES

Figure 2.1:	Composite domain Ω	8
Figure 2.2:	Domain of Example 1.	11
Figure 2.3:	Green functions of Example 1 with $K_1 = 1$ and (a) $K_2 = 10$, (b) $K_2 = 0.1$	12
Figure 2.4:	Domain of Example 2. K_1 (blue) = $\exp(-0.1)$, K_2 (red) = $\exp(7.0)$, $L = 12$	13
Figure 2.5:	Green functions of Example 2 for (a) $\mathbf{x}_0 = (9, 6)$, (b) $\mathbf{x} = (5, 10)$	14
Figure 2.6:	Mean (a) and variance (b) of hydraulic conductivity of Example 3.	15
Figure 3.1:	Relative error in the estimation of $\bar{u}(z = 0)$, $\sigma_Y^2 = \sigma_\beta^2 = 0.1$	27
Figure 3.2:	Relative error in the estimation of $\bar{u}(z = 0)$, $\sigma_Y^2 = \sigma_\beta^2 = 0.5$	27
Figure 3.3:	Response surface of $u(z = 0)$ using 1-term K-L expansions.	28
Figure 3.4:	Estimated mean of u profile for different times using level 7 quadrature rules.	30
Figure 3.5:	Estimated variance of u profile for different times using level 7 quadrature rules.	31
Figure 3.6:	Response surface for $t = 0.6$ at locations (a) $z = 0.4$ and (a) $z = 0.5$ along the profile using 1-term K-L expansions.	31
Figure 4.1:	Absolute value mean (a) and variance (b) of Q_L and Y_L for the elliptic boundary-value problem (3.10)– (3.12) and $\sigma_Y^2 = 0.1$	43
Figure 4.2:	Absolute value mean (a) and variance (b) of Q_L and Y_L for the elliptic boundary-value problem (3.10)– (3.12) and $\sigma_Y^2 = 1.0$	44
Figure 4.3:	Standardized cost of the simple MC and MLMC estimators for different target values of the standard deviations of the estima- tor and $\sigma_Y^2 = 0.1$. Final discretization level is fixed at $M_L = 256$	45
Figure 4.4:	Standardized cost of the simple MC and MLMC estimators for different target values of the standard deviations of the estima- tor and $\sigma_Y^2 = 1.0$. Final discretization level is fixed at $M_L = 256$	46
Figure 5.1:	Observations (solid line) and release distribution to be recon- structed (dotted line).	58
Figure 5.2:	Actual (solid line) and reconstructed (dotted line) release his- tories using HMC.	59
Figure 5.3:	Breakthrough curves of contaminant at observation locations along the transport domain.	61
Figure 5.4:	Actual (solid line) and reconstructed (dotted line) release his- tories using HMC.	62
Figure 5.5:	Autocorrelation functions for q_j , $j = \{0, \dots, 63\} \setminus \{47, 63\}$, and $\mathbf{A} = \mathbf{A}_1$ (top), $\mathbf{A} = \mathbf{A}_2$ (middle) and $\mathbf{A} = \mathbf{I}$ (bottom).	70

Figure 6.1:	Synthetic log-hydraulic conductivity field.	83
Figure 6.2:	Reconstructed log-hydraulic conductivity field for Test 1.	84
Figure 6.3:	Reconstructed log-conductivity fields for Test 2.	86
Figure 6.4:	Observed (continuous) and reconstructed (dotted) hydraulic heads at observation locations for Test 2(d).	87
Figure 6.5:	Reconstructed log-conductivity fields for Test 3.	88
Figure 6.6:	Observed (continuous) and reconstructed (dotted) hydraulic heads at observation locations for Test 3(d).	89

LIST OF TABLES

Table 2.1:	Absolute differences \mathcal{E}_p and \mathcal{E}_h for various grid sizes and $K_1 = 1$.	13
Table 3.1:	Model parameters for steady problem (3.10)–(3.12).	25
Table 3.2:	Model parameters for unsteady problem (3.14)–(3.18).	29
Table 6.1:	Properties of point loading in three test cases.	83
Table 6.2:	Inversion parameter sets	84
Table 6.3:	Penalty results for Test 2.	87
Table 6.4:	Penalty results for Test 3.	89

ACKNOWLEDGEMENTS

I would like to express my deepest gratitude to my advisor and friend, Prof. Daniel M. Tartakovsky, without whom I simply wouldn't be writing these words. I can't think of higher praise for a man than the extent to which he has changed for the best another man's life. Thank you for all your guidance and support. All the best.

I would like to thank Frank J. Alexander, Marian Anghel, Brendt Wohlberg and Velimir Vesselinov at Los Alamos National Laboratory, and Prof. Robert Scheichl at University of Bath, UK. I would also like to thank Peter Quinlan and all the staff at the Hidrogeology Division of Dudek in Encinitas. It has been an excellent experience to work with you and learn from you.

Special thanks to Professors Sutanu Sarkar, Juan Carlos del Álamo, Yuri Bazilevs and Michael Holst for serving in my committee. Also, thanks to the staff at the department of Mechanical and Aerospace Engineering, the International Center and the University at large for their support throughout my time at UCSD.

Thanks to my parents, Ruby and Crisóstomo, and my siblings Crisóstomo, Andrés, Paula and Steffanny for all their love and support. Gracias por siempre creer en mí. Los quiero mucho.

To all my friends: Vern and Patty, thank you for everything. Baldo, Bego, Manu, Silvana, Javier, Marco, Genaro, Jesús, Mercedes, Matt, Yin and Francesca, you are San Diego to me. Koby, Ben, TJ, Sveta and Prof. Ilenia Battiato, EBU I misses you. Viviana, Cristian, Khalay, Manti, Andrés, Nelson, Milena and Lina, thank you for always believing in me. Thank you very much Viktoría, my venture partner. Thank you Katie, I miss you.

Finally, thank you Kaitlyn Suski, my love and my rock. May we go places together.

Chapter 2, in part, is a reprint of the material as it appears in International Journal for Uncertainty Quantification: Barajas-Solano, D. A., Tartakovsky, D. M., "Computing Green's functions for flow in heterogeneous composite media", vol. 3, no. 1, pp. 39-46, 2013. The dissertation author was the primary investigator and author of the paper.

Chapter 3, in part, has been submitted for publication of the material: Barajas-Solano, D. A., Tartakovsky, D. M., “Stochastic collocation methods for nonlinear parabolic equations with uncertain parameters”. The dissertation author was the primary investigator and author of the paper.

Chapter 4, in part, is currently being prepared for submission for publication of the material. Barajas-Solano, D. A., Tartakovsky, D. M. The dissertation author was the primary investigator and author of the paper.

Chapter 5, in part, has been submitted for publication of the material: Barajas-Solano, D. A., Alexander, F. J., Anghel, M., Tartakovsky, D. M., “Efficient Reconstruction of Contaminant Release History”. The dissertation author was the primary investigator and author of the paper.

Chapter 6, in part, is currently being prepared for submission for publication of the material: Barajas-Solano, D. A., Wohlberg, B., Vesselinov, V., Tartakovsky, D. M., “Linear Functional Minimization for Inverse Modeling”. The dissertation author was the primary investigator and author of the paper.

VITA

- 2008 B. S. in Civil Engineering *summa cum laude*,
Industrial University of Santander, Bucaramanga, Colombia
- 2009-2013 Graduate Research Assistant, Department of Mechanical and
Aerospace Engineering, University of California, San Diego
- 2010 M. S. in Engineering Science (Applied Mechanics),
University of California, San Diego
- 2013 Ph. D. in Engineering Science (Applied Mechanics),
University of California, San Diego

PUBLICATIONS

Barajas-Solano, D. A., Tartakovsky, D. M., “Computing Green’s functions for flow in heterogeneous composite media”, *International Journal for Uncertainty Quantification*, vol. 3, no. 1, pp. 39-46, 2013.

Barajas-Solano, D. A., Alexander, F. J., Anghel, M., Tartakovsky, D. M. “Efficient Reconstruction of Contaminant Release History”, Submitted for publication.

Barajas-Solano, D. A., Tartakovsky, D. M., “Stochastic collocation methods for nonlinear parabolic equations with uncertain parameters”, Submitted for publication.

Barajas-Solano, D. A., Wohlberg, B., Vesselinov, V., Tartakovsky, D. M., “Linear Functional Minimization for Inverse Modeling”, Under preparation.

ABSTRACT OF THE DISSERTATION

Data-driven Simulations of Distributed Systems

by

David Alonso Barajas-Solano

Doctor of Philosophy in Engineering Science (Applied Mechanics)

University of California, San Diego, 2013

Professor Daniel M. Tartakovsky, Chair

This dissertation deals with mathematical modeling of complex distributed systems whose parameters are heterogeneous and heavily under-specified by data. Such problems are ubiquitous in every field of science and engineering, where one or more deterministic models exist to describe a given phenomenon but only a limited number of measurements of a model's parameters and its state variables are available. The main focus of this dissertation is on parameter identification (PI) and uncertainty quantification (UQ). The first part of this dissertation deals with development and numerical implementation of an algorithm to compute accurately and efficiently Green's functions, which are often used in both PI and UQ analyses of linear systems with piece-wise continuous parameters. The second part of this

dissertation explores the propagation of parametric uncertainty through a modeling process, in which quantities of interest are described by nonlinear elliptic and parabolic partial differential equations. We demonstrate that the variance of uncertain parameters (a measure of their uncertainty) strongly affects the regularity of a system's stochastic response, restricting the use of modern probabilistic UQ methods (e.g, polynomial chaos expansions and stochastic collocation methods) to low distributed parameters with low noise-to-signal ratios. High ratios adversely affect the stability and scalability of such methods. The third part of this dissertation deals with this issue by developing a multi-level Monte Carlo algorithm that outperforms direct Monte Carlo and allows for systematic treatment of different sources of bias in the computed estimators. In the final part of this dissertation we explore two PI strategies based on a Bayesian framework. The first strategy is to sample a posterior distribution using a generalized hybrid Monte Carlo (gHMC) method. We develop acceleration schemes for improving the efficiency of gHMC, and use them to estimate parameters in reactive transport systems with sparse concentration measurements. The second strategy is to compute the maximum a posteriori estimator of the configuration of spatially distributed, piece-wise continuous parameters by using a linearized functional minimization algorithm. Total variation regularization (TV) is employed as a prior on the parameter distribution, which allows one to capture large-scale features of system behavior from sparse measurements of both system parameters and transient system states.

Chapter 1

Introduction

1.1 Epistemic uncertainty in distributed systems

Analysis and predictions of physical phenomena require the formulation of mathematical models that capture various mechanisms of interest. Such models take the form of partial differential equations (PDEs) if a phenomenon of interest is defined over a space-time continuum. Parameters of these PDEs can vary in space and time, reflecting a system's heterogeneity that can manifest itself with different intensity on different scales. Hydrogeological applications provide a penultimate example of multiscale heterogeneity [26]. On the meter scale, properties of a geological formation can vary from one point to another in a way that maintains a constant average value over the whole structure. On the kilometer scale, one might encounter multiple geological structures whose vastly different properties give rise to spatially varying averaged parameters.

Once the required scale of parameter heterogeneity has been ascertained, one can proceed with assigning parameter values. This task involves measuring both system parameters (coefficients in the corresponding PDEs) and system states (unknowns described by these PDEs). The latter set of measurements serves to identify and calibrate model parameters, as well as to validate the model. In general, data collection is an expensive exercise (in terms of both time and money). Consequently, a typical problem in science and engineering involves system parameters that are underspecified by data and, hence, are inherently *uncertain*. This

then must be taken into account in order to produce valuable estimators of the relevant quantities of interest.

We assume that the phenomena under consideration are well understood, i.e., that their PDE-based models capture the “real” behavior of these phenomena with desired accuracy. Our focus is on quantification of *epistemic* (parametric) uncertainty. More specifically, we develop and analyze tools for epistemic uncertainty quantification and parameter estimation in highly under-characterized distributed system. Chapters 2–4 deal with uncertainty propagation, i.e., with estimating the statistics of quantities of interest in terms of the statistics of the input parameters. This is done by adopting a probabilistic framework that is outlined in Section 1.2. Chapters 5 and 6 contain our implementation of two alternative strategies for parameter identification, both of which are based on a Bayesian framework outlined in Section 1.3.

1.2 Uncertainty propagation

A probabilistic framework for the characterization of parametric uncertainty treats uncertain parameters as random fields. Governing PDEs and, correspondingly, system states become stochastic. Given a certain probabilistic characterization of the input parameters one would like to compute the probabilistic characterization of the system state. This procedure amounts to propagation of parametric uncertainty through a modeling process.

Various approaches to uncertainty quantification can be subdivided into three groups: direct methods, sampling methods, and spectral stochastic methods. Direct methods derive and solve a set of deterministic PDEs governing either relevant statistics or full probability density functions (PDFs) of the quantities of interest [64, 65, 82, 81, 83]. These methods generally require one to introduce assumptions and constraints on the stochastic characterization of input parameters, which limits their range of applicability.

Sampling methods, such as Monte Carlo simulations and its various variants, consist of generating multiple realizations of input parameters, solving deter-

ministic PDEs for each of these realizations, and computing the ensemble statistics of their solutions. Variants of sampling methods differ in the way they draw realizations of input parameters, but share the requirement that the number of realizations be large. If solving a deterministic PDE for each realization of input parameters is computationally expensive, the overall computational burden of sampling methods might become prohibitive. On the other hand, these methods are easy to implement, highly robust, and versatile.

Spectral stochastic methods use a finite number of random variables to approximate the probability space of input parameter fields, and employ spectral decompositions of the stochastic response (random system states) over the finite-dimensional probability space [34, 94, 29, 57, 91, 100]. Two hypotheses underpin these methods. First, it is assumed that the probability space can be characterized with a relatively small number of random variables, so that the dimensionality of the problem is not intractably large. Second, it is assumed that the stochastic response is sufficiently regular to allow a spectral representation. If any of these two assumptions is violated, stochastic spectral methods can become computationally more expensive than Monte Carlo simulations.

This dissertation deals with all three approaches to uncertainty quantification. Chapter 2 presents a computationally efficient strategy to compute Green's functions that appear in various direct methods, and proposes a regularization strategy that improves the accuracy of their approximation for problems with piece-wise continuous coefficients. Chapter 3 contains an implementation of the method of stochastic collocation (SC), a spectral stochastic method, to solve a class of nonlinear elliptic and parabolic problems. Analysis of the performance of SC for these problems sheds new light on the limitations of spectral stochastic methods, and establishes the requirements that an uncertainty quantification technique appropriate for these problems should satisfy. Finally, in Chapter 4 a multi-level Monte Carlo approach is proposed in order to address some of the issues outlined above for nonlinear elliptic and parabolic problems.

1.3 Parameter identification

As mentioned in Section 1.1, one can gain information about uncertain system parameters from measurements of the state of a system. This *inverse modeling* procedure is known as calibration within a deterministic framework, and as probabilistic updating of parameter distributions within a probabilistic framework. Bayes' theorem is typically used to condition the probability distribution of the system parameters to available measurements of the system states. Such an approach allows for both assimilation of various types of measurements as they become available and incorporation of the modeler's previous knowledge or expectations about the system parameters.

Chapters 5 and 6 introduce novel Bayesian techniques for parameter identification. A generalized hybrid Monte Carlo (gHMC) method for sampling the posterior probability distribution is presented in Chapter 5. This approach aims to reduce the correlation between samples in a generated Markov chain in order to reduce the variance of the estimators of the statistics of quantities of interest. It is applied to various identification problems in reactive transport. In Chapter 6 we present a linear functional minimization algorithm to compute maximum posterior estimators of heterogeneous, piece-wise continuous distributed parameters. This algorithm proposes the total variation (TV) regularization as a prior capturing the piece-wise continuity hypothesis. It is used to identify a composite conductivity distribution for saturated steady and transient problems with deterministic external loadings.

Chapter 2

Green functions for flow in heterogeneous composite media

2.1 Introduction

Green's functions are often used to quantify parametric uncertainty in physical systems described by partial differential equations (PDEs). They allow for direct analysis of the effects of uncertain forcings (source functions, initial and boundary conditions), whose effect is additive; and facilitate quantification of uncertainty in system parameters (conductivity, porosity, etc.), whose effect is multiplicative. Green's functions were employed to handle random parameters in a variety of fields, including dispersion of passive scalars in turbulent flows [49], flow and reactive transport in porous media [64, 65, 87, 88, 95, 82, 83], subsurface imaging [12] and parameter estimation and source identification [44].

Our analysis is motivated by the nonlocal formalism [64, 65], which employs moment differential equations (MDEs) to quantify uncertainty in predictions of steady-state flow in heterogeneous porous media with uncertain conductivity. The approach we propose relies on Green's functions to represent the nonlocal nature of the ensemble statistics of a system's response (i.e., hydraulic heads and fluxes). This methodology, coupled with asymptotic expansions in (small) variances of system parameters, was used to model linear [87], nonlinear [85] and

free-surface [82, 81] flows, as well as transport of chemically inert [25, 61] and active [83] solutes in porous media with statistically homogeneous and inhomogeneous [95] uncertain (random) parameters.

A typical example of the use of Green's functions in the context of uncertainty quantification is the computation of statistics of the hydraulic head h for steady saturated flow, governed by an elliptic equation [87]

$$\nabla \cdot [k(\mathbf{x})\nabla h(\mathbf{x})] + f(\mathbf{x}) = 0, \quad \mathbf{x} \in \Omega, \quad (2.1)$$

where uncertain hydraulic conductivity $k(\mathbf{x})$ and source function $f(\mathbf{x})$ are modeled as random fields. Let $Y(\mathbf{x}) = \ln k(\mathbf{x})$ be multivariate Gaussian, with mean $\langle Y \rangle$ and variance σ_Y^2 , and correlation function $C_Y(\mathbf{x}, \mathbf{y})$. For mildly heterogeneous media with small variances σ_Y^2 , the mean hydraulic head $\langle h(\mathbf{x}) \rangle$ can be expanded into a perturbation series in powers of σ_Y^2 [87],

$$\langle h(\mathbf{x}) \rangle = \langle h^{(0)}(\mathbf{x}) \rangle + \langle h^{(1)}(\mathbf{x}) \rangle + \dots \quad (2.2)$$

The zeroth-order approximation satisfies

$$\nabla \cdot [K\nabla \langle h^{(0)}(\mathbf{x}) \rangle] + \langle f(\mathbf{x}) \rangle = 0, \quad (2.3)$$

where K is the geometric mean of the field $k(\mathbf{x})$. The first-order approximation can then be computed as

$$\langle h^{(1)}(\mathbf{x}) \rangle = - \int_{\Omega} \left[K \frac{\sigma_Y^2}{2} \nabla_{\mathbf{y}} \langle h^{(0)}(\mathbf{y}) \rangle - \hat{\mathbf{r}}^{(1)}(\mathbf{y}) \right] \cdot \nabla_{\mathbf{y}} G(\mathbf{y}; \mathbf{x}) d\mathbf{y}. \quad (2.4)$$

In this expression, $G(\mathbf{y}; \mathbf{x})$ is the Green's function associated with (2.3) and $\hat{\mathbf{r}}^{(1)}(\mathbf{x})$ is given by

$$\hat{\mathbf{r}}^{(1)}(\mathbf{y}) = K^2 \int_{\Omega} C_Y(\mathbf{x}, \mathbf{y}) \nabla_{\mathbf{y}} \nabla_{\mathbf{x}}^{\top} G(\mathbf{y}; \mathbf{x}) \nabla_{\mathbf{y}} \langle h^{(0)}(\mathbf{y}) \rangle d\mathbf{y}. \quad (2.5)$$

Other statistics of $h(\mathbf{x})$ can be expressed in terms of the Green's function G as well.

When Green's functions are derived analytically, such approaches can be computationally more efficient than Monte Carlo simulations and other numerical techniques for solving stochastic PDEs. They also provide physical insight into

how uncertainty in parameters and/or driving forces affects predictive uncertainty. However, Green’s functions for many problems of practical significance cannot be obtained analytically. The need to compute Green’s functions numerically is by far the largest computational expense in the uncertainty quantification approaches that rely on them.

Presence of the Dirac delta function in a Green’s function PDE compromises the accuracy and convergence of regular finite element methods [19]. Further contributing to the loss of solution regularity are discontinuous coefficients in the governing PDEs [41] that describe, for example, flow in heterogeneous composite media [95]. The ability to compute efficiently Green’s functions in such a setting is of crucial importance to many uncertainty quantification efforts. Meeting this goal for uncertainty quantifications in elliptic PDEs is the major goal of the present analysis.

Green’s functions of some elliptic operators decay with the distance from the location of the Dirac forcing (point source). It is therefore possible to speak of a “support domain” for the Green’s function, defined as a portion of the computational domain in which the function is non-zero with a given degree of accuracy. Outside of the support domain the Green’s function is small enough not to affect a global quantity, e.g., an integral or another quantity dependent on the function in a weak sense. In such a case, homogeneous boundary conditions of a Green’s function PDE are transferred to the boundary of the support domain. Being able to solve a Green’s function PDE on the support domain, a (small) subset of the computational domain, leads to a significant reduction in computational time.

2.2 Regularized formulation of Green’s function problem

Let $\Omega \subset \mathbb{R}^d$ ($d = 1, 2, 3$) denote a convex domain with the Lipschitz-continuous boundary Γ . The domain Ω is composed of two disjoint units Ω_1 and Ω_2 such that $\Omega = \Omega_1 \cup \Omega_2$ and Γ_{12} being the boundary between them (Figure 2.1). Consider an elliptic PDE for the Green’s function $G(\mathbf{x}; \mathbf{x}_0)$,




Figure 2.1: Composite domain Ω .

$$-\nabla \cdot [K(\mathbf{x})\nabla G(\mathbf{x}; \mathbf{x}_0)] = \delta(\mathbf{x} - \mathbf{x}_0) \quad \text{in} \quad \mathbf{x}, \mathbf{x}_0 \in \Omega, \quad (2.6)$$

subject to the boundary conditions

$$G(\mathbf{x}) = 0 \quad \text{on} \quad \mathbf{x} \in \Gamma_D \quad (2.7)$$

$$\nabla G(\mathbf{x}) \cdot \mathbf{n} = 0 \quad \text{on} \quad \mathbf{x} \in \Gamma_N. \quad (2.8)$$

The coefficient $K(\mathbf{x})$ is piecewise constant in Ω_1 and Ω_2 , $\delta(\mathbf{x} - \mathbf{x}_0)$ is the Dirac delta function centered at point \mathbf{x}_0 , and \mathbf{n} is the unit normal vector pointing outward of Ω . The Dirichlet and Neumann boundary conditions are prescribed in the boundary segments Γ_D and Γ_N ($\Gamma = \Gamma_D \cup \Gamma_N$), respectively. Along the interface Γ_{12} , the Green's function and its flux satisfy the continuity conditions

$$[G(\mathbf{x})]_{\Gamma_{12}} = 0, \quad [K\nabla G(\mathbf{x}) \cdot \mathbf{n}_{12}]_{\Gamma_{12}} = 0, \quad (2.9)$$

where $[\cdot]$ denotes the jump of an enclosed quantity across the interface, and \mathbf{n}_{12} is the unit vector normal to Γ_{12} , pointing from Ω_1 to Ω_2 . The Green's function G is singular at $\mathbf{x} = \mathbf{x}_0$ due to the Dirac forcing.

Let $H^1(\Omega)$ be the space of square-integrable functions with square-integrable weak derivatives up to the first order, and V its subspace given by

$$V = \{v \in H^1(\Omega) : v|_{\Gamma_D} = 0\}. \quad (2.10)$$

Multiplying both sides of (2.6) with $v \in V$ and applying the Green's theorem, one

obtains the following weak formulation: Find $G_d \in V$ such that

$$\int_{\Omega} K(\mathbf{x}) \nabla G_d \cdot \nabla v \, d\mathbf{x} = v(\mathbf{x}) \quad \text{for all } v \in V. \quad (2.11)$$

The problem with this standard formulation is that $G \notin H^1(\Omega)$ due to the singularity at \mathbf{x}_0 . This implies that the solution G of (2.6) is not globally smooth enough for G_d to converge to G in the limit of mesh refinement [19]. In order to overcome this issue, we follow the finite element strategy [21, 19, 41] for solving the Poisson-Boltzmann equation (2.6). The approach divides the solution into a singular part and a regular part [21]. The regular problem admits a unique solution $\tilde{G} \in H^1$ and thus can be computed using a standard finite element approach [41].

To illustrate the regularization methodology, consider a domain Ω such that the boundary Γ_1 of Ω_1 contains a portion of the domain boundary Γ , and let $\mathbf{x}_0 \in \Omega_1$. Following [41], the total solution G can be decomposed into three components G_s , G_h and \tilde{G} such that

$$G = \begin{cases} G_s + G_h + \tilde{G} & \text{if } \mathbf{x} \in \Omega_1, \\ \tilde{G} & \text{if } \mathbf{x} \in \Omega_2. \end{cases} \quad (2.12)$$

We define G_s , the *singular* component, as a solution of

$$-\nabla \cdot (K_1 \nabla G_s) = \delta(\mathbf{x} - \mathbf{x}_0), \quad \mathbf{x} \in \mathbb{R}^d. \quad (2.13)$$

That is, G_s is the free-space Green's function for Laplace's equation with forcing δ/K_1 . For $d = 2$, $G_s = -(2\pi K_1^{-1} \ln(|\mathbf{x} - \mathbf{x}_0|))$ where $|\cdot|$ is the Euclidean norm. We also define the component G_h , the harmonic extension of the trace of the singular component G_s on Γ_1 into Ω_1 , as a solution of

$$-\nabla \cdot (K_1 \nabla G_h) = 0 \quad \text{in } \mathbf{x} \in \Omega_1, \quad G_s + G_h = 0 \quad \text{on } \mathbf{x} \in \Gamma_1. \quad (2.14)$$

Substituting (2.12)–(2.14) into (2.6) and enforcing the boundary conditions (2.7)

and (2.8), one obtains the following problem for the *regular* component \tilde{G} :

$$\begin{aligned}
-\nabla \cdot (K(\mathbf{x})\nabla\tilde{G}) &= 0 & \text{in } \mathbf{x} \in \Omega, \\
\tilde{G} &= 0 & \text{on } \mathbf{x} \in \Gamma_D, \\
\nabla(\tilde{G} + G_s + G_h) \cdot \mathbf{n} &= 0 & \text{on } \mathbf{x} \in \Gamma_N, \\
[\tilde{G}] &= 0 & \text{on } \mathbf{x} \in \Gamma_{12}, \\
[K(\mathbf{x})\nabla\tilde{G} \cdot \mathbf{n}_{12}] &= K_1\nabla(G_s + G_h) \cdot \mathbf{n}_{12} & \text{on } \mathbf{x} \in \Gamma_{12}.
\end{aligned} \tag{2.15}$$

It is proved in [41] that \tilde{G} is unique and lies in $H^1(\Omega)$, thus completing a methodology for approximating numerically the solution of (2.6).

2.2.1 Extensions of the regularization formulation

The approach described above can be generalized to more elaborate Green's function problems as long as a singular component analogous to (2.13) can be computed in closed form. Consider for example an equation

$$-\nabla \cdot [K(\mathbf{x})\nabla G] + \alpha \frac{\partial}{\partial x_3} [K(\mathbf{x})G] = \delta(\mathbf{x} - \mathbf{x}_0), \tag{2.16}$$

which arises in the context of unsaturated flow in porous media [88]. For constant $K = K_1$, the free-space solution of this problem is

$$G_s = -\frac{1}{2\pi K_1} e^{-\alpha(x_3 - x_{0,3})/2} \mathcal{K}_0(\alpha|\mathbf{x} - \mathbf{x}_0|/2),$$

where \mathcal{K}_0 is the zeroth-order modified Bessel function of the second kind. A harmonic expansion and a regular component can be defined in a fashion similar to (2.14) and (2.15).

Another possible generalization is the relaxation of the piecewise constant condition on $K(\mathbf{x})$ to the weaker piecewise Lipschitz continuity [20]. Such $K(\mathbf{x})$ fields capture small-scale variability inside each domain Ω_i , superimposed on large-scale variability between domains. Consider a setting similar to (2.6), but with the piecewise Lipschitz continuous $K(\mathbf{x})$. While K_1 is no longer constant, we nevertheless define the singular component G_s as a solution of

$$-\nabla \cdot (K^*\nabla G_s) = \delta(\mathbf{x} - \mathbf{x}_0), \quad \mathbf{x} \in \mathbb{R}^d, \tag{2.17}$$

where $K^* = K(\mathbf{x}_0)$. This solution is often available analytically, and from it one can define the harmonic extension in a weak sense as a solution $G_h \in V_1$ of

$$\int_{\Omega_1} K(\mathbf{x}) \nabla G_h \cdot \nabla v \, d\mathbf{x} = \int_{\Omega_1} [K^* - K(\mathbf{x})] \nabla G_s \cdot \nabla v \, d\mathbf{x} \quad \text{for all } v \in V_1, \quad (2.18)$$

with $V_1 = \{v \in H^1(\Omega) : v|_{\Gamma_1} = 0\}$. The remaining regular component is defined in a manner analogous to (2.15).

2.3 Numerical examples

The regularized formulation of the Green's function problem discussed in the previous section makes it explicit that differing values of K within subdomains Ω_i ($i = 1, 2$) affect the global behavior of the Green's function. This difference can be of one order of magnitude as in the Poisson-Boltzmann problem, or of many orders of magnitude as in subsurface applications [98]. To study the effect this difference has on a solution of the Green's function, we solve (2.6) on the domain shown in Figure 2.2, with $K_1 = 1.0$ and $K_2 = 10.0$ or 0.1 . The solutions G for both cases, computed using linear triangular finite elements and a grid size $h = 0.02$, are shown in Figure 2.3.




Figure 2.2: Domain of Example 1.

To assess the accuracy of the regularization approach compared to a standard discrete approximation, we computed solutions to the previous problem, using both the regularization approach and the discrete approximation G_d of (2.11), for different grid sizes h . For each value of h , let \mathcal{E}_p and \mathcal{E}_h denote the absolute difference between G and G_d at point $(0.5, 0.5)$ and at a distance h in the x_1 direction

from the point source, respectively. Table 2.1 shows that \mathcal{E}_p decays with h as h^2 , while \mathcal{E}_h remains constant independent of grid size; this confirms that G_d doesn't converge uniformly to G in the limit of mesh refinement, as stated in [19]. This result is to be expected as it is not possible for an approximation in H^1 to capture the behavior of the Green's function in the vicinity of the singularity. The regularization approach is more accurate than the standard discrete approximation, especially for low-resolution (large grid size h) simulations in which the vicinity of the singularity can be rather large.

Next we evaluate the size of the support domains of the Green's functions in Figure 2.3. Figure 2.3(a) shows that if the Dirac impulse is located in the low- K region, then the support domain lies within Ω_1 and a portion of Ω_2 in the vicinity of Γ_{12} . If the Dirac impulse lies within the high- K region, then the jump conditions (2.9) at the interface Γ_{12} can be approximated with the homogeneous Neumann (no-flow) boundary condition for the Green's function PDE defined on Ω_1 , and the Green's function PDE defined on Ω_2 is subject to the continuity (Dirichlet) condition $[G] = 0$ for $\mathbf{x} \in \Gamma_{12}$. The two PDEs become decoupled, and the support domain of G becomes larger (Figure 2.3(b)).





Figure 2.3: Green functions of Example 1 with $K_1 = 1$ and (a) $K_2 = 10$, (b) $K_2 = 0.1$.

To illustrate these points further, we compute solutions of (2.6) defined on the $L \times L$ square domain shown in Figure 2.4. Its internal geometry is reconstructed in [98] from synthetic geostatistical data by means of indicator kriging with $K_1 = \exp(-0.1)$ and $K_2 = \exp(7.0)$. The Green's function, for the Dirac forcing placed in both subdomains, is presented in Figure 2.5.

Table 2.1: Absolute differences \mathcal{E}_p and \mathcal{E}_h for various grid sizes and $K_1 = 1$.

h	\mathcal{E}_p		\mathcal{E}_h	
	$K_2/K_1 = 10$	$K_2/K_1 = 0.1$	$K_2/K_1 = 10$	$K_2/K_1 = 0.1$
0.02	5.98×10^{-5}	4.58×10^{-5}	7.3×10^{-3}	7.3×10^{-3}
0.025	9.39×10^{-5}	7.17×10^{-5}	7.3×10^{-3}	7.3×10^{-3}
0.05	3.96×10^{-4}	3.02×10^{-4}	7.3×10^{-3}	7.1×10^{-3}
0.10	1.89×10^{-3}	1.44×10^{-3}	7.3×10^{-3}	7.7×10^{-3}

Figure 2.5(a) shows that, as expected from earlier discussions, the Green's function with \mathbf{x}_0 in the low- K region has the support domain restricted to that subdomain. A sizable portion of the rest of the domain can then be disregarded, with the extent of that reduction depending on the level of accuracy required. On the other hand, the Dirac forcing located in the high- K region significantly extends the support domain, reducing the aforementioned numerical advantages.

**Figure 2.4:** Domain of Example 2. K_1 (blue) = $\exp(-0.1)$, K_2 (red) = $\exp(7.0)$, $L = 12$.

Finally, we use the Green's functions in Figure 2.5 to solve a stochastic problem

$$\nabla \cdot [K(\mathbf{x})\nabla h(\mathbf{x})] + f(\mathbf{x}) = 0, \quad \mathbf{x} \in \Omega, \quad (2.19)$$

in the composite domain Ω of Figure 2.4 subjected to homogeneous Dirichlet boundary conditions. The source function $f(\mathbf{x})$ is a stationary random field with




Figure 2.5: Green functions of Example 2 for (a) $\mathbf{x}_0 = (9, 6)$, (b) $\mathbf{x} = (5, 10)$.

unit mean and variance, and exponential correlation function

$$C_f(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|_2/l)$$

with correlation length $l = L/2$, where $\|\cdot\|_2$ denotes the Euclidean norm. In terms of the Green's function for problem (2.6), the mean and variance of h can be written as

$$\langle h(\mathbf{x}) \rangle = \int_{\Omega} \langle f(\mathbf{y}) \rangle G(\mathbf{y}; \mathbf{x}) d\mathbf{y}, \quad (2.20)$$

$$\sigma_h^2(\mathbf{x}) = \iint_{\Omega \times \Omega} C_f(\mathbf{y}, \mathbf{z}) G(\mathbf{y}; \mathbf{x}) G(\mathbf{z}; \mathbf{x}) d\mathbf{y} d\mathbf{z}. \quad (2.21)$$

Results computed using (2.20) and (2.21) are shown in Figure 2.6.

2.4 Conclusions

A regularization methodology for the numerical computation of Green's functions of elliptic boundary value problems in heterogeneous composite media has been studied. Green's functions are routinely used in uncertainty quantification, particularly in the moment differential equations approach to solving stochastic partial differential equations. Numerical experiments confirm that the regularization methodology allows for accurate and computationally efficient computation of Green's functions compared to standard methods that do not take into account




Figure 2.6: Mean (a) and variance (b) of hydraulic conductivity of Example 3.

the lack of regularity of the solution stemming from the Dirac forcing. In the case of composite media with discontinuous coefficients, the strength of the jump and the location of the forcing have a damping or amplifying effect on the value of the Green's function throughout the domain. If damping occurs, the opportunity arises for reducing the support domain of the Green's function, thus reducing the problem size and cutting the computation cost.

Chapter 2, in part, is a reprint of the material as it appears in International Journal for Uncertainty Quantification: Barajas-Solano, D. A., Tartakovsky, D. M., "Computing Green's functions for flow in heterogeneous composite media", vol. 3, no. 1, pp. 39-46, 2013. The dissertation author was the primary investigator and author of the paper.

Chapter 3

Stochastic collocation methods for nonlinear systems

3.1 Introduction

Nonlinear parabolic partial differential equations (PDEs), and their elliptic (steady-state) counterparts, describe a wide range of physical phenomena that range from heat conduction in solids to multiphase flow in porous media to electrodynamics [85, and the references therein]. The ability of these equations to predict underlying phenomena is often, if not always, compromised by uncertainty in their parameterization. This uncertainty arises from ubiquitous heterogeneity of ambient environments in which such phenomena occur, scarcity of parametric data, and imprecise knowledge of forcings (sources and initial and boundary conditions). Quantification of the impact of parametric uncertainty on the veracity of model predictions is an integral part of modern scientific computing.

A standard approach to quantification of parametric uncertainty relies on the probabilistic framework. The latter treats uncertain system parameters and forcings as random fields. Solutions of the corresponding PDEs with random coefficients are given in terms of probability density functions or, equivalently, statistical moments of dependent variables or other quantities of interest (QoI). Computing these statistics is the major focus of uncertainty quantification.

Monte Carlo simulations (MCS)—a method that consists of i) generating multiple equally probable realizations of input parameters, ii) solving deterministic PDEs for each realization, and iii) evaluating ensemble statistics of these solutions—provide the most robust and straightforward way to solve PDEs with random coefficients. Since MCS impose no limitations on statistical properties of input parameters, they usually serve as a yardstick against which the performance of other approaches is compared. MCS entail no modifications of existing deterministic solvers and therefore is often referred to as a nonintrusive technique. Yet MCS require a large number of realizations to converge, especially when used to solve nonlinear stochastic PDEs (SPDEs). This renders them computationally expensive (often, prohibitively so). Much of the research in the field of SPDEs is driven by the goal of designing numerical techniques that are computationally more efficient than MCS.

One alternative to MCS is to derive deterministic PDEs that govern the evolution of either statistical moments [47, 96] or probability density functions [18, 84] of dependent variables. The reliance on new governing equations implies that these approaches are intrusive, even though their underlying structure often remains the same and existing solvers can be used. In most implementations, these methods do not rely on finite-term approximate representations of random parameter fields, e.g., on truncated Karhunen-Loève (K-L) expansions, and thus do not suffer from “the curse of dimensionality”. When random parameters enter SPDEs as multiplicative noise, these approaches require a closure approximation. A systematic approach to obtaining such closures is based on perturbation expansions of relevant quantities into series in the powers of variances of the input parameters (e.g., [85, and the references therein]). This formally limits the applicability of such techniques to PDEs whose random coefficients exhibit low noise-to-signal ratios (coefficients of variation), even though in linear diffusion-type problems they might remain accurate for variances of log-transformed input fields as high as 4 [101].

Various flavors of stochastic finite element methods (FEMs) provide another alternative to MCS. The stochastic FEMs start by constructing (e.g., by means of truncated Karhunen-Loève expansions) a finite-dimensional probability space on

which an SPDE solution is defined. The Galerkin FEM [34, 8], often equipped with h -type and p -type adaptivity, approximates such solutions in the resulting composite probability-physical space. Stochastic Galerkin and collocation methods [100] employ orthogonal basis expansions of an SPDE solution in the chosen finite-dimensional probability space. These methods are often referred to as non-perturbative, even though their applications to systems whose parameters exhibit large coefficients of variation are scarce. They outperform MCS when random parameter fields exhibit long correlations and, therefore, can be accurately represented by, e.g., a few terms in their K-L expansions. As the correlation length of an input parameter decreases, its K-L expansion requires more terms to maintain the same accuracy, thus increasing the dimensionality of the probability space on which solution is defined. Once the number of random variables exceeds a certain threshold the stochastic FEMs become computationally less efficient than MCS, a phenomenon known as the curse of dimensionality.

Nonlinear problems of the kind considered in the present analysis pose additional challenges to stochastic Galerkin and collocation methods (and other spectral representations). In particular, the convergence rate of the associated expansions and quadrature rules is limited by the smoothness of the state variable manifold in the probability space. A problem's nonlinearity degrades the solution's regularity in probability space, undermining the performance of stochastic Galerkin and collocation methods. Another issue is that of scalability with the dimensionality of the chosen probability space, which must be taken into account to control the computation cost while ensuring an accurate solution. Computational approaches for dealing with the occurrence of discontinuities and/or the loss of regularity include direction-adaptive quadrature rules [33, 31], multi-element generalized polynomial chaos [94, 91], multi-element probabilistic collocation [29], locally-adaptive stochastic collocation [1, 57] and wavelet expansions [51, 52]. These methods have been shown to be efficient for problems with low-dimensional probability space. Their applicability is curtailed by the exponential cost increase associated with adaptability in high-dimensional spaces [57].

We investigate the limits of applicability of collocation techniques for solving

nonlinear advection-diffusion equations with random coefficients, a class of problems that includes the Richards and viscous Burgers equations. Section 3.2 contains a problem formulation and its statistical parameterization. Stochastic collocation approaches for solving this problem are discussed in Section 3.3. In Section 3.4 we test these strategies on elliptic and parabolic nonlinear Richards equations. Due to the similarities between these approaches, observations drawn from the application of stochastic collocation can be generalized to stochastic Galerkin methods. Section 3.5 presents recommendations on the appropriateness of stochastic Galerkin and collocation for this type of problems.

3.2 Problem statement

Consider a nonlinear parabolic equation

$$\frac{\partial \theta(\mathbf{x}, u)}{\partial t} = \nabla \cdot [K(\mathbf{x}, u) \nabla (u - x_3)], \quad \mathbf{x} = (x_1, \dots, x_d)^\top \in D \quad (3.1)$$

defined on a d -dimensional domain $D \subset \mathbb{R}^d$ with $1 \leq d \leq 3$. In the context of two-phase fluid flow in porous media, this equation is known as the Richards equation wherein the state variable $u(\mathbf{x}, t)$ is the fluid pressure, $\theta(\mathbf{x}, u)$ is the fluid saturation of a porous medium, $K(\mathbf{x}, u)$ is the saturation-dependent hydraulic conductivity of the medium, and the explicit presence of the vertical coordinate x_3 accounts for gravity [85]. Equation (3.1) is supplemented with constitutive relations

$$\theta = \theta[\Lambda(\mathbf{x}), u], \quad K = K[\Lambda(\mathbf{x}), u], \quad (3.2)$$

where $\Lambda(\mathbf{x})$ is the set of heterogeneous model parameters. A problem formulation is completed by specifying boundary conditions

$$\mathbb{B}(\mathbf{x}, t; u) = g(\mathbf{x}, t), \quad \mathbf{x} \in \partial D, \quad (3.3)$$

where \mathbb{B} is the boundary operator representing, e.g., to the Dirichlet and/or Neumann boundary conditions on various segments of the domain's boundary ∂D ; and $g(\mathbf{x}, t)$ denotes corresponding boundary functions.

Heterogeneity of the input parameters $\Lambda(\mathbf{x})$, practical impossibility of measuring their values at every point $\mathbf{x} \in D$, and measurement errors at points where

data are available render values of these parameters uncertain. To quantify parametric uncertainty, we treat these parameters as random fields $\Lambda(\mathbf{x}, \omega)$, where $\omega \in \Omega$ is the “coordinate” in the suitable probability space Ω . This renders the state variable $u(\mathbf{x}, t)$ random as well, $u(\mathbf{x}, t, \omega)$.

3.2.1 Probabilistic model parameterization

We assume that two parameters (α and K_s) from the parameter set are uncertain. They are modeled as second-order stationary, log-normal, mutually uncorrelated random fields. The log-transforms $Y = \ln K_s$ and $\beta = \ln \alpha$ are multivariate Gaussian fields, which we represent via truncated Karhunen-Loève (K-L) expansions

$$Y(\mathbf{x}, \omega) = \langle Y \rangle + \sum_{j=1}^P \sqrt{\lambda_{1j}} u_{1j}(\mathbf{x}) \xi_{1j}(\omega) \quad (3.4a)$$

$$\beta(\mathbf{x}, \omega) = \langle \beta \rangle + \sum_{j=1}^P \sqrt{\lambda_{2j}} u_{2j}(\mathbf{x}) \xi_{2j}(\omega). \quad (3.4b)$$

Here ξ_{ij} are i.i.d. normal random variables with zero mean and unit variance, λ_{ij} are the eigenvalues of the Karhunen-Loève expansions, and u_{ij} their eigenfunctions. Retaining the leading P terms in the infinite Karhunen-Loève series introduces a truncation error, manifested as a bias in the estimators of the statistics of the state variables. This truncation error can be reduced by increasing P .

This finite-term truncation enables one to characterize the event space by a random vector

$$\boldsymbol{\xi} = (\xi_{11}, \dots, \xi_{1P}, \xi_{21}, \dots, \xi_{2P})^\top, \quad (3.5)$$

with multi-Gaussian joint probability distribution function (PDF)

$$\rho(\boldsymbol{\xi}) = \frac{1}{(2\pi)^{(N/2)}} e^{-\boldsymbol{\xi}^\top \boldsymbol{\xi} / 2} \quad (3.6)$$

and support $\Gamma = \mathbb{R}^N$, where $N = 2P$ is the dimensionality of the problem. In other words, the random solution $u(\mathbf{x}, t, \omega)$ is approximated by a random solution $u(\mathbf{x}, t, \boldsymbol{\xi})$.

3.3 Stochastic collocation

Statistics of the state variable $u(\mathbf{x}, t, \boldsymbol{\xi})$ can now be computed as integrals weighted with the PDF (3.6). For example, the ensemble mean of $u(\mathbf{x}, t, \boldsymbol{\xi})$ is defined by

$$\mathbb{E}[u] = \int_{\Gamma} u(\mathbf{x}, t, \boldsymbol{\xi}) \rho(\boldsymbol{\xi}) d\boldsymbol{\xi}. \quad (3.7)$$

Stochastic collocation methods compute such multivariate integrals using quadrature rules of the form

$$\mathbb{E}[u] \approx \sum_{i=1}^M u(\mathbf{x}, t, \boldsymbol{\xi}_i) \omega_i, \quad (3.8)$$

where $\boldsymbol{\xi}_i$ and ω_i are the nodes and weights of the quadrature rule, respectively, and $u(\mathbf{x}, t, \boldsymbol{\xi}_i)$ is the solution of the deterministic problem (3.1)-(3.4) for a particular realization $\boldsymbol{\xi}_i$ of the random vector $\boldsymbol{\xi}$. This approach is a sampling technique whose *raison d'être* is the claim that it requires a smaller number of terms, M , to achieve the same accuracy as MCS. Recall that an estimator of $\mathbb{E}[u]$ computed with MCS—a penultimate sampling technique that consists of drawing equally probable realizations $\boldsymbol{\xi}_i$ from the distribution (3.6) and setting $w_i = 1/M$ for all $i = 1, \dots, M$ —converges as $1/\sqrt{M}$.

Quadrature rules of the form (3.8) are commonly based on a polynomial interpolation of the integrand. The extent to which they can provide a certain level of accuracy in the estimation of moments at a cost lower than MCS depends on how closely can polynomials approximate the response $u(\mathbf{x}, t, \boldsymbol{\xi}_i)$.

3.3.1 Selection of quadrature rules

Construction of an appropriate interpolatory quadrature rule is a key part of the stochastic collocation approach. Let $A(q, N)$ denote an N -dimensional rule with a construction parameter (“order”) q . The number of nodes M in this rule depends on q and N , i.e., $M = M(q, N)$. A set (“grid”) of nodes $\{\boldsymbol{\xi}_i\}_{i=1}^M$ corresponding to the N -dimensional rule $A(q, N)$ is denoted by $H(q, N)$.

For selected dimensionality N and weight $\rho(\boldsymbol{\xi})$, one can design a rule that is exact for integrands of total polynomial degree equal or less than a target value

$l(q, N)$ [72], known as *degree of exactness*, so that l increases with q for fixed N . However, in general, response surfaces $u(\mathbf{x}, t, \boldsymbol{\xi})$ derived from stochastic PDE problems cannot be represented exactly by polynomials of any order, therefore, the use of a quadrature rule of polynomial degree of exactness l introduces an interpolation error on top of the truncation error of the Karhunen-Loève expansion. This interpolation error decreases as l increases, so for practical applications it is advisable to use quadrature rule constructions that allow for variable and/or adaptive degree of exactness to adequately compute quadratures like (3.8).

In order to be able to control both the error introduced by truncating the Karhunen-Loève expansions (3.4) and the interpolation error of a quadrature rule $A(q, N)$, the number of terms $M(q, N)$ in (3.8) must be adequately *scalable* with increasing dimensionality N and/or order q . In other words, the growth rate of M with N and q must be as slow as possible. Additionally, a quadrature rule $A(q, N)$ must be *stable*, i.e., the quantity $R = \sum_{i=1}^M |w_i| \geq 1$ must be reasonably small and not scale poorly with dimensionality N and q , this in order to avoid the appearance of non-physical artifacts during the quadrature evaluation (3.8) due to numerical cancellation.

Quadrature rules can be either *global* (with fixed quadrature nodes) or *locally adaptive* (with node locations determined by features of the stochastic response). Global constructions are defined chiefly by the quadrature weight $\rho(\boldsymbol{\xi})$ for which they are designed, and allow for a certain level of adaptivity. For example if the smoothness of the stochastic response $u(\mathbf{x}, t, \boldsymbol{\xi})$ exhibits a directional dependance, one can introduce non-isotropic quadrature rules that locate quadrature nodes preferentially along more critical directions [33]. One can also choose quadrature rule constructions that produce an embedded series of grids for increasing order,

$$H(q, N) \subset H(q + 1, N). \quad (3.9)$$

This enables one to increase the rule order q without wasting previous work performed, thus allowing for order adaptivity.

3.3.2 Quadrature rules for multinormal weights

The probability density function of the weight $\rho(\boldsymbol{\xi})$ in (3.7) affects the construction of a quadrature rule $A(q, N)$. If the weight $\rho(\boldsymbol{\xi})$ is multi-Gaussian, one can use either appropriate tensor products or sparse grids.

The former takes advantage of the multi-Gaussianity of $\rho(\boldsymbol{\xi})$ to construct N -dimensional quadrature rules as tensor products of one-dimensional Gauss-Hermite (GH) quadratures in each of the N directions. If the same number m of nodes is used in all directions, then the total number of quadrature points is $M = m^N$. This indicates an unacceptable exponential growth of the number of nodes with the dimensionality N of the problem. Moreover, standard GH quadratures and their tensor products are not embedded. While in some cases one can construct embedded extensions of GH quadratures of the Kronrod type, general embedded rules of arbitrary order do not exist [48]. On the other hand, all the weights w_i in GH quadratures are positive, which ensures that the resulting quadrature rule has ideal stability ($R = 1$).

Sparse grids constructed with the Smolyak algorithm [79] produce quadrature rules that have significantly fewer nodes than their counterparts based on tensor products. They also exhibit favorable scalability properties with the dimensionality N and order q . A straightforward way to construct a multi-Gaussian sparse grid is to combine the Smolyak algorithm with GH quadratures (the SGH rules); more sophisticated constructs include the Genz-Keister (GK) rules [32]. The latter rely on possible Gauss-Kronrod extensions of GH quadratures, and yield embedded rules that are generally more stable and have fewer nodes than SGH rules. On the negative side, since the Smolyak algorithm introduces negative weights, it leads to quadrature rules whose stability number R increases with q and N .

When sparse quadrature rules are to be used, GK rules are preferable over SGH rules due to their better stability and lower node count for increasing q and N . The stability of quadrature rules plays an important role in the accuracy of estimators of moments for nonlinear stochastic PDEs, as will be illustrated in Section 3.4.

3.3.3 Quadrature rules for other UQ methods

The procedure described above enables one to compute ensemble moments of quantities of interest in a non-intrusive way, i.e., by utilizing an existing deterministic solver for the PDE under consideration. Generalized polynomial chaos expansions (gPC) [100] provide an intrusive alternative (requiring modifications of the deterministic PDE solver), which allows computing a full probabilistic description of quantities of interest. To regain non-intrusiveness, one can approximate the coefficients in a gPC expansion with the weighted quadrature rules described above by employing the projection integral definition of these coefficients [100]. The error of the resulting gPC approach is that of the truncated gPC expansion and that of the quadrature rule.

3.4 Numerical experiments

The major goal of the numerical experiments performed below is to investigate the accuracy and robustness of stochastic collocation (SC) methods for solving nonlinear elliptic (steady-state) and parabolic (time-dependent) advection-diffusion equations. Our focus is on the effect of statistical characteristics (especially variance) of random coefficients on the performance of stochastic collocation methods. The impact of the stability of quadrature rules is also considered.

3.4.1 Elliptic problems

We start with an elliptic problem that represents a steady one-dimensional version of (3.1),

$$\frac{d}{dz} \left[K(z, u, \omega) \frac{d(u-z)}{dz} \right] = 0, \quad 0 < z < 1 \quad (3.10)$$

subject to boundary conditions

$$K \frac{du}{dz}(z=0, \omega) = -Q, \quad u(z=1, \omega) = 0 \quad (3.11)$$

where Q is the prescribed (deterministic) boundary flux. We employ an exponential Gardner’s constitutive relation

$$K(z, u, \omega) = \begin{cases} K_s(z, \omega) e^{\alpha(z, \omega)u(z, \omega)}, & u \leq 0, \\ K_s(z, \omega), & u > 0. \end{cases} \quad (3.12)$$

Theoretical considerations and experimental evidence (see, e.g., [86] and the references therein) suggest that the random system parameters $K_s(z, \omega)$ and $\alpha(z, \omega)$ have lognormal distributions. We assume their Gaussian counterparts $Y = \ln K_s$ and $\beta = \ln \alpha$ to be mutually uncorrelated and stationary (statistically homogeneous), each possessing an exponential correlation function.

We use SC to estimate the ensemble mean of the state variable, $\mathbb{E}[u(z, \omega)] \equiv \bar{u}(z)$. This estimate is compared with a Monte Carlo solution, $\bar{u}_{\text{ref}}(z)$, which is treated as ground truth or reference. The MC solution is obtained by averaging solutions corresponding to 10^8 MC realizations generated at the same resolution of the discrete spatial scheme used for the deterministic solution of (3.10)–(3.11). The results are presented in the form of a relative error between $\bar{u}(0)$ and $\bar{u}_{\text{ref}}(0)$,

$$\epsilon_r = \frac{|\bar{u}(0) - \bar{u}_{\text{ref}}(0)|}{|\bar{u}_{\text{ref}}(0)|}. \quad (3.13)$$

The estimators and reference value of the ensemble mean are computed up to 3 significant digits. Our goal is to analyze this error a function of a measure of the associated computational cost.

We conduct two numerical experiments, which are characterized by different levels of parametric uncertainty (variance of the random parameters), $\sigma_Y^2 = \sigma_\beta^2 = 0.1$, and $\sigma_Y^2 = \sigma_\beta^2 = 0.5$. Both cases assume unit correlation lengths for Y and β . These and other parameter values are summarized in Table 3.1.

Table 3.1: Model parameters for steady problem (3.10)–(3.12).

Test	Q	$K_G = \exp \langle Y \rangle$	$\alpha_G = \exp \langle \beta \rangle$	σ_Y^2	σ_β^2	l_Y	l_β
A	0.1	1.0	7.5	0.1	0.1	1.0	1.0
B	0.1	1.0	7.5	0.5	0.5	1.0	1.0

Each experiment computes $\bar{u}(z = 0)$ using both the GK and SGH rules of order $q = 1, 2$ and 3 , together with the corresponding Karhunen-Loève expansions of $N = 5, 10$ and 15 terms per random variable. This results in a total of $10, 20$ and 30 dimensions of the stochastic problem, respectively. All deterministic equations (in SC and MC) are solved using uniform linear finite elements with 100 elements. The nonlinear systems of equations resulting from the spatial discretization of (3.10) together with (3.11) are solved using the modified Newton's method with line search as implemented by the KINSOL library [40].

Computational cost for non-intrusive methods such as SC and MC is driven by the properties of the solution scheme of the deterministic subproblems, and the total number of samples. For the finite element scheme described above, the computational cost of the t -th sample is of the form $k^{(t)}E$, where E is the number of elements of the discretization, and $k^{(t)}$ is a coefficient that includes the number of evaluations of the nonlinear functional stemming from the discretization of (3.10)–(3.11), the approximation of its Jacobian, and the solution of the associated (tridiagonal) linear systems of equations. The partial cost of each of these operations is proportional to E , and also the number of times these operations are carried per sample is approximately independent of E , as numerical experience has shown. Therefore it can be said that $k^{(t)}$ is constant for a given E on average over M stochastic samples. The measure of the relative computational cost can be expressed as $M\bar{k}E$, where \bar{k} is the average of the $k^{(t)}$ over those M samples. We use this relative computational cost for comparing the different estimators.

Figures 3.1 and 3.2 show the results of the numerical experiments. It can be seen that, when the variance of the model parameters is low (Test A, Figure 3.1), both SC implementations are more accurate for less cost than simple MC. The error ϵ_r decreases as the rule order q increases, saturating to the estimator bias introduced by the K-L expansion, which itself drops with increasing number of dimensions N . The differences in cost and accuracy between the GK and SGH rules are minor. On the other hand, for the higher variance of Test B (Figure 3.2), the advantage of SC over MC is less clear, with MC outperforming both sparse grid constructions. This drop in performance can be attributed to the way the




Figure 3.1: Relative error in the estimation of $\bar{u}(z=0)$, $\sigma_Y^2 = \sigma_\beta^2 = 0.1$.

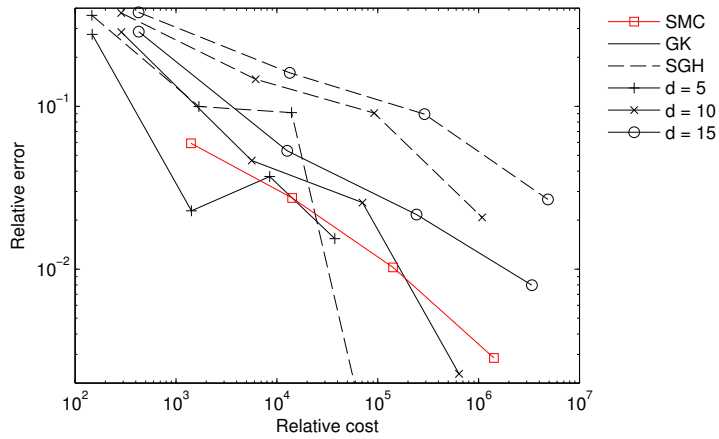


Figure 3.2: Relative error in the estimation of $\bar{u}(z=0)$, $\sigma_Y^2 = \sigma_\beta^2 = 0.5$.

variance affects the smoothness of the state variable manifold in probability space.

To illustrate this idea, Fig. 3.3 presents the stochastic response of u at $z=0$ as a function of the first stochastic terms of the decomposition of Y and β (all other dimensions set to zero), and for the two levels of variance considered. Comparison of Figs. 3.3(a) and 3.3(b) reveals that the response is less likely to be interpolated accurately using the same sparse grid with increasing variance of the input random fields, therefore increasing the interpolation error. The rescaling that occurs with increasing variance brings fringe behavior closer to the origin of the probability space, thus lowering the regularity of the response and requiring

higher level quadrature rules for an accurate estimation of moments. This effect can render quadrature-based approaches for the estimation of moments less efficient than the simpler Monte Carlo approach, as seen in Figure 3.2.




Figure 3.3: Response surface of $u(z = 0)$ using 1-term K-L expansions.

3.4.2 Parabolic problems

Consider the one-dimensional form of (3.1)

$$\frac{\partial \theta}{\partial t} = \frac{\partial}{\partial z} \left[K \frac{\partial}{\partial z} (u - z) \right], \quad 0 < z < 1 \quad (3.14)$$

with boundary conditions

$$u(z = 0) = u_t, \quad u(z = 1) = u_b, \quad (3.15)$$

and an initial condition

$$u(t = 0) = u_b. \quad (3.16)$$

The constitutive relations $\theta = \theta(u)$ and $K = K(u)$ are given by the commonly used van Genuchten and Mualem relations,

$$K = \begin{cases} K_s(z, \omega) S_e^{1/2} [1 - (1 - S_e^{1/m})^m]^2, & u \leq 0 \\ K_s(z, \omega), & u > 0 \end{cases}, \quad (3.17)$$

$$S_e \equiv \frac{\theta - \theta_r}{\theta_s - \theta_r} = \begin{cases} (1 + |\alpha(z, \omega) u|^n)^{-m}, & u \leq 0, \\ 1, & u > 0 \end{cases}. \quad (3.18)$$

The parameter values used in the simulations reported below are provided in Table 3.2.

Deterministic solutions of this problem develop a sharp moving front across which the state variable u drops from u_t to the background value u_b . The dependence of the sharp front's dynamics on the random input parameters is expected to translate into a lack of regularity in the probability space. We analyze this phenomenon by comparing the performance of stochastic collocation and simple Monte Carlo simulations.

As before, we assume that the random system parameters K_s and α have log-normal distributions, and that their Gaussian counterparts Y and β are mutually uncorrelated and stationary, each possessing an exponential correlation function. The model parameters in Table 3.2 are chosen so that the deterministic solution of the pressure field develops sharp moving fronts. Such a parameter set is representative of the scenarios encountered in stochastic analyses of nonlinear transport problems.

Table 3.2: Model parameters for unsteady problem (3.14)–(3.18).

Parameter values		
$\theta_s = 0.368$	$\alpha_G = \exp(\langle\beta\rangle)$	$= 2.01$
$\theta_r = 0.102$	$K_G = \exp(\langle Y\rangle)$	$= 0.5532$
$n = 2.0$	σ_Y^2	$= 1.0 \times 10^{-3}$
$m = 0.5$	σ_β^2	$= 1.0 \times 10^{-3}$
$u_t = -1.25$	l_Y	$= 1.0$
$u_b = -16.67$	l_β	$= 1.0$

The reported expectation and variance of $u(z, t, \omega)$ are evaluated using either MCS with 10^5 samples or stochastic collocation with the GK and SGH quadrature rules. Deterministic solutions are computed using the method of lines, with 100 uniform linear finite elements and the 3-stage implicit Runge-Kutta scheme RADAU IIA for temporal integration. The number of dimensions and rule order of the quadrature rules are chosen to approximate the MC estimator while keeping the cost lower than that of MC. Two K-L terms per random field, for a total of 4

stochastic dimensions, and rules of order 7 (polynomial degree of exactness of 15) were required in order to approximate the MC estimator. Simulation results are presented in Figures 3.4 and 3.5.

Even though the input parameters have low variances, high-level quadrature rules (order 7) were required to accurately interpolate the stochastic response. The GK rule gives accurate results, while the SGH rule produces undesirable oscillations along the profiles; these oscillations are artifacts introduced by the lack of stability of the high-level SGH quadrature rules. These results suggest that the more favorable stability properties of the GK rule makes it preferable to the SGH rule.

The input parameters in this example were assigned correlation lengths that equal the domain size. Therefore, a relatively small number of terms in the K-L expansions (2 per random field) was necessary to build reliable quadrature rules. If the correlation lengths were to decrease, the number of terms in the K-L expansions would have to increase to keep the same truncation error. In this scenario of high-order quadrature rules, the price to be paid for an increase in the dimensionality of the problem is very high.




Figure 3.4: Estimated mean of u profile for different times using level 7 quadrature rules.

Another issue of interest is the effect of increasing variance on the performance of stochastic collocation methods. To evaluate this, we plot the response sur-




Figure 3.5: Estimated variance of u profile for different times using level 7 quadrature rules.

face of $u(z, t)$ at a fixed time and location, as a function of the first terms of the K-L expansions of Y and β (Fig. 3.6). To emphasize the nonlinearity of the response surfaces, the input parameters' variance was increased to $\sigma_Y^2 = \sigma_\beta^2 = 1.0 \times 10^{-2}$.




Figure 3.6: Response surface for $t = 0.6$ at locations (a) $z = 0.4$ and (a) $z = 0.5$ along the profile using 1-term K-L expansions.

Figure 3.6 illustrates pitfalls in using SC methods to compute estimators of moments. Even for small variances, the response surface develops very steep regions, which necessitates the use of high-level quadrature rules to accurately interpolate the response. As the variance increases, the required rule order increases

as well, undermining the advantage of stochastic collocation as compared to MCS.

A possible approach to addressing the presence of a clear transition in the probability space of the response from u_t to u_b is to employ an adaptive construction of a quadrature rule that accurately captures this transition [1, 4, 46, 57]. However, this strategy is problematic for this class of problems, because the location of the irregularity in the probability space changes with both time t and space x . This raises the computational cost of building a quadrature rule that captures the dynamic transition along the profile as time progresses. Therefore the use of adaptive schemes to estimating moments of quantities of interest is computationally feasible, i.e., competitive with MCS, only when response surfaces exhibit no such dynamic behavior.

These limitations do not imply that SC methods are unusable for all transient problems. In some applications, such as the transition from two steady states of problem (3.10)-(3.11) for different values of the boundary flux q , response surfaces in the probability space can be integrated accurately using a global fixed sparse grid for appropriate values of variance and correlation lengths of input random fields [54], and therefore can benefit from the use of SC methods. It is therefore of much importance for the modeler to properly identify the features of the transient problem studied in order to determine the applicability of different UQ methods.

3.5 Conclusions

We have presented a numerical implementation of global stochastic collocation strategies for the estimation of first moments of quantities of interest in nonlinear parabolic and elliptic problems. It has been shown that the performance of such strategies, and in general, strategies based on a spectral decomposition of state variables in probability space, is tied strongly to the way the stochastic properties of the input random fields affect the regularity of the systems' stochastic response in the probability space. If random input fields have low variance and large correlation lengths, stochastic collocation strategies are competitive against

alternative uncertainty quantification methods, such as Monte Carlo simulations (MCS). Increasing variance affects the regularity of the stochastic response, requiring higher-order quadrature rules to accurately approximate the moments of interest, thus increasing the overall computational cost beyond that of MCS.

The effect of short correlation lengths is more straightforward to understand and has been discussed thoroughly: for a fixed number of stochastic dimensions, a bias introduced by a truncated K-L expansion of random fields increases as their correlation length decreases; the bias is reduced only by increasing the number of stochastic dimensions and, thus, the dimensionality of the problem. Resulting high-dimensional formulations restrict usable quadrature rules to low orders, and by extension to regular problems. Given the nonlinear nature of the parabolic and elliptic problems studied, this becomes an important limitation to be taken into account when deciding on an uncertainty quantification framework.

It is important to note that the estimator bias introduced by the K-L expansion is not the only source of bias in the computation of estimators of statistics of quantities of interest associated to nonlinear parabolic and elliptic problems. Additionally, one must take into account the bias introduced by the spatiotemporal discretization of the governing partial differential equations. A robust uncertainty propagation framework must be able to deal with these sources of bias in a systematic fashion and produce estimators with the level of accuracy desired by the modeler.

Chapter 3, in part, has been submitted for publication of the material: Barajas-Solano, D. A., Tartakovsky, D. M., “Stochastic collocation methods for nonlinear parabolic equations with uncertain parameters”. The dissertation author was the primary investigator and author of the paper.

Chapter 4

Multi-level Monte Carlo for nonlinear systems

4.1 Introduction

The computational examples of Chapter 3 served to demonstrate that, when used to quantify parametric uncertainty in nonlinear elliptic and parabolic PDEs, the state-of-the-art spectral stochastic approaches are restricted to low-dimensional and low-variance scenarios. Both restrictions are not met in many practical applications.

Monte Carlo (MC) simulations and its variants can deal with high-dimensional and high-variance problems, but are computationally expensive. Properties of simple MC, e.g., the independence of the convergence rate of the estimator from the number of dimensions are highly desirable. Yet this convergence rate is so slow that computing an estimator with a certain target standard deviation might not be feasible in practice. This is particularly so if the computational cost associated with each sample of the quantities of interest is large, as is the case for nonlinear systems such as those described above.

These shortcomings motivate the development of approaches that preserve the robustness of MC and have a reasonable computational cost. Additionally, an ideal approach must be able to systematically deal with various sources of bias in

the estimators. Specifically, it must be able to compute estimators for a target level of error, which includes not only the estimator’s variance but also errors introduced by the truncated representation of the stochastic response in space, time and probability.

We argue that multi-level Monte Carlo (MLMC) satisfies the requirements outlined above to a satisfactory extent. The idea of MLMC was first introduced to estimate quantities of interest described by stochastic ordinary differential equations [35]. It has since been applied to stochastic partial differential equations [23, 27, 28]. The MLMC formalism [23] introduces a sequence of random variables that converge in the mean to the quantity of interest. The sequence is derived from a corresponding sequence of sequentially refined spatial discretizations of stochastic PDEs. An estimator of the relevant statistics is computed from the statistics of the sequence of random variables, which are estimated using standard techniques (e.g., simple MC). This approach allows for systematic quantification of the estimator bias introduced by a spatial discretization of stochastic PDEs. The reliance on simple MC to solve the sub-problems enables one to use an arbitrary number of terms for these discrete representations. This sidesteps the issue of bias introduced by the finite representation of the input random fields altogether.

We use the MLMC variant [23] to solve the family of nonlinear problems studied in Chapter 3. A general theory of MLMC is presented in Section 4.2. Its numerical implementation is discussed in Section 4.3. In Section 4.4 we use MLMC to solve the steady problem formulated in Chapter 3. Our analysis demonstrates that MLMC outperforms simple Monte Carlo even when input parameters are high-dimensional and exhibit high variances. Conclusions and directions of further research are outlined in Section 4.5.

4.2 Multi-level Monte Carlo

Consider a infinite-dimensional probability space $(\Omega, \mathcal{F}, \mathbb{P})$. A random variable $Q : \Omega \rightarrow \mathbb{R}$ is defined via a functional \mathcal{G} as $Q = \mathcal{G}(\omega)$, where $\omega \in \Omega$. We assume that the quantity of interest Q is in general inaccessible, i.e., cannot be

computed exactly given a certain $\omega \in \Omega$. Instead, we introduce a sequence of functionals \mathcal{G}_M , $M \in \mathbb{N}$ that approximate the quantity of interest Q with increasing accuracy as M increases for a fixed ω . Specifically, we assume that a sequence of random variables $Q_M = \mathcal{G}_M(\omega)$ converges in the mean to Q as $M \rightarrow \infty$, that is

$$\mathbb{E}[Q_M - Q] = \mathcal{O}(M^{-\alpha}) \text{ as } M \rightarrow \infty, \quad (4.1)$$

for a certain constant α independent of M . Instead of trying to approximate the statistics of Q , we compute unbiased estimators of the statistics of Q_M . By virtue of (4.1), these estimators converge to the inaccessible statistics of Q as M increases and the variance of the estimators decreases.

To illustrate this procedure, let us consider an unbiased simple Monte Carlo (SMC) estimator of $\mathbb{E}[Q_M]$,

$$\hat{Q}_{M,N}^{\text{SMC}} = \frac{1}{N} \sum_{t=1}^N Q_M^{(t)}, \quad (4.2)$$

where N is the number of MC samples. The variance of this estimator is given by

$$\text{Var}[\hat{Q}_{M,N}^{\text{SMC}}] = \frac{1}{N} \text{Var}[Q]. \quad (4.3)$$

A goal is to use $\mathbb{E}[Q_M]$ to approximate $\mathbb{E}[Q]$. To quantify how these quantities differ from each another, we introduce the mean square error

$$MSE\{\hat{Q}_{M,N}^{\text{SMC}}\} = \mathbb{E}[(\hat{Q}_{M,N}^{\text{SMC}} - \mathbb{E}[Q])^2], \quad (4.4)$$

which, by virtue of the unbiased nature of $\hat{Q}_{M,N}^{\text{SMC}}$, can be rewritten as

$$\begin{aligned} MSE\{\hat{Q}_{M,N}^{\text{SMC}}\} &= \mathbb{E} \left[(\hat{Q}_{M,N}^{\text{SMC}} - \mathbb{E}[\hat{Q}_{M,N}^{\text{SMC}}] + \mathbb{E}[\hat{Q}_{M,N}^{\text{SMC}}] - \mathbb{E}[Q])^2 \right] \\ &= \mathbb{E} \left[(\hat{Q}_{M,N}^{\text{SMC}} - \mathbb{E}[\hat{Q}_{M,N}^{\text{SMC}}])^2 \right] + \mathbb{E} \left[(\mathbb{E}[\hat{Q}_{M,N}^{\text{SMC}}] - \mathbb{E}[Q])^2 \right] \\ &= \text{Var}[\hat{Q}_{M,N}^{\text{SMC}}] + (\mathbb{E}[Q_M - Q])^2. \end{aligned} \quad (4.5)$$

This expression shows that MSE of the SMC estimator is given by both the estimator's standard deviation and its bias. One can compute an estimator with a root MSE below a certain threshold ϵ in two steps. First, (4.1) is used to determine the value of M for which $(\mathbb{E}[Q_M - Q])^2 \leq \epsilon^2/2$. Second, the value of N for which

$\text{Var}[\hat{Q}_{M,N}^{\text{SMC}}] \leq \epsilon^2/2$ is determined from (4.3). The drawback of this approach is that it is very computationally expensive. The condition $\text{Var}[\hat{Q}_{M,N}^{\text{SMC}}] \leq \epsilon^2/2$ requires $N = O(\epsilon^{-2})$, which can be rather large if ϵ is small, i.e., if a highly accurate estimator is required. Moreover, if ϵ or α or both are small, then a large M might be required to satisfy $(\mathbb{E}[Q_M - Q])^2 \leq \epsilon^2/2$. Since computing each $Q_M^{(t)}$ becomes more expensive as M increases, this procedure requires computing a large number of expensive samples.

Two alternative strategies can be deployed to improve this situation. The first is to build a sequence of functionals \mathcal{G}_M that maximizes the rate of decay of the bias in (4.1), while keeping the cost per sample reasonable. The second is to compute an estimator whose variance satisfies the required bound at a cost smaller than that of simple Monte Carlo. The first approach is not discussed in this work, i.e., we assume that the sequence of functionals \mathcal{G}_M is given and therefore that α cannot be modified. On the other hand, we argue that multi-level Monte Carlo (MLMC) can be used to reduce the computational cost associated with a certain bound on the estimator's variance.

Let $\{M_l\}_{l=0}^L$ be a sequence in \mathbb{N} , such that $M_L = M$ and $M_l = sM_{l-1}$ for $1 \leq l < L$ and an $s \in \mathbb{N} \setminus \{1\}$. Each value of the index l represents a “level”, and for each level we define an associated random variable Q_{M_l} . This sequence of random variables is used to build a telescopic expansion for the expectation of Q ,

$$\begin{aligned} \mathbb{E}[Q_M] &= \mathbb{E}[Q_{M_L}] - \mathbb{E}[Q_{M_{L-1}}] + \cdots + \mathbb{E}[Q_{M_1}] - \mathbb{E}[Q_{M_0}] + \mathbb{E}[Q_{M_0}] \\ &= \sum_{l=1}^L \mathbb{E}[Q_{M_l} - Q_{M_{l-1}}] + \mathbb{E}[Q_{M_0}] = \sum_{l=0}^L \mathbb{E}[Y_l]. \end{aligned} \quad (4.6)$$

Here $Y_0 \equiv Q_{M_0}$ and $Y_l \equiv Q_{M_l} - Q_{M_{l-1}}$ for $1 \leq l < L$, i.e., the expectation of Q is written as the sum of the expectations of random variables Y_l . For each of these expectations, we introduce an SMC estimator

$$\hat{Y}_{l,N_l}^{\text{SMC}} = \frac{1}{N_l} \sum_{t=1}^{N_l} [Q_{M_l}^{(t)} - Q_{M_{l-1}}^{(t)}],$$

where both $Q_{M_l}^{(t)}$ and $Q_{M_{l-1}}^{(t)}$ are computed for the same $\omega^{(t)} \in \Omega$ independently from each other. Once all estimators of $\mathbb{E}[Y_l]$ have been computed, an MLMC estimator

of Q is computed as

$$\hat{Q}_M^{\text{MLMC}} = \sum_{l=0}^L \hat{Y}_{l, N_l}^{\text{SMC}}. \quad (4.7)$$

The variance of this estimator is the sum of the variances of $\hat{Y}_{l, N_l}^{\text{SMC}}$,

$$\text{Var}[\hat{Q}_M^{\text{MLMC}}] = \sum_{l=0}^L \frac{1}{N_l} \text{Var}[Y_l]. \quad (4.8)$$

An advantage of MLMC is that the MLMC estimator for a certain M (i.e., a certain permissible level of estimator bias) can be cheaper to compute than the SMC estimator for that M [23]. That is because the variances $\text{Var}[Y_l]$ decay to zero as $l \rightarrow \infty$ if Q_M converges to Q in mean square, which implies that one needs progressively fewer samples to reduce $\text{Var}[Y_l]$ as l increases. In other words, most samples are required at the coarsest levels where samples are cheaper to compute, and only a few samples are needed at the more expensive finer levels.

According to Theorem 1 in [23] (see also [35]), these considerations imply that, for a given M , MLMC is significantly cheaper than SMC for a given M under certain assumptions on the decays of the estimator bias, the sequence of variances in (4.8) and the computational cost per sample w_l . Specifically, condition (4.1) is supplemented with the following assumptions:

- a) $\text{Var}[Y_l] = \mathcal{O}(M^{-\beta})$ as $M \rightarrow \infty$ for $\beta > 0$ independent of M ,
- b) $w_l = \mathcal{O}(M^\gamma)$ as $M \rightarrow \infty$ for $\gamma > 0$ independent of M .

For $\alpha \geq \min\{\beta, \gamma\}/2$ the overall cost of computing the MLMC estimator is shown to be much smaller than the cost of computing the SMC estimator. These assumptions hold for linear diffusion problems [23]. We investigate the validity of these assumptions for the nonlinear diffusion problems considered in Chapter 3. While both (4.1) and Assumption a) are likely to hold even for nonlinear problems, the validity of Assumption b) is less clear because solving a deterministic nonlinear PDE necessitates the use of an iterative scheme to solve an associated system of nonlinear algebraic equations. The numerical experiments of section 4.4 show that Assumptions a) and b) are satisfied (albeit in a weaker sense) for the steady nonlinear diffusion problem considered.

4.3 Implementation of MLMC

A general outline of our implementation of MLMC is as follows:

Require: Zeroth-level M_0 and bound on root MSE ϵ .

- 1: Determine the smallest L such that $M = M_L$ satisfies

$$|\mathbb{E}[Q_M - Q]| \leq \epsilon_{\text{bias}}, \quad \epsilon_{\text{bias}} = \epsilon/\sqrt{2}. \quad (4.9)$$

- 2: For each level l , compute the number of samples $\{N_l\}_{l=1}^L$ in

$$\text{Var}[\hat{Q}_M^{\text{MLMC}}] = \sum_{l=0}^L \frac{1}{N_l} \text{Var}[Y_l] \leq \epsilon_{\text{est}}, \quad \epsilon_{\text{est}} = \epsilon/\sqrt{2}, \quad (4.10)$$

that minimizes the total computational cost

$$W[\hat{Q}_M^{\text{MLMC}}] = \sum_{l=0}^L \sum_{t=1}^{N_l} w_l^{(t)}, \quad (4.11)$$

where $w_l^{(t)}$ is the cost per sample at the l -th level.

- 3: Compute the MLMC estimator.

Determination of L and $M = M_L$ that satisfy (4.9) requires the knowledge not only of the asymptotic behavior of the estimator bias, but also of its actual value. Since the latter is generally not known *a priori*, a guess of L has to be sequentially refined. Let M' denote the level at which the asymptotic regime of (4.1) takes hold, i.e, $|\mathbb{E}[Q_M - Q]| \approx CM^{-\alpha}$ for $M > M'$. Since $M_l = s^\alpha M_{l-1}$, it follows from (4.6) that for $M_{l-1} > M'$

$$\begin{aligned} |\mathbb{E}[Q_{M_{l-1}} - Q]| &= s^\alpha |\mathbb{E}[Q_{M_l} - Q]| = |\mathbb{E}[Q_{M_l} - Q] - \mathbb{E}[Q_{M_l} - Q_{M_{l-1}}]| \\ &\leq |\mathbb{E}[Q_{M_l} - Q]| + |\mathbb{E}[Y_l]|. \end{aligned}$$

That allows us to replace the condition $|\mathbb{E}[Q_{M_L} - Q]| \leq \epsilon_{\text{bias}}$ in (4.9) with

$$(s^\alpha - 1)|\mathbb{E}[Q_{M_L} - Q]| \leq |\mathbb{E}[Y_L]| \leq (s^\alpha - 1)\epsilon_{\text{bias}}. \quad (4.12)$$

Instead of determining L in advance, we progressively increase L until (4.12) is satisfied.

Once L is determined, the next task is to compute $\{N_l\}_{l=1}^L$. We assume that for a certain level l the cost per sample $w_l^{(t)}$ has a well defined stationary average \bar{w}_l . Then the total simulation cost is

$$W[\hat{Q}_M^{\text{MLMC}}] = \sum_{l=0}^L N_l \bar{w}_l. \quad (4.13)$$

By treating N_l as continuous variables, minimizing (4.13) subject to the constraint (4.10), and rounding off the solution of this optimization problem to the next largest integer yields [28]

$$N_l = \text{ceil} \left[\frac{C_{N_L}}{\epsilon_{\text{est}}^2} \sqrt{\frac{\text{Var}[Y_l]}{\bar{w}_l}} \right], \quad C_{N_L} = \sum_{k=0}^L \sqrt{\text{Var}[Y_k] \bar{w}_k}. \quad (4.14)$$

Equation (4.14) expresses each N_l in terms of the corresponding variance $\text{Var}[Y_l]$ and cost-per-sample \bar{w}_l . These latter quantities are computed by embedding the estimation of N_l into the sequential algorithm that progressively increases L :

Require: Zeroth-level M_0 , bound on root MSE ϵ and maximum permissible number of levels L_{max}

- 1: **for** $L = 0$ to L_{max} **do**
- 2: Estimate $\text{Var}[Y_L]$ and \bar{w}_L from an arbitrary number of samples N_{est}
- 3: Compute C_{N_L} and N_L using (4.14)
- 4: Update $\text{Var}[Y_L]$, $\mathbb{E}[Y_L]$ and \bar{w}_L using $N_L - N_{\text{est}}$ samples
- 5: $N_L^* \leftarrow N_L$
- 6: **for** $l = 0$ to $L - 1$ **do**
- 7: Compute N_l using C_{N_L} and (4.14)
- 8: Update $\text{Var}[Y_L]$, $\mathbb{E}[Y_l]$ and \bar{w}_L using $N_l - N_l^*$ samples
- 9: $N_l^* \leftarrow N_l$
- 10: **end for**
- 11: **if** $L \geq 1$ and $(s^\alpha - 1)\epsilon_{\text{bias}} \geq \mathbb{E}[Y_L]$ **then**
- 12: Finish
- 13: **end if**
- 14: **end for**

Note that this algorithm doesn't require knowledge of the decay rate of $\text{Var}[Y_l]$ and the growth rate of \bar{w}_l . These quantities are estimated on-the-fly as the

number of levels L is increased. The interplay between these rates determines the level at which most of the computational effort occurs [23]. It follows from (4.14) and (4.13) that $\sqrt{W} \sim \sum_l \sqrt{\text{Var}[Y_l] \bar{w}_l}$. Therefore, if $\text{Var}[Y_l]$ decreases faster than \bar{w}_l increases, most of the work occurs at the zeroth-level. Conversely, if $\text{Var}[Y_l]$ decreases slower than \bar{w}_l increases, most of the work occurs in the final level.

4.4 Numerical experiments

We use the elliptic boundary-value problem (3.10)–(3.12) to illustrate our MLMC. The quantity of interest in this problem is the expectation (ensemble mean) of the pressure head $u(z)$ at $z = 0$, i.e., $Q \equiv u(0)$ in the notation of this chapter. We assume that the saturated conductivity $K_s(z)$ is a second-order stationary random field with log-normal point PDF, and that log-conductivity $Y(z) = \ln K_s(z)$ has an exponential covariance function

$$C_Y(x, x') = \sigma_Y^2 \exp\left(-\frac{|x - x'|}{\lambda}\right)$$

with correlation length $\lambda = 0.3$.

To facilitate comparison with the stochastic collocation technique used in Chapter 3, we use the same finite element (FE) discretization with first-order uniform elements (see Section 3.4 for details). The hierarchy of MC levels is constructed by using different levels of spatial discretization, with the number M of FE elements given by $M = s^l M_0$ with level index $l \geq 0$ and $s = 2$. The spatial discretization produces a set of nonlinear algebraic equations, which is solved using the KINSOL library [40] implementation of the modified Newton’s method with line search.

The MLMC estimator of the $\mathbb{E}[Q]$ is computed for two levels of variance, $\sigma_Y^2 = 0.1$ and 1.0. Each repetition of the log-conductivity field $Y(z)$ is generated using its K-L decomposition with 800 terms. The use of a high number of the K-L terms allows us to sidestep altogether the problem of the estimator bias introduced by a truncated K-L expansion mentioned in Section 3.2. In the experiments reported below we assume that the estimator bias due to the finite K-L representation is negligible (3.10)–(3.12).

The computational cost at each Monte Carlo level and its behavior with increasing level are worthwhile discussing. For the t -th realization of parameter field $Y(x)$ at the l -th MC level, the computational cost associated with approximating $Y_l^{(t)}$ is that of approximating $Q_l^{(t)}$ and $Q_{l-1}^{(t)}$. Per the discussion in Section 3.4, for the numerical scheme described above these costs are respectively $k_l^{(t)}M_l$ and $k_{l-1}^{(t)}M_{l-1}$, where $k_l^{(t)}$ and $k_{l-1}^{(t)}$ are coefficients associated with the numerical solution of the respective systems of nonlinear algebraic equations. The coefficient $k_l^{(t)}$ includes information about the costs of repeated evaluations of the nonlinear functional and an approximation of its Jacobian, as well as of repeated solutions of the corresponding tridiagonal linear systems. Each of these partial costs is proportional to M_l in the case of the one-dimensional boundary-value problem. Finally, we assume that other costs, such as that of generating the random samples and evaluating the K-L expansion, are negligible compared to the cost of solving the deterministic PDE problems.

Given that $M_l = s^{-1}M_{l-1}$, we approximate the cost associated with the t -th sample at the l -th level by $(k_l^{(t)} + s^{-1}k_{l-1}^{(t)})M_l$. The cost associated with the l -th level is $N_l\bar{k}_lM_l$, where N_l is the number of samples for the l -th level and \bar{k}_l is the average of $(k_l^{(t)} + s^{-1}k_{l-1}^{(t)})$ over the N_l samples. Our numerical experiments revealed that for this problem \bar{k}_l is approximately constant across all levels. Therefore we introduce a “standardized cost” associated with each MC level by disregarding the constant \bar{k}_l and only taking into account the scaling factor M_l and the number of repetitions. This standardized cost is N_lM_l for each level, and the total standardized cost is the sum of each cost per level for all MC levels. We use this standardized cost to compare simple and multi-level MC estimators.

Figure 4.1 exhibits estimators of the (absolute value of) mean and variance of both Q_l and Y_l for $\sigma_Y^2 = 0.1$. These results were computed using the MLMC algorithm with 4 levels and the standard deviation target $\epsilon_{\text{est}} = \sqrt{2} \times 10^{-5}$. Figure 4.1 (a) demonstrates that $|\mathbb{E}[Y_l]|$ decays as $M^{-\alpha}$ with $\alpha \approx 0.89$. This finding experimentally verifies the hypothesis that the discretization bias $|\mathbb{E}[Q_{M_L} - Q]|$ decays as $M^{-\alpha}$. Figure 4.1 (b) shows that $\text{Var}[Y_l]$ decays as $M^{-\beta}$ with $\beta \approx 2.0$.

Since the variance of Y_l decays faster ($\beta \approx 2.0$) than the computational

cost per sample ($\gamma \approx 1.0$), the concluding remarks of Section 4.3 suggest that the majority of the computational cost is incurred at the zeroth MC level. In fact, for $\sigma_Y^2 = 0.1$, $\epsilon_{\text{est}} = \sqrt{2} \times 10^{-5}$, $M_0 = 16$ and $L = 4$ the zeroth level accounts for 46 % of the total computational cost.

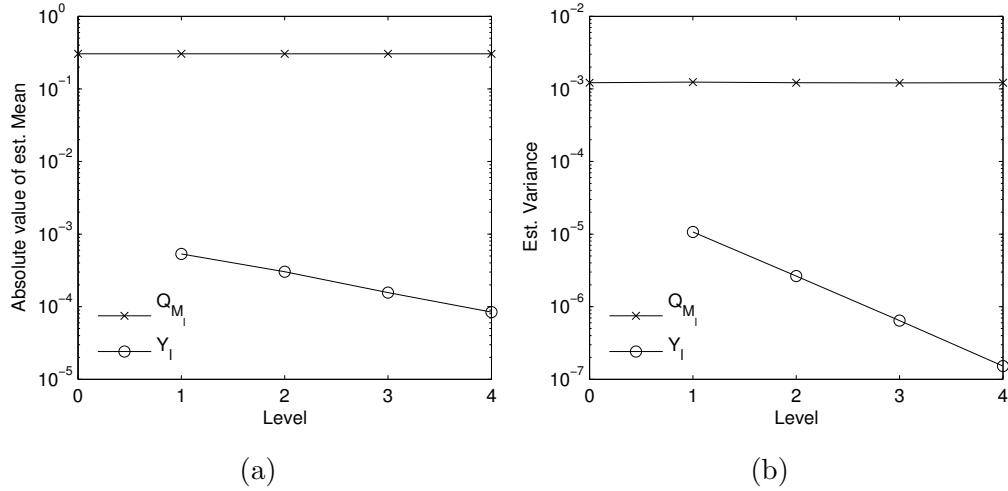


Figure 4.1: Absolute value mean (a) and variance (b) of Q_L and Y_L for the elliptic boundary-value problem (3.10)– (3.12) and $\sigma_Y^2 = 0.1$.

Figure 4.2 shows estimators of the mean and variance of both Q_l and Y_l , this time for $\sigma_Y^2 = 1.0$. These results were computed using the MLMC algorithm with 4 levels and a standard deviation target $\epsilon = 5\sqrt{2}/2 \times 10^{-5}$. The behavior is similar to that observed for $\sigma_Y^2 = 0.1$. First, the absolute value of the expectation of Y_l decays approximately as $M^{-\alpha}$ but with a smaller coefficient $\alpha \approx 0.84$, and the first level expectation $|\mathbb{E}[Y_1]|$ is larger for $\sigma_Y^2 = 1.0$. In other words, $|\mathbb{E}[Y_l]|$ not only decays (slightly) slower, but also starts from a higher value, so that reducing the discretization bias $|\mathbb{E}[Q_{M_L} - Q]|$ to an acceptable level, i.e., verifying the condition $(s^\alpha - 1)|\mathbb{E}[Y_L]| \leq \epsilon_{\text{est}}$, takes more levels for larger variances of the model parameters. Second, the variance of Y_l decays as $M^{-\beta}$ with $\beta \approx 2.0$, but both $\text{Var}[Y_l]$ and $\text{Var}[Q_l]$ are an order of magnitude larger for $\sigma_Y^2 = 1.0$ than for $\sigma_Y^2 = 0.1$. Hence the number of samples required to satisfy a certain bound on the standard deviation of the estimator increases. We conclude that increasing the variance of the model parameters affects the computation of the MLMC estimator

satisfying a certain bound on its RMSE in two ways. It decreases the rate of decay of the discretization bias, requiring more levels to satisfy the bound; and it increases the number of samples per level required to satisfy the bound on the standard deviation of the estimator.

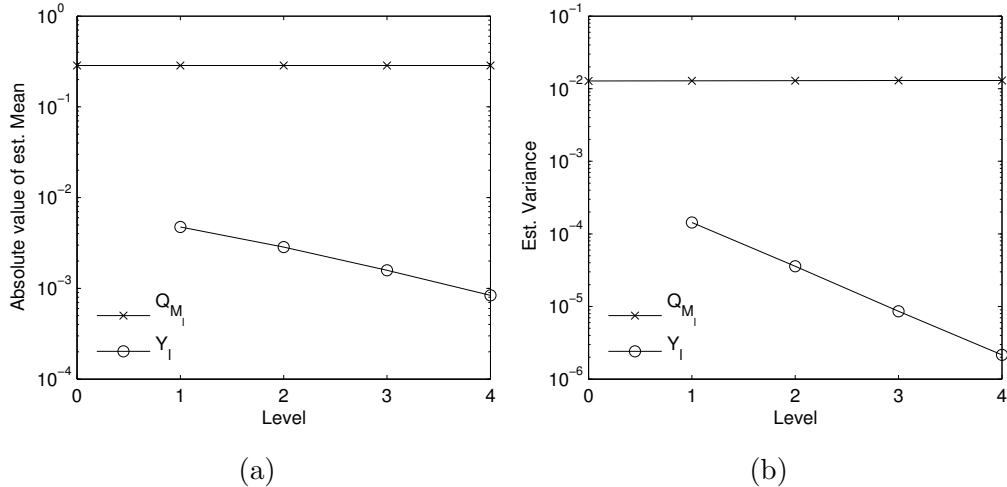


Figure 4.2: Absolute value mean (a) and variance (b) of Q_L and Y_L for the elliptic boundary-value problem (3.10)– (3.12) and $\sigma_Y^2 = 1.0$.

The dependence of the scaling exponent α on the variance of the random input parameters is determined not only by the properties of the governing PDE but also by a numerical scheme used to approximate the quantity of interest at a certain level. It is expected that the use of a numerical scheme different from the one used in our analysis would affect the value and behavior of α . The selection of a numerical scheme that maximize the decay of the discretization bias with increasing MLMC levels is a subject for future studies.

Next, we examine the relative performance of our MLMC approach and simple Monte Carlo in terms of the computational cost associated with a given bound on the RMSE of the estimator. Since the discretization bias depends on the discretization level but not on the method used to compute the estimator, we specify the final level as $M_L = 256$ and compute the estimator of $\mathbb{E}[Q_L]$ using simple MC, MLMC with 2 levels ($M_0 = 64$, $M_1 = 128$, $M_2 = 256$) and MLMC with 4 levels ($M_0 = 16$, $M_1 = 32$, \dots , $M_4 = 256$) for different target values of the

standard deviations of the estimator ϵ_{est} .

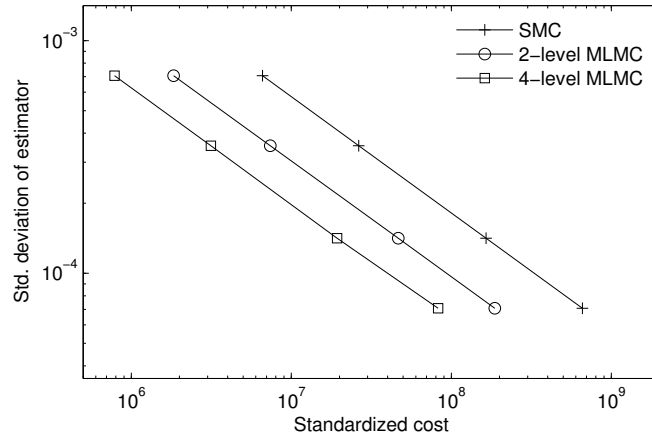


Figure 4.3: Standardized cost of the simple MC and MLMC estimators for different target values of the standard deviations of the estimator and $\sigma_Y^2 = 0.1$. Final discretization level is fixed at $M_L = 256$.

Figure 4.3 shows the standardized cost as a function of ϵ_{est} for $\sigma_Y^2 = 0.1$. The cost curves of the different estimators follow the same power-law behavior, albeit with different leading coefficients. The 2-level MLMC estimator and 4-level MLMC estimators have a cost approximately 3.5 times and 8.5 times lower than the simple MC estimator, respectively. The fact that MLMC provides significant savings relative to simple MC is a direct consequence of its general properties discussed earlier. Since the variance of Y_l decays faster than the standardized cost per sample at the l -th level, the majority of the computational effort of MLMC takes place in the lowest MC level. This explains why increasing the number of levels in MLMC (e.g., going from 2 to 4 levels) reduces the cost of the MLMC estimator with respect to the equivalent simple MC estimator.

Although not shown in Figure 4.3, the addition of an extra level (i.e., starting at $M_0 = 8$) yields the cost savings of 9.8 relative to simple MC. This further increase in efficiency provides an evidence of diminishing returns associated with introducing further levels in the lower end. This phenomenon occurs because decreasing M_0 not only decreases the cost per sample but also increases the variance of Y_1 [23]. At a certain point, the cost at level 1 would be comparable to the cost

at level 0 (due to an increase in the number of samples), undermining the efficiency of MLMC over simple Monte Carlo.




Figure 4.4: Standardized cost of the simple MC and MLMC estimators for different target values of the standard deviations of the estimator and $\sigma_Y^2 = 1.0$. Final discretization level is fixed at $M_L = 256$.

Figure 4.4 shows the standardized cost as a function of ϵ_{est} , this time for $\sigma_Y^2 = 1.0$. As expected, given a certain standard deviation target, the cost of the corresponding estimators is higher than in the case of $\sigma_Y^2 = 0.1$. Yet the cost ratios between the different estimators remain approximately the same: the 2-level and 4-level MLMC estimators outperform the simple MC estimator (in terms of the computational cost) by the factor of 3.5 and 8.5, respectively. Hence the efficiency gained by bending the cost curves via the use of the MLMC algorithm transfers to high variance cases.

4.5 Conclusions and further work

The multi-level Monte Carlo (MLMC) framework for computing statistics of quantities of interest offers a systematic way to deal with the estimator bias introduced by replacing the quantities of interest Q with a sequence of approximations Q_M that converges to Q as $M \rightarrow \infty$. The use of MLMC also reduces the computational cost with respect to simple Monte Carlo. The systematic treatment of the

discretization bias and the robustness of MLMC estimators in high-dimensional probability spaces alleviate the difficulties encountered by stochastic collocation techniques in Chapter 3.

In the follow-up MLMC-related investigations we will apply it to non-linear transient problems similar to those described in Section 3, and explore the behavior of $\mathbb{E}[Y_l]$ and $\text{Var}[Y_l]$ with increasing MC level and computational cost per sample. We will also study the relationship between the decay of the discretization bias and the discretization scheme used to solve the governing PDEs. This will enable us to select an efficient scheme that maximizes the rate of decay of the bias.

Chapter 4, in part, is currently being prepared for submission for publication of the material. Barajas-Solano, D. A., Tartakovsky, D. M. The dissertation author was the primary investigator and author of the paper.

Chapter 5

Efficient Reconstruction of Contaminant Release History

5.1 Introduction

An efficient and accurate reconstruction of contaminant release history is essential to regulatory and remedial efforts. Most such efforts rely on measurements of pollutant concentration to identify sources and/or release history of a pollutant. Quite often available concentration data are corrupted by measurement errors. Additional complications arise from heterogeneity and insufficient site characterization, although we do not consider these effects here.

A detailed review of the state-of-the-art in the field of inverse modeling as related to contaminant source identification is presented in [7]. The existing approaches can be subdivided into two broad classes: deterministic and probabilistic. Deterministic approaches include, but are not limited to, Tikhonov regularization of convolution integrals [77, 56], least-square estimation from analytical approximations [15], least-square solution of an optimal control problem [37], the method of quasi-reversibility [78], and the backward beam equation method [6, 10]. These approaches provide estimates of the release history from a source of known locations, without a clear way to quantify the uncertainty associated with such estimates. The robustness of these methods is highly sensitive to measurement errors, and

more often than not their mathematical formulations are fundamentally ill-posed.

While existing probabilistic approaches, such as random walk particle tracking for the backward transport equation [11], the minimum relative entropy [99] and adjoint methods (e.g., [67]), alleviate some of these problems, others remain. For example, these and similar methods do not take advantage of the regularizing nature of the measurement noise and, hence, are often ill-posed. Thus, the minimum relative entropy method treats concentration measurements as ensemble averages. Additionally, there are some outstanding issues with quantifying uncertainty [66] and the inability of many existing approaches to handle more than one observation point [70].

Finally, most existing approaches to the reconstruction of release history are restricted to linear transport phenomena, i.e., to migration of contaminants that are either conservative (all the references above) or exhibit first-order (linear) reaction rates [68, 69]. This is because such approaches are based on either Green's functions [77, 99] or analytically derived adjoint equations [67, 70]. The use of Kalman filters for source identification [39] is formally limited to linear transport phenomena and Gaussian errors. While both limitations can be somewhat relaxed by employing various generalizations of the Kalman filter (e.g., the extended Kalman filter and the ensemble Kalman filter), their use is suboptimal and is known to fail if the nonlinearity is too strong.

Purely statistical approaches to history reconstruction, such as the geostatistical inversion with Bayesian updating proposed in [80], are applicable to nonlinear transport. Since this is achieved by ignoring governing equations, the reconstructed release histories could have non-physical characteristics, including negative concentrations. These problems have been alleviated by introducing additional constraints into an optimization functional and requiring the reconstructed field to be Gaussian [59, 58]. Combining these geostatistical approaches with analytically derived adjoint equations [60, 76] however brings back the linearity requirement.

We present an optimal reconstruction of contaminant release history that fully utilizes all available information and requires neither the linearity of governing

transport equations nor the Gaussianity of the underlying fields. In Section 5.2 we formulate the problem of reconstructing the contaminant release history from noisy observations. Section 5.3 introduces our general computational framework, which is further implemented in Section 5.4 for various examples.

5.2 Problem formulation

We begin with the premise that subsurface migration of a chemically active contaminant in a porous medium Ω is adequately described by an advection-dispersion-reaction equation with a reaction term $R(c)$,

$$\frac{\partial c}{\partial t} = \nabla \cdot (\mathbf{D}\nabla c) - \nabla \cdot (\mathbf{u}c) - R(c) + r(\mathbf{x}, t), \quad \mathbf{x} \in \Omega. \quad (5.1)$$

Here $c = c(\mathbf{x}, t)$ is the solute concentration at point \mathbf{x} and time t , \mathbf{u} is the average (macroscopic) pore velocity, and \mathbf{D} is the dispersion coefficient tensor. Both the location and duration of the contaminant release, i.e., the source function $r(\mathbf{x}, t)$, can be unknown, but only the former source of uncertainty is treated in the computational examples of section 5.4). Our goal is to reconstruct the release history $r(\mathbf{x}, t)$ from concentration data $\bar{c}_{mi} = \bar{c}(\mathbf{x}_m, t_i)$ collected at points $\{\mathbf{x}_m\}$ ($m = 1, \dots, M$) at times t_i ($i = 1, \dots, I$).

Concentration measurements are corrupted by measurement errors. We assume that the measured concentrations \bar{c}_{mi} differ from the true concentration by an additive measurement noise, so that

$$\bar{c}_{mi} = c(\mathbf{x}_m, t_i) + \epsilon_{mi}, \quad (5.2)$$

where the errors ϵ_{mi} are zero-mean Gaussian random variables described by the covariance matrix $\mathbb{E}[\epsilon_{mi}\epsilon_{nj}] = \delta_{ij}R_{mn}$, where $\mathbb{E}[\cdot]$ is the expectation operator, δ_{ij} denotes the Kronecker delta function, and R_{mn} are components of the spatial covariance matrix \mathbf{R} of measurement errors. This treatment of measurement noise assumes that the measurements are well separated in time to neglect any temporal correlations, but the model can be easily extended to include temporal correlations. We use the additive error model (5.2) of [99], rather than the multiplicative

error model of [77], for the purpose of illustration only. Both models have similar effects on the accuracy of history reconstruction [66] and can be handled by our approach.

In a typical situation, one has prior information (or a belief) about potential sources of contamination (a region Ω_c within the flow domain Ω) and a time period $[T_l, T_u]$ during which the release has occurred. Examples of Ω_c include spatially distributed zones of contamination (e.g., landfills) and a collection of point sources (e.g., localized/small industrial sites or storage facilities) some of which have contributed to contamination. The lower (T_l) and upper (T_u) bounds of the release interval might represent the time when a landfill became operational and the time when contamination has first been detected, respectively. In the absence of prior information about the release occurrence, one can assume a uniform random distribution of the release in $[T_l, T_u] \times \Omega$. We allow for an arbitrary number of measurement points and for either discrete or continuous-in-time measurements.

5.2.1 Likelihood function

To simplify the exposition, we assume a spatially distributed chemical release at time $t = 0$ only, i.e., $r(\mathbf{x}, t) = c(\mathbf{x}, 0)\delta(t)$. Given the measurements $\bar{c}(\mathbf{x}_m, t_i)$ and the noise model (5.2) we would like to determine the likelihood of a given release configuration $c(\mathbf{x}, 0)$. Unfortunately, the measurements, generally taken at later times, do not estimate directly the likelihood of a release configuration. Nevertheless, because the transport equation (5.1) is deterministic, we can implicitly assess the likelihood of a given release configuration $P[c(\mathbf{x}, 0)]$, from the probability (likelihood) of a given (computed) concentration history $c(\mathbf{x}, t)$. This likelihood can be expressed as [2]

$$P[c(\mathbf{x}, 0)] \sim \exp\{-\tilde{H}_{\text{obs}}[c(\mathbf{x}, t)]\}, \quad (5.3)$$

where $\tilde{H}_{\text{obs}}[c(\mathbf{x}, t)]$ is the so-called ‘‘Hamiltonian’’ or log-likelihood function,

$$\tilde{H}_{\text{obs}}[c(\mathbf{x}, t)] = \frac{1}{2} \sum_{m,n=1}^{M,I} \Delta_{mi}(\mathbf{R}^{-1})_{mn} \Delta_{ni}, \quad (5.4)$$

with $\Delta_{mi} \equiv c(\mathbf{x}_m, t_i) - \bar{c}(\mathbf{x}_m, t_i)$. Since (5.1) uniquely determines the evolution of the solute concentration from its initial state $c(\mathbf{x}, 0)$, the Hamiltonian (5.4) is a non-linear functional of the initial conditions $c(\mathbf{x}, 0)$, i.e., $\tilde{H}_{\text{obs}}[c(\mathbf{x}, t)] = H_{\text{obs}}[c(\mathbf{x}, 0)]$.

This formulation assumes that the measurement errors ϵ_{mi} are Gaussian and uncorrelated with the state of the system. Other distributions of the measurement noise and the stochasticity of governing equations can be handled as well [2]. The Hamiltonian for stochastic systems (which can represent, e.g., uncertain hydraulic conductivity and flow velocity that are treated as random fields) can be reformulated to explicitly include the dynamical equations [2, 3].

The contribution of highly fluctuating, unphysical initial conditions is reduced by adding a regularization term $H_{\text{reg}}[c(\mathbf{x}, 0)]$ to the observation Hamiltonian (5.4) and replacing the likelihood function (5.3) with

$$P[c(\mathbf{x}, 0)] \sim \exp\{-H[c(\mathbf{x}, 0)]\}, \quad (5.5a)$$

where

$$H[c(\mathbf{x}, 0)] = H_{\text{obs}}[c(\mathbf{x}, 0)] + \gamma H_{\text{reg}}[c(\mathbf{x}, 0)] \quad (5.5b)$$

and the weight γ is a tuning hyperparameter. The regularization term H_{reg} is equivalent to a Bayesian prior on the initial condition. The selection of an appropriate regularization Hamiltonian is particularly important for problems in which the observation Hamiltonian doesn't define a proper probability distribution for $c(\mathbf{x}, 0)$ due to a lack of measurements. For a one-dimensional source profile, the squared gradient of the initial spatial profile can play a role of the regularization Hamiltonian. In higher dimensions, one can use a thin-plate penalty functional [93].

A conceptual difference between our approach and maximum likelihood methods is worthwhile emphasizing. Rather than sampling the Gibbs distribution $\exp\{-H[c(\mathbf{x}, 0)]\}$, as we do here, maximum likelihood methods minimize the Hamiltonian (5.5b) over $c(\mathbf{x}, 0)$. While standard maximum likelihood methods determine the mode and variance of the posterior distribution under a Gaussian approximation, the approach described below can be used to determine the mean and higher-order statistics and it is valid even when the posterior distribution is multi-modal.

5.3 Monte Carlo sampling

In principle, one can sample the Gibbs distribution by using Markov-chain Monte Carlo (MCMC) (e.g., [58]). However, quite often, the disadvantage of local MCMC-based methods is their slow convergence. To improve convergence, we apply a Generalized Hybrid Monte Carlo (gHMC), which enables one to efficiently sample release configurations $c(\mathbf{x}, 0)$ with probability given by (5.5a)

5.3.1 Hybrid Monte Carlo (HMC)

Hybrid Monte Carlo (HMC) refers to a class of methods that combine Hamiltonian molecular dynamics with Metropolis-Hastings Monte Carlo simulations (see [62] for an introductory survey). Specifically, a time-discretized integration of the molecular dynamics equations is used to propose a new configuration, which is then accepted or rejected by the standard Metropolis-Hastings Monte Carlo criteria. The change in total energy serves as the acceptance/rejection criteria.

In HMC one treats the log-likelihood function H in (5.5b) as the configurational Hamiltonian for a system of N “particles”, each of which has unit mass and generalized coordinates q_1, q_2, \dots, q_N . Each of these generalized positions corresponds to the solute concentration $c(\mathbf{x}, t)$ at a space-time point (\mathbf{x}, t) . In the following, the particle positions correspond to the initial concentration at time $t = 0$, e.g., $q_i = c(x_i, 0)$, $x_i = iL/(N - 1)$, $i = 0, \dots, N - 1$ for a contaminant release over the one-dimensional domain $[0, L]$.

At any given time, the state of the system is completely described by (\mathbf{q}, \mathbf{p}) , where $\mathbf{q} = \{q_i\}_{i=1}^N$ and $\mathbf{p} = \{p_i\}_{i=1}^N$. Here, the momentum of the i -th particle, p_i , is $dp_i/d\tau = p_i$, where τ is the fictitious time of the molecular dynamics. The kinetic energy of the system of N particles is given by

$$H_K(\mathbf{p}) = \frac{1}{2} \sum_{i=1}^N |p_i|^2, \quad (5.6)$$

and the total Hamiltonian of the system is

$$\hat{H}(\mathbf{q}, \mathbf{p}) = H(\mathbf{q}) + H_K(\mathbf{p}). \quad (5.7)$$

The Hamiltonian dynamics of such a system is

$$\frac{dq_i}{d\tau} = p_i, \quad \frac{dp_i}{d\tau} = F_i, \quad F_i = -\frac{\partial H}{\partial q_i}, \quad (5.8)$$

where F_i is the force acting on the i -th “molecule” that is to be computed from the governing transport equation. During the time interval $\Delta\tau$, the system evolves from its current state (\mathbf{q}, \mathbf{p}) to a new state $(\tilde{\mathbf{q}}, \tilde{\mathbf{p}})$, which can be computed by discretizing the Hamiltonian dynamics (5.8). An example of such a discretization is the standard leapfrog method, which may be written as

$$\tilde{q}_i = q_i + \Delta\tau p_i + \frac{\Delta\tau^2}{2} F(\mathbf{q}) \quad (5.9a)$$

$$\tilde{p}_i = p_i + \frac{\Delta\tau}{2} \{F(\mathbf{q}) + F(\tilde{\mathbf{q}})\}, \quad (5.9b)$$

for $i = 1, \dots, N$. Multiple dynamic updates, i.e., multiple applications of equation (5.9), can be performed. For the hybrid Monte Carlo method, the number of updates M is larger than one. For $M = 1$ we obtain the Langevin Monte Carlo method [62]. This completes the “proposal part” of HMC.

The remaining part of HMC consists of deciding whether to accept or reject the new state $(\tilde{\mathbf{q}}, \tilde{\mathbf{p}})$. This is done by the Metropolis-Hastings sampling strategy, according to which the new state $(\tilde{\mathbf{q}}, \tilde{\mathbf{p}})$ is accepted with probability

$$Q = \min \left\{ 1, \exp\{\hat{H}(\mathbf{q}, \mathbf{p}) - \hat{H}(\tilde{\mathbf{q}}, \tilde{\mathbf{p}})\} \right\}. \quad (5.10)$$

The momenta variables $\tilde{\mathbf{p}}$ are resampled after each acceptance/rejection of the position variables according to a Gaussian distribution of independent variables $\exp(-H_K)$. The time-marching and acceptance/rejection process represents one step in the Markov chain, and therefore one Monte Carlo sample. It is important to note that the update from (\mathbf{q}, \mathbf{p}) to $(\tilde{\mathbf{q}}, \tilde{\mathbf{p}})$ does not conserve energy as a result of the time discretization. The extent to which energy is not conserved is controlled by the time step $\Delta\tau$. Detailed balance is achieved if the configuration obtained after evolving several steps is accepted with probability Q in (5.10). Thus, the Metropolis step corrects for time discretization errors.

As we have noted before, the method samples from the multivariate target distribution, $\sim \exp(-\hat{H})$, by generating a Markov chain. Sampling from this

density allows us to estimate the mean state (initial configuration) and its variance. Markov chain sampling from the posterior distribution involves a *transient phase*, in which we start from some initial state and simulate the Markov chain for a period long enough to reach its stationary density, followed by a *sampling phase*, in which we assume that the Markov chain visits states from this stationary density. If the chain has converged, and the sampling phase is long enough to cover the entire posterior distribution, accurate inferences about any quantity of interest are made by computing the sample mean, variance and other desired statistics [50].

5.3.2 Generalized Hybrid Monte Carlo (gHMC)

In many cases the generalized hybrid Monte Carlo (gHMC) of [90] can improve the efficiency of standard HMC via a nonlocal sampling strategy described in some detail below. For $\mathbf{q}, \mathbf{p} \in \mathbb{R}^N$, gHMC replaces the Hamiltonian dynamics in (5.8) with a more general formulation,

$$\frac{d\mathbf{q}}{d\tau} = \mathbf{A}\mathbf{p}, \quad \frac{d\mathbf{p}}{d\tau} = \mathbf{A}^\top \mathbf{F}, \quad (5.11)$$

where \mathbf{A} is a linear operator represented by a $\mathbb{R}^{N \times N}$ matrix. The corresponding leapfrog discretization is

$$\mathbf{q}(\delta\tau) = \mathbf{q} + \delta\tau \mathbf{A}\mathbf{p} + \frac{\delta\tau^2}{2} \mathbf{A}\mathbf{A}^\top \mathbf{F}[\mathbf{q}], \quad (5.12a)$$

$$\mathbf{p}(\delta\tau) = \mathbf{p} + \mathbf{A}^\top \frac{\delta\tau}{2} \{\mathbf{F}[\mathbf{q}] + \mathbf{F}[\mathbf{q}(\delta\tau)]\}. \quad (5.12b)$$

The two formulations, (5.8) and (5.11), are identical if \mathbf{A} is the identity matrix. The goal is to find a matrix \mathbf{A} that leads to a significant reduction of the temporal correlations of the Markov chain without appreciably increasing the cost of the update due to matrix-vector multiplications.

In order to illustrate how the introduction of the matrix \mathbf{A} can help reducing the correlations of the Markov chain, consider the problem with $\mathbf{q} \in \mathbb{R}^N$ and Hamiltonian

$$H(\mathbf{q}) = -\frac{1}{2}(\mathbf{q} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1}(\mathbf{q} - \boldsymbol{\mu}),$$

so that the forcing is given by $\boldsymbol{\Sigma}^{-1}(\mathbf{q} - \boldsymbol{\mu})$. For the case $\mathbf{A} = \mathbf{I}$, it can be seen from (5.12) that the different components q_i are updated at different rates, given

by the covariance matrix Σ . For a given $\delta\tau$ some components would be updated with long steps, while others would be updated with shorter steps.

The disadvantage of such a configuration is that too long of a step for a certain component might increase the total Hamiltonian enough to produce a rejection according to (5.10). If the rejection rate of the chain is too large, one would have to reduce $\delta\tau$, which affects all components. The issue of the rejection rate would be addressed, but then some components would be updated with very short steps, increasing their correlation. To solve this issue, one can remove the appearance of Σ altogether by choosing \mathbf{A} such that $\mathbf{A}\mathbf{A}^\top\Sigma = \mathbf{I}$. If Σ is a valid covariance matrix, this is trivially accomplished by choosing \mathbf{A} as the Cholesky factor of Σ .

Unfortunately, in general, the Hamiltonian (5.5b) for our problem doesn't have a simple bilinear form for which an appropriate selection of the matrix \mathbf{A} can be derived. Nevertheless, it stands to reason that one can build a matrix \mathbf{A} for more complex systems that still reduces to an extent the correlation of the Markov chain.

5.4 Numerical applications

In this section we study various applications that show how the framework outlined above can be applied to source identification problems. In the first two test cases we study the implementation of HMC to contaminant transport problems with a nonlinear reaction term for different configurations of observations. In the third case we study a linear advection-dispersion problem and explore possible selections of the gHMC acceleration matrix.

5.4.1 Continuous in space, discrete in time measurements

We consider a one-dimensional source profile and a regularization operator, which is the ℓ_2 -norm of the gradient of the initial spatial distribution. Furthermore, we assume that the time of the contaminant release is precisely known and we impose no constraints on the total mass of the released contaminant. We also assume

that measurements are taken at time T at all space points in the domain and that the observations have independent, Gaussian-distributed errors with variance $R_{mn} = 2\sigma_\epsilon^2\delta_{mn}$. Choosing a continuous formulation, the Hamiltonian $H[c(\mathbf{x}, 0)]$ in (5.5b) with the continuous version of (5.4) takes the form

$$H = \frac{1}{2\sigma_\epsilon^2} \int_{\Omega} [c(\mathbf{x}, T) - \bar{c}(\mathbf{x}, T)]^2 dx + \gamma \int_{\Omega} |\nabla c(\mathbf{x}, 0)|^2 dx. \quad (5.13)$$

Introducing a method-of-lines discretization of c in terms of basis functions ψ_i such that $c(t) = \sum c_i(t)\psi_i$, the Hamiltonian can be computed as

$$H = \frac{1}{2\sigma_\epsilon^2} \Delta \mathbf{c}^\top \mathbf{M} \Delta \mathbf{c} + \gamma \mathbf{q}^\top \mathbf{M}' \mathbf{q}, \quad (5.14)$$

where $\mathbf{c}(t) = (c_1(t), \dots, c_N(t))^\top$, $\mathbf{q} = \mathbf{c}(0)$ and $\Delta c_i = c_i(T) - \bar{c}_i(T)$. The matrices \mathbf{M} and \mathbf{M}' are symmetric with coefficients

$$M_{ij} = \int_{\Omega} \psi_i \psi_j dx, \quad M'_{ij} = \int_{\Omega} \psi_i \psi_j dx.$$

The sensitivity of H in the discrete form (5.14) can be computed by using the adjoint method for sensibility analysis [16, 53]. A possible approach is to discretize the problem (5.1) into a set of ordinary differential equations (ODEs), and compute the solution $\mathbf{v}(t)$ of the adjoint of the forward ODE problem [53]. This strategy leads to

$$-\mathbf{F} = \nabla_{\mathbf{q}} H = 2\gamma \mathbf{M}' \mathbf{q} + \mathbf{v}(0), \quad (5.15)$$

from which the force vector in (5.8) driving the Hamiltonian dynamics is derived. Details of the derivation are presented in Appendix A.1. For each step of the leapfrog scheme (5.9) we need to solve forward and backward systems of ODEs. Those systems must be integrated accurately enough in time to ensure that the rejection rate of the Markov chain is not negatively affected.

We test this formulation with a one-dimensional transport application with uniform velocity u and dispersion coefficient D . The unknown contaminant release distribution at $t = 0$ and corresponding concentration measurements at a later time, $T = 100\Delta t$, are shown in Figure 5.1. This spatial release was selected to mimic a localized source with additional weak concentration fluctuations superposed.




Figure 5.1: Observations (solid line) and release distribution to be reconstructed (dotted line).

Transport is modeled over a unit domain with no-flow boundary conditions on both ends, and the reaction model [55]

$$R(c) = 2k(c^2 - c_{eq}^2), \quad (5.16)$$

corresponding to a nonlinear heterogeneous (precipitation/dissolution) reaction with equilibrium concentration c_{eq} . Here k denotes the kinetic rate constant normalized by porosity. The parameter values are set $D = 1.0$, $u = 50.0$, $c_{eq} = 0.4$ and $k = 1.0$. The standard deviation of the measurements is set to $\sigma_\epsilon = 0.02$. A space-time discretization with 129 spatial points and $\Delta t = 1 \times 10^{-5}$ is employed.

Observations are available at all spatial points at time T (see Figure 5.1). The release history is then inferred using the HMC scheme described above, carried out with hybrid time-step $\Delta\tau = 0.18$, $M = 5$, and regularizing parameter $\gamma = 0.025$. A total of 4×10^6 samples of the release configuration were computed, and the first half of the chain rejected as the “burn-in” period. Our method reconstructs the configuration that mimics the actual release as shown in Figure 5.2. The variance of the estimate is too small to be noticeable on the scale of Figure 5.2.




Figure 5.2: Actual (solid line) and reconstructed (dotted line) release histories using HMC.

5.4.2 Discrete in space, continuous in time measurements

The second case deals with the same one-dimensional reactive transport problem, and employs the ℓ_2 -norm of the gradient of the initial condition as the regularization term. As in the first case, we assume that the time of the contaminant release is precisely known and no constraints on the total mass of the release contaminant is imposed. In contrast to the first case, the measurements are taken continuously over the time interval $(0, T]$ at a discrete subset J of locations in the spatial domain. We assume that the measurement errors are uncorrelated in space and time, and have the same variance σ_ϵ^2 at every point. This case represents observations of contaminant breakthrough curves at a number of sampling locations.

The Hamiltonian corresponding to this setup is

$$H = \frac{1}{2\sigma_\epsilon^2} \int_0^T \sum_{j \in J} [c(\mathbf{x}_j, t) - \bar{c}(\mathbf{x}_j, t)]^2 dt + \gamma \int_\Omega |\nabla c(\mathbf{x}, 0)|^2 dx. \quad (5.17)$$

It is evaluated, together with its sensitivity with respect to the initial condition, by using a method-of-lines discretization of the concentration field $c(\mathbf{x}, t)$. Once the governing PDE has been discretized into a system of ODEs, one can compute the sensitivity $\nabla_{\mathbf{q}} H$ via the adjoint sensitivity method used in the previous case. The disadvantage of this approach, as mentioned earlier, is that it incurs two levels

of numerical error: the integration error of the forward problem, which affects the initial condition of the adjoint problem; and the integration error of the backward problem. If these errors are significant, both the quality of the estimator and the rejection rate of the Markov chain can be compromised. Reducing the error requires one to decrease the time step used for integration in both directions, which would increase the computational cost per leapfrog step.

To partially alleviate this problem, we use a single-step ODE integration scheme for the forward problem and compute the sensitivity of H with respect to the initial condition via multiple applications of the chain rule [24]. Let \mathbf{c}^i ($i = 0, \dots, I$) be a vector of discretized states evaluated at time $t = iT/I$, and $\tilde{\mathbf{c}}^i$ be a vector of the measurements at time t_i in the elements corresponding to the J measurements locations and zeros in the other elements. Let \mathbf{C} be a diagonal matrix with ones on the diagonal elements corresponding to the J subset of measurement locations, and zeros in all other locations. We use this notation to rewrite the observation Hamiltonian and the sensitivity as

$$H_{\text{obs}}(\mathbf{c}, \mathbf{q}) = \frac{1}{2\sigma_\epsilon^2} \Delta t \sum_{i=1}^I (\mathbf{c}^i - \tilde{\mathbf{c}}^i)^\top \mathbf{C} (\mathbf{c}^i - \tilde{\mathbf{c}}^i),$$

and

$$\nabla_{\mathbf{q}} H_{\text{obs}} = \frac{1}{\sigma_\epsilon^2} \Delta t \sum_{i=1}^I \left(\frac{d\mathbf{c}^i}{d\mathbf{q}} \right)^\top \mathbf{R}^{-1} (\mathbf{c}^i - \tilde{\mathbf{c}}^i),$$

respectively. Using the chain rule, this gives

$$\nabla_{\mathbf{q}} H_{\text{obs}} = \frac{1}{\sigma_\epsilon^2} \Delta t \left(\frac{d\mathbf{c}^1}{d\mathbf{q}} \right)^\top \left[\mathbf{C} (\mathbf{c}^1 - \tilde{\mathbf{c}}^1) + \left(\frac{d\mathbf{c}^2}{d\mathbf{c}^1} \right)^\top [\mathbf{C} (\mathbf{c}^2 - \tilde{\mathbf{c}}^2) + \dots] \right].$$

This implies that the sensitivities can be evaluated by repeatedly computing products of the form $(d\mathbf{c}^{i+1}/d\mathbf{c}^i)^\top \mathbf{u}$. If these products can be computed exactly, then this approach provides the exact sensitivity of the (space-time discretized) system, which is useful for problems with costly forward and backward solutions. The disadvantage of this approach is that it is highly application-specific and restricts the selection of ODE solvers to a specific family. Details of the implementation of this ad-hoc sensitivity analysis approach are presented in Appendix A.2.

We test this formulation on a one-dimensional transport problem defined in the domain $[0, 1]$ with constant velocity u and dispersion coefficient D . The reaction model (5.16) with $c_{eq} = 0.4$ and $k = 1.0$ is used. Boundary conditions are $dc/dx = 0$ at $x = 0, 1$. The transport equation is discretized with finite-volumes scheme consisting of 128 cells of size Δx . Concentration measurements are taken at the cells with centers $x = L/2 - \Delta x/2$, $x = 3L/4 - \Delta x/2$ and $x = L - \Delta x/2$ over the time period $(0, 2.5 \times 10^{-2}]$ (Figure 5.3). The standard deviation of these measurements is set to $\sigma_\epsilon = 0.02$.




Figure 5.3: Breakthrough curves of contaminant at observation locations along the transport domain.

The release configuration is inferred using the HMC scheme, carried out with hybrid time-step $\Delta\tau = 0.18$, $M = 5$, and regularizing parameter $\gamma = 0.025$. A total of 2×10^6 samples of the release configuration were computed, and the first half of the chain rejected as the burn-in period. Figure 5.4 shows that, for the initial condition shown, the HMC scheme is able to infer the main features of the initial condition, namely the location of the release and the total mass of contaminant released. The estimator obtained is of high variance given the dearth of data available and the relatively high measurement error.




Figure 5.4: Actual (solid line) and reconstructed (dotted line) release histories using HMC.

5.5 Application of gHMC to linear transport

In order to study the construction of an acceleration matrix A appropriate for contaminant transport, we consider a 1-D advection-dispersion (no reaction) problem

$$\frac{\partial c}{\partial t} + u \frac{\partial c}{\partial x} = D \frac{\partial^2 c}{\partial x^2}, \quad x \in [0, 2\pi], \quad t \in (0, T], \quad (5.18)$$

with uniform coefficients u and D . This equation is subject to the periodic boundary condition

$$c(0, t) = c(2\pi, t),$$

and (unknown) initial condition

$$c(x, 0) = c_0(x).$$

Similar to the case studied in Section 5.4.2, available concentration data consist of a set of measurements continuous in time on the interval $(0, T]$ collected at a subset J of the discrete locations x_j , $c_{\text{obs},j} = c_{\text{obs}}(x_j, t)$. The data come with space-time uncorrelated additive errors of equal variance σ_c^2 .

Observation Hamiltonian. The state variable $c(x, t)$ is discretized into N functions $c_j(t) = c(x_j, t)$, where the $x_j = 2\pi j/N$, $j = 0, \dots, N-1$ are N equi-distant

nodes along the domain $[0, 2\pi)$. We define the measurement Hamiltonian as

$$H_{\text{obs}} = \frac{1}{2\sigma_\epsilon^2} \sum_{j \in J} \int_0^T [c_j(t) - c_{\text{obs},j}(t)]^2 dt. \quad (5.19)$$

It defines the probability distribution of the initial release vector \mathbf{q} with components $q_j = c_{0,j}$.

A solution of (5.18) is represented by a discrete Fourier transform (DFT),

$$\hat{c}_k = \frac{1}{N} \sum_{j=0}^{N-1} c_j e^{-ikx_j}, \quad k = -N/2, \dots, N/2 - 1, \quad (5.20)$$

which defines the N Fourier modes \hat{c}_k . The backward or inverse transform is given by

$$c_j = \sum_{k=-N/2}^{N/2-1} \hat{c}_k e^{ikx_j}, \quad j = 0, \dots, N-1. \quad (5.21)$$

Let \mathbf{c} denote a vector of discrete values c_j , and $\hat{\mathbf{c}}$ denote the vector of Fourier modes \hat{c}_k . Then (5.20) and (5.21) are rewritten as

$$\hat{\mathbf{c}} = \frac{1}{N} \mathcal{F} \mathbf{c}, \quad \mathbf{c} = \mathcal{F}^* \hat{\mathbf{c}}, \quad (5.22)$$

where \mathcal{F} is the DFT matrix whose elements are

$$\mathcal{F}_{pq} = \omega^{(p-N/2)q}, \quad \omega = e^{-2\pi i/N} \quad (5.23)$$

and $(\cdot)^*$ denotes the Hermitian adjoint. By projection, (5.18) is discretized into a set of uncoupled ODEs for the Fourier modes

$$\frac{\partial \hat{c}_k}{\partial t} = -(Dk^2 + iku)\hat{c}_k, \quad k = -N/2, \dots, N/2 - 1.$$

Their solutions are

$$\hat{c}_k = \hat{q}_k \exp \{-(Dk^2 + iku)t\}. \quad (5.24)$$

Substituting (5.22) and (5.24) into (5.19) yields an observation Hamiltonian

$$\begin{aligned} H_{\text{obs}} &= \frac{1}{2\sigma_\epsilon^2} (\hat{\mathbf{q}} - \hat{\mathbf{q}}_{\text{obs}})^* \left(\int_0^T \mathbf{B}^* \mathcal{F}_J \mathcal{F}_J^* \mathbf{B} dt \right) (\hat{\mathbf{q}} - \hat{\mathbf{q}}_{\text{obs}}), \\ &= \frac{1}{2\sigma_\epsilon^2} (\hat{\mathbf{q}} - \hat{\mathbf{q}}_{\text{obs}})^* \hat{\mathbf{G}}_{\text{obs}} (\hat{\mathbf{q}} - \hat{\mathbf{q}}_{\text{obs}}), \end{aligned} \quad (5.25)$$

where $\hat{\mathbf{q}} \equiv \hat{\mathbf{c}}_0$, $\hat{\mathbf{q}}_{\text{obs}} = \hat{\mathbf{c}}_{\text{obs},0}$ is the vector of Fourier modes of the unknown underlying release configuration, \mathcal{F}_J corresponds to the J columns of \mathcal{F} , $\mathbf{B}(t)$ is a diagonal matrix with elements $B_{kk} = \exp[-(Dk^2 + iku)t]$, and $\hat{\mathbf{G}}_{\text{obs}}$ is the Hermitian (semi)positive definite matrix. The observation Hamiltonian specifies a multivariate normal distribution for the Fourier modes of the release configuration. The mean of this distribution corresponds to the underlying configuration. Given that \mathbf{q} and $\hat{\mathbf{q}}$ are related via a linear transformation, it follows that the observation Hamiltonian specifies a multivariate normal distribution for \mathbf{q} .

If the measurements are available in every node of the computational domain, i.e., if $\mathcal{F}_J = \mathcal{F}$, then $\mathcal{F}\mathcal{F}^* = N\mathbf{I}$ and (5.25) simplifies to

$$H_{\text{obs}} = \frac{N}{2\sigma_\epsilon^2} (\hat{\mathbf{q}} - \hat{\mathbf{q}}_{\text{obs}})^* \left(\int_0^T \mathbf{B}^* \mathbf{B} dt \right) (\hat{\mathbf{q}} - \hat{\mathbf{q}}_{\text{obs}}).$$

That is equivalent to

$$H_{\text{obs}} = \frac{N}{2\sigma_\epsilon^2} \sum_{k=-N/2}^{N/2-1} |\hat{q}_k - \hat{q}_{\text{obs},k}|^2 \hat{g}_k \quad (5.26)$$

where the coefficients \hat{g}_k are given by

$$\hat{g}_k = \int_0^T |\exp\{-(Dk^2 + iku)t\}|^2 dt = \int_0^T e^{-2Dk^2t} dt = \frac{1 - \exp(-2Dk^2T)}{2Dk^2}.$$

Note that all coefficients \hat{g}_k are real, symmetric ($\hat{g}_k = \hat{g}_{-k}$), and depend only on the dispersion coefficient D .

It follows from (5.25) that

$$\nabla_{\mathbf{q}} H_{\text{obs}} = \frac{1}{\sigma_\epsilon^2} \mathcal{F}^* \hat{\mathbf{G}}_{\text{obs}} (\hat{\mathbf{q}} - \hat{\mathbf{q}}_{\text{obs}}) = \mathcal{F}^* \hat{\mathbf{G}}_{\text{obs}} \mathcal{F} (\mathbf{q} - \mathbf{q}_{\text{obs}}) = \mathbf{G}_{\text{obs}} (\mathbf{q} - \mathbf{q}_{\text{obs}})$$

where $\mathbf{G}_{\text{obs}} = \mathcal{F}^* \hat{\mathbf{G}}_{\text{obs}} \mathcal{F}$. That brings the forcing into the form $\mathbf{F}[\mathbf{q}] = \Sigma^{-1}(\mathbf{q} - \boldsymbol{\mu})$ required by our analysis in Section 5.3.2, which suggests a possibility of computing the acceleration matrices as $\mathbf{A}\mathbf{A}^* = \mathbf{G}_{\text{obs}}^{-1}$. Unfortunately this is not generally feasible, because the matrix \mathbf{G}_{obs} is singular unless the measurements are taken at every node of the domain. The singularity of \mathbf{G}_{obs} implies the singularity of the multivariate normal distribution of $\hat{\mathbf{q}}$ given by H_{obs} , which means the distribution is concentrated in a r -dimensional subspace of \mathbb{C}^N . Since \mathbf{q} results from a linear

transformation of $\hat{\mathbf{q}}$, the multivariate normal distribution of \mathbf{q} is also degenerate. This implies that there are linear combinations of $q_j = c_{0,j}$ that cannot be sampled with the distribution stemming from the measurements. In other words, there is an unobservable subset of initial configurations.

The SVD decomposition $\hat{\mathbf{G}}_{\text{obs}} = \mathbf{USV}^*$ provides some insight into features of the distribution of \mathbf{q} that can be sampled from the distribution of the observations. Specifically, the vectors forming a basis for $\ker \hat{\mathbf{G}}_{\text{obs}}$ have negligible terms associated with the lower Fourier modes of \mathbf{q} , i.e., $|V_{k,j}| \approx 0$ for small $|k|$ and $\text{rank} \hat{\mathbf{G}}_{\text{obs}} < j$. This implies that the lower frequency components of \mathbf{q} fall mostly on the observable subspace. In general, $|V_{k,j}| \neq 0$ for high $|k|$ and $\text{rank} \hat{\mathbf{G}}_{\text{obs}} < j$, which implies that high frequency features are mostly unobservable.

Regularization Hamiltonian. The regularization Hamiltonian extends the distributions of \mathbf{q} and $\hat{\mathbf{q}}$, which makes them well defined. After a real-space discretization, the regularization Hamiltonian takes the form

$$H_{\text{reg}} = \gamma \mathbf{q}^\top \mathbf{G}_{\text{reg}} \mathbf{q}, \quad (5.27)$$

where γ is a regularization hyperparameter and \mathbf{G}_{reg} is the circulant matrix

$$\mathbf{R} = \frac{1}{\Delta x} \begin{bmatrix} 2 & -1 & & & -1 \\ -1 & 2 & -1 & & \\ & \ddots & \ddots & \ddots & \\ & & -1 & 2 & -1 \\ -1 & & & -1 & 2 \end{bmatrix},$$

or $\mathbf{G}_{\text{reg}} = \text{circ}\{\mathbf{r}\}^\top$, where $\mathbf{r} = (2, -1, 0, \dots, 0, -1)^\top / \Delta x$ and $\Delta x = 2\pi/N$. The matrix \mathbf{G}_{reg} extends the probability distribution by assigning a high energy (low probability) to configurations with large high-frequency components. To demonstrate this, we rewrite \mathbf{G}_{reg} as

$$\mathbf{G}_{\text{reg}} = \text{circ}\{\mathbf{r}\}^\top = \mathcal{F}^* \text{diag}\{\hat{\mathbf{r}}\} \mathcal{F}, \quad \hat{r}_k = \frac{1}{\pi} \left[1 - \cos \left(\frac{2\pi k}{N} \right) \right],$$

where $\hat{\mathbf{r}}$ is the DFT of \mathbf{r} . The components \hat{r}_k of vector $\hat{\mathbf{r}}$ increase with frequency k , with the zeroth frequency giving rise to $\hat{r}_0 = 0$. The latter is to be expected

since the regularization operator does not affect the observability of the zeroth frequency, which corresponds to the average of the initial release.

Note that a Fourier-space discretization of the regularization Hamiltonian leads to a similar bilinear form for \mathbf{G}_{reg} , with $\hat{r}_k = 2\pi k^2/N^2$. Indeed, these \hat{r}_k have a similar asymptotic behavior as $k \rightarrow 0$.

Acceleration matrix. For the full Hamiltonian $H = H_{\text{obs}} + H_{\text{reg}}$, the forcing is given by

$$\mathbf{F}[\mathbf{q}] = -\hat{\mathbf{G}}_{\text{obs}}(\hat{\mathbf{q}} - \hat{\mathbf{q}}_{\text{obs}}) - \hat{\mathbf{G}}_{\text{reg}}\hat{\mathbf{q}} = -\mathbf{G}_{\text{obs}}(\mathbf{q} - \mathbf{q}_{\text{obs}}) - \mathbf{G}_{\text{reg}}\mathbf{q} \quad (5.28)$$

where $\mathbf{G}_{\text{obs}} = \mathcal{F}^*\hat{\mathbf{G}}_{\text{obs}}\mathcal{F}$ and $\mathbf{G}_{\text{reg}} = \mathcal{F}^*\hat{\mathbf{G}}_{\text{reg}}\mathcal{F}$. This suggests that choosing the acceleration matrix \mathbf{A} , such that $\mathbf{A}\mathbf{A}^*(\mathbf{G}_{\text{obs}} + \mathbf{G}_{\text{reg}}) = \mathbf{I}$, would reduce the correlation of the Markov chain. Since \mathbf{G}_{obs} and \mathbf{G}_{reg} are Hermitian (semi)positive definite, their sum $\mathbf{G} = \mathbf{G}_{\text{obs}} + \mathbf{G}_{\text{reg}}$ is at least Hermitian (semi)positive definite. In fact, \mathbf{G} is a full rank matrix and, therefore, it can be factorized via a Cholesky decomposition $\mathbf{G} = \mathbf{R}^\top\mathbf{R}$. Hence the matrix \mathbf{A} , defined by $\mathbf{A}\mathbf{A}^*\mathbf{G} = \mathbf{I}$, is given by

$$\mathbf{A} = \mathbf{A}_1 = \mathbf{R}^{-1}, \quad \mathbf{R}^\top\mathbf{R} = \mathbf{G}. \quad (5.29)$$

The added cost of computing the acceleration matrix \mathbf{A} is the Cholesky factorization cost and the cost of computing four matrix-vector products. For dense matrices, these costs are $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$, respectively.

An advantage of the Cholesky factorization is that the vector of momenta \mathbf{p} in (5.12) can be chosen as real and multivariate normal, with zero mean and identity covariance matrix. A drawback is its relatively high cost per step in the Markov chain. Moreover, the matrix \mathbf{G} becomes more poorly conditioned as $\gamma \rightarrow 0$, which might affect the accuracy of the Cholesky decomposition.

A computationally efficient alternative for the construction of \mathbf{A} is to employ the following heuristic. Instead of using the full correlation matrices in Fourier space, $\hat{\mathbf{G}}_{\text{obs}}$ and $\hat{\mathbf{G}}_{\text{reg}}$, to define \mathbf{G} , we approximate it as $\mathbf{G} \approx \mathcal{F}^*\text{diag}\{\bar{\mathbf{g}}\}\mathcal{F}$, where $\bar{\mathbf{g}}$ is the vector with components $\bar{g}_i = \{\hat{\mathbf{G}}_{\text{obs}} + \hat{\mathbf{G}}_{\text{reg}}\}_{ii}$. This approximation allows one to factorize \mathbf{G} as $\mathbf{G} \approx \bar{\mathbf{G}} = \mathcal{F}^*\mathbf{D}\mathbf{D}^*\mathcal{F}$, where $\mathbf{D} = \text{diag}\{(\bar{\mathbf{g}})^{1/2}\}$ with the

square root understood as element-wise. This argument suggests that the acceleration matrix \mathbf{A} can be constructed as

$$\mathbf{A} = \mathbf{A}_2 = \frac{1}{N} \mathcal{F}^* \mathbf{D}^{-1}, \quad \mathbf{D} = \text{diag}\{(\bar{\mathbf{g}})^{1/2}\}, \quad \bar{\mathbf{g}} = \text{diag}\{\hat{\mathbf{G}}_{\text{obs}} + \hat{\mathbf{G}}_{\text{reg}}\}, \quad (5.30)$$

which gives

$$\mathbf{A}\mathbf{p} = \frac{1}{N} (\mathcal{F}^* \mathbf{D}^{-1} \mathbf{p}), \quad \mathbf{A}^* \mathbf{F} = \mathbf{D}^{-1} \left(\frac{1}{N} \mathcal{F} \mathbf{F} \right).$$

Note that we have replaced the transpose of \mathbf{A} with its Hermitian transpose due to the complex nature of the DFT. This implies that the transpose in (5.11)–(5.12) must be replaced with an Hermitian transpose, and that in order to guarantee that $\mathbf{q} \in \mathbb{R}^N$ we must generalize the momenta such that $\mathbf{p} \in \mathbb{C}^N$.

Once the acceleration matrix \mathbf{A} in (5.30) is constructed, products of the form $\mathbf{A}\mathbf{p}$ and $\mathbf{A}^* \mathbf{F}$ can be computed using DFTs. The computational cost per leapfrog step is reduced from four matrix-vector products of cost $\sim \mathcal{O}(N^2)$ to four of cost $\sim \mathcal{O}(N \log N)$, and no Cholesky decomposition is necessary.

Since $\bar{\mathbf{g}} = \partial^2 H / \partial \hat{\mathbf{q}}^2$, the approximation (5.30) can be thought as building \mathbf{A} from the diagonal of the Hessian of H with respect to $\hat{\mathbf{q}}$ (a similar heuristic is employed in [63] for Bayesian learning). This observation begs the following question: Why do we take $\bar{\mathbf{g}} = \partial^2 H / \partial \hat{\mathbf{q}}^2$ instead of $\bar{\mathbf{g}} = \partial^2 H / \partial \mathbf{q}^2$, which would produce a similar acceleration matrix \mathbf{A} without the Fourier transforms? The answer is that the matrix $\hat{\mathbf{G}}_{\text{obs}} + \hat{\mathbf{G}}_{\text{reg}}$ is more concentrated along its diagonal than \mathbf{G} is. Hence more information about the observation operator is conserved by taking the diagonal of $\hat{\mathbf{G}}_{\text{obs}} + \hat{\mathbf{G}}_{\text{reg}}$ than the diagonal of \mathbf{G} .

Sampling of momentum vector. In order to retain the validity of the leapfrog method with generalized momenta, we require said momenta to be associated with a kinetic energy following a bilinear form. We can achieve this by assuming that $\mathbf{p} \in \mathbb{C}^N$ has a general complex normal distribution $\mathcal{CN}(0, \mathbf{\Gamma}, \mathbf{C})$ with unit covariance $\mathbf{\Gamma}$ and

$$\mathbf{C} = \mathbf{A}^* \bar{\mathbf{G}} (\text{conj } \mathbf{A}) = \frac{1}{N} \mathbf{D} \mathcal{F} \mathcal{F}^\top \mathbf{D}^{-1}. \quad (5.31)$$

For \mathcal{F}_{pq} given by (5.23),

$$\frac{1}{N}(\mathcal{F}\mathcal{F}^\top)_{p,q} = \sum_{r=0}^{N-1} \omega^{r(p+q-N)} = \mathbf{P} = \begin{cases} 1 & \text{if } p+q = kN, k = \dots, -1, 0, 1, \dots, \\ 0 & \text{otherwise.} \end{cases}$$

Since \mathbf{D} is diagonal and \mathbf{P} is a permutation matrix, (5.23) yields $\mathbf{C} = \mathbf{P}$.

Let $\mathbf{p} = \mathbf{X} + i\mathbf{Y}$ with $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^N$. The vector $\mathbf{V}^\top = [\mathbf{X}^\top \ \mathbf{Y}^\top]$ is multivariate normal with zero mean. Given $\mathbf{\Gamma}$ and \mathbf{C} , the cross-covariance matrix of this vector is

$$\mathbb{E}[\mathbf{X}\mathbf{Y}^\top] = \frac{1}{2}\Im\{\mathbf{\Gamma} + \mathbf{C}\} = 0, \quad \mathbb{E}[\mathbf{Y}\mathbf{X}^\top] = \frac{1}{2}\Im\{-\mathbf{\Gamma} + \mathbf{C}\} = 0,$$

since both $\mathbf{\Gamma}$ and \mathbf{C} are real. In other words, the real and imaginary parts of \mathbf{p} are mutually uncorrelated. The covariance matrix of this vector is

$$\mathbb{E}[X_i X_j] = \frac{1}{2}\Re\{\Gamma_{ij} + C_{ij}\} = \begin{cases} 1 & \text{if } i = j = 0, -N/2 \\ 1/2 & \text{if } i = -j, i, j \neq 0 \\ 0 & \text{otherwise,} \end{cases} \quad (5.32a)$$

$$\mathbb{E}[Y_i Y_j] = \frac{1}{2}\Re\{\Gamma_{ij} - C_{ij}\} = \begin{cases} 1/2 & \text{if } i = -j, i, j \neq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (5.32b)$$

It follows from (5.32) that only the components $p_k = X_k + iY_k$ with $k = -k$ are correlated. Their covariances are $\mathbb{E}[X_k X_{-k}] = 0.5$, $\mathbb{E}[Y_k Y_{-k}] = -0.5$. Since \mathbf{p} must be complex-symmetric to guarantee that \mathbf{q} remains real, we generate \mathbf{p} as

$$\begin{aligned} X_{-N/2} &\sim \mathcal{N}(0, 1), \\ X_{-N/2+1} &\sim \mathcal{N}(0, 1/2), & Y_{-N/2+1} &\sim \mathcal{N}(0, 1/2) \\ &\vdots \\ X_0 &\sim \mathcal{N}(0, 1), \\ X_1 = X_{-1}, & & Y_1 = -Y_{-1}, \\ &\vdots \\ X_{N/2-1} = X_{-N/2+1}, & & Y_{N/2-1} = -Y_{-N/2+1}. \end{aligned}$$

Hence the vector \mathbf{p} is generated with N independent identically distributed normal random variables.

Computational examples. Next we illustrate the effect of selecting the acceleration matrix \mathbf{A} on the Markov chains produced by the generalized Hybrid Monte Carlo algorithm. Three alternatives for \mathbf{A} are considered: $\mathbf{A} = \mathbf{A}_1$ in (5.29), $\mathbf{A} = \mathbf{A}_2$ in (5.30) and $\mathbf{A} = \mathbf{I}$ (no acceleration). We apply these selections to the model problem (5.18) with parameters $D = 1.0$, $u = 10.0$ and $N = 64$. Observations are taken at locations x_J , $J = \{47, 63\}$. For each case, $P = 10$ chains with 2×10^4 samples are run, and the last half samples are kept. Using these samples over P runs, we compute the autocorrelation function $\rho(s)$ for each q_j , $j = 0, \dots, N - 1$, except for $j = 47, 63$, which are included in the observations.

Figures 5.5 show that $\mathbf{A} = \mathbf{A}_1$ produces highly uncorrelated chains for each of the q_j studied, $\mathbf{A} = \mathbf{A}_2$ produces more correlated chains, and $\mathbf{A} = \mathbf{I}$ produces the most correlated chains. As expected, \mathbf{A}_2 provides a compromise between the low autocorrelation / high expense of the full Cholesky decomposition and the high autocorrelation / low cost of $\mathbf{A} = \mathbf{I}$.

For problems with reaction terms, the forcing $\mathbf{F} = -\nabla_{\mathbf{q}}H$ is not a linear function of \mathbf{q} as in (5.28). In such cases, the selection of the acceleration matrix \mathbf{A} is not straightforward. The challenge is to find an approximation to the forcing that is linear in \mathbf{q} , i.e., preserves the form (5.28) with \mathbf{G}_{obs} and \mathbf{G}_{reg} independent of \mathbf{q} . This is required to guarantee the reversibility of the Hamiltonian dynamics.

Such an approximation can be obtained by disregarding the nonlinear reaction term and using \mathbf{G}_{obs} in (5.25) and \mathbf{G}_{reg} in (5.27), which are functions only of the temporal and spatial domain properties and the hyperparameters σ_ϵ^2 and γ . This selection is equivalent to taking $\mathbf{F} \approx -\mathbf{G}_{\text{lin}}(\mathbf{q} - \mathbf{q}_{\text{obs}})$, where \mathbf{G}_{lin} is the Hessian of the advection-diffusion (linear) portion of the Hamiltonian. It gives the acceleration matrix \mathbf{A}_1 of (5.29). This choice is justified for non-periodic boundary conditions if the contaminant plume does not reach the domain's boundaries during the simulation time. An alternative is to take only diagonal portion of the Hessian of an advection-diffusion portion of the Hamiltonian. This would produce the acceleration matrix \mathbf{A}_2 in (5.30).




Figure 5.5: Autocorrelation functions for q_j , $j = \{0, \dots, 63\} \setminus \{47, 63\}$, and $\mathbf{A} = \mathbf{A}_1$ (top), $\mathbf{A} = \mathbf{A}_2$ (middle) and $\mathbf{A} = \mathbf{I}$ (bottom).

5.6 Conclusions and further work

We presented a computationally efficient and accurate algorithm for identification of sources and release histories of (geo)chemically active solutes. The algorithm is based on a generalized hybrid Monte Carlo approach, in which MC sampling is accelerated by the use of discrete adjoint equations. Some of the salient features of our approach are: 1) its ability to handle nonlinear systems, since it requires no linearizations; and 2) its compatibility with various regularization strategies.

The introduction of an acceleration matrix to the gHMC scheme was tested for a advection-dispersion problem. Our analysis demonstrated that the proposed acceleration matrices improve upon basic HMC. The generalization of these constructions to problems with nonlinear reaction terms and non-periodic boundary conditions is the subject of future work.

Chapter 5, in part, has been submitted for publication of the material: Barajas-Solano, D. A., Alexander, F. J., Anghel, M., Tartakovsky, D. M. “Efficient Reconstruction of Contaminant Release History”. The dissertation author was the primary investigator and author of the paper.

Chapter 6

Inverse Modeling via Linear Functional Minimization

6.1 Introduction

The spatially distributed properties of groundwater flow and transport domains, such as conductivity, transmissivity, dispersivity, etc., are often highly heterogeneous. Heterogeneity is a product of multiple geological, mechanical and physico-chemical processes that form subsurface environments [71, 17]. It occurs over a large range of scales, from well-defined large geological structures with significantly different hydrogeological properties to smaller variations of these properties inside each structures. Taking into account multiple scales of heterogeneity in complex geologic environments is a must in order to build useful predictive models [96, 26].

A variety of approaches have been proposed to deal with multiscale heterogeneity of flow and transport fields [96, 26]. Generally, they start with building a conceptual model and identifying a set of relevant parameters. The next step is to select parameter values via an appropriate inversion process that leverages both available field observations and prior knowledge and expectations, such as delineation of large-scale structures from geological information [17, 30]. Most modern inversion techniques are designed to handle only one scale of heterogeneity [97], fo-

cusing on either single-facies heterogeneous systems [30] or multiple homogeneous facies [45, 98, 89]. A two-scale alternative [42, 43] combines a zonation approach for large-scale facies delineation with a stochastic description of within-facies variability. This strategy is strongly dependent on the selection of an adequate stochastic model.

We develop an inversion methodology that uses observations of both system parameters and transient system states to compute an estimator of the underlying parameter field. No strong assumptions, such as prescribed zonation or stochastic models for small-scale heterogeneity, are introduced. Instead we adopt a Bayesian framework and choose a prior for parameter fields, which reflects the knowledge of their spatial distribution (i.e., that parameters are piecewise continuous). The total variation (TV) regularization [74] is used to define the prior with desired properties. A maximum *a posteriori* (MAP) estimator of the system parameters is obtained by solving the associated high-dimensional nonlinear optimization problem. The linearized functional minimization algorithm [97] is used in this task. Our algorithm splits the nonlinearities of the log-posterior into two parts, each of which is treated separately in an iterative scheme: The nonlinearity of the data fidelity terms associated with system states is treated using a linearization approach based on the Levenberg-Marquardt method [14, 9]; and the nonlinearity associated with the TV regularization operator is dealt with using the alternating direction method of multipliers (ADMM) [13].

TV regularization has been employed in previous works for inverse modeling of spatially distributed parameters of elliptic equations [5, 22, 9], proposing different algorithms to the one employed in this study. Furthermore, our proposed methodology allows incorporation of transient information of various system states for which there are available measurements. This feature is particularly useful from a history matching perspective, where information of the response of measured system states can be leveraged to improve the reconstruction of the model parameters. The implementation presented in this work uses measurements of a model parameter and a single state, but can be extended to include multiple states.

Section 6.2 presents the inversion methodology as applied to a saturated

flow problem. The linearized functional minimization algorithm is described in Section 6.3. A numerical scheme used both to solve the forward problem and to compute the sensitivities is outlined in Section 6.4. The computational cost of the overall inversion procedure is analyzed in Section 6.5. We apply our inversion methodology to a synthetic two-dimensional saturated flow problem with a highly heterogeneous, piecewise continuous conductivity field. Error-free measurements of the synthetic conductivity field are taken over a grid of discrete sampling locations, together with measurements of the hydraulic head response of the system at a number of discrete locations under point loadings. The results obtained for different selections of the inversion coefficients are presented in Section 6.6. Conclusions and implications of this work are presented in Section 6.7.

6.2 Formulation of optimization problem

Problem formulation. We consider an n -dimensional groundwater flow equation

$$\frac{\partial h}{\partial t} = \nabla \cdot (K(\mathbf{x})\nabla h) + q(\mathbf{x}, t), \quad (6.1)$$

subject to appropriate boundary and initial conditions. Here $K(\mathbf{x})$ is the spatially varying saturated hydraulic conductivity (a system parameter), $h(\mathbf{x}, t)$ is the hydraulic head (a system state), and the source term $q(\mathbf{x}, t)$ represents, e.g., pumping wells. A spatio-temporal discretization of (6.1) replaces continuous functions $K(\mathbf{x})$ and $h(\mathbf{x}, t)$ with their discrete counterparts arranged in vectors \mathbf{k} and \mathbf{h} , respectively. The size of these matrices depends on the numerical method used to solve (6.1). We employ a method-of-lines (MOL) discretization, in which $K(\mathbf{x})$ is evaluated at N_k discrete grid points in a computational domain and $h(\mathbf{x}, t)$ is computed at N_h discrete grid points in space and N_t discrete points in time. As a result, vectors \mathbf{k} and \mathbf{h} have N_k and $N_h N_t$ components, respectively. Combined with the initial and boundary conditions, (6.1) defines a nonlinear map $\mathbf{h} = \mathbf{f}(\mathbf{k})$.

Suppose that hydraulic conductivity $K(\mathbf{x})$ is measured at $N_{\text{obs},k}$ grid points ($N_{\text{obs},k} \ll N_k$) of the numerical mesh used to discretize (6.1). Suppose further that hydraulic head $h(\mathbf{x}, t)$ is measured at $N_{\text{obs},h}$ grid points at each of the N_t discrete

time points. The support volume of the conductivity measurements is considered representative of the cell size in the numerical model. These measurements are assembled into vectors $\hat{\mathbf{k}}$ and $\hat{\mathbf{h}}$ whose dimensions are $N_{\text{obs},k}$ and $N_{\text{obs},h}N_t$, respectively. Linear operators \mathbf{M}_k and \mathbf{M}_h provide spatio-temporal maps between \mathbf{k} and $\hat{\mathbf{k}}$ and between \mathbf{h} and $\hat{\mathbf{h}}$, such that $\mathbf{M}_k\mathbf{k} = \hat{\mathbf{k}}$ and $\mathbf{M}_h\mathbf{h} = \hat{\mathbf{h}}$. The measurement operator matrices \mathbf{M}_k and \mathbf{M}_h have dimensions $N_{\text{obs},k} \times N_k$ and $N_{\text{obs},h}N_t \times N_hN_t$, respectively.

Our goal is to compute estimators of \mathbf{k} and \mathbf{h} using the nonlinear map $\mathbf{h} = \mathbf{f}(\mathbf{k})$ and the available measurements $\hat{\mathbf{k}}$ and $\hat{\mathbf{h}}$.

Optimization problem. Since the measurement operator \mathbf{M}_k is often highly undetermined, it is in general not possible to use direct inversion or to compute an ordinary least squares estimator. The inclusion of the state variable data $\hat{\mathbf{h}}$ is expected to reduce the indeterminacy of the inversion problem, but it is generally not enough to define a unique solution of the inversion problem. Instead, our goal is to compute maximum *a posteriori* (MAP) estimators of \mathbf{k} and \mathbf{h} . These MAP estimators combine the measurements $\hat{\mathbf{k}}$ and $\hat{\mathbf{h}}$ with the available general knowledge of the properties of $K(\mathbf{x})$ and $h(\mathbf{x}, t)$.

Following [97], we express the inversion problem in variational form as the minimization of a (negative) likelihood functional. This likelihood is composed of data fidelity terms, which penalize the difference between model estimates and measurements, and regularization terms, which express the prior knowledge of the properties of the (discretized) fields \mathbf{k} and \mathbf{h} . Enforcing the constraint $\mathbf{h} = \mathbf{f}(\mathbf{k})$, that gives rise to an optimization problem over \mathbf{k} ,

$$\arg \min_{\mathbf{k}} \frac{\alpha}{2} \|\mathbf{M}_k\mathbf{k} - \hat{\mathbf{k}}\|_2^2 + \frac{\beta}{2} \|\mathbf{M}_h\mathbf{f}(\mathbf{k}) - \hat{\mathbf{h}}\|_2^2 + \gamma R_k(\mathbf{k}) + \frac{\delta}{2} R_h(\mathbf{f}(\mathbf{k})), \quad (6.2)$$

where $R_k(\mathbf{k})$ and $R_h(\mathbf{h} = \mathbf{f}(\mathbf{k}))$ are the regularization terms for \mathbf{k} and \mathbf{h} , respectively. From a Bayesian point of view, the data fidelity terms provide the likelihood of observations given a certain configuration of \mathbf{k} , and the regularization terms provide the priors on \mathbf{k} and \mathbf{h} . The inversion coefficients α , β , γ and δ determine the extent to which low probability is assigned to a certain \mathbf{k} configuration, i.e., the relative contributions of each term in (6.2). For example if one sets $\alpha = 1$ as

a reference, then β would determine how the data fidelity of $\mathbf{h} = \mathbf{f}(\mathbf{k})$ affects probability of \mathbf{k} relative to the data fidelity of \mathbf{k} .

In subsurface environments consisting of multiple facies (see Section 6.1) the hydraulic conductivity field $K(\mathbf{x})$ is piecewise continuous. To reflect this fact, we select the regularization for \mathbf{k} to be given by the ℓ_1 norm of the gradient of the conductivity field ∇K . In n -dimensional domains this regularization, which is often referred to as Total Variation (TV) [74], penalizes general heterogeneous fields and promotes fields with low small-scale variation inside facies and large jumps across the $(n - 1)$ -dimensional surfaces separating the facies. Since \mathbf{h} is continuous and defined by the map $\mathbf{h} = \mathbf{f}(\mathbf{k})$, we decided not to add an additional regularity requirement on \mathbf{h} . Instead, we add an additional regularization term for \mathbf{k} , which penalizes the ℓ_2 norm of ∇K . This increases stability during the minimization of (6.2). To avoid interference with the regularization properties of the TV norm, the ℓ_2 regularization term is multiplied by a coefficient δ that satisfies the inequality $\delta \ll \gamma$.

Two sources contribute to high nonlinearity of the objective functional (6.2): the data fidelity penalty on the deviation of the modeled head from observations, $\|\mathbf{M}_h \mathbf{f}(\mathbf{k}) - \hat{\mathbf{h}}\|_2^2$, and the TV norm. The latter also renders the optimization functional non-differentiable with respect to \mathbf{k} . Without the TV norm, the problem is equivalent to a nonlinear least squares estimation, which can be solved by a variety of standard techniques (e.g., Levenberg-Marquardt). On the other hand, the problem without the non-quadratic penalty is equivalent to linear minimization equipped with the TV norm, for which many efficient numerical approaches have been proposed [73, 36, 13]. We pursue a hybrid approach [97] that combines the strengths of the methods designed for dealing with each type of nonlinearity. This strategy separates the two sources of nonlinearity by first linearizing the map $\mathbf{h} = \mathbf{f}(\mathbf{k})$, and then minimizing a sequence of the linearized models equipped with the TV norm.

For a current estimate $\mathbf{k}^{(c)}$ of the minimizer, we replace $\mathbf{f}^{(c)}$ with the affine model [14, 9]

$$\mathbf{f}(\mathbf{k}) \approx \mathbf{f}(\mathbf{k}^{(c)}) + \mathbf{J}_f(\mathbf{k}^{(c)})(\mathbf{k} - \mathbf{k}^{(c)}),$$

where $\mathbf{J}_f(\mathbf{k}^{(c)}) = \partial \mathbf{f} / \partial \mathbf{k}^{(c)}$ is the Jacobian of function \mathbf{f} evaluated at $\mathbf{k}^{(c)}$. Next, we approximate the head data fidelity term with

$$\|\mathbf{M}_h \mathbf{f}(\mathbf{k}) - \hat{\mathbf{h}}\|_2^2 \approx \|\mathbf{M}_h \mathbf{J}_f \mathbf{k} + \mathbf{M}_h (\mathbf{f}(\mathbf{k}) - \mathbf{J}_f(\mathbf{k}^{(c)}) \mathbf{k}^{(c)}) - \hat{\mathbf{h}}\|_2^2.$$

The modified minimization problem associated with the iterate $\mathbf{k}^{(c)}$ is a quadratic form, except for the TV norm. Any of the previously mentioned TV optimization schemes can be used to solve this optimization problem. A solution of this linearized sub-problem yields a new iterate $\mathbf{k}^{(c+1)}$, for which a new linearized minimization problem is solved, and so forth. This linearized functional minimization algorithm is described in detail in the following section.

6.3 Linearized functional minimization

In this section we describe the linearized functional minimization algorithm [97], which we use to solve (6.2) together with a TV regularization term. In broad terms, this algorithm is a combination of a Levenberg trust-region algorithm, which is employed to propose linearized minimization sub-problems, and an Alternating Direction Method of Multipliers (ADMM) approach [13], which is used to solve these sub-problems.

Consider a generalized problem with a parameter vector $\mathbf{u} \in \mathbb{R}^N$ and a state vector $\mathbf{v} \in \mathbb{R}^M$, stemming from a suitable discretization of a partial-differential equation (PDE) on a two-dimensional domain Ω and time interval $(0, T]$. The discretized PDE provides a map $\mathbf{v} = \mathbf{f}(\mathbf{u})$. Let \mathbf{P} and \mathbf{Q} denote measurement operators, such that $\mathbf{P}\mathbf{u} = \mathbf{s}$ and $\mathbf{Q}\mathbf{v} = \mathbf{t}$, with $\mathbf{s} \in \mathbb{R}^{N'}$ and $\mathbf{t} \in \mathbb{R}^{M'}$, and $M' \ll M$ and $N' \ll N$. The cost functional (6.2) with a TV regularization term takes the form

$$\arg \min_{\mathbf{u}} \frac{\alpha}{2} \|\mathbf{P}\mathbf{u} - \mathbf{s}\|_2^2 + \frac{\beta}{2} \|\mathbf{Q}\mathbf{f}(\mathbf{u}) - \mathbf{t}\|_2^2 + \gamma \|\mathcal{D}(\mathbf{u})\|_1 + \frac{\delta}{2} \|\mathcal{D}(\mathbf{u})\|_2^2, \quad (6.3)$$

where $\mathcal{D}(\cdot) = \sqrt{[\mathbf{D}_x(\cdot)]^2 + [\mathbf{D}_y(\cdot)]^2}$ is a discretized gradient magnitude operator, with the operations $[\cdot]^2$ and $\sqrt{\cdot}$ understood as element-wise, and \mathbf{D}_x and \mathbf{D}_y are linear discrete differentiation operators in the x and y directions respectively.

We use an iterative trust region approach to minimize (6.3). Starting from an iterate $\mathbf{u}^{(c)}$, we compute the next iterate as $\mathbf{u}^{(+)} = \mathbf{u}^{(c)} + \mathbf{w}$, where \mathbf{w} is chosen so as to minimize a linearization of (6.3). For a \mathbf{w} with small $\|\mathbf{w}\|_2^2$, one can approximate \mathbf{v} as $\mathbf{v} \approx \mathbf{v}^{(c)} + \mathbf{A}\mathbf{w}$, where \mathbf{A} is the Jacobian of \mathbf{f} whose elements are $A_{pq} = \partial f_p / \partial u_q$, and $\mathbf{v}^{(c)} = \mathbf{f}(\mathbf{u}^{(c)})$. The step \mathbf{w} is chosen as the minimizer of

$$\begin{aligned} \arg \min_{\mathbf{w}} \frac{\alpha^*}{2} \|\mathbf{P}(\mathbf{u}^{(c)} + \mathbf{w}) - \mathbf{s}\|_2^2 + \frac{\beta^*}{2} \|\mathbf{Q}(\mathbf{v}^{(c)} + \mathbf{A}\mathbf{w}) - \mathbf{t}\|_2^2 \\ + \gamma^* \|\mathcal{D}(\mathbf{u}^{(c)} + \mathbf{w})\|_1 + \frac{\delta^*}{2} \|\mathcal{D}(\mathbf{u}^{(c)} + \mathbf{w})\|_2^2 + \frac{1}{2} \|\mathbf{w}\|_2^2 \end{aligned} \quad (6.4)$$

where the penalty on $\|\mathbf{w}\|_2^2$ is added to provide regularity (convexity) of the solution. The inversion coefficients of the linearized problem (α^* , β^* , γ^* and δ^*) are chosen initially to coincide with their counterparts in (6.3). They can also be allowed to vary in conjunction in order to control the trust region size throughout the minimization process.

The ADMM approach [13] is used to minimize the linearized functional (6.4). This allows us to split the regularization terms from the deviation penalties. Introducing auxiliary variables \mathbf{d}_x and \mathbf{d}_y , we recast (6.4) as

$$\begin{aligned} \arg \min_{\mathbf{w}} \frac{\alpha^*}{2} \|\mathbf{P}(\mathbf{u}^{(c)} + \mathbf{w}) - \mathbf{s}\|_2^2 + \frac{\beta^*}{2} \|\mathbf{Q}(\mathbf{v}^{(c)} + \mathbf{A}\mathbf{w}) - \mathbf{t}\|_2^2 \\ + \gamma^* \left\| \sqrt{\mathbf{d}_x^2 + \mathbf{d}_y^2} \right\|_1 + \frac{\delta^*}{2} \left\| \sqrt{\mathbf{d}_x^2 + \mathbf{d}_y^2} \right\|_2^2 + \frac{1}{2} \|\mathbf{w}\|_2^2, \end{aligned}$$

subject to $\mathbf{d}_x = \mathbf{D}_x(\mathbf{u}^{(c)} + \mathbf{w})$ and $\mathbf{d}_y = \mathbf{D}_y(\mathbf{u}^{(c)} + \mathbf{w})$. This constrained problem is solved iteratively via the splitting scheme

$$\begin{aligned} \mathbf{w}^{(k+1)} \equiv \arg \min_{\mathbf{w}} \frac{\lambda}{2} \|\mathbf{D}_x \mathbf{w} - (\mathbf{d}_x^{(k)} - \mathbf{D}_x \mathbf{u}^{(c)} - \mathbf{b}_x^{(k)})\|_2^2 \\ + \frac{\lambda}{2} \|\mathbf{D}_y \mathbf{w} - (\mathbf{d}_y^{(k)} - \mathbf{D}_y \mathbf{u}^{(c)} - \mathbf{b}_y^{(k)})\|_2^2 + \frac{1}{2} \|\mathbf{w}\|_2^2 \\ + \frac{\alpha^*}{2} \|\mathbf{P}\mathbf{w} - (\mathbf{s} - \mathbf{P}\mathbf{u}^{(c)})\|_2^2 + \frac{\beta^*}{2} \|\mathbf{Q}\mathbf{A}\mathbf{w} - (\mathbf{t} - \mathbf{Q}\mathbf{v}^{(c)})\|_2^2, \end{aligned} \quad (6.5)$$

$$\begin{aligned} (\mathbf{d}_x^{(k+1)}, \mathbf{d}_y^{(k+1)}) \equiv \arg \min_{\mathbf{d}_x, \mathbf{d}_y} \gamma^* \left\| \sqrt{\mathbf{d}_x^2 + \mathbf{d}_y^2} \right\|_1 + \frac{\delta^*}{2} \left\| \sqrt{\mathbf{d}_x^2 + \mathbf{d}_y^2} \right\|_2^2 \\ + \frac{\lambda}{2} \|\mathbf{D}_x \mathbf{w}^{(k+1)} - (\mathbf{d}_x - \mathbf{D}_x \mathbf{u}^{(c)} - \mathbf{b}_x^{(k)})\|_2^2 + \frac{\lambda}{2} \|\mathbf{D}_y \mathbf{w}^{(k+1)} - (\mathbf{d}_y - \mathbf{D}_y \mathbf{u}^{(c)} - \mathbf{b}_y^{(k)})\|_2^2, \end{aligned} \quad (6.6)$$

where $\lambda > 0$. The dual variables \mathbf{b}_x and \mathbf{b}_y are updated as

$$\mathbf{b}_x^{(k+1)} = \mathbf{b}_x^{(k)} + \mathbf{D}_x(\mathbf{u}^{(c)} + \mathbf{w}^{(k+1)}) - \mathbf{d}_x^{(k+1)}, \quad (6.7)$$

$$\mathbf{b}_y^{(k+1)} = \mathbf{b}_y^{(k)} + \mathbf{D}_y(\mathbf{u}^{(c)} + \mathbf{w}^{(k+1)}) - \mathbf{d}_y^{(k+1)}. \quad (6.8)$$

Initial values for the iteration scheme are set as $\mathbf{b}_x^{(0)} = 0$, $\mathbf{b}_y^{(0)} = 0$, $\mathbf{d}_x^{(0)} = \mathbf{D}_x \mathbf{u}^{(c)}$ and $\mathbf{d}_y^{(0)} = \mathbf{D}_y \mathbf{u}^{(c)}$.

The linear minimization problem (6.5) is obtained by differentiating (6.5) and setting equal to 0, and then solving the resulting $N \times N$ linear algebraic system

$$\begin{aligned} & (\mathbf{I} + \alpha^* \mathbf{P}^\top \mathbf{P} + \beta^* \mathbf{A}^\top \mathbf{Q}^\top \mathbf{Q} \mathbf{A} + \lambda \mathbf{D}_x^\top \mathbf{D}_x + \lambda \mathbf{D}_y^\top \mathbf{D}_y) \mathbf{w}^{(k+1)} \\ & = \alpha^* \mathbf{P}^\top (\mathbf{s} - \mathbf{P} \mathbf{u}^{(c)}) + \beta^* \mathbf{A}^\top \mathbf{Q}^\top (\mathbf{t} - \mathbf{Q} \mathbf{v}^{(c)}) \\ & + \lambda \mathbf{D}_x^\top (\mathbf{d}_x^{(k)} - \mathbf{D}_x \mathbf{u}^{(c)} - \mathbf{b}_x^{(k)}) + \lambda \mathbf{D}_y^\top (\mathbf{d}_y^{(k)} - \mathbf{D}_y \mathbf{u}^{(c)} - \mathbf{b}_y^{(k)}). \end{aligned} \quad (6.9)$$

The problem for $\mathbf{d}_x^{(k+1)}$ and $\mathbf{d}_y^{(k+1)}$ is solved component-wise to yield

$$d_{xi}^{(k+1)} = h_{xi} \sigma(h_{xi}, h_{yi}, \gamma^*/\lambda, \delta^*/\lambda) \quad (6.10)$$

$$d_{yi}^{(k+1)} = h_{yi} \sigma(h_{xi}, h_{yi}, \gamma^*/\lambda, \delta^*/\lambda) \quad (6.11)$$

where $\mathbf{h}_x = \mathbf{D}_x(\mathbf{u}^{(c)} + \mathbf{w}^{(k+1)}) + \mathbf{b}_x^{(k)}$, $\mathbf{h}_y = \mathbf{D}_y(\mathbf{u}^{(c)} + \mathbf{w}^{(k+1)}) + \mathbf{b}_y^{(k)}$, and σ is the shrinkage function [36] defined as

$$\sigma(a, b, \xi, \eta) = \frac{\max(0, \sqrt{a^2 + b^2} - \xi)}{(1 + \eta) \sqrt{a^2 + b^2}}. \quad (6.12)$$

The splitting scheme can be generalized to 3-dimensional domains. In such a case, one introduces an additional variable \mathbf{d}_z subject to $\mathbf{d}_z = \mathbf{D}_z(\mathbf{u}^{(c)} + \mathbf{w})$ into the splitting. The $(\mathbf{d}_x, \mathbf{d}_y, \mathbf{d}_z)$ sub-problem is solved via the ansatz $d_{xi} = h_{xi} \sigma$, $d_{yi} = h_{yi} \sigma$, $d_{zi} = h_{zi} \sigma$, which leads to a shrinkage formula similar to (6.12) except with $\sqrt{a^2 + b^2 + c^2}$ instead of $\sqrt{a^2 + b^2}$.

The iteration procedure (6.5)–(6.8) is repeated until the relative change in the functional (6.4) between iterations falls below a certain tolerance. Completing this iteration procedure produces a candidate \mathbf{w} , which is accepted or rejected as follows. (I.) The candidate is accepted if it produces a reduction of the functional (6.3) and rejected otherwise. (II.) If the candidate is rejected, the coefficients α^* ,

β^* , γ^* and δ^* are reduced by a factor of 2 (thus reducing the trust region) and a new candidate is computed. This process is repeated until a satisfactory candidate is found. (III.) If the first candidate \mathbf{w} for a certain starting point is satisfactory, the coefficients α^* , β^* , γ^* and δ^* are increased by a factor of 2 (thus expanding the trust region).

The minimization of (6.3) is carried until the relative change in the functional falls below a certain tolerance. The vector $\mathbf{u}^{(+)} = \mathbf{u}^{(c)} + \mathbf{w}$ of the last iteration is taken as the estimator of the parameter vector.

6.4 Forward solution

The methodology described in the previous section requires computing the state and the Jacobian associated with each iterate of the model parameter vector \mathbf{u} . In this study, we have employed a method-of-lines (MOL) strategy to solve the flow equation (6.1). An MOL discretization of a linear PDE produces a system of N_h linear ODEs

$$\left[\mathbf{M} \frac{d}{dt} + \mathbf{S}(\mathbf{k}) \right] \tilde{\mathbf{h}}(t) = \mathbf{g}(\mathbf{k}) + \mathbf{q}, \quad (6.13)$$

where $\tilde{\mathbf{h}}(t)$ is the hydraulic head vector at time t , \mathbf{M} and \mathbf{S} are the corresponding ‘‘capacitance’’ and ‘‘conductivity’’ matrices, \mathbf{g} is the forcing vector representing the boundary conditions, and \mathbf{q} is the vector representing the system loading. The set of ODEs (6.13) is integrated in time using an appropriate time-stepping scheme, and the state vector is stored for the N_t discrete observation times t_1, t_2, \dots, t_{N_t} . The full state vector is assembled as $\mathbf{h} = [\tilde{\mathbf{h}}(t_1)^\top, \tilde{\mathbf{h}}(t_2)^\top, \dots, \tilde{\mathbf{h}}(t_{N_t})^\top]^\top$.

Assuming that the system load \mathbf{q} is independent of the conductivity field \mathbf{k} , and using the chain rule with respect to the i -th element of \mathbf{k} in (6.13), we obtain a set of linear ODEs governing the evolution of the i -th column of the sensitivity $\mathbf{J}_{,i} = \partial \tilde{\mathbf{h}} / \partial k_i$:

$$\left[\mathbf{M} \frac{d}{dt} + \mathbf{S}(\mathbf{k}) \right] \mathbf{J}_{,i} = \frac{\partial \mathbf{g}(\mathbf{k})}{\partial k_i} - \frac{\partial \mathbf{S}(\mathbf{k})}{\partial k_i} \tilde{\mathbf{h}}. \quad (6.14)$$

The computation of the full sensitivity requires the integration in time of N_k sets of N_h ODEs. For each forward simulation, the sensitivities at the N_t observation

times are stored, and the Jacobian \mathbf{A} of the functional $\mathbf{h} = \mathbf{f}(\mathbf{k})$ is assembled as $\mathbf{A} = [\mathbf{J}(t_1)^\top, \mathbf{J}(t_2)^\top, \dots, \mathbf{J}(t_{N_t})^\top]^\top$.

The sets of ODEs (6.13) and (6.14) are similar, differing only in their arguments and right-hand-sides. Therefore the computation of the state vector and Jacobian can be understood as the simultaneous time-stepping of $N_h(N_k + 1)$ linear ODEs. Given that one of the leading costs of the time-stepping of linear systems of ODEs using implicit schemes is the factorization of the iteration matrix, integrating state and sensitivities together in time leads to significant computational savings.

6.5 Computational cost

The feasibility of this inversion strategy depends on one’s ability to efficiently solve the forward problems and the iterative optimization sub-problems. The most expensive operations of this method are the integration of the forward equations (6.13)–(6.14) and the linear algebraic system (6.9). The time-stepping of the ODEs using implicit methods requires the solution of $m(N_k + 1)$ linear algebraic systems of size $N_h \times N_h$ per time step, with m an integer associated with the scheme. We therefore recommend the use of high-order, variable step ODE integrators for their solution. The use of high-order methods with variable time steps allows for accurate time-stepping with long time steps, reducing the total amount of time steps and the overall cost.

Several robust high-order implicit schemes are known to perform well on sets of ODEs stemming from MOL discretizations. These include the CVODES package [40, 75], which uses high-order backward differences formulae, and Implicit Runge-Kutta (IRK) methods [38]. In the simulations reported below we used the 3-stage, order 5 Radau IIA IRK formula, which allows for time step selection via a posteriori error control.

Our implementation allows one to control both the error of the Jacobian and the state. Our experience shows that controlling for the Jacobian tends to produce smaller steps than just controlling for the step, without any appreciable

impact in the inversion procedure. Hence we used only state error control. The accuracy of the Jacobian's approximation affects chiefly the computation of $\mathbf{w}^{(k+1)}$ in (6.5) and (6.9). If the approximation of \mathbf{A} were to be insufficiently accurate to negatively affect the computation of \mathbf{w} , the candidate would still be rejected if it does not decrease the objective functional (6.3). Nevertheless, our experience shows that this is not the case: controlling only for the error of the state produces adequate approximations of \mathbf{A} .

When written in terms of $\mathbf{A}^\top \mathbf{Q}^\top \mathbf{Q} \mathbf{A}$, the size of the system (6.9) is $N_k \times N_k$. By repeated applications of the matrix inversion lemma one can rewrite this system in terms of $\mathbf{Q} \mathbf{A} \mathbf{A}^\top \mathbf{Q}^\top$, which reduces its size to $(N_{\text{obs},h} N_t) \times (N_{\text{obs},h} N_t)$. This reduction is possible if the terms $\alpha \mathbf{P}^\top \mathbf{P}$ and $\lambda \mathbf{D}_x^\top \mathbf{D}_x + \lambda \mathbf{D}_y^\top \mathbf{D}_y$ are either diagonalizable or trivially invertible. For example, if one uses a circulant discretization of the differentiation operators, the term $\lambda \mathbf{D}_x^\top \mathbf{D}_x + \lambda \mathbf{D}_y^\top \mathbf{D}_y$ is a block circulant with circulant blocks (BCCB) matrix. It can be diagonalized and inverted via the 2D discrete Fourier transform. This approach can lead to nontrivial savings in the solution of (6.9) if $N_{\text{obs},h} N_t \ll N_k$.

6.6 Application

To test our inversion strategy, we consider flow in a square domain $[0, 12] \times [0, 12]$ that is discretized into 30 cells in each spatial direction. The reference log conductivity field, shown in Figure 6.1, represents two heterogeneous facies with spatial averages differing by several orders of magnitude. Error-free measurements of hydraulic conductivity are taken at the $N_{\text{obs},k} = 25$ cells labeled with an 'x'. Note that only two observations fall into the low conductivity region. A first-order finite-element scheme is used to compute the head response at the nodes of each conductivity cell. Error-free measurements of hydraulic head are taken at $N_t = 6$ discrete times ($10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1$) in the $M_{\text{obs},h} = 25$ locations shown with white circles and at $N_t = 6$ discrete times ($10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 10^0, 10^1$) [T].

Boundary conditions are of constant hydraulic heads $h = 20$ along the left boundary and $h = 10$ along the right boundary. No-flow boundary conditions are




Figure 6.1: Synthetic log-hydraulic conductivity field.

imposed along the upper and lower boundaries. The initial condition is the steady configuration of heads in the absence of external loadings.

The algorithm is tested for three system loading (pumping) scenarios. The first scenario is of zero loadings, so the hydraulic head field is steady. The other two scenarios consist of a single steady point loading of strength and location indicated in Table 6.1 representing synthetic pumping tests. Additionally, the inversion is performed for 4 combinations of the inversion parameters α , β , γ and δ of equation (6.2), indicated in Table 6.2. The parameter sets are chosen to evaluate the effect of different degrees of penalization of the deviation of the reconstructed heads, TV regularization and ℓ_2 regularization.


Table 6.1: Properties of point loading in three test cases.

Scenario	Location	Strength
1	No loading	
2	(6.0, 6.0)	5×10^0
3	(2.8, 6.0)	5×10^3

Figure 6.2 exhibits the estimated conductivity fields obtained for the first (loading-free) scenario, together with the set (a) of inversion parameters (as listed in Table 6.2). Our inversion algorithm is capable of reconstructing the overall features of the original field, namely, the location, orientation and extent of the low conductivity intrusion, even though the intrusion itself is highly undersampled (2

Table 6.2: Inversion parameter sets

Set	α	β	γ	δ
a	1.0	1.0	1.0×10^{-6}	2.0×10^{-8}
b	1.0	0.5	1.0×10^{-6}	2.0×10^{-8}
c	1.0	1.0	1.0×10^{-4}	2.0×10^{-8}
d	1.0	0.5	1.0×10^{-4}	2.0×10^{-8}

**Figure 6.2:** Reconstructed log-hydraulic conductivity field for Test 1.

out of 25 conductivity measurements). Additionally, the intrusion is reconstructed as a region of conductivity higher than in the original field.

The satisfactory reconstruction is a consequence of the additional information provided by the hydraulic head observations, together with the piece-wise a priori information provided by the TV operator. One can think of the hydraulic head data fidelity term, after the linearization of the functional $\mathbf{h} = \mathbf{f}(\mathbf{k})$ around the solution, as a quadratic form

$$\frac{1}{2} \mathbf{w}^\top \mathbf{A}^\top \mathbf{Q}^\top \mathbf{Q} \mathbf{A} \mathbf{w} + (\text{other terms}).$$

The rank of the symmetric positive semi-definite matrix $\mathbf{A}^\top \mathbf{Q}^\top \mathbf{Q} \mathbf{A}$ gives an idea of the dimensionality of the parameter subspace observable via observations of the system state. In the first loading scenario the matrix $\mathbf{A}^\top \mathbf{Q}^\top \mathbf{Q} \mathbf{A}$, evaluated around the estimated field shown in Figure 6.2, has an effective rank of 25 (the number of head observations). Its eigenvectors show that the introduction of state measurements allows for observation of features that extend spatially beyond the

observation points. This added information, together with the piecewise regularity assumptions encapsulated into the TV norm, allow for the reconstruction shown in Figure 6.2.

While the reconstruction in Figure 6.2 captures the overall trend in hydraulic conductivity, it treats the actual observed cells as isolated, point-wise features. This is an effect of the TV operator, which penalizes discontinuities along $(n - 1)$ -dimensional manifolds separating n -dimensional facies, in favor of point discontinuities. Point discontinuities imply only n jumps per discontinuity and therefore introduce less total variation than a jump across a manifold of large surface area. This effect is more noticeable in the observational cells with values that are significantly different from the overall conductivity field.

Figure 6.3 shows the estimated conductivity fields obtained for the second loading scenario. The location of the point load is marked with a rhombus. The reconstructed hydraulic head at each observation point compared to observations for the inversion test (d) are presented in Figure 6.4. The conductivity estimator is similar to the one obtained for the steady case, albeit with an additional feature. In the immediate vicinity of the head and conductivity observation locations along the low-conductivity intrusion, the conductivity field is more accurately estimated as being of lower conductivity. This effect is more noticeable in the vicinity of the head observation point (5.6, 6.0), for which there is a significant observed decay in head due to the point load (Figure 6.4). The extent of the low-conductivity patch around observation point (5.6, 6.0) is limited by the TV regularization. As mentioned earlier, it has a tendency to reduce the area of the surface separating the dark blue patch around (5.6, 6.0) and the surrounding low-conductivity intrusion. Indeed, the extent of this patch is reduced as the TV coefficient γ increases (Figures 6.3 (c) and (d)).

Comparison of Figures 6.2 and 6.3 reveals that the addition of transient information has a clear effect on the reconstruction results. This is to be expected because more information is available in Test 2. The rank of $\mathbf{A}^\top \mathbf{Q}^\top \mathbf{Q} \mathbf{A}$ evaluated around the estimated field in Figure 6.3 (a) increases from 25 (Test 1) to 76 (Test 2), which implies that more information from the underlying field can be recovered.




Figure 6.3: Reconstructed log-conductivity fields for Test 2.

Figure 6.5 presents the estimated conductivity fields obtained for the third loading scenario, with the location of the point load marked with a rhombus. The reconstructed and observed hydraulic heads at each observation point for Test 3 (d) are shown in Figure 6.6. The latter figure demonstrates that if the point loading is located in the high conductivity region, it produces a significant variation of heads in observation locations both along the low conductivity intrusion and in the center and upper left portions of the high conductivity region. As expected, the addition of more transient data increases the amount of recoverable information. The rank of $\mathbf{A}^T \mathbf{Q}^T \mathbf{Q} \mathbf{A}$ (with \mathbf{A} evaluated around the reconstructed field shown in Figure 6.5 (a)) increases from 25 (no transient data) to 121.

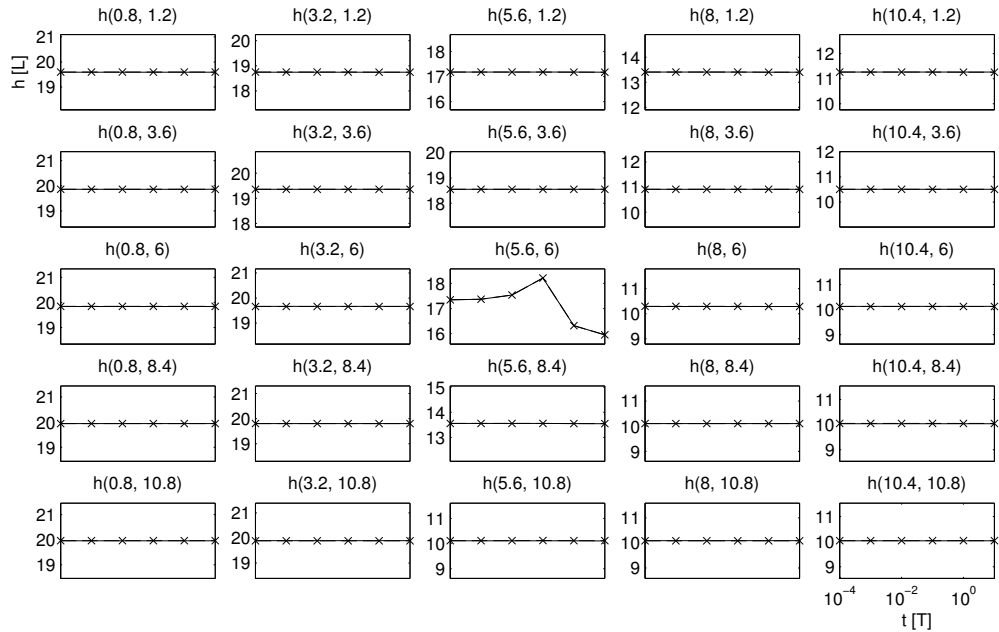


Figure 6.4: Observed (continuous) and reconstructed (dotted) hydraulic heads at observation locations for Test 2(d).

Table 6.3: Penalty results for Test 2.

Set	$\ \mathbf{P}\mathbf{u} - \mathbf{s}\ _2^2/2$	$\ \mathbf{Q}\mathbf{v} - \mathbf{t}\ _2^2/2$	TV $\ \mathcal{D}(\mathbf{u})\ _1$	$\ \mathcal{D}(\mathbf{u})\ _2^2$
a	9.69×10^{-11}	4.25×10^{-6}	5.92×10^2	3.59×10^3
b	2.81×10^{-10}	1.68×10^{-5}	6.10×10^2	3.75×10^3
c	1.03×10^{-6}	7.04×10^{-4}	5.97×10^2	3.48×10^3
d	9.80×10^{-7}	1.73×10^{-3}	5.94×10^2	3.64×10^3

Figure 6.5 (a) reveals that for a stringent penalty on the reconstruction of heads ($\beta = 1$) and a relatively lax TV requirement ($\gamma = 1 \times 10^{-6}$) this information increase produces a reconstruction of the intrusion that includes two fingers of high conductivity connecting the two observation locations. This implies that in order to accurately match the head observations (see Table 6.4 with heads data fidelity $\frac{1}{2}\|\mathbf{Q}\mathbf{v} - \mathbf{t}\|_2^2 = 2.13 \times 10^{-5}$) without introducing either the required heterogeneity into the field (which would significantly increase TV) or the well delineated boundary between facies (which is not observable), our reconstruction approach prefers to pay a small price in TV by introducing the high-conductivity fingers. A more regular reconstruction can be obtained either by relaxing the




Figure 6.5: Reconstructed log-conductivity fields for Test 3.

penalty on reconstruction of heads (Figure 6.5 (b), which increases $\frac{1}{2}\|\mathbf{Q}\mathbf{v} - \mathbf{t}\|_2^2$ from 2.13×10^{-5} to 1.85×10^{-3}) or by increasing the TV coefficient (Figure 6.5 (c), which increases $\frac{1}{2}\|\mathbf{Q}\mathbf{v} - \mathbf{t}\|_2^2$ to 1.68×10^{-3} , but decreases TV from 1.06×10^3 to 6.38×10^2). A proper selection of inversion parameters reflecting the desired balance between reconstruction and regularization is the responsibility of the modeler.

6.7 Conclusions and future work

Our numerical experiments show that the linearized functional minimization algorithm is a feasible and promising approach for inverse modeling of geophys-

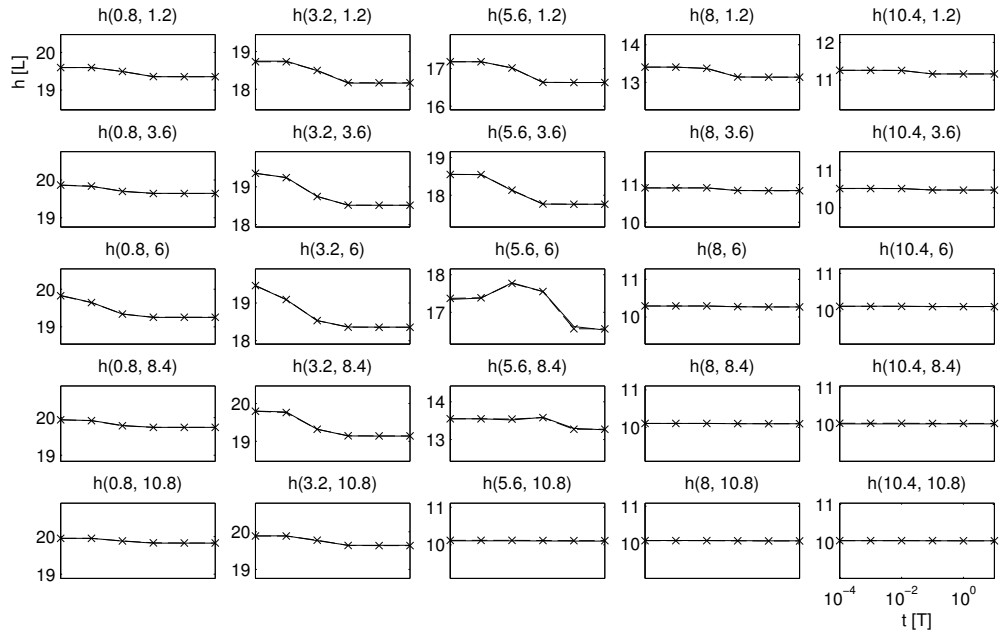


Figure 6.6: Observed (continuous) and reconstructed (dotted) hydraulic heads at observation locations for Test 3(d).

Table 6.4: Penalty results for Test 3.

Set	$\ \mathbf{P}\mathbf{u} - \mathbf{s}\ _2^2/2$	$\ \mathbf{Q}\mathbf{v} - \mathbf{t}\ _2^2/2$	TV $\ \mathcal{D}(\mathbf{u})\ _1$	$\ \mathcal{D}(\mathbf{u})\ _2^2$
a	2.24×10^{-10}	2.13×10^{-5}	1.06×10^3	1.52×10^4
b	6.01×10^{-8}	1.85×10^{-3}	8.72×10^2	9.18×10^3
c	5.95×10^{-7}	1.68×10^{-3}	6.38×10^2	5.78×10^3
d	1.07×10^{-6}	5.35×10^{-3}	6.06×10^2	4.29×10^3

ical systems. The application to a steady flow inversion problem shows that the strategy can be used to detect large-scale features of the parameter field, provided it is equipped with appropriate regularization terms reflecting the prior knowledge of the field (e.g., the TV norm). Incorporation of transient information about the state variables increases the amount of information that can be recovered (Tests 2 and 3 shown in Figures 6.3–6.6). The modeler must pay attention to a possible appearance of extraneous features in the reconstructed field. These are introduced by the algorithm in order to match additional head observations, while keeping the total variation of the field low (Figure 6.5).

The approach is not limited to systems described by linear governing partial

differential equations, such as the groundwater flow equation (6.1). All that is required is for the Jacobian of the functional that relate the system parameters to the system states to be computable. This implies that the strategy can be easily expanded to incorporate measurements of multiple state variables and both measurements and previous knowledge of multiple spatially distributed parameters.

Our linearized functional minimization approach to parameter identification presents a promising venue of research. Future work will be focused on designing a regularization operator that improves over TV by addressing its observed limitations. These include its tendency to penalize large discontinuity surfaces in favor of smaller surfaces or even point discontinuities. The trust region approach used to minimize (6.3) can be improved for robustness and efficiency, together with the selection of the inversion parameters.

Chapter 6, in part, is currently being prepared for submission for publication of the material: D. A. Barajas-Solano, D. A., Wohlberg, B., Vesselinov, V., Tartakovsky, D. M., “Linear Functional Minimization for Inverse Modeling”. The dissertation author was the primary investigator and author of the paper.

Chapter 7

Conclusions

This dissertation leads to the following major conclusions.

1. We have proposed a regularization methodology for the numerical computation of Green's functions for elliptic boundary value problems with variable piecewise continuous coefficients. These Green's functions can be used in uncertainty quantification to solve deterministic partial differential equations (PDEs) governing the statistics of the quantities of interest that can be derived from the original elliptic stochastic PDEs. Our regularization methodology accounts for the lack of a solution's regularity due to the Dirac forcing to the Green's functions' problems. If not addressed, the lack of regularity can significantly affect the accuracy of the computed Green's functions. Numerical experiments show that the proposed regularization methodology improves the accuracy of the approximated Green's function over standard methods that do not take into account the issue of regularity.
2. We have compared the performance of sampling techniques (e.g., Monte Carlo simulations) and global stochastic collocation strategies for the estimation of first moments of quantities of interest in nonlinear stochastic parabolic and elliptic PDEs. The performance of global stochastic collocation methods, and in general strategies based on a spectral decomposition of state variables in probability space, is significantly affected by the statistical properties of uncertain model parameters. As the variance of random

input fields increases it becomes necessary to use methods with high-order accuracy. At some point these methods are computationally less efficient than Monte Carlo simulations and its derivatives. These methods also scale poorly with increasing number of dimensions under high-order of accuracy requirements. A decrease in the correlation length of random input fields can significantly increase the computational cost of the estimators. Finally, different quantities of interest derived from the same nonlinear system can have markedly different behaviors in probability space. Our numerical experiments demonstrate the need to develop uncertainty quantification strategies that can outperform simple Monte Carlo under high variance requirements in high dimensional spaces.

3. To address the shortcomings of stochastic collocation techniques, we have developed a Multi-level Monte Carlo (MLMC) framework for computing the statistics of quantities of interest described by stochastic PDEs. This framework has been applied to a family of nonlinear elliptic boundary value problems with uncertain coefficients. It has yielded a significant reduction in the computational cost relative to that of simple Monte Carlo simulations. Additionally, MLMC allows for the quantification of the estimator bias introduced by the spatial discretization of the stochastic PDEs. We have performed numerical experiments with a high number of stochastic dimensions, effectively removing the bias introduced by the truncation of the stochastic representation of the random input fields. The selection of spatio-temporal discretization schemes, which increase the rate of decay of the various biases with increasing Monte Carlo level, is a critical area for future research.
4. We have presented a computationally efficient algorithm for identification of sources and release histories of (geo)chemically active solutes. It relies on a Bayesian formulation of the inversion problem, and employs a generalized hybrid Monte Carlo (gHMC) method to sample the posterior distribution. The latter accelerates the Markov chain Monte Carlo (MCMC) sampling by using information about the gradient of the likelihood to reduce the correlation of the Markov chain. While our inversion algorithm can in principle

handle nonlinear problems, the computation of the sensitivities for these problems is a challenge. Robust and efficient sensitivity analysis techniques must be employed to keep the overall computational cost reasonably low. We have constructed acceleration matrices for the gHMC scheme for a linear advection-dispersion problem. These acceleration matrices significantly reduce the correlation of the Markov chain over HMC without acceleration. The generalization of these constructions to problems with nonlinear reaction terms and non-periodic boundary conditions is the subject of further work.

5. We have introduced a novel strategy for estimation of parameters in nonlinear PDEs from discrete measurements of system parameters and system states. This approach computes the maximum *a posteriori* estimator based on a likelihood functional. The latter reflects available discrete (in space) measurements of system parameters and discrete (in space and time) measurements of transient system states. The approach makes use of regularization terms (e.g., the total variation (TV) regularization norm) that capture the prior information about system parameters, e.g., the fact that they are piecewise continuous. Our numerical experiments show that the proposed strategy is capable of detecting large-scale features and values of the uncertain parameter field. Further work is required to design better regularization operators that overcome some of the shortcomings of the TV norm. The incorporation of more informative data, such as tracer measurements. will also be explored.

Appendix A

Computation of sensitivities for the hybrid Monte Carlo method

A.1 Adjoint sensitivity analysis

The forward problem (5.1) can be written in discrete form as a regular ODE problem

$$\mathbf{f}(t, \mathbf{c}, \mathbf{c}_0) = \mathbf{c}_t - \mathbf{A}\mathbf{c} + \mathbf{R}(\mathbf{c}) - \mathbf{r} = 0, \quad \mathbf{c}(0) = \mathbf{q}, \quad (\text{A.1})$$

where \mathbf{A} is the discretized form of the linear advection-diffusion operator $\mathcal{D}(c) = \nabla \cdot (\mathbf{D}\nabla c) - \nabla \cdot (\mathbf{u}c)$ plus boundary conditions, \mathbf{r} is the discretized source vector (which we set to zero without loss of generality), and \mathbf{R} is a vector with components $R_i = R(c_i)$.

In the following discussion, the subindices of \mathbf{f} , \mathbf{c} and scalar functions denote partial derivatives. The sensitivity of H in (5.13) with respect to vector \mathbf{c}_0 is computed following the procedure outlined in [16]. We begin by defining the functionals

$$g(t, \mathbf{c}, \mathbf{q}) = \frac{1}{2\sigma_\epsilon^2} \Delta \mathbf{c}^\top \mathbf{M} \Delta \mathbf{c} + \gamma \mathbf{q}^\top \mathbf{M}' \mathbf{q}, \quad G(\mathbf{c}, \mathbf{q}) = \int_0^T g(t, \mathbf{c}, \mathbf{q}) dt. \quad (\text{A.2})$$

Then $H = g(T) = dG/dT$, and $\nabla_{\mathbf{q}} H = \nabla_{\mathbf{q}} (dG/dT)$. Next, for $\mathbf{f} = \mathbf{0}$ we introduce a Lagrange multiplier $\mathbf{v}(t)$ such that

$$G(\mathbf{c}, \mathbf{q}) = G(\mathbf{c}, \mathbf{q}) - \int_0^T \mathbf{v}^* \mathbf{f}(t, \mathbf{c}, \mathbf{q}) dt$$

where $(\cdot)^*$ denotes the Hermitian adjoint. The sensitivity of G is given by

$$\nabla_{\mathbf{q}}G = \int_0^T [g_{\mathbf{c}}\mathbf{c}_{\mathbf{q}} + g_{\mathbf{q}}] dt - \int_0^T \mathbf{v}^*[\mathbf{f}_{\mathbf{c}}\mathbf{c}_{\mathbf{q}} + (\mathbf{c}_{\mathbf{q}})_t] dt. \quad (\text{A.3})$$

Integrating by parts,

$$\int_0^T \mathbf{v}^*(\mathbf{c}_{\mathbf{q}})_t dt = [\mathbf{v}^*\mathbf{c}_{\mathbf{q}}]_0^T - \int_0^T \mathbf{v}_t^*\mathbf{c}_{\mathbf{q}} dt. \quad (\text{A.4})$$

Substituting this into (A.3),

$$\nabla_{\mathbf{q}}G = \int_0^T g_{\mathbf{q}} dt - \int_0^T \mathbf{c}_{\mathbf{q}}[-g_{\mathbf{c}} + \mathbf{v}^*\mathbf{f}_{\mathbf{c}} - \mathbf{v}_t^*] dt - [\mathbf{v}^*\mathbf{c}_{\mathbf{q}}]_0^T. \quad (\text{A.5})$$

Choosing \mathbf{v} as a solution of the reverse-time ODE problem

$$-\mathbf{v}_t^* + \mathbf{v}^*\mathbf{f}_{\mathbf{c}} - g_{\mathbf{c}} = 0, \quad \mathbf{v}(T) = 0,$$

greatly simplifies (A.5).

For the sensitivity of $H = g(T)$, we take the total derivative of (A.5) with respect to T to obtain

$$\nabla_{\mathbf{q}}H = g_{\mathbf{q}} - \int_0^T \mathbf{c}_{\mathbf{q}}[-(\mathbf{v}_T^*)_t + \mathbf{v}_T^*\mathbf{f}_{\mathbf{c}}] dt + (\mathbf{v}_T^*\mathbf{c}_{\mathbf{q}})|_{t=0} - \frac{d}{dT}(\mathbf{v}^*\mathbf{c}_{\mathbf{q}})|_{t=T}$$

Taking \mathbf{v}_T to be a solution of the ODE problem

$$-(\mathbf{v}_T^*)_t + \mathbf{v}_T^*\mathbf{f}_{\mathbf{c}} = 0, \quad \mathbf{v}_T^*(T) = -\mathbf{v}_t(T) = g_{\mathbf{c}},$$

the sensitivity of H reduces to

$$\nabla_{\mathbf{q}}H = \nabla_{\mathbf{q}}g + \mathbf{v}_T(0) \quad (\text{A.6})$$

Since \mathbf{M} and \mathbf{M}' are symmetric, $g_{\mathbf{c}} = \mathbf{M}\Delta\mathbf{c}/\sigma_c^2$, and $g_{\mathbf{q}} = 2\gamma\mathbf{M}'\mathbf{q}$. Also, $\mathbf{f}_{\mathbf{c}} = \mathbf{A} + \mathbf{R}_{\mathbf{c}}$. Finally, writing \mathbf{v} for \mathbf{v}_T , we obtain (5.15).

A.2 Ad-hoc sensitivity analysis

For the problem in section 5.4.2 we use the linearized Runge-Kutta (Rosenbrock) method ROS2 of [92] for time stepping of the forward ODE problem. The

advantage of using this method is that it allows for a linear implicit treatment of the dispersion operator and a linearization of the reaction operator, while the advection operator is treated explicitly.

We assume that the advection-dispersion-reaction equation can be discretized into an autonomous system of ODEs

$$\mathbf{c}_t = \mathbf{f}(\mathbf{c}) = (\mathbf{A}_D + \mathbf{A}_A)\mathbf{c} - \mathbf{R}(\mathbf{c}),$$

where \mathbf{c} is the state vector, \mathbf{A}_D is the discretized linear dispersion operator, \mathbf{A}_A is the discretized linear advection operator, and $\mathbf{R}(\mathbf{c})$ is the reaction vector. Time stepping is performed via a scheme

$$\mathbf{c}^{n+1} = \mathbf{c}^n + (2 - b)\Delta t \mathbf{k}_1 + b\Delta t \mathbf{k}_2, \quad (\text{A.7})$$

$$(\mathbf{I} - \theta\Delta t \mathbf{J})\mathbf{k}_1 = \mathbf{f}(\mathbf{c}^n), \quad (\text{A.8})$$

$$(\mathbf{I} - \theta\Delta t \mathbf{J})\mathbf{k}_2 = \mathbf{f}\left(\mathbf{c}^n + \frac{1}{2b}\Delta t \mathbf{k}_1\right) - \frac{1}{b}\mathbf{k}_1, \quad (\text{A.9})$$

where $\mathbf{J} = \mathbf{f}_c(\mathbf{c}^n)$ is the Jacobian of \mathbf{f} with respect to the state. The coefficients θ and b are taken for this application as $\theta = 1 - \sqrt{2}/2$ and $b = 1/2$, respectively. The left-hand side operators of (A.8)–(A.9) are approximated via approximate matrix factorization (AMF) to obtain the split form

$$(\mathbf{I} - \theta\Delta t \mathbf{J}) \approx (\mathbf{I} - \theta\Delta t \mathbf{A}_D)(\mathbf{I} + \theta\Delta t \mathbf{R}_c(\mathbf{c}^n)).$$

The discussion in section 5.4.2 led us to conclude that it is necessary to compute products of the form $(d\mathbf{c}^{i+1}/d\mathbf{c}^i)^\top \mathbf{u}$ in order to apply the ad-hoc sensitivity technique of [24]. The formulae for the computation of these products is derived from the time-stepping scheme (A.7)–(A.9). In particular, differentiating (A.7) with respect to the state and multiplying by a test vector \mathbf{u} gives the single-step sensitivity product as

$$\left(\frac{d\mathbf{c}^{n+1}}{d\mathbf{c}^n}\right)^\top \mathbf{u} = \mathbf{u} + \frac{3}{2}\Delta t \left(\frac{d\mathbf{k}_1^n}{d\mathbf{c}^n}\right)^\top \mathbf{u} + \frac{1}{2}\Delta t \left(\frac{d\mathbf{k}_2^n}{d\mathbf{c}^n}\right)^\top \mathbf{u}.$$

The next task is to derive formulas for the Jacobians of the stage derivatives \mathbf{k}_1 and \mathbf{k}_2 . Let \mathbf{M} be the AMF-ed left-hand-side matrix of (A.8)–(A.9). Differentiating (A.8)–(A.9) with respect to the state and multiplying by the test vector \mathbf{u}

gives the formulae

$$\left(\frac{d\mathbf{k}_1^n}{d\mathbf{c}^n}\right)^\top \mathbf{u} = \left[\mathbf{J}_0 - \left(\frac{d\mathbf{M}}{d\mathbf{c}^n}\mathbf{k}_1^n\right)\right] \mathbf{v}, \quad \mathbf{M}^\top \mathbf{v} = \mathbf{u}$$

and

$$\left(\frac{d\mathbf{k}_2^n}{d\mathbf{c}^n}\right)^\top \mathbf{u} = \mathbf{J}_1^\top \mathbf{v} + \left(\frac{d\mathbf{k}_1^{n+1}}{d\mathbf{c}^n}\right)^\top (\Delta t \mathbf{J}_1 - 2\mathbf{I}) \mathbf{v} - \left(\frac{d\mathbf{M}}{d\mathbf{c}^n}\mathbf{k}_2^n\right)^\top \mathbf{v},$$

with $\mathbf{J}_0 = \mathbf{f}_c(\mathbf{c}^n)$, $\mathbf{J}_1 = \mathbf{f}_c(\mathbf{c}^n + \Delta t \mathbf{k}_1^n)$.

The computation of the products $(d\mathbf{M}/d\mathbf{c}^n)\mathbf{k}_i^n$, $i = 1, 2$ is highly problem-specific. It depends on the structure of the second-order derivatives of the reaction vector with respect to the state. For the reaction model (5.16) and a method-of-lines discretization, the Jacobian \mathbf{R}_c is diagonal and so the computation of these products is straightforward. For different reaction models and more sophisticated discretization schemes the computation might be more involved.

Bibliography

- [1] N. Agarwal and N. R. Aluru. A domain adaptive stochastic collocation approach for analysis of MEMS under uncertainties. *J. Comput. Phys.*, 228(20):7662–7688, 2009.
- [2] F. J. Alexander, G. L. Eyink, and J. M. Restrepo. Accelerated Monte Carlo for optimal estimation of time series. *J. Stat. Phys.*, 119:1331–1345, 2005.
- [3] C. Archambeau, M. Opper, Y. Shen, D. Cornford, and J. Shawe-Taylor. Variational inference for diffusion processes. In J. C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20, pages 1–8. MIT Press, 2007.
- [4] R. Archibald, A. Gelb, R. Saxena, and D. Xiu. Discontinuity detection in multivariate space for stochastic simulations. *J. Comput. Phys.*, 228(7):2676–2689, 2009.
- [5] U. Ascher and E. Haber. Computational methods for large distributed parameter estimation problems with possible discontinuities. In *Proceedings of the Inverse Problems, Design & Optimization Symposium*, pages 201–208, 2004.
- [6] J. Atmadja and A. C. Bagtzoglou. Pollution source identification in heterogeneous porous media. *Water Resour. Res.*, 37(8):2113–2125, 2001.
- [7] J. Atmadja and A. C. Bagtzoglou. State of the art report on mathematical methods for groundwater pollution source identification. *Environ. Forensics*, 2(3):205–214, 2001.
- [8] I. Babuška, R. Tempone, and G. E. Zouraris. Solving elliptic boundary value problems with uncertain coefficients by the finite element method: the stochastic formulation. *Comput. Meth. Appl. Mech. Eng.*, 194(12-16):1251–1294, 2005.
- [9] M. Bachmayr and M. Burger. Iterative total variation schemes for nonlinear inverse problems. *Inverse Prob.*, 25(10):105004, 2009.

- [10] A. C. Bagtzoglou and J. Atmadja. Marching-jury backward beam equation and quasi-reversibility methods for hydrologic inversion: Application to contaminant plume spatial distribution recovery. *Water Resour. Res.*, 39(2), 2003. doi:10.1029/2001WR001021.
- [11] A. C. Bagtzoglou, D. E. Dougherty, and A. F. B. Tompson. Application of particle methods to reliable identification of groundwater pollution sources. *Water Resour. Manage.*, 6:15–23, 1992.
- [12] L. Borcea, G. Papanicolaou, C. Tsogka, and J. Berryman. Imaging and time reversal in random media. *Inverse Prob.*, 18(5):1247, 2002.
- [13] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learning*, 3(1):1–122, 2010.
- [14] K. P. Bube and R. T. Langan. Hybrid ℓ^1/ℓ^2 minimization with applications to tomography. *Geophysics*, 62(4):1183–1195, 1997.
- [15] J. B. Butcher and T. D. Gauthier. Estimation of residual dense NAPL mass by inverse modeling. *Ground Water*, 32(1):71–78, 1994.
- [16] Y. Cao, S. Li, L. Petzold, and R. Serban. Adjoint sensitivity analysis for differential-algebraic equations: The adjoint DAE system and its numerical solution. *SIAM J. Sci. Comput.*, 24(3):1076–1089, 2003.
- [17] J. Carrera, A. Alcolea, A. Medina, J. Hidalgo, and L. J. Slooten. Inverse problem in hydrogeology. *Hydrogeol. J.*, 13(1):206–222, 2005.
- [18] H. D. Chen, S. Y. Chen, and R. H. Kraichnan. Probability distribution of a stochastically advected scalar field. *Phys. Rev. Lett.*, 63(24):2657–2660, 1989.
- [19] L. Chen, M. J. Holst, and J. Xu. The finite element approximation of the nonlinear Poisson-Boltzmann equation. *SIAM J. Numer. Anal.*, 45(6):2298–2320, 2007.
- [20] Z. Chen and X. Yue. Numerical homogenization of well singularities in the flow transport through heterogeneous porous media. *Multiscale Model. Simul.*, 1(2):260–303, 2003.
- [21] I.-L. Chern, J.-G. Liu, and W.-C. Wang. Accurate evaluation of electrostatics for macromolecules in solutions. *Methods Appl. Anal.*, 10(2):309–327, 2003.
- [22] E. T. Chung, T. F. Chan, and X.-C. Tai. Electrical impedance tomography using level set representation and total variational regularization. *J. Comput. Phys.*, 205(1):357 – 372, 2005.

- [23] K. Cliffe, M. Giles, R. Scheichl, and A. Teckentrup. Multilevel monte carlo methods and applications to elliptic pdes with random coefficients. *Comput. Visual. Sci.*, 14(1):3–15, 2011.
- [24] D. Daescu, G. R. Carmichael, and A. Sandu. Adjoint implementation of rosenbrock methods applied to variational data assimilation problems. *J. Comput. Phys.*, 165(2):496 – 510, 2000.
- [25] G. Dagan. Theory of solute transport by groundwater. *Ann. Rev. Fluid Mech.*, 19:183–215, 1987.
- [26] G. de Marsily, F. Delay, J. Gonalvs, P. Renard, V. Teles, and S. Violette. Dealing with spatial heterogeneity. *Hydrogeol. J.*, 13(1):161–183, 2005.
- [27] Y. Efendiev, O. Iliev, and C. Kronsbein. Multi-level monte carlo methods using ensemble level mixed msfem for two-phase flow and transport simulations. Technical Report 217, Fraunhofer Inst. für Techno-und Wirtschaftsmathematik, ITWM, 2012.
- [28] P. J. F. Müller and D. Meyer. Multilevel monte carlo for two phase flow and transport in random heterogeneous porous media. Technical Report 2012-12, Eidgenössische Technische Hochschule, may 2012.
- [29] J. Foo, X. Wan, and G. E. Karniadakis. The multi-element probabilistic collocation method (ME-PCM): Error analysis and applications. *J. Comput. Phys.*, 227(22):9572–9595, 2008.
- [30] H. J. H. Franssen, A. Alcolea, M. Riva, M. Bakr, N. van der Wiel, F. Stauffer, and A. Guadagnini. A comparison of seven methods for the inverse modelling of groundwater flow. application to the characterisation of well catchments. *Adv. Water Resour.*, 32(6):851 – 872, 2009.
- [31] B. Ganapathysubramanian and N. Zabararas. Sparse grid collocation schemes for stochastic natural convection problems. *J. Comput. Phys.*, 225(1):652–685, 2007.
- [32] A. Genz and B. D. Keister. Fully symmetric interpolatory rules for multiple integrals over infinite regions with Gaussian weight. *J. Comput. Appl. Math.*, 71(2):299–309, 1996.
- [33] T. Gerstner and M. Griebel. Numerical integration using sparse grids. *Numer. Algorithms*, 18:209–232, 1998.
- [34] R. G. Ghanem and P. D. Spanos. *Stochastic finite elements: a spectral approach*. Springer-Verlag New York, Inc., New York, NY, USA, 1991.

- [35] M. B. Giles. Multilevel monte carlo path simulation. *Oper. Res.*, 56(3):607–617, 2008.
- [36] T. Goldstein and S. J. Osher. The Split Bregman method for l1-regularized problems. *SIAM J. Imag. Sci.*, 2(2):323–343, 2009.
- [37] M. Gugat. Contamination source determination in water distribution networks. *SIAM J. Appl. Math.*, 72(6):1772–1791, 2012.
- [38] E. Hairer and G. Wanner. Stiff differential equations solved by radau methods. *J. Comput. Appl. Math.*, 111(12):93 – 111, 1999.
- [39] G. S. Herrera and G. F. Pinder. Space-time optimization of groundwater quality sampling networks. *Water Resour. Res.*, 41(12), 2005. doi:10.1029/2004WR003626.
- [40] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward. SUNDIALS: Suite of nonlinear and differential/algebraic equation solvers. *ACM Trans. Math. Softw.*, 31(3):363–396, Sept. 2005.
- [41] M. Holst, J. A. McCammon, Z. Yu, Y. Zhou, and Y. Zhu. Adaptive finite element modeling techniques for the Poisson-Boltzmann equation. *Commun. Comput. Phys.*, 11(1):179–214, 2012.
- [42] L. Hu. Combination of dependent realizations within the gradual deformation method. *Math. Geol.*, 34(8):953–963, 2002.
- [43] L. Hu and M. Le Ravalec-Dupin. An improved gradual deformation method for reconciling random and gradient searches in stochastic optimizations. *Math. Geol.*, 36(6):703–719, 2004.
- [44] C. Huang, T. Hsing, N. Cressie, A. R. Ganguly, V. A. Protopopescu, and N. S. Rao. Bayesian source detection and parameter estimation of a plume model based on sensor network measurements. *Appl. Stochastic Models Bus. Ind.*, 26(4):331–348, 2010.
- [45] M. A. Iglesias and D. McLaughlin. Level-set techniques for facies identification in reservoir modeling. *Inverse Prob.*, 27:035008, 2011.
- [46] J. D. Jakeman, R. Archibald, and D. Xiu. Characterization of discontinuities in high-dimensional stochastic problems on adaptive sparse grids. *J. Comput. Phys.*, 230(10):3977–3997, 2011.
- [47] K. D. Jarman and T. F. Russell. Eulerian moment equations for 2-D stochastic immiscible flow. *Multiscale Model. Simul.*, 1(4):598–608, 2003.

- [48] D. K. Kahaner and G. Monegato. Nonexistence of extended Gauss-Laguerre and Gauss-Hermite quadrature rules with positive weights. *Z. Angew. Math. Mech. (ZAMM)*, 29:983–986, 1978.
- [49] R. H. Kraichnan. Eddy viscosity and diffusivity: Exact formulas and approximations. *Complex Systems*, 1:805–820, 1987.
- [50] D. P. Landau and K. Binder. *A Guide to Monte Carlo Simulations in Statistical Physics*. Cambridge University Press, 2009.
- [51] O. P. Le Maître, O. M. Knio, H. N. Najm, and R. G. Ghanem. Uncertainty propagation using Wiener-Haar expansions. *J. Comput. Phys.*, 197(1):28–57, 2004.
- [52] O. P. Le Maître, H. N. Najm, R. G. Ghanem, and O. M. Knio. Multi-resolution analysis of Wiener-type uncertainty propagation schemes. *J. Comput. Phys.*, 197(2):502–531, 2004.
- [53] S. Li and L. Petzold. Adjoint sensitivity analysis for time-dependent partial differential equations with adaptive mesh refinement. *J. Comput. Phys.*, 198(1):310–325, 2004.
- [54] W. Li, Z. Lu, and D. Zhang. Stochastic analysis of unsaturated flow with probabilistic collocation method. *Water Resour. Res.*, 45(8), 2009. doi:10.1029/2008WR007530.
- [55] P. C. Lichtner and D. M. Tartakovsky. Upscaled effective rate constant for heterogeneous reactions. *Stochas. Env. Res. Risk Assess.*, 17(6):419–429, 2003.
- [56] C. Liu and W. P. Ball. Application of inverse methods to contaminant source identification from aquitard diffusion profiles at Dover AFB, Delaware. *Water Resour. Res.*, 35(7):1975–1985, 1999.
- [57] X. Ma and N. Zabarar. An adaptive hierarchical sparse grid collocation algorithm for the solution of stochastic differential equations. *J. Comput. Phys.*, 228(8):3084–3113, 2009.
- [58] A. M. Michalak and P. K. Kitanidis. A method for enforcing parameter nonnegativity in Bayesian inverse problems with an application to contaminant source identification. *Water Resour. Res.*, 39(2), 2003. doi:10.1029/2002WR001480.
- [59] A. M. Michalak and P. K. Kitanidis. Application of geostatistical inverse modeling to contaminant source identification at Dover AFB, Delaware. *J. Hydraul. Res.*, 42:9–18, 2004.

- [60] A. M. Michalak and P. K. Kitanidis. Estimation of historical groundwater contaminant distribution using the adjoint state method applied to geostatistical inverse modeling. *Water Resour. Res.*, 40(8), 2004. doi:10.1029/2004WR003214.
- [61] E. Morales-Casique, S. P. Neuman, and A. Guadagnini. Nonlocal and localized analyses of nonreactive solute transport in bounded randomly heterogeneous porous media: Theoretical framework. *Adv. Water Resour.*, 29(8):1238–1255, 2006.
- [62] R. M. Neal. Probabilistic inference using Markov chain Monte Carlo methods. Technical Report CRG-TR-93-1, University of Toronto, 1993.
- [63] R. M. Neal. *Bayesian learning for neural networks*. PhD thesis, University of Toronto, 1995.
- [64] S. P. Neuman and S. Orr. Prediction of steady state flow in nonuniform geologic media by conditional moments: Exact nonlocal formalism, effective conductivities, and weak approximation. *Water Resour. Res.*, 29:341–364, 1993.
- [65] S. P. Neuman, D. Tartakovsky, T. C. Wallstrom, and C. L. Winter. Correction to prediction of steady state flow in nonuniform geologic media by conditional moments: Exact nonlocal formalism, effective conductivities, and weak approximation by shlomo p. neuman and shlomo orr. *Water Resour. Res.*, 32(5):1479–1480, 1996.
- [66] R. M. Neupauer, B. Borchers, and J. L. Wilson. Comparison of inverse methods for reconstructing the release history of a groundwater contamination source. *Water Resour. Res.*, 36(9):2469–2475, 2000.
- [67] R. M. Neupauer and J. L. Wilson. Adjoint method for obtaining backward-in-time location and travel time probabilities of a conservative groundwater contaminant. *Water Resour. Res.*, 35(11):3389–3398, 1999.
- [68] R. M. Neupauer and J. L. Wilson. Backward location and travel time probabilities for a decaying contaminant in an aquifer. *J. Contam. Hydrol.*, 66(12):39 – 58, 2003.
- [69] R. M. Neupauer and J. L. Wilson. Forward and backward location probabilities for sorbing solutes in groundwater. *Adv. Water Resour.*, 27(7):689–705, 2004.
- [70] R. M. Neupauer and J. L. Wilson. Backward probability model using multiple observations of contamination to identify groundwater contamination sources at the Massachusetts Military Reservation. *Water Resour. Res.*, 41(2), 2005. doi:10.1029/2003WR002974.

- [71] B. Noetinger, V. Artus, and G. Zargar. The future of stochastic and upscaling methods in hydrogeology. *Hydrogeol. J.*, 13(1):184–201, 2005.
- [72] E. Novak and K. Ritter. Simple cubature formulas with high polynomial exactness. *Constr. Approx.*, 15(4):499–522, 1999.
- [73] P. Rodríguez and B. Wohlberg. Efficient minimization method for a generalized total variation functional. *IEEE Trans. Image Process.*, 18(2):322–332, Feb. 2009.
- [74] L. Rudin, S. J. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60(1-4):259–268, November 1992.
- [75] R. Serban and A. C. Hindmarsh. CVODES, the sensitivity-enabled ode solver in SUNDIALS. In *Proceedings of the 5th International Conference on Multibody Systems, Nonlinear Dynamics and Control, Long Beach, CA, Long Beach, CA, USA, Sept. 2005*.
- [76] S. Shlomi and A. M. Michalak. A geostatistical framework for incorporating transport information in estimating the distribution of a groundwater contaminant plume. *Water Resour. Res.*, 43(3), 2007. doi:10.1029/2006WR005121.
- [77] T. H. Skaggs and Z. J. Kabala. Recovering the release history of a groundwater contaminant. *Water Resour. Res.*, 30(1):71–79, 1994.
- [78] T. H. Skaggs and Z. J. Kabala. Recovering the history of a groundwater contaminant plume: Method of quasi-reversibility. *Water Resour. Res.*, 31(11):2669–2673, 1995.
- [79] S. A. Smolyak. Quadrature and interpolation formulas for tensor products of certain classes of functions. *Dokl. Akad. Nauk SSSR*, 4:240–243, 1963.
- [80] M. F. Snodgrass and P. K. Kitanidis. A geostatistical approach to contaminant source identification. *Water Resour. Res.*, 33:537–546, 1997.
- [81] A. M. Tartakovsky, P. Meakin, and H. Huang. Stochastic analysis of immiscible displacement of the fluids with arbitrary viscosities and its dependence on support scale of hydrological data. *Adv. Water Resour.*, 27(12):1151–1166, 2004.
- [82] A. M. Tartakovsky, S. P. Neuman, and R. J. Lenhard. Immiscible front evolution in randomly heterogeneous porous media. *Phys. Fluids*, 15(11):3331–3341, 2003.
- [83] D. M. Tartakovsky and S. Broyda. PDF equations for advective-reactive transport in heterogeneous porous media with uncertain properties. *J. Contam. Hydrol.*, 120-121:129–140, 2011.

- [84] D. M. Tartakovsky, M. Dentz, and P. C. Lichtner. Probability density functions for advective-reactive transport in porous media with uncertain reaction rates. *Water Resour. Res.*, 45, 2009. doi:10.1029/2008WR007383.
- [85] D. M. Tartakovsky, A. Guadagnini, and M. Riva. Stochastic averaging of nonlinear flows in heterogeneous porous media. *J. Fluid Mech.*, 492:47–62, 2003.
- [86] D. M. Tartakovsky, Z. Lu, A. Guadagnini, and A. M. Tartakovsky. Unsaturated flow in heterogeneous soils with spatially distributed uncertain hydraulic parameters. *J. Hydrol.*, 275(3-4):182–193, 2003.
- [87] D. M. Tartakovsky and S. P. Neuman. Transient flow in bounded randomly heterogeneous domains 1. Exact conditional moment equations and recursive approximations. *Water Resour. Res.*, 34(1):1–12, 1998.
- [88] D. M. Tartakovsky, S. P. Neuman, and Z. Lu. Conditional stochastic averaging of steady state unsaturated flow by means of Kirchhoff transformation. *Water Resour. Res.*, 35(3):731–745, 1999.
- [89] D. M. Tartakovsky and B. Wohlberg. Delineation of geologic facies with statistical learning theory. *Geophys. Res. Lett.*, 31(18), 2004. doi:10.1029/2004GL020864.
- [90] R. Toral and A. Ferreira. A general class of hybrid Monte Carlo methods. In R. Gruber and M. Tomasi, editors, *Proceedings of Physics Computing*, volume 94, pages 265–268. EPS, 1994.
- [91] J. Tryoen, O. P. Le Maître, M. Ndjinga, and A. Ern. Intrusive Galerkin methods with upwinding for uncertain nonlinear hyperbolic systems. *J. Comput. Phys.*, 229(18):6485–6511, 2010.
- [92] J. G. Verwer, E. J. Spee, J. G. Blom, and W. Hundsdorfer. A second-order rosenbrock method applied to photochemical dispersion problems. *SIAM J. Sci. Comput.*, 20(4):1456–1480, Feb. 1999.
- [93] G. Wahba. *Spline models for observational data*, volume 59. Society for Industrial Mathematics, 1990.
- [94] X. Wan and G. Karniadakis. Multi-element generalized polynomial chaos for arbitrary probability measures. *SIAM J. Sci. Comput.*, 28(3):901–928, 2006.
- [95] C. L. Winter and D. M. Tartakovsky. Groundwater flow in heterogeneous composite aquifers. *Water Resour. Res.*, 38(8), 2002. doi:10.1029/2001WR000450.

- [96] C. L. Winter, D. M. Tartakovsky, and A. Guadagnini. Moment differential equations for flow in highly heterogeneous porous media. *Surv. Geophys.*, 24(1):81–106, 2003.
- [97] B. Wohlberg, D. M. Tartakovsky, and M. Dentz. Linearized functional minimization for inverse modeling. In *XIX International Conference on Water Resources (CMWR)*, Urbana-Champaign, IL, USA, June 2012.
- [98] B. Wohlberg, D. M. Tartakovsky, and A. Guadagnini. Subsurface characterization with Support Vector Machines. *IEEE Trans. Geosci. Remote Sens.*, 44(1):47 – 57, 2006. doi:10.1109/TGRS.2005.859953.
- [99] A. D. Woodbury and T. J. Ulrych. Minimum relative entropy inversion: Theory and application to recovering the release history of a groundwater contaminant. *Water Resour. Res.*, 32(9):2671–2681, 1996.
- [100] D. Xiu. *Numerical Methods for Stochastic Computations: A Spectral Method Approach*. Princeton Univ. Press, Princeton, NJ, 2010.
- [101] M. Ye, S. P. Neuman, A. Guadagnini, and D. M. Tartakovsky. Nonlocal and localized analyses of conditional mean transient flow in bounded, randomly heterogeneous porous media. *Water Resources Research*, 40(5), 2004. doi:10.1029/2003WR002099.