

**Forward Modelling and Inversion of Time-Domain
Electromagnetic Geophysical Surveys in the Presence of
Chargeable Materials**

by

Patrick Belliveau

B.Sc., Simon Fraser University, 2010

M.Sc., Memorial University of Newfoundland, 2014

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL
STUDIES
(Geophysics)

The University of British Columbia
(Vancouver)

April 2019

© Patrick Belliveau, 2019

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation entitled:

Forward Modelling and Inversion of Time-Domain Electromagnetic Geophysical Surveys in the Presence of Chargeable Materials

submitted by Patrick Belliveau in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Geophysics

Examining Committee:

Eldad Haber, Earth, Ocean and Atmospheric Sciences

Supervisor

Douglas Oldenburg, Earth, Ocean and Atmospheric Sciences

Supervisory Committee Member

Roger Beckie, Earth, Ocean and Atmospheric Sciences

University Examiner

Chen Greif, Computer Science

University Examiner

Additional Supervisory Committee Members:

Christian Schoof, Earth, Ocean, and Atmospheric Sciences

Supervisory Committee Member

Abstract

The work described in this thesis examines transient geophysical electromagnetic forward modelling and inversion in the presence of induced polarization (IP) effects. The thesis introduces a new method of modelling IP using stretched exponential relaxation. A three-dimensional (3D) forward modelling algorithm taking full account of the coupling of IP and electromagnetic induction is developed. The stretched exponential modelling algorithm has been implemented using efficient numerical methods that allow it to tackle large-scale problems and are amenable to use in inversion. In particular, a parallel time-stepping technique has been developed that allows transient electric fields at multiple time steps to be computed simultaneously. The behavior of the stretched exponential model is demonstrated by applying it to synthetic numerical examples that simulate a grounded source IP survey with significant electromagnetic induction effects and a concentric-loop airborne electromagnetic sounding over a polarizable body.

An inversion algorithm using the stretched exponential model was developed that is able to recover the 3D structure of physical properties of the earth related to IP from transient geophysical electromagnetic data. The method is tested on a simple synthetic example problem. The thesis finishes with the development of a novel stochastic parametric level-set inversion algorithm, which could be useful in applying stretched exponential inversion to real world problems in the future. The algorithm addresses some of the shortcomings of the simple inversion approach used for stretched exponential inversion earlier in the thesis. The stochastic parametric inversion algorithm is used to solve shape reconstruction inverse problems in which the object of interest is embedded in a heterogeneous background medium that is known only approximately. Shape reconstruction is posed as a

stochastic programming problem, in which the background medium is treated as a random field with a known probability distribution. It is demonstrated that by using accelerated stochastic gradient descent the method can be applied to large-scale problems. The capabilities of the method are demonstrated on a simple 2D model problem and in a more demanding application to a 3D inverse conductivity problem in geophysical imaging.

Lay Summary

This thesis describes research in the field of electromagnetic geophysical imaging. The work led to the development of novel algorithms and software to facilitate three-dimensional imaging of the electromagnetic properties of the earth's subsurface. Electromagnetic imaging results are useful to geologists and other geoscience professionals in their attempts to assess subsurface geology for applications such as natural resource exploration and environmental assessment. The particular focus of this thesis was on improving the quality and efficiency of imaging algorithms that consider an electromagnetic phenomenon known as induced polarization. Induced polarization effects can be important indicators of geological features of interest to practitioners but they are difficult to model and standard methods either ignore them or treat them in some simplified form. This thesis has developed the computational methods necessary to fully consider the effects of induced polarization in electromagnetic imaging.

Preface

This thesis is my original work. It has resulted in one peer reviewed journal publication, one peer reviewed journal submission currently under review, one expanded abstract publication, and seven conference presentations.

The parallel time-stepping algorithm in chapter 2 originated in discussions with Dr. Haber and is based on his original idea. I developed the algorithm and implemented it. It was published in [11]. Chapter 2 provides more detailed analysis and more extensive background material than was included in the published paper.

The stretched exponential forward modelling algorithm in chapter 3 also originated in discussions with Dr. Haber and is based on his original idea. I developed, implemented and tested the algorithm based on the discussions with Dr. Haber. A version of chapter 3 was published as [11]. Chapter 3 contains a more extensive comparison of the stretched exponential model of induced polarization with the Cole-Cole model than was included in the published paper.

The stretched exponential inversion algorithm and workflow described in chapter 4 is my original work. It was carried out independently but benefited from discussions with Dr. Haber and Roman Shekhtman.

The stochastic inversion algorithm described in chapter 5 originated in discussions with Dr. Haber based on my original idea. I implemented and analyzed the algorithms. A version of this chapter has been submitted for publication and is currently under review.

The work in this thesis entailed significant software development. All development was done within the open source jInv software framework, which was created in Dr. Haber's research group at UBC and is available online at <https://github.com/juliaInv>. Four of the conference presentations I gave over the course

of my Ph.D. focused on the jInv software package.

Table of Contents

Abstract	iii
Lay Summary	v
Preface	vi
Table of Contents	viii
List of Tables	xi
List of Figures	xiii
Glossary	xx
Acknowledgments	xxii
Dedication	xxiii
1 Introduction	1
1.1 Research Motivation	1
1.2 Geophysical background	3
1.2.1 Maxwell's equations	5
1.2.2 IP models	7
1.3 Thesis outline	10
2 Efficient modelling of non-dispersive TEM data	12
2.1 Discretization of the time-domain quasi-static Maxwell equations	13

2.1.1	Spatial discretization	14
2.1.2	Discretization in time	17
2.2	Parallel time-stepping	26
2.2.1	Evaluating accuracy	28
2.2.2	Evaluating performance	30
2.3	Conclusions	35
3	Forward modelling of the dispersive time-domain Maxwell Equations using the stretched exponential function	36
3.1	The stretched exponential formulation	36
3.1.1	Comparison of stretched exponential and Cole-Cole models	41
3.2	Discretization	45
3.2.1	Backward Euler Algorithm	45
3.2.2	Solving the linear system	47
3.2.3	BDF2	52
3.2.4	Implementation	53
3.3	Synthetic modelling examples	54
3.3.1	Grounded source example	54
3.3.2	Inductive source example	58
3.4	Conclusion	61
4	Stretched exponential inversion	63
4.1	The gradient array example	64
4.2	Inversion methodology	65
4.3	Inversion results	67
4.4	Conclusions	73
5	Randomized background stochastic inversion	75
5.1	Introduction	75
5.2	Method	77
5.2.1	Computing Background Samples	81
5.2.2	Implementation	82
5.3	Application to a Simple Model Problem	83
5.3.1	SAA inversion	86

5.3.2	Optimization by stochastic gradient descent	91
5.4	Application to non-linear 3D problem	93
5.5	Conclusions	99
6	Conclusions	101
	Bibliography	105

List of Tables

Table 1.2	Maxwell's equations	6
Table 2.1	Parallel time-stepping scheme. Six parallel forward modelling simulations with constant step-sizes shown in this table were run in parallel to simulate the receiver voltages at the same observation times used in the reference backward Euler simulation. The number of observation times per parallel process varied but the total number of time-steps was the same in all the simulations.	32
Table 2.2	Backward Euler solution run times.	33
Table 2.3	Parallel time-stepping run times using 6 parallel time-stepping processes run across 4 compute nodes of a computer cluster. . .	34
Table 5.1	Starting guesses and parameter bounds for straight ray tomography inversion parameters. ϕ is the ellipse orientation angle in $^\circ$, x and y are the horizontal and vertical coordinates of the centre position, and s the slowness inside the ellipse.	87
Table 5.2	Marginal statistics of SAA inversion results. The sample means of the recovered values of each parameter are shown for the 1, 10, and 100 background inversion experiments in a), with the corresponding sample standard deviations shown in b). Recall that the true parameter values are $\phi = 30^\circ$, $x = 0.5$, $y = 0.5$, and $s = 4$. See table 5.1 for parameter bounds and initial guesses. . .	89

Table 5.3 Marginal statistics of SAA inversion results using the ADAM optimization algorithm. Recall that the true parameter values are $\phi = 30^\circ$, $x = 0.5$, $y = 0.5$, and $s = 4$. See table 5.1 for parameter bounds and initial guesses. 93

List of Figures

Figure 1.1	Illustration of typical grounded and inductive source surveys. Left image shows current electrodes connected to ground and energy source, and receiver measuring the potential difference between separate measurement electrodes. Right image shows a concentric loop airborne electromagnetic (EM) system, where a current run through the outer loop creates a time-varying magnetic field that induces a response from the earth that is measured by the inner coil.	4
Figure 1.2	DC experimental setup. From Tomoquest.com	5
Figure 2.1	Yee grid. Components of the electric field are discretized on mesh edges in corresponding Cartesian coordinate directions. Magnetic fields are discretized on face centres. Figure adapted from [51].	16
Figure 2.2	Typical example of realistic time-domain electromagnetic (TDEM) transmitter current time-dependence plotted in black, with dots showing times at which solution is computed. The solid blue line shows the variation in the step-size over the full time range of the simulation. There are six unique step-sizes.	20
Figure 2.3	Transmitter current step-off waveform. Assumes transmitter current has been on long enough that fields are steady-state at $t = 0$, where forward modelling begins. Shutoff takes place in one time-step.	22

Figure 2.4	Comparison of time-stepping methods in concentric loop example. In the legend, BDF2 1 step init refers to the solution computed using backward Euler for the first step only, using second order backward differentiation formula (BDF2) thereafter with no special treatment of the discontinuity. BDF2 3 step init refers to the solution computed using backward Euler time-stepping for the first three steps and BDF2 stepping thereafter. Dashed lines indicate negative values. a) shows the overall behaviour of each method over the full 200 time-step simulation. c) illustrates the large oscillation observed from the second time-step when using BDF2 with one step initialization.	25
Figure 2.5	Parallel time-stepping accuracy on grounded source test problem.	29
Figure 2.6	Time-stepping scheme for reference backward Euler solution . There are 10 steps each of sizes 2×10^{-5} s, 4×10^{-5} s, 8×10^{-5} s, 1.6×10^{-4} s, and 14 steps of size 3.2×10^{-4} s.	31
Figure 2.7	Pardiso parallel scaling results when using a single sequential backward Euler time-stepping process. Parallel efficiency for N threads is $T_1/(N \cdot T_N)$ where T_1 is the sequential run time and T_N is the run time using N threads.	34
Figure 3.1	a) Stretched exponential impulse response for various values of β . b) Cole-Cole impulse responses for various values of c . For both plots $\eta = 0.1$, $\tau = 0.1$ and $\sigma_\infty = 0.1$	42
Figure 3.2	Cole-Cole impulse response with $c = 0.5$, $\tau = 0.1$, $\eta = 0.1$, $\sigma_\infty = 0.1$ plotted in blue alongside stretched exponential impulse responses with parameters chosen by non-linear least squares fitting. The green line attempted to fit the $0.001 < t < 100$ s time range and the magenta line the $0.1 < t < 100$ s range.	44
Figure 3.3	Error in PCG computed example problem electric fields for loose and tight PCG relative residual stopping tolerances. . . .	51

Figure 3.4	Gradient array example survey layout. a) Cartoon showing transmitter, receiver array bounding box, and sample cartoon receiver. b) Plan view layout of transmitter, receiver array and two block synthetic model.	55
Figure 3.5	Visualization of OcTree mesh discretization of the conductivity model.	56
Figure 3.6	DC inline electric field at the earth's surface for two block example model. The locations of the buried blocks are shown by black squares.	57
Figure 3.7	In-line component of surface electric field for gradient array example. a) Plan-view in-line field at 0.01 s after current shut-off. Left and right black squares show the boundaries of the chargeable and conductive blocks, respectively. b) Plan view field at 0.14 s. c) Electric field decays above the centre of the chargeable block, with stretched exponential field in blue and corresponding non-dispersive field ($\eta = 0$) in red. Dashed lines indicate negative values. d) Electric fields above the centre of the conductive block, showing that the stretched exponential and non-dispersive responses are almost identical.	58
Figure 3.8	Section view of block in a half-space with chargeable overburden model. The overburden depth is 10 m and the block extends from 40-90 m depth with a plan view cross-section of 100×100 m. The background halfspace conductivity was 0.005 S/m and the conductivity of the block was 0.1 S/m. The overburden had $\sigma_{\infty} = 0.01$ S/m $\eta = 0.3$, $\tau = 1 \times 10^{-3}$ s and $\beta = 0.8$	59

Figure 3.9	Effect of chargeable overburden on ATEM response from buried conductor. All plots show the vertical component of $d\mathbf{b}/dt$. The background halfspace response is shown in green in each plot. Dashed lines indicate negative values. a) Block with no overburden. The response with the conductive block at a depth of 40 m is shown in blue. The magenta line shows the response obtained when the top of the block is moved to a depth of 70 m. b) Stretched exponential response with 10 m of overburden ($\sigma_\infty = 0.005$, $\eta = 0.3$, $\tau = 1 \times 10^{-3}$, $\beta = 0.8$) is shown in black, along with the overburden free and background responses. c) Stretched exponential response with the top of the block at 70 m in black, along with the overburden free and background responses. d) Stretched exponential response with chargeable overburden and no conductive block in black, along with background EM response.	60
Figure 4.1	Gradient array example survey layout. a) Cartoon showing transmitter, receiver array bounding box, and sample cartoon receiver. b) Plan view layout of transmitter, receiver array and two block synthetic model.	64
Figure 4.2	Selected slices of true conductivity model. Left hand block had $\sigma_\infty = 0.018$ S/m and right block $\sigma_\infty = 0.1$ S/m. Left block had $\eta = 0.25$, $\beta = 0.5$ and $\tau = 0.5$ s. Right block was not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.	65
Figure 4.3	Selected slices of recovered conductivity model from DC and early off-time inversion. Thick black lines show true block boundaries. True conductivities are 0.01 S/m for the background, 0.018 S/m for the left block and 0.1 S/m for the right block. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.	68
Figure 4.4	Misfit convergence curve for σ_∞ inversion.	69

Figure 4.5	Selected slices of recovered chargeability model from DC and mid to late off-time data inversion. Thick black lines show true block boundaries. Left hand block has $\eta = 0.25$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.	69
Figure 4.6	Selected slices of recovered time constant model from late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\tau = 0.5$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.	70
Figure 4.7	Selected slices of recovered exponent model from late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\beta = 0.5$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.	70
Figure 4.8	Selected slices of recovered chargeability model from stage 4 mid to late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\eta = 0.25$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.	71
Figure 4.9	Selected slices of recovered time constant model from stage 4 mid to late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\tau = 0.5$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.	72
Figure 4.10	Selected slices of recovered exponent model from stage 4 mid to late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\beta = 0.5$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.	72
Figure 4.11	Voltage data for the x-directed receivers due to the x-directed transmitter at 0.02 s after source shutoff.	73
Figure 4.12	Voltage data for the x-directed receivers due to the x-directed transmitter at 0.02 s after source shutoff.	73

Figure 4.13	true and predicted voltage data above the centre of the chargeable block from 0.02 s after source shutoff onward plotted on log scale. Blue line shows observed data and red line shows predicted data. Dashed lines indicate negative values.	74
Figure 5.1	Cartoon schematic of straight ray tomography experiment showing straight ray paths from the transmitters to the receivers. . .	83
Figure 5.2	True model used for straight ray tomography inversions. The mean model a) consists of homogeneous horizontal layers. The random perturbation shown in b) is added to the mean model to form c), used as the background in the true model shown in d).	84
Figure 5.3	Results of 100 independent single background inversions. The recovered values for each inversion parameter are plotted as histograms with accompanying kernel density estimates. The red vertical lines show the true parameter values. Note that the estimated densities may have support outside the parameter bounds. They are included in order to give general visual impressions of the spread of inversion results and these effects are not of concern.	88
Figure 5.4	Kernel density plots showing inversion results for single background, 10 background SAA, and 100 background SAA inversions. 100 background angle results are not shown because they took the upper bound value of 90° for all but two of these inversions rendering the corresponding kernel density plot meaningless. The vertical axes of the plots were clipped in order to show the single and 10 background results at reasonable scales.	90
Figure 5.5	Kernel density plots of inversion results for 100 background and 500 background ADAM inversions with mini-batch size 1. Horizontal axis limits are the same as in fig. 5.4	94
Figure 5.6	Mean background conductivity in the x-z plane at $y = 0$ in the core region of the mesh.	97
Figure 5.7	True model slices	98

Figure 5.8 Results from single background and SAA ADAM inversions of DC data. Blue dots show the recovered values of each parameter from the 20 single background inversions and the stars show the parameter values recovered by the SAA inversion. In each plot, the vertical axis limits are the parameter bounds, the dashed black line shows the initial guess and the red line the true value of the parameter. 100

Glossary

2D two-dimensional

3D three-dimensional

ATEM airborne time-domain electromagnetic

BDF backward differentiation formula

BDF2 second order backward differentiation formula

DC direct current

DCIP direct current induced polarization, an IP experiment conducted alongside a DC survey. EM induction effects are not considered.

EM electromagnetic

EMIP coupled electromagnetic induction and induced polarization

FLCBDF2 fixed leading coefficient second order backward differentiation formula

FDTD finite difference time domain

IP induced polarization

ODE ordinary differential equation

PDE partial differential equation

RMS root mean squared

TDEM time-domain electromagnetic

Acknowledgments

First and foremost I would like to acknowledge the patient and thoughtful guidance of my supervisor, Eldad Haber. He helped shape my research direction while always encouraging me to follow my curiosity and believe in my abilities. I was also very lucky to have the chance to learn from Doug Oldenburg, who taught me much about electromagnetics and about teaching.

Seogi Kang and Dave Marchant were invaluable sources of sound advice on IP and EM as well as being warm, friendly, and collegial folks to chat with around the office. Christoph Schwarzbach was also an invaluable source of insight on EM. These folks, along with fellow office mates and geophysics students Sanna Tyrväinen, Xiaojin Tan, Michelle Liu, Luz-Angelica Caudillo Mata, Wenke Wilhelms, Devin Cowan, Daniel Bild-Enkin, Mike MacMillan, Mike Mitchell, Thibaut Astic, Dom Fournier, Lindsey Heagy, and Klara Steklova, made UBC a collegial and downright fun place to work.

I would be remiss not to mention the two pre-UBC mentors in academia that introduced me to research, gave me much guidance and shared their love of learning: my undergraduate research supervisor Gwenn Flowers and my M.Sc. supervisor Colin Farquharson.

Most importantly, I could not have made it through my Ph.D. without the constant love and support of my parents, my sister Rachael, and my partner Caroline.

Dedication

To my family

Chapter 1

Introduction

1.1 Research Motivation

At a broad level, this thesis is concerned with numerical modelling of transient electromagnetic (EM) geophysical surveys and associated parameter estimation problems. The particular emphasis is on understanding how to do this efficiently when the surveys are affected by the phenomenon of induced polarization (IP). Geophysical methods designed to detect IP have a long history in mineral exploration. They have been used in many exploration contexts but are particularly important in the search for disseminated mineralization. Some important types of disseminated mineralization exhibit strong IP anomalies but are not sensitive to other geophysical methods—see e.g. Ward 96, Zonge et al. 106. IP is the most important geophysical method in exploration for porphyry copper deposits, for example [106]. It is also increasingly important in environmental applications [54].

IP is a dynamic EM phenomenon. However, in traditional processing and interpretation of geophysical surveys designed to detect it, it is often treated as a static electrical phenomenon—e.g. as in [67]. Such treatments neglect two key factors that affect IP data. They ignore the nature of the time or frequency dependence of IP relaxations and the coupling between IP and EM induction. A full treatment of the problem requires the ability to solve the quasi-static Maxwell equations—in either the frequency or time-domain—using a constitutive law that models the EM response of polarizable materials.

Traditional IP surveys excite the earth using energy sources connected to the earth by grounded electrodes. IP effects can also be generated by closed loop systems that excite the earth by EM induction. The increasingly common observation of IP effects in airborne transient inductive source EM data has generated renewed interest in understanding and modelling IP in the context of transient electromagnetics. In this thesis I develop tools that are general and efficient for modelling grounded and inductive source transient electromagnetic surveys in the presence of polarizable media.

Simulation tools that can accurately model realistic three-dimensional (3D) geometries are very useful both for developing physical understanding and in designing surveys. Modelling the coupling of EM and IP effects can help a geophysicist understand in detail how EM coupling might affect a grounded source IP survey and allow them to use inductive responses as useful signal rather than just noise. Some work has been done on modelling the quasi-static time-domain Maxwell equations in the presence of polarizable media (e.g. [25, 62]) but it remains an under-researched area.

The first two chapters of the thesis are focused on forward modelling and the last two on geophysical inversion. Inversion, the process of estimating a subsurface physical property model from data, is a key step in geophysical data processing/interpretation. We want to be able to invert time-domain EM data for both conductivity and IP parameters in 3D, whether the data comes from traditional grounded source surveys designed to detect IP or from inductive source surveys. After developing an algorithm for simulating IP effects in transient electromagnetic data, described in chapter 3 of this thesis, I describe the corresponding inversion algorithm in chapter 4.

In addition to the need for appropriate forward modelling tools, a major difficulty in inverting EM data for IP parameters is underdeterminedness in the inverse problem. Not having enough data to uniquely define a model of the subsurface is standard in geophysical inversion but the problem is exacerbated in the case of multi-parameter problems like IP inversion where it may be difficult to distinguish whether some features in the data are due to IP or to variations in conductivity. Also, in multi-parameter inverse problems, a geophysicist is often only interested in a subset of the parameters but is forced to estimate them all in order to create a

model that matches observed geophysical data. I address these issues in the final chapter of this thesis by developing a stochastic inversion methodology that improves the precision of estimates of parameters of interest due to uncertainty in the nuisance parameter estimates.

1.2 Geophysical background

This section will place IP in the wider context of electromagnetic geophysics, describe some general principles of EM surveying, and how IP fits into the picture by way of constitutive relations. By electromagnetic geophysics, here we mean low frequency (or long time) methods that are sensitive primarily to the electrical conductivity of the subsurface. This excludes high frequency techniques such as ground penetrating radar which are, of course, electromagnetic in nature but fundamentally different from the range of low-frequency methods normally encompassed by the term electromagnetic geophysics.

These methods are part of the field of applied geophysics, whose major goal is the mapping of variations in the material properties of the earth's subsurface using remote measurements. Such information can be useful in characterizing subsurface geology or locating buried objects without labour intensive and expensive direct methods such as drilling [92].

The family of geophysical EM methods may be divided into natural source methods, which sense the earth response to natural variations in atmospheric EM fields, and controlled source methods that measure the response due to artificial excitations of the earth controlled by a geophysicist. This thesis will focus exclusively on controlled source methods. These may be further divided into grounded and inductive source methods. In grounded source surveys the energy source is directly connected to the earth by means of electrodes, allowing electric currents to flow from the source into the earth. By contrast, inductive source surveys use closed coils of wire to create time-varying magnetic fields that excite responses in the earth through EM induction. Typical source and receiver arrangements for these survey types are shown in fig. 1.1.

These methods have been designed primarily to sense the electrical conductivity of the subsurface but they can all be influenced by the IP phenomenon. When




Figure 1.1: Illustration of typical grounded and inductive source surveys. Left image shows current electrodes connected to ground and energy source, and receiver measuring the potential difference between separate measurement electrodes. Right image shows a concentric loop airborne EM system, where a current run through the outer loop creates a time-varying magnetic field that induces a response from the earth that is measured by the inner coil.

an electric field is applied to a conductive material, ions will move freely through it. This is the flow of electric current. When the movement is impeded, charges will build up, whether that be at the interface between two homogeneous regions of differing conductivities or due to more complex microscopic effects. This creates electrical polarization in the subsurface. Normally, when the applied field is removed, the charges will return to an electrically neutral state extremely quickly—far too quickly to be detected in an ordinary geophysical experiment. However, in some geological materials large polarization effects are observed that build up slowly under an applied current and then relax slowly enough to be measured in a geophysical EM survey. This is the IP phenomenon. The effect has complex origins in the microscopic properties of the rocks. At the macroscopic level polarization acts to impede current flow due to an externally applied field. Then, when the applied field is removed, the relaxation of the polarized charges to equilibrium will result in a small but measurable current flow. This effect is modelled as reduced steady-state conductivity in the static case and as frequency dependent conductivity for harmonic sources. In the time-domain the current flow at a given time is modelled as depending on the history of the electric field driving the current flow and not just its instantaneous value.

In the canonical IP experiment, a direct current is driven into the ground using a pair of electrodes connected to an energy source. The current remains switched on

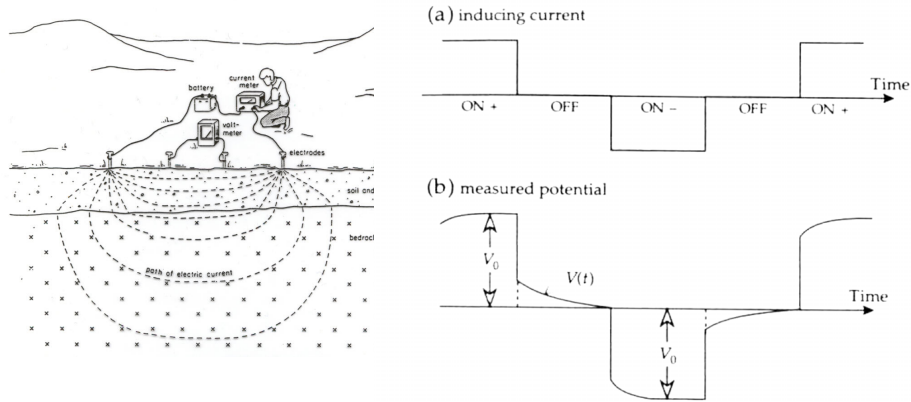


Figure 1.2: DC experimental setup. From Tomoquest.com

until the polarization has reached a steady state, at which point potential differences between electrodes at various locations on the earth's surface are measured. Then the current is abruptly switched off. The potential differences resulting from the current created by the polarized charges returning to equilibrium is then measured. This is illustrated in fig. 1.2

1.2.1 Maxwell's equations

To describe this process quantitatively, we turn to Maxwell's equations. Macroscopic EM fields are governed by the macroscopic Maxwell equations, the fundamental laws of classical electromagnetism. These may be stated as a set of partial differential equations (PDEs) describing the behaviour of EM fields, augmented by a set of empirical rules called constitutive relations that define how EM fields are modified by materials. Maxwell's equations are collected in table 1.2. The equations may be expressed in the time-domain, where the fields are functions of space and time. What are known as the frequency-domain equations are obtained by taking the Fourier transforms with respect to time of the time-domain equations.

The macroscopic Maxwell equations relate several vector fields: the electric field \mathbf{e} , the magnetic induction \mathbf{b} , the magnetic field \mathbf{h} , the electric displacement \mathbf{d} and the electric current \mathbf{j} . There are three constitutive laws. The first relates the

Name	Time	Frequency
Faraday's law	$\nabla \times \mathbf{e} = -\frac{\partial \mathbf{b}}{\partial t}$ (1.1.1)	$\nabla \times \mathbf{E} = -i\omega \mathbf{B}$ (1.1.2)
Ampère-Maxwell	$\nabla \times \mathbf{h} - \mathbf{j} = \frac{\partial \mathbf{d}}{\partial t}$ (1.1.3)	$\nabla \times \mathbf{h} - \mathbf{j} = i\omega \mathbf{D}$ (1.1.4)
Gauss's law	$\nabla \cdot \mathbf{d} = \rho_f$ (1.1.5)	$\nabla \cdot \mathbf{D} = \rho_f$ (1.1.6)
Unnamed	$\nabla \cdot \mathbf{b} = 0$ (1.1.7)	$\nabla \cdot \mathbf{B} = 0$ (1.1.8)

Table 1.2: Maxwell's equations

electric field to the electric displacement, the second, the magnetic induction to the magnetic field, and the third the electric field to the electric current. In general these may be non-linear functions

$$\mathbf{d} = \mathbf{d}(\mathbf{e}) \quad (1.1)$$

$$\mathbf{h} = \mathbf{h}(\mathbf{b}) \quad (1.2)$$

$$\mathbf{j} = \mathbf{j}(\mathbf{e}) \quad (1.3)$$

but the EM properties of most earth materials may be described by linear constitutive relations. These are typically written

$$\mathbf{d} = \boldsymbol{\varepsilon} \mathbf{e} \quad (1.4)$$

$$\mathbf{h} = \boldsymbol{\mu}^{-1} \mathbf{b} \quad (1.5)$$

$$\mathbf{j} = \boldsymbol{\sigma} \mathbf{e}, \quad (1.6)$$

where $\boldsymbol{\varepsilon}$ is called the electrical permittivity, $\boldsymbol{\mu}$ the magnetic permeability, and $\boldsymbol{\sigma}$ the electrical conductivity. The constants of proportionality may be scalars or symmetric rank-2 tensors but depend only on position, not on time or frequency.

In this thesis I will consider Maxwell's equations in the quasi-static approximation, in which the electric displacement term in Ampère-Maxwell equation is

ignored. This is a standard assumption in geophysical EM [97]. Most materials have ϵ on the order of the permittivity of free space, $\epsilon_0 = 8.85 \times 10^{-12}$ F/m. The least conductive common earth materials have conductivities on the order of 10^{-6} S/m. Larger values in the range 10^{-3} -1 S/m are most often encountered and some important geological materials such as nickel sulphides can have conductivities as high as 10^5 S/m [70]. Frequency-domain geophysical experiments typically use frequencies under 10^5 Hz and time-domain experiments operate on time-scales longer than 10^{-6} s. Consider the Ampère-Maxwell equation in frequency using linear constitutive relations

$$\nabla \times \mu^{-1} \mathbf{b} - (\sigma + i\omega\epsilon)\mathbf{e} = 0. \quad (1.7)$$

It is clear from the above paragraph that in typical geophysical parameter regimes $\sigma \gg \omega\epsilon$, making it safe to ignore the electric displacement.

1.2.2 IP models

The constitutive law that really concerns us here is that between electric field and electric current, which we use to model the IP effect. For common materials that don't exhibit IP, Ohm's law is typically used to describe how currents flow in response to applied electric fields

$$\mathbf{j} = \sigma \mathbf{e} \quad (1.8)$$

where the electrical conductivity σ may be a scalar or symmetric rank-2 tensor that varies in space but not in time or frequency. When IP effects are present, the flow of electric current will depend on the frequency of the applied electric field. These effects can be modelled by replacing the standard Ohm's law with a dispersive constitutive relation.

Early IP modelling treated the phenomenon statically, as a perturbation of the DC resistivity problem. This was characterized mathematically by Seigel [85]. He described the polarization of a subsurface body as a perturbation of the conductivity, writing Ohm's law as

$$\mathbf{j} = \sigma(1 - \eta)\mathbf{e} \quad (1.9)$$

where η is a unit-less number on the interval $[0, 1)$ called chargeability. The chargeability quantifies the magnitude of the polarization effect, given sufficient time for charge accumulation. This provided a powerful framework for the interpretation of IP data from the standard time-domain IP experiment described above. However, Seigel himself noted that chargeability says nothing about how polarization potentials decay in time. In the 1970s, several authors showed how mineral types with similar conductivities and chargeabilities could be distinguished based on the spectral characteristics of their IP responses—see e.g. [71, 94, 107]. Pelton et al. [71] described this spectral IP effect using an Ohm’s law with a frequency dependent complex conductivity, where the frequency dependence is given by the phenomenological Cole-Cole model

$$\sigma(\omega) = \sigma_{\infty} \left(1 - \frac{\eta}{1 + (1 - \eta)(i\omega\tau)^c} \right) \quad (1.10)$$

where ω is angular frequency, η is Seigel’s chargeability, τ is a time constant that governs the overall time-scale of the relaxation, the exponent c controls the shape of the spectrum, and σ_{∞} is the conductivity at infinite frequency or in the absence of IP effects. This phenomenological model was introduced by Cole and Cole [22] to describe dispersive dielectric permittivity. Pelton et al. [71] had considerable success in discriminating between minerals based on their Cole-Cole parameters in laboratory experiments. The Cole-Cole model has become the de facto standard model for characterizing the IP spectrum of geological materials. Zhdanov [105] proposed a more general model of induced polarization based on effective medium theory, known as GEMTIP, that includes the Cole-Cole model as a special case but it does not seem to be widely used. This may be because the Cole-Cole model has been seen as sufficient to fit IP data seen in the field or because it has not improved the ability to correlate IP with geology. Given one of these models, or any other frequency dependent conductivity model, frequency domain IP effects can be modelled using standard methods for solving the frequency domain Maxwell equations designed without IP in mind. There are several examples of this approach in the geophysical literature, including [24, 100, 102].

It should be noted that in the analysis of data collected in frequency domain IP surveys, it is not standard to model inductive effects by solving the full quasi-

static Maxwell equations. Typical experiments have traditionally been designed to reduce or eliminate the effects of EM induction. What induction did occur was considered noise and had to be removed from field data using imperfect pre-processing procedures. Classical examples of this approach include [71] and [79].

Grounded source frequency domain surveys designed to detect IP effects are still in use, especially for near-surface environmental work (see e.g. [54, 76]) but time-domain surveys remain the more common way to collect data in the mineral exploration setting. Being able to extract information on the ground based on the shape and time-scale of the IP decay in a time-domain IP experiment, as well as being able to fully model EM coupling effects motivates us to study how to efficiently model IP effects in the time-domain.

A larger motivation for the EM geophysics community at large to study IP in the time domain (and to study coupled IP and EM induction effects) is the increasing awareness of IP effects in airborne time-domain electromagnetic (ATEM) data. The well known phenomenon of negative transients in concentric loop ATEM data (see e.g. [56, 99]) cannot be modelled using non-dispersive ATEM modelling algorithms. More subtly, IP signals may mask conductive features at depth without the smoking gun of negative transients. Some groups (e.g. [56, 87]) have focused on separating induction and IP effects in ATEM data through data processing rather than capturing both phenomena in earth model based simulations. Uniquely, Kang and Oldenburg [53] have taken a compromise approach between data driven decoupling and full time domain EM-IP simulation. They simulated IP affected ATEM data as a linear perturbation of non-dispersive EM data, then used those perturbations as data to invert for chargeability.

Unfortunately, modelling IP in the time domain is more complex than in frequency. The simple notion of Ohm's law with a frequency dependent conductivity is lost and the constitutive law becomes a convolution. Most work in modelling time-domain dispersive EM in time has attempted to avoid treating the convolutional Ohm's law directly by performing the modelling in frequency domain at many frequencies and then transforming the results to the time-domain. This was done in the suite of software released through AMIRA project P223F [100]. More recently, Fiandaca et al. [32] used this approach in a scheme to invert ATEM data for Cole-Cole parameters. Their inversion algorithm was computationally

tractable because they used one-dimensional forward modelling. The approach of modelling in frequency then transforming to time becomes very computationally expensive for 3D problems. Zaslavsky and Druskin [104] were able to more efficiently compute time-domain responses from frequency domain information using a Krylov subspace reduction technique.

By contrast, there have been very few methods described that model time-domain IP by direct time-stepping. Recently, Commer et al. [25] modelled dispersive TEM fields using an explicit time-stepping finite difference time-domain (FDTD) algorithm, approximating the Cole-Cole convolutional Ohm's law using a spectrum of Debye decays. Marchant [62] used an implicit time-stepping approach, directly treating the convolutional Ohm's law using numerical quadrature, with the convolution kernel transformed to the time-domain by the digital filtering technique. The only other approach we are aware of in the literature was developed by Marchant et al. [63]. In this approach the transformation of Ohm's law to time is handled using rational function approximations, which allow it to be described by an ordinary differential equation in time. This approach was recently extended by Cai et al. [20] to make it more computationally efficient. In chapter 3 I will discuss the details of time domain EM and IP modelling in more detail and describe our new approach to the problem.

1.3 Thesis outline

The first part of this thesis (chapters 2 and 3) considers forward simulation of transient EM survey data, while the second part (chapters 4 and 5) discusses the inversion of such data. Chapter 2 focuses on the development of efficient time discretization methods for the quasi-static Maxwell equations. IP is ignored in this chapter. Modelling IP effects adds additional computational complexity to the standard EM forward simulation problem. I develop efficient methods for the non-dispersive case before expanding to consider IP effects in chapter 3. The main result is a novel parallel time-stepping algorithm. I give a brief overview of the spatial discretization approach, followed by a discussion of discretization in time and the development of the parallel time-stepping algorithm.

Chapter 3 presents a novel approach to modelling IP effects in time-domain

electromagnetic data developed by my supervisor Dr. Haber and myself, based on stretched exponential relaxation. Stretched exponential relaxation has been used to model relaxations in disordered systems in other fields such as the study of glassy systems in solid-state physics. It has been suggested by Everett [31] and by Haber [37] as a way to model IP effects but has not been studied in detail in the geophysical community. I compare the stretched exponential model of IP to the widely used Cole-Cole model and develop an efficient forward modelling algorithm using the stretched exponential approach. The capabilities of the algorithm are demonstrated using synthetic examples based on typical scenarios in grounded source IP surveying and airborne EM surveying.

Chapter 4 of the thesis describes preliminary investigations into inverting transient EM data to recover the parameters that describe polarizable materials in the stretched exponential model. This is a difficult, multi-parameter, highly under-determined inverse problem. The work in chapter 4 presents a proof of concept demonstration that recovery of stretched exponential IP related parameters from transient EM data is possible but further work is needed in order to develop robust inversion algorithms and workflows suitable for use by practising geophysicists.

Chapter 5 begins to address these issues. One approach to addressing the non-uniqueness in the stretched exponential inverse problem is to use parametric inversion. In a parametric inversion one assumes that the target of the inversion is a homogeneous body embedded in a heterogeneous background medium and attempts to recover its shape, position and physical properties—rather than trying to recover physical property values for each cell in a computational mesh. In chapter 5 I develop a stochastic parametric inversion algorithm that accounts for uncertainty in estimation of the heterogeneous background medium. I test the algorithm on a simple linear two-dimensional (2D) inverse problem and on a 3D direct current (DC) resistivity inverse problem. Applying this approach to stretched exponential inversion is beyond the scope of this thesis but presents an excellent opportunity for future work.

Chapter 2

Efficient modelling of non-dispersive TEM data

Modelling the quasi-static time-domain Maxwell equations is a computationally expensive problem, even in the non-dispersive case. Adding induced polarization (IP) effects further increases the complexity. Therefore efficient numerical methods are essential to the successful application of coupled modelling of electromagnetic (EM) and IP effects to large scale real world problems. A significant amount of research has investigated spatial discretization of Maxwell's equations over large domains and at high spatial resolutions—e.g domain decomposition ([29, 75]) and multigrid methods ([38]). Advancing solutions in time, however, remains a bottleneck. This chapter addresses that bottleneck by developing a novel parallel in time forward modelling algorithm for the quasi-static Maxwell equations. I consider only the non-dispersive case in this chapter before considering IP effects in chapter 3.

I start this chapter by describing the formulation of Maxwell's equations I will solve and giving a brief overview of how the equations are discretized in space. I use established methods for spatial discretization. I then discuss some general considerations for time-discretization of Maxwell's equations before describing the parallel time-stepping algorithm for geophysical time-domain electromagnetic (TDEM) simulations.

2.1 Discretization of the time-domain quasi-static Maxwell equations

This section will describe how I've discretized the quasi-static Maxwell equations to simulate time-domain electromagnetic geophysical surveys over non-dispersive earth models. These surveys are governed by Faraday's law and Ampère's law supplemented by appropriate initial and boundary conditions. Recall Faraday's law and Ampère's law

$$\nabla \times \mathbf{e} = -\frac{\partial \mathbf{b}}{\partial t} \quad (2.1)$$

$$\nabla \times \mu^{-1} \mathbf{b} - \mathbf{j} = \mathbf{j}^s, \quad (2.2)$$

where \mathbf{e} is the electric field, \mathbf{b} the magnetic flux density, \mathbf{j} the electric current in the earth and \mathbf{j}^s the source current. In this chapter I will neglect IP effects and assume that currents in the earth are governed by Ohm's law,

$$\mathbf{j} = \sigma \mathbf{e}, \quad (2.3)$$

where σ is the isotropic conductivity, a scalar function of space. Substituting Ohm's law into Ampère's law eliminates explicit dependence on \mathbf{j} . I eliminate explicit dependence on \mathbf{b} by combining eq. (2.1) and eq. (2.2) to give the single vector-valued parabolic partial differential equation (PDE)

$$\nabla \times (\mu^{-1} \nabla \times \mathbf{e}) + \sigma \frac{\partial \mathbf{e}}{\partial t} = -\frac{\partial \mathbf{j}^s}{\partial t}. \quad (2.4)$$

This electric field equation is the preferred formulation of Maxwell's equations for modelling grounded source surveys as it provides the most natural setting for discretizing the sources and computing observable quantities from the electric fields.

I consider this equation on a rectangular spatial domain $\Omega \in \mathbb{R}^3$, on the closed time interval $[0, t^*]$. I assume steady-state initial conditions. For grounded source surveys the initial electric field is computed by solving the steady-state problem

$$\nabla \cdot \sigma \nabla \phi = \nabla \cdot \mathbf{j}^s, \quad (2.5)$$

where $\mathbf{e}(t = 0) = -\nabla\phi$. The initial electric field is zero in inductive source problems. I impose natural boundary conditions on the electric field

$$\hat{\mathbf{n}} \times (\mu^{-1} \nabla \times \mathbf{e}) = \mathbf{0}, \quad (2.6)$$

where $\hat{\mathbf{n}}$ is the unit normal vector to the boundary. The diffusive electric fields in quasi-static electromagnetics decay asymptotically toward zero away from energy sources. In practice the boundary conditions are implemented approximately by choosing a domain large enough such that the boundaries are far enough from any sources that 2.6 is approximately satisfied.

Collecting the information from the last paragraphs, I can now state the initial boundary value problem that is to be solved

$$\nabla \times (\mu^{-1} \nabla \times \mathbf{e}) + \sigma \frac{\partial \mathbf{e}}{\partial t} = -\frac{\partial \mathbf{j}_s}{\partial t} \quad \text{on} \quad \Omega \times (0, t^*] \quad (2.7a)$$

$$\nabla \cdot \sigma \nabla \phi = \nabla \cdot \mathbf{j}_s \quad \text{on} \quad \Omega \times (-\infty, 0] \quad (2.7b)$$

$$\hat{\mathbf{n}} \times (\mu^{-1} \nabla \times \mathbf{e}) = \mathbf{0} \quad \text{on} \quad \partial\Omega. \quad (2.7c)$$

I take a method of lines approach to discretizing the problem. The main focus of this chapter is on modelling and utilizing the temporal information in TDEM geophysical surveys. I have used existing numerical methods and software for spatial discretization. I will give a brief overview of the spatial discretization I have employed before moving on to describe time-discretization methods and the application to grounded source TDEM inversion.

2.1.1 Spatial discretization

I perform spatial discretization using the mimetic finite volume method on locally refined rectilinear meshes called OcTrees. The general approach I use is described in [37]. It is based on the mimetic finite difference algorithms developed by Hyman and Shashkov [49]. It was extended for use in modelling the quasi-static Maxwell equations on OcTree meshes in work by Haber and Heldmann [38] and by Horesh and Haber [45].

I discretize the weak form of eq. (2.7a),

$$(\mu^{-1}\nabla \times \mathbf{e}, \nabla \times \mathbf{v}) + (\sigma \frac{\partial \mathbf{e}}{\partial t}, \mathbf{v}) = -(\frac{\partial \mathbf{j}_s}{\partial t}, \mathbf{v}), \quad (2.8)$$

where the parentheses indicate the $L_2(\Omega)$ inner product, \mathbf{v} is an arbitrary test function from the Sobolev space $H_0(\text{curl}, \Omega)$, and the curl operator is to be understood in the weak sense. The definition of $H_0(\text{curl}, \Omega)$ is

$$H_0(\text{curl}, \Omega) = \{\mathbf{u} \in L_2(\Omega)^3 : \nabla \times \mathbf{u} \in L_2(\Omega)^3, \mathbf{n} \times \mathbf{u} = 0 \in \partial\Omega\}, \quad (2.9)$$

where $L_2(\Omega)^3$ is the set of 3-component vector valued functions whose components are square integrable on Ω . A solution of eq. (2.8), called a weak solution, is a function in $H_0(\text{curl}, \Omega)$ that satisfies eq. (2.8) for arbitrary test functions in $H_0(\text{curl}, \Omega)$.

The weak formulation of the electric field curl-curl equation admits a larger class of solutions than the strong form. In particular it provides a natural setting to study non-smooth solutions. It also provides a convenient setting in which to formulate mimetic finite volume discretization. For a detailed and rigorous treatment of the weak formulation of Maxwell's equations the reader is referred to [65].

The discretization of eq. (2.8) follows three steps. First the curl operator is discretized, followed by the inner products, giving a system of ordinary differential equations (ODEs). Finally, the time derivatives are discretized to give an implicit time-stepping scheme to solve for the time evolution of the electric field.

The electric field is approximated on mesh cell edges and its curl on face centres, following the staggered grid approach introduced by Yee [103]. This is illustrated in fig. 2.1. The use of staggered grids is standard in the numerical solution of Maxwell's equations. See e.g. [4, 23, 30] for examples from the geophysical literature using staggered grids to approximate Maxwell's equations using spectral, finite difference, and finite element methods, respectively.

I will briefly sketch the discretization of the curl operator. The curl of a vector field \mathbf{u} may be defined as

$$(\nabla \times \mathbf{u}) \cdot \hat{\mathbf{n}} = \lim_{A \rightarrow 0} \left(\frac{1}{|A|} \oint_{\partial A} \mathbf{u} \cdot d\mathbf{r} \right) \quad (2.10)$$




Figure 2.1: Yee grid. Components of the electric field are discretized on mesh edges in corresponding Cartesian coordinate directions. Magnetic fields are discretized on face centres. Figure adapted from [51].

where A is an area and the line integral is around the boundary of A . Let a_f be the area of a face f_i of the finite volume mesh. The component of $\nabla \times \mathbf{u}$ normal to f_i is approximated by

$$(\nabla \times \mathbf{u}) \cdot \hat{\mathbf{n}}_f \approx \frac{1}{a_f} \text{mid}(\mathbf{u}) \quad (2.11)$$

where $\text{mid}(\mathbf{u})$ is the midpoint rule approximation of the line integral of \mathbf{u} around the boundary of the face f_i . Applying this over the entire mesh gives the discrete curl operator \mathbf{C} which maps an edge grid vector to a face grid vector.

The inner product integrals are approximated using the three-dimensional (3D) trapezoidal rule. For a detailed derivation see [39]. This leads to the implicit system of ODEs

$$\mathbf{C}^T \mathbf{M}_{\mu^{-1}} \mathbf{C} \mathbf{e} + \mathbf{M}_{\sigma} \frac{\partial \mathbf{e}}{\partial t} = -\frac{\partial \mathbf{s}}{\partial t} \quad (2.12)$$

where \mathbf{e} is now understood to be an edge grid vector and \mathbf{s} is the discretization of the source current. The \mathbf{M} matrices are symmetric positive definite mass matrices arising from discretization of the inner products in eq. (2.8). Their subscripts refer to the PDE coefficients involved in the respective inner products.

I have used the julia software package jInv ([80]) to implement the spatial dis-

cretization. Julia is a relatively new dynamic language designed for high performance technical computing [15]. It allows rapid prototyping, with a syntax similar to Matlab, but gives high performance and has built in distributed memory parallelism that allows for code that scales to large distributed memory systems.

jInv is a software package and framework for PDE constrained parameter estimation problems. It provides finite volume discretizations of differential operators, other building blocks for developing software for the numerical solution of partial differential equations, and optimization routines designed specifically for PDE constrained parameter estimation problems. It uses Julia's builtin distributed memory parallelism and is designed to scale to large scale industrial problems. I was not involved in the initial development of jInv but have become a contributor to the core jInv package, a primary maintainer of the JOcTree package which provides support for finite volume discretizations on OcTree meshes, and the primary developer of the open source transient EM module used to run transient EM forward modelling and inversions in the jInv framework.

2.1.2 Discretization in time

Having approximated eq. (2.8) by the system of ODEs eq. (2.12), I now study its discretization in time using backward differentiation formulas (BDFs). The quasi-static Maxwell equations in time are extremely stiff [40]. It is difficult to give a precise and rigorous definition of stiffness of ODEs. Ascher and Greif [8] give the following intuitive definition:

The initial-value ODE problem is stiff if the step size needed to maintain absolute stability of the forward Euler method is much smaller than the step size needed to represent the solution accurately.

Stiffness is a common property of systems of ODEs that result from the discretization of parabolic PDEs such as the curl-curl equations derived from the quasi-static Maxwell equations. This seems to occur because it is often the case that the solutions of these equations have secondary features that decay extremely quickly while their primary features, that are typically of interest, evolve more slowly.

Because of stiffness the choice of time-stepping method is dictated as much by stability considerations as accuracy considerations I limit my attention to time-

stepping methods that maintain absolute stability for all step-sizes. Such methods are called A-stable. All A-stable methods are implicit, meaning that a system of linear or non-linear equations must be solved at each time-step. This makes them much more expensive per step than explicit methods, in which the solution at a given time can be computed explicitly in terms of known quantities. However, for stiff problems, the much larger step-sizes that can be taken with implicit methods more than compensate for the increased cost of each step.

In addition to A-stability, time-stepping methods for some very stiff problems, including the quasi-static Maxwell equations, must possess a property known as L-stability. Roughly speaking, L-stable methods are methods that when applied to ODEs with solutions that decay in time, will produce approximate solutions that decay arbitrarily close to zero in a single step as the step-size grows toward infinity [6, 7]. A-stable methods that are not L-stable will produce decaying solutions for all step-sizes but may introduce spurious solution components that decay very slowly. I have seen in practice that such methods (e.g. trapezoidal rule, TR-BDF2) fail for transient electromagnetic geophysics problems.

I use the class of A and L-stable methods known as BDFs. I have used the first and second order BDF methods. Higher order methods could be used but in the transient EM geophysics use case the benefit is not likely to be large enough to outweigh the additional code complexity incurred, especially when considering sensitivity computations in inverse problems. The simplest BDF method is the backward Euler algorithm. It approximates the solution of the initial value problem

$$\dot{\mathbf{y}}(t) = \mathbf{f}(t, \mathbf{y}) \quad (2.13)$$

$$\mathbf{y}(t_0) = \mathbf{y}_0 \quad (2.14)$$

at discrete times $\{t_{n+1}\}$ by

$$\mathbf{y}_{n+1} = \mathbf{y}_n + \Delta t_n \mathbf{f}(t_{n+1}, \mathbf{y}_{n+1}), \quad (2.15)$$

where $\Delta t_n = t_{n+1} - t_n$. Applying backward Euler to eq. (2.12) gives a system of

linear equations to be solved to update the electric field at each time-step

$$(\mathbf{C}^T \mathbf{M}_{\mu^{-1}} \mathbf{C} + \frac{1}{\Delta t_n} \mathbf{M}_{\sigma}) \mathbf{e}_{n+1} = \frac{1}{\Delta t_n} (\mathbf{M}_{\sigma} \mathbf{e}_n + \mathbf{s}_n - \mathbf{s}_{n+1}). \quad (2.16)$$

For second order time-stepping I used the variable step-size and fixed leading coefficient forms of the second order backward differentiation formula (BDF2). Defining $\rho_n = \Delta t_n / \Delta t_{n-1}$, the variable step-size BDF2 method is

$$\frac{1 + 2\rho_n}{1 + \rho_n} \mathbf{y}_{n+1} = (1 + \rho_n) \mathbf{y}_n - \frac{\rho_n^2}{1 + \rho_n} \mathbf{y}_{n-1} + \Delta t f(\mathbf{y}_{n+1}, t_{n+1}). \quad (2.17)$$

The method maintains stability and second order accuracy for $\rho < 1 + \sqrt{2}$ [43]. Applying this method to eq. (2.12) gives the linear time-stepping system of equations

$$\left(\mathbf{C}^T \mathbf{M}_{\mu^{-1}} \mathbf{C} + \frac{1 + 2\rho}{\Delta t_n (1 + \rho)} \mathbf{M}_{\sigma} \right) \mathbf{e}_{n+1} = \frac{1}{\Delta t_n} \left(\mathbf{M}_{\sigma} \left[(1 + \rho) \mathbf{e}_n - \frac{\rho^2}{1 + \rho} \mathbf{e}_{n-1} \right] - \frac{1 + 2\rho}{1 + \rho} \mathbf{s}_{n+1} + (1 + \rho) \mathbf{s}_n - \frac{\rho^2}{1 + \rho} \mathbf{s}_{n-1} \right). \quad (2.18)$$

Two main factors complicate the use of BDF2 in TDEM simulations. The first difficulty is that the source current in TDEM surveys is usually discontinuous in time. Secondly, one is usually interested in simulating TDEM fields on logarithmic time scales. Previous work [68] has shown that simulating TDEM geophysical surveys over such time-scales is most efficiently accomplished by using sparse direct matrix factorization methods to solve eq. (2.16). The time-range of interest is divided into regions, with the step-size kept constant in each region. For example, fig. 2.2 shows a synthetic transmitter waveform—the current in the transmitter as a function of time—similar to that employed in the commercial SkyTEM TDEM system [88]. A single step-size is used to discretize the on-time, during which magnetic fields in the earth are developed. The response to the shutoff of the transmitter is simulated over a time range from 0.01 ms to 6 ms after the end of the transmitter shutoff, using six distinct step-sizes. The backward Euler system matrix must be factored once for each unique step-size.

Like the backward Euler system, the variable step-size BDF2 system matrix




Figure 2.2: Typical example of realistic TDEM transmitter current time-dependence plotted in black, with dots showing times at which solution is computed. The solid blue line shows the variation in the step-size over the full time range of the simulation. There are six unique step-sizes.

eq. (2.19) remains constant while the step-size is constant. However, the system matrix at a transitional step between a time region where a constant step size Δt_1 was and Δt_2 is different from the matrix for subsequent steps of size Δt_2 . Matrix factorization is computationally expensive and for TDEM simulation methods to be efficient, it is essential to minimize the number of factorizations. There are two main approaches in the ODE literature to avoid factoring an additional matrix when the step-size changes (see e.g. [6, 43]). In the first approach, the fields at previous times may be interpolated to compute the solution at the time that will allow the constant step-size BDF2 method for the new step size to be used at the step-size boundary. In the second approach, which I have used, the BDF2 method is reformulated in so-called fixed leading coefficient form. This is more robust than the interpolation approach. In the fixed leading coefficient form, the coefficient of the mass term at the leading time is replaced by one that depends only on the current step-size, at the cost of an additional term in the right hand side of the time-stepping linear system. Using the fixed leading coefficient second order backward differentiation formula (FLCBDF2) in the discretization of eq. (2.16) gives the

time-stepping system of equations

$$\begin{aligned} (\mathbf{C}^T \mathbf{M}_{\mu^{-1}} \mathbf{C} + \frac{3}{2\Delta t_n} \mathbf{M}_{\sigma}) \mathbf{e}_{n+1} = \frac{1}{\Delta t_n} \mathbf{M}_{\sigma} \left[\frac{3}{2} (1 + \rho^2/3) \mathbf{e}_n - \frac{\rho^2}{2} \mathbf{e}_{n-1} \right] - \\ - \mathbf{d}\mathbf{s}_{n+1} + \frac{\rho - 1}{2} (\mathbf{C}^T \mathbf{M}_{\mu^{-1}} \mathbf{C} \mathbf{e}_n + \mathbf{d}\mathbf{s}_n). \end{aligned} \quad (2.19)$$

where $\mathbf{d}\mathbf{s}_{n+1}$ denotes the second order backward approximation to the time derivative of the transmitter current at time t_{n+1} .

For efficiency, it is crucial that BDF2 time-stepping is initialized without performing an additional factorization. BDF2 is a multi-step method, with computation of each step requiring the approximate solution at the two previous steps. Therefore the first step must be computed using a different algorithm. Standard practice is to initialize BDF2 with a single step of backward Euler. Unfortunately, this requires an additional factorization, since the system matrix for backward Euler time-stepping is different from the BDF2 matrix. However, the constant step-size and fixed leading coefficient BDF2 system matrices for step-size Δt are equal to the backward Euler system matrix for step-size $\Delta t_{be} = 2\Delta t/3$. I have initialized BDF2 time-stepping by taking two backward Euler steps of size Δt_{be} and linearly interpolating between them to compute the fields at time $t_0 + \Delta t$, where t_0 is the time at which initial conditions are specified. The additional backward Euler step needed for interpolation is much cheaper to compute than an additional factorization and the interpolation does not cause a degradation in the overall order of accuracy of the BDF2 method.

Now I address the treatment of discontinuities in the electric field. BDF2, and indeed all higher order time-stepping methods, lose their convergence properties in the presence of discontinuities in the solution [89]. In TDEM surveying, electrical currents in the earth are induced by rapid changes in the magnetic field of the transmitter. The change in magnetic field is achieved by shutting off the transmitter current as quickly as possible. In cases where the transmitter on-time is short, and data will be collected in early times after the transmitter shutoff, accurate modelling of the full transmitter current time-dependence, such as that depicted in fig. 2.2 is necessary. In other cases, the transmitter current may be modelled as $I(t) = 1 - H(t)$, where $H(t)$ is the Heaviside step function. Such a waveform is




Figure 2.3: Transmitter current step-off waveform. Assumes transmitter current has been on long enough that fields are steady-state at $t = 0$, where forward modelling begins. Shutoff takes place in one time-step.

illustrated in fig. 2.3. Regardless of the exact waveform, the time-derivative of the source current, which appears in the forward modelling equation eq. (2.16), is typically discontinuous at the point where the current reaches 0. Fortunately the point of discontinuity depends only on the transmitter waveform and so is known a priori. In simulation of actual geophysical surveys the waveform will be known only approximately but the approximate waveform will be taken as a fixed input to the forward modelling algorithm—thereby making the point of discontinuity a fixed input to the algorithm. This makes it straightforward to take steps to minimize its effect. I have found empirically that when no effort is made to address discontinuities, much error is introduced into the solution near the time of the discontinuity but that the accuracy at late times is not severely affected.

Because BDF2 is a backward looking two step method, it will use information from both sides of a discontinuity when taking the first step immediately after the discontinuity, giving spurious results. For example, in the approximate step-off waveform in fig. 2.3 the true time derivative of the current at 1×10^{-4} s is clearly

0 A/s. However, the BDF2 approximation of the derivative is

$$\frac{dI}{dt} \approx \frac{3}{2}I(1 \times 10^{-4}) - 2I(5 \times 10^{-5}) + \frac{1}{2}I(0) = \frac{1}{2}. \quad (2.20)$$

I address this problem by reverting to backward Euler time-stepping for two steps after crossing a discontinuity so that the solution at the time of the discontinuity is not used in any BDF2 steps. In the case of a step-off transmitter I simply take the first three steps using backward Euler.

I use a simple concentric loop survey configuration above a homogeneous earth—for which an analytic solution is available—to empirically study the effectiveness of this approach. This model problem consists of a circular loop transmitter of radius r , placed at the surface of a homogeneous half space of conductivity σ_h . Assuming that a steady current runs through the transmitter for times $t < 0$, and that the current is instantaneously terminated at $t = 0$, the time derivative of the vertical component of the magnetic field for $t > 0$ is

$$\frac{\partial h_z}{\partial t} = \frac{1}{\mu_0 \sigma_h r^3} \left(3 \operatorname{erf}(\tau) - \frac{2}{\sqrt{\pi}} \tau (3 + 2\tau^2) e^{-\tau^2} \right), \quad (2.21)$$

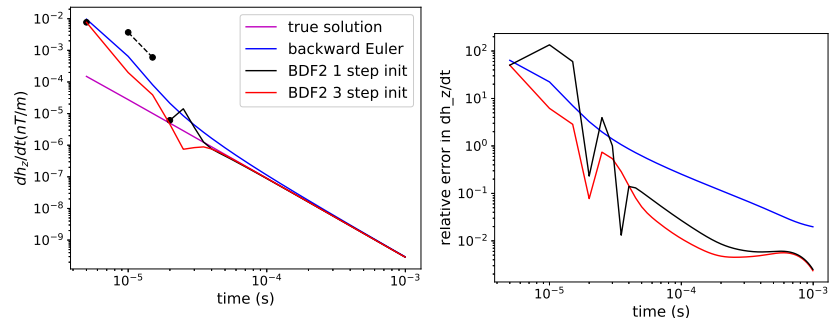
where $\tau = r \sqrt{\frac{\sigma_h \mu_0}{4t}}$, erf is the error function, and μ_0 is the magnetic permeability of free space [97]. This idealized problem is based on the typical airborne transient EM survey in which a horizontal induction coil at the centre of a horizontal loop transmitter measures the $\frac{\partial h_z}{\partial t}$ response when the current in the loop is abruptly shut off.

I simulated the electric fields due to a loop with radius 13.5 m at the surface of a half-space with a conductivity of 0.01 S/m using backward Euler time-stepping, BDF2 initialized with a single step of backward Euler, and BDF2 initialized with three steps of backward Euler. The spatial discretization was the same in all three cases. It used an OcTree mesh consisting of a core region of 5 m^3 fine cells covering a rectangular volume of side length 200 m centred at the loop centre. The fields were simulated from 0.005-1 ms using 200 steps of length 0.005 ms. The transmitter current was set to 1 A at the initial time $t = 0$ (immediately prior to the instantaneous shutoff) and 0 A for all subsequent times. The z-component of the magnetic field dh_z/dt was computed from the electric fields at each time-step by

numerically computing the z -component of $\nabla \times \mathbf{e}$ at the centre of the transmitter loop. Because spatial discretization adds additional error to the numerical solution, the analytic solution eq. (2.21) cannot be used to perform a rigorous convergence study of my implementations of the backward Euler and BDF2 time-stepping algorithms. However, by using a relatively fine spatial mesh much can be learned from comparisons with the analytic solution.


The true solution and computed dh_z/dt from all three time-stepping methods are shown in fig. 2.5a. All three methods are wildly inaccurate in the first few time-steps. This is the expected behaviour. In these early steps the solution is dominated by the approximation of the transmitter step-off. Fortunately these artefacts are quickly damped and at later times the solution is less affected by the accuracy of the current waveform approximation. This is consistent with the general behaviour of BDF methods applied to parabolic problems. In these problems it is not necessary to resolve the details of early time behaviour in order to capture the late time behaviour [6]. In practice, when approximating a real transmitter waveform by a step-off, one must simply choose a small enough step-size such that any artefacts from inaccurate approximation of the step-off are sufficiently damped before the earliest time of interest.

After these early steps dominated by artefacts, both BDF2 solutions are much more accurate than the backward Euler solution. This can be seen by eye in fig. 2.5a and in the relative errors in the approximate solutions plotted in fig. 2.4b. Figure 2.4b also shows that initializing the time-stepping with three backward Euler steps in order to avoid using BDF2 across the discontinuity in the solution does improve the accuracy at later times, as well as eliminating the large oscillation seen in the naive solution where BDF2 is applied across the discontinuity. This oscillation can be more clearly seen in fig. 2.4c, in which the solutions are plotted on a linear, rather than a logarithmic scale. At very late times the two BDF2 solutions do converge but this is likely because other sources of error, such as those due to spatial discretization, are starting to dominate the discrepancy with the true solution.



(a) dh_z/dt , plotted on log scale

(b) Relative error



(c) dh_z/dt for first 5 time-steps, plotted on linear scale

Figure 2.4: Comparison of time-stepping methods in concentric loop example. In the legend, BDF2 1 step init refers to the solution computed using backward Euler for the first step only, using BDF2 thereafter with no special treatment of the discontinuity. BDF2 3 step init refers to the solution computed using backward Euler time-stepping for the first three steps and BDF2 stepping thereafter. Dashed lines indicate negative values. a) shows the overall behaviour of each method over the full 200 time-step simulation. c) illustrates the large oscillation observed from the second time-step when using BDF2 with one step initialization.

2.2 Parallel time-stepping

The computational cost of transient geophysical EM forward modelling is dominated by the solution of the linear system of equations needed to advance the electric fields at each time-step. The process is inherently sequential. As discussed above, sparse direct methods have proven to be more efficient than iterative methods for geophysical transient EM modelling. It is possible to use parallel sparse factorization software but the factorization computations have limited parallel scaling.

With these considerations in mind I seek alternative methods for the parallel numerical solution of the stretched exponential forward modelling problem and other parabolic initial-boundary value problems. Parallelism across time has received much less attention from the numerical analysis community than parallelism in space but a body of literature on the topic does exist. The greatest effort has been in the acceleration of high order time-stepping methods. The prototypical algorithm of this class is the parareal method [59]. These methods are likely of little use here since they require serial execution of a low order time-stepping method to facilitate parallelization of the high-order method. Here I am seeking to parallelize a first or second order method. Other methods such as space-time multigrid [46] and multigrid waveform relaxation [47] treat the entire space-time problem all at once, which is difficult for large scale problems. Additionally, it is not clear how to reuse information from forward modelling in derivative computations with these methods. Relatively recently, parallel exponential integrator methods, such as the method presented by Börner et al. [16], have shown promising performance and potential for reuse of information during derivative computations but they are difficult to differentiate. It may be that their parallelism can scale to large numbers of processors but this is not clear to us.

Here I propose a parallel time-stepping method for time-domain electromagnetic geophysics that is a small adjustment of the methods already discussed in this chapter. The key insight leading to the method is that due to the diffusive nature of quasi-static transient EM fields it is possible to use independent time stepping processes to simulate electric fields at different time-scales of interest. For parabolic PDEs, with sources that step off or ramp off in time, it is possible to model late

time behaviour in a stable fashion without resolving early-time behaviour [6]. Features of the fields decay in magnitude and become spatially smoother as the time after source shutoff increases. For example, to simulate a survey collecting data over the time interval 0.001-1 s, a minimum step-size of 1×10^{-4} or smaller would be required but a step-size of 0.05 s might be sufficient to roughly approximate the electric field 1 s after current shutoff.

In my parallel time-stepping algorithm I consider modelling the electric field at each time of interest as an independent forward problem. I then choose a single time step size and number of steps adequate to accurately model the field at that time. Then each time-stepping process has a single step-size. Recall that the matrix of the time-stepping linear system remains constant while the step-size remains constant, so each independent forward problem in the parallel time-stepping scheme requires only a single factorization, and the time-stepping processes are trivially parallel.


I have tested this parallel time stepping approach only for ideal step-off sources. As just described it should also work for more complex transmitter current time-dependence (waveform), provided that the source waveform can be adequately discretized in time using the same step size required for simulating late time fields. In cases where the source waveform varies on a shorter time-scale than the step-size required for simulating late time fields, a more sophisticated scheme would need to be used in order to achieve good parallel scaling. A possible approach would be that described by Gander and Güttel [34]. The reader is directed to that work for a detailed explanation of the method. In their method, an inhomogeneous linear initial value problem is split into two independent sub-problems that can be solved in parallel, one homogeneous and one inhomogeneous. In this manner a small step time-stepping process incorporating the source behaviour may be run in parallel with a large step process run until the time of interest. The solution of the true problem is then given by the sum of the independent sub-problems. I have not implemented it but I believe my parallel time-stepping approach could be augmented by such a scheme in order to extend it to more complex transmitter waveforms.

The main determining factor of the parallel efficiency of the method is the number of time steps, which I call spin-up steps, needed to simulate the electric field at a given time. If a very small step-size—and therefore an excessive number


of steps—are needed to ensure accuracy at late times then the benefit of performing multiple factorizations in parallel will be greatly reduced. Although knowledge of early time-behaviour is not required to resolve late-time behaviour, the accuracy of the electric field at any time t will be proportional to the step-size or square of the step-size for backward Euler and BDF2, respectively. I have found by numerical experiment that the number of backward Euler time steps required to resolve late time fields with sufficient accuracy was large enough to significantly dissipate the speedup acquired through parallelization. The benefit of parallel time-stepping is greatly improved by the use of second order time-stepping methods.

2.2.1 Evaluating accuracy

The parallel BDF2 time-stepping algorithm was tested with a range of example problems to get an empirical sense of the typical number of spin-up iterations required to achieve sufficient accuracy over a wide range of time-scales. For each test problem I computed a reference backward Euler solution using a very small step-size and then computed the relative error between the reference solution and BDF-2 solutions with various constant step-sizes. Results for a typical test problem using a reference solution with BE step-size of 1×10^{-5} s are shown in figure fig. 2.5. The test problem used to generate those figures was that of computing the electric field at the surface of the earth due to a 2 km long straight wire grounded transmitter over a homogeneous half-space. The transmitter waveform was a simple step-off. The mesh consisted of a core region of $50 \times 50 \times 50$ m cells covering the transmitter and a 1.2×1.2 km area centred about the centre of the transmitter. fig. 2.5a shows the component of the electric field parallel to the transmitter wire, at the point half way between the transmitter electrodes. It shows that for the first two to three time-steps for each BDF2 step size, the fields differ significantly from the reference backward Euler solution, as expected. A step-size of 5×10^{-4} s for example is not small enough to accurately resolve the electric field at 10^{-3} s. However, the fields for all BDF2 step sizes start to agree with the backward Euler field after four steps. Figure 2.5b shows the relative difference between the reference



(a) Electric field in-line component over the time interval $[10^{-4}, 10^{-1}]$ s for reference backward Euler solution computed using step-size of 1×10^{-5} s and BDF-2 solutions at various step-sizes shown in the legend.



(b) Relative error between reference backward Euler solution and BDF-2 solutions for each BDF-2 step-size.

Figure 2.5: Parallel time-stepping accuracy on grounded source test problem.

backward Euler solution and the BDF2 solutions, which is computed point-wise as

$$\text{rel. dif.} = \frac{d - d_{\text{ref}}}{d_{\text{ref}}}. \quad (2.22)$$

For all step-sizes the BDF2 solutions have a relative difference with the reference solutions of less than 1×10^{-2} after 6 time steps. In the context of inverse modelling such accuracy is normally sufficient. Errors in the physical property models will overwhelm discrepancies from this parallel time-stepping scheme. Testing on more complex conductivity models did not show a degradation in accuracy.

2.2.2 Evaluating performance

I tested the performance of the algorithm on a synthetic problem with a more realistic time-stepping scheme. This problem is based on the gradient array style of survey used in induced polarization experiments. It includes two perpendicular long grounded wire transmitters that intersect at their centres. An array of electric dipole receivers is laid out on a regular grid centred at the point of intersection of the transmitter wire paths. The volume enclosing the survey area was meshed with a core volume of $50 \times 50 \times 25$ m cells, with the cell size gradually expanding from there to fill the domain. With the domain being made large enough to approximately fulfill the zero-flux boundary conditions. In total the mesh had 41120 cells and 135233 edges. Constraining the so-called hanging edges of the mesh (a detail of spatial discretization not discussed in this chapter) then gave a linear system with 112849 unknowns to be solved at each time-step.

This single mesh was used to simulate the data from both transmitters. The linear systems of equations at each time-step of the forward modelling process were solved using the Pardiso sparse-direct linear solver software package [84]. As discussed earlier in this chapter, factorization of the system matrix is the most expensive element of transient EM simulations when using direct methods. The system matrix depends on the mesh, earth model and time-step size but not on the transmitter or the time itself. The transmitter enters the system only through the right hand side. Therefore multiple transmitters and multiple time-steps of equal size can be handled very efficiently by a direct solver—each additional transmitter and equally sized time-step adds a forward and backward substitution step but not




Figure 2.6: Time-stepping scheme for reference backward Euler solution . There are 10 steps each of sizes 2×10^{-5} s, 4×10^{-5} s, 8×10^{-5} s, 1.6×10^{-4} s, and 14 steps of size 3.2×10^{-4} s.

an additional factorization. Previous surveys of sparse matrix factorization software such as those by Gould et al. [36] and by Davis et al. [26] have found Pardiso to be among the best sparse factorization software in terms of performance and usability. My own previous work ([9]) and the experience of my research group have found it to be the fastest such software for factorizing matrices and solving linear systems arising in geophysical EM simulations.

The current waveforms of both transmitters were modelled as step-offs. Receiver voltages were computed at 22 time steps in the interval $[2.8 \times 10^{-4}, 7.5 \times 10^{-3}]$ s. In the reference backward Euler solution this was modelled using 54 steps of 5 unique sizes, as illustrated in fig. 2.6. Due to its lower accuracy, a larger number of steps before the first observation time are needed to eliminate initialization artefacts in backward Euler time-stepping than with BDF2. In this example 12 steps were taken to reach the first observation time.

I performed parallel BDF2 time-stepping on the same mesh to compute the receiver voltages at the same 22 observation times. I did not have sufficient computational resources to test using a separate time-stepping process for each observation time but I was able to divide the 22 observation times into 6 parallel time-stepping processes that each used could be performed in parallel, with each using a constant

step-size	first obs. time	last obs. time	# of obs. times	total # of steps
3.5×10^{-5}	2.8×10^{-4}	4.4×10^{-4}	3	13
6.5×10^{-5}	5.2×10^{-4}	7.6×10^{-4}	3	13
1.15×10^{-4}	9.2×10^{-4}	1.4×10^{-3}	4	13
2.15×10^{-4}	1.72×10^{-3}	2.68×10^{-3}	4	13
3.75×10^{-4}	3×10^{-3}	4.28×10^{-3}	3	13
6.15×10^{-4}	4.92×10^{-3}	7.48×10^{-3}	5	13

Table 2.1: Parallel time-stepping scheme. Six parallel forward modelling simulations with constant step-sizes shown in this table were run in parallel to simulate the receiver voltages at the same observation times used in the reference backward Euler simulation. The number of observation times per parallel process varied but the total number of time-steps was the same in all the simulations.

step-size and thus only a single factorization. For each parallel time-stepping process, 8 steps were taken to reach the first observation time. The step sizes used and observation times covered by each parallel time-stepping process are summarized in table 2.1. These tests were performed on a small computer cluster containing four twelve core compute nodes, with two nodes running one time-stepping process each and the other two nodes each running two time-stepping processes. Attempting to run more than two parallel time-stepping processes per node resulted in significantly reduced performance. Each time-stepping process is both compute and memory intensive and this poor performance was presumably due to low memory bandwidth on the cluster nodes causing the computations to become memory bound rather than compute bound. The specific time-stepping scheme detailed in table 2.1 was chosen to minimize the number of time-steps per parallel process while balancing their computational loads.

The reference backward was computed in approximately 89 s on a single core of a single node of the cluster. This time, as with all timings in this performance test, did not include initial setup computations such as the construction of the mesh and differential operator matrices. This method of measuring performance was chosen because it is most representative of the typical cost of forward modelling during an inversion. The setup computations need only be done once at the beginning of an inversion and subsequent simulations with updated earth models need

# of threads	run time (s)
1	89.
2	62.
4	46.
6	40.
12	36.

Table 2.2: Backward Euler solution run times.

only update the mass matrices that depend on the earth model.

As previously mentioned in this chapter, aside from running parallel time-stepping processes, the linear solver itself can be parallelized. The Pardiso solver uses openMP shared memory parallelism. Using shared memory limits the parallel scaling to a single node but provides efficient parallelism within a node. Distributed memory parallel sparse matrix factorization software that can scale across multiple computers, such as the MuMPS package [3], are available. Previous experience with MuMPS seems to indicate that good parallel scaling across multiple machines will only occur for the largest of geophysical EM simulation problems.

The parallel time-stepping algorithm provides a way to expand the parallelism of transient EM forward modelling beyond what a parallel sparse linear solver can provide. In this example the best performance was achieved when combining the two modes of parallelism. To achieve a full comparison, I ran both the reference backward Euler solution and the parallel time-stepping BDF2 solutions with Pardiso in both single and multi-threaded modes.

I ran the reference backward Euler solution with 1, 2, 4, 6, and 12 threads. Recall that this simulation involved 54 time-steps of 5 unique sizes, thus requiring 5 factorizations. The run times for these tests are shown in table 2.2 and the corresponding parallel efficiency plot is shown in fig. 2.7. Using Pardiso gives a non-trivial speedup but the parallel efficiency drops off precipitously as the number of threads increases.

The parallel time-stepping scheme achieved lower run times than was possible by parallelizing the linear solver. This was the case despite the fact that the parallel time-stepping solution actually performed more floating point operations than the backward Euler solution, using a total of 78 time-steps and six factoriza-




Figure 2.7: Pardiso parallel scaling results when using a single sequential backward Euler time-stepping process. Parallel efficiency for N threads is $T_1/(N \cdot T_N)$ where T_1 is the sequential run time and T_N is the run time using N threads.

threads per process	run time (s)
1	20
2	15
4	11
6	11

Table 2.3: Parallel time-stepping run times using 6 parallel time-stepping processes run across 4 compute nodes of a computer cluster.

tions across all six processes in comparison with the 54 steps and 5 factorizations used by the backward Euler solution. The parallel time-stepping algorithm is able to take advantage of distributed memory high performance computing hardware, while the sequential time-stepping algorithm is not. With each process running single threaded, the parallel time-stepping run time was 20 s—already a substantial improvement over single time-stepping process running with 12 threads. This was improved by running Pardiso with multiple threads in each time-stepping process. The results are summarized in table 2.3. These results show that the parallel time-stepping method is capable of significantly improving upon the parallelism present in matrix factorization software.

2.3 Conclusions

The work presented in this chapter has shown that the speed and accuracy of geophysical transient EM forward modelling can be significantly improved by the use of more sophisticated numerical methods than are typically used in the temporal discretization of the problem. I have investigated the use of second order time-stepping methods in initial boundary value problems involving the time-domain quasi-static Maxwell equations. I found that as long as care is taken in initializing the time-stepping and handling discontinuities, then these methods provide superior performance, relative to first order methods. The use of second order methods allowed the development of a parallel time-stepping algorithm that removed a major bottleneck in the parallel scaling of quasi-static transient EM forward modelling. I demonstrated that the parallel time-stepping method could achieve significantly greater speedup than parallelization of the linear algebra of the time-stepping alone could achieve.

Chapter 3

Forward modelling of the dispersive time-domain Maxwell Equations using the stretched exponential function

This chapter will describe the algorithm I have developed for modelling time-domain quasi-static electromagnetic (EM) fields in the presence of polarizable earth materials. I start by stating the forward modelling problem precisely, discussing the difficulties introduced by the consideration of induced polarization (IP), and motivating the choice to use stretched exponential relaxation to model IP decays. I then discuss the similarities and differences between the stretched exponential model of IP and the Cole-Cole model, the de facto standard frequency domain geophysical IP model, before describing the implementation of the stretched exponential forward modelling algorithm. The chapter finishes with synthetic modelling examples illustrating the capabilities of the stretched exponential approach.

3.1 The stretched exponential formulation

I am interested in solving the quasi-static time-domain Maxwell equations to determine the electric field due to either grounded or inductive sources, in the pres-

ence of polarizable media. The polarization effect is contained in the constitutive relation between the electric current \mathbf{j} and the electric field \mathbf{e} . I formulate the quasi-static Maxwell equations in the time-domain in terms of the electric field \mathbf{e} , magnetic flux density \mathbf{b} , current in the earth \mathbf{j} and source current \mathbf{j}^s . I use zero-flux boundary conditions. Other boundary conditions could be used but zero-flux were sufficient for the work considered in this thesis. I will assume steady state initial conditions. Formally, the forward modelling problem is described by the following initial boundary value problem on a cuboidal domain $\Omega \in \mathbb{R}^3$ and time interval $t \in [0, t^*]$.

$$\nabla \times \mathbf{e} = -\frac{\partial \mathbf{b}}{\partial t} \quad \text{on} \quad \Omega \times (0, t^*) \quad (3.1a)$$

$$\nabla \times \mu^{-1} \mathbf{b} - \mathbf{j} = \mathbf{j}^s \quad \text{on} \quad \Omega \times (0, t^*) \quad (3.1b)$$

$$\frac{\partial \mathbf{e}}{\partial t} = 0 \quad \text{on} \quad (-\infty, 0] \quad (3.1c)$$

$$\hat{\mathbf{n}} \times (\mu^{-1} \nabla \times \mathbf{e}) = \mathbf{0} \quad \text{on} \quad \partial\Omega \times (0, t^*). \quad (3.1d)$$

For inductive sources, I assume the electric field is zero initially and for grounded sources, I compute the initial electric field by solving the DC problem for the given electrode configuration. The current in the earth \mathbf{j} is the primary concern here. In the non-dispersive case, explicit dependence on \mathbf{j} may be eliminated by the use of the standard constitutive relation $\mathbf{j} = \sigma \mathbf{e}$, where σ might be a scalar or tensor quantity but depends only on position.

As briefly discussed in the introduction to this thesis, the relationship between electric field and current is less straightforward when IP effects are considered. The IP phenomenon is most often described in the frequency domain, as a frequency dependent conductivity. Here Ohm's law remains a simple linear relationship

$$\mathbf{J} = \sigma(\omega) \mathbf{E}. \quad (3.2)$$

When transformed to the time-domain, this becomes a convolution

$$\mathbf{j}(t) = \int_{-\infty}^{\infty} \hat{\sigma}(t - \tilde{\tau}) \mathbf{e}(\tilde{\tau}) d\tilde{\tau} \quad (3.3)$$

where $\hat{\sigma}(t)$ is the impulse response of the conductivity. Due to causality and the initial condition on the electric field for $t \leq 0$, this may be reduced to [62]

$$\mathbf{j}(t) = \sigma_{\infty} \mathbf{e}(t) - \int_0^t \hat{\sigma}(t - \tilde{\tau}) \mathbf{e}(\tilde{\tau}) d\tilde{\tau}. \quad (3.4)$$

Note that I have not assumed any particular frequency dependence here. The Cole-Cole model is the most popular choice but eq. (3.4) applies generally to any frequency dependence. Making a different choice will simply lead to a different impulse response function.

The question then becomes, how to deal with this convolution. This question has seen some recent attention in the geophysical EM community and in the electrical engineering finite difference time domain (FDTD) modelling community. As noted by, e.g. Commer et al. [25], Takayama and Klaus [91], the efforts fall into three main categories: direct numerical evaluation of the impulse response and convolution integral (e.g. [62, 101]), auxiliary differential equation methods (e.g. [63, 74, 91]), and methods based on the Z-transform (e.g. [90, 98]).

These methods are all based on translating frequency domain characterizations of dispersive electromagnetic media to the time-domain and they all have advantages and disadvantages whose importance depends on the application. For example, Z-transform methods can be made highly efficient only when constant time-stepping is employed, making them inappropriate for geophysics when the variation in EM fields over logarithmic time scales is typically of interest. Auxiliary differential equation methods work well when the frequency dependence has a simple enough form but require the use of fractional derivatives for more complex scenarios. Direct convolution methods are quite general but can be computationally expensive and difficult to differentiate with respect to model parameters.

Another approach is to consider a parametrization of the IP effect defined naturally in the time-domain. This is the approach I take in this chapter. I use stretched exponential relaxation to describe the decay of IP currents. The stretched exponential function

$$f(t) = e^{-t^{\beta}/\tau}, \quad (3.5)$$

where $\beta \in (0, 1]$ and τ is a time-constant, describes the relaxation of complex sys-

tems whose multiple components exhibit exponential decays of different time constants [2]. It has been used widely in solid state physics to describe relaxations in glassy systems [2], and also in other areas of physics such as time-resolved luminescence spectroscopy [13] and the diffusion of water in biological tissues [12]. In the application to glassy systems it is often thought of as a time-domain analogue to frequency domain representations similar to the Cole-Cole model [2, 44]. This experience in the physics literature motivates the choice to explore the use of the stretched exponential to describe geophysical IP effects. Everett [31] and Haber [37] have noted how the stretched exponential may be used to describe the decay of IP currents but to my knowledge no electromagnetic simulation algorithms based on the idea have been developed.

The algorithm developed in this chapter incorporates stretched exponential relaxation into Maxwell's equations using an auxiliary differential equation approach. The implementation originated from an auxiliary differential equation algorithm for the special case of Debye dispersion (Cole-Cole with $c = 1$) by Marchant et al. [63]. Here I generalize the approach and make key changes to improve the stability of the resulting numerical algorithm. I will show that formulating the constitutive law as an ordinary differential equation (ODE) is mathematically equivalent to the convolution representation.

Before describing the implementation of the stretched exponential approach, I will discuss the general approach to defining the $\mathbf{j}(\mathbf{e})$ constitutive law via an ODE and how it relates to convolutions. I start with two definitions. First, I introduce the effective DC conductivity $\sigma_0 = \sigma_\infty(1 - \eta)$. As noted by Seigel [85], this is the effective conductivity observed in a DC survey of a material with conductivity σ_∞ in the absence of IP effects and chargeability η . Next I write the total earth current $\mathbf{j}(t)$ as

$$\mathbf{j}(t) = \sigma_0 \mathbf{e}(t) - \mathbf{r}(t). \quad (3.6)$$

I call $\mathbf{r}(t)$ the residual current. It contains any deviations from standard non-dispersive current flow. This decomposition of the current will be important for achieving a stable and robust discretization of the time-domain dispersive Maxwell equations.

Before specializing to the stretched exponential case, consider defining the con-

stitutive law relating \mathbf{e} and \mathbf{r} by the following ordinary differential equation (ODE) in \mathbf{r} :

$$\alpha\sigma_0\frac{\partial\mathbf{e}}{\partial t} = \frac{\partial\mathbf{r}}{\partial t} + \lambda(\alpha, \theta, \beta, \mathbf{x}, t)\mathbf{r}, \quad (3.7)$$

where λ is an arbitrary smooth function of position and possibly time, and α , θ , and β are spatially varying parameters. They will be related to the Cole-Cole parameters below. I will also associate the initial condition $\mathbf{r}(0) = 0$ with the ODE. This condition gives the correct initial current for both grounded and inductive source time-domain EM experiments.

The ODE (3.7) has the solution

$$\mathbf{r}(t) = \alpha\sigma_0 \int_0^\infty G(\lambda, t - \tilde{\tau}) \frac{\partial\mathbf{e}}{\partial\tilde{\tau}} d\tilde{\tau} \quad (3.8)$$

where G is a Green's function. The $\mathbf{r}(0) = 0$ conditions implies $G = 0$ for $\tilde{\tau} > t$, which is consistent with the physical notion that \mathbf{r} is caused by an applied electric field \mathbf{e} . Thus the integration may be truncated at t , giving

$$\mathbf{r}(t) = \alpha\sigma_0 \int_0^t G(\lambda, t - \tilde{\tau}) \frac{\partial\mathbf{e}}{\partial\tilde{\tau}} d\tilde{\tau}. \quad (3.9)$$

Considering the derivative of the Green's function in the distributional sense, one may use a generalized integration by parts to give

$$\mathbf{r}(t) = -\alpha\sigma_0 \int_0^t \frac{\partial G(\lambda, \mathbf{x}, t - \tilde{\tau})}{\partial\tilde{\tau}} \mathbf{e} d\tilde{\tau}. \quad (3.10)$$

Comparing this result with eq. (3.4) it is seen that $-\alpha\sigma_0 \frac{\partial G(\mathbf{x}, t - \tilde{\tau})}{\partial\tilde{\tau}}$ is equivalent to a conductivity impulse response that arises when transforming a frequency dependent Ohm's law to the time domain.

In principle one could choose the function λ in eq. (3.7) to generate the Cole-Cole impulse response but except for the special cases of $c = 1, 0.5$, this would likely still lead to an integral that would need to be evaluated numerically. Stretched exponential relaxation offers a similar range of behaviour to the Cole-Cole model and has a simple functional form in the time-domain. Alternatively, I will choose λ such that the residual currents decay with a stretched exponential time dependence in the absence of induction effects. One may ignore induction when $\frac{\partial\mathbf{e}}{\partial t}$ is negli-

bly small. Thus we choose λ such that $\mathbf{r} = Ae^{-t^\beta/\theta}$ satisfies the homogeneous version of eq. (3.7)

$$\frac{\partial \mathbf{r}}{\partial t} + \lambda(\alpha, \theta, \mathbf{x}, t)\mathbf{r} = 0. \quad (3.11)$$

This implies $\lambda = \theta^{-1}\beta t^{\beta-1}$. In the general case when induction and IP effects occur simultaneously, \mathbf{r} will depend on the electric field and the constitutive law relating current and electric field will be defined by eq. (3.6) and the inhomogeneous ODE

$$\alpha\sigma_0 \frac{\partial \mathbf{e}}{\partial t} = \frac{\partial \mathbf{r}}{\partial t} + \theta^{-1}\beta t^{\beta-1}\mathbf{r}. \quad (3.12)$$

The forward modelling initial boundary value problem is then fully specified by Maxwell's equations ((3.1)), the constitutive relation defined by eqs. (3.6) and (3.12), and the initial condition $\mathbf{r}(0) = 0$.

3.1.1 Comparison of stretched exponential and Cole-Cole models

It is instructive to compare the properties of the stretched exponential to the Cole-Cole model. In the case of Debye dispersion the two models are entirely equivalent, which may be shown with the appropriate choice of parameters. Following Marchant et al. [63] and Haber [37], I write Ohm's law in frequency for the Debye model as

$$\frac{1}{\sigma_0} \left(1 - \frac{i\omega\tau\eta}{1+i\omega\tau} \right) \mathbf{J} = \mathbf{E}. \quad (3.13)$$

In this case, Ohm's law may be analytically transformed to the time domain, where it becomes the differential equation

$$(1-\eta) \frac{\partial \mathbf{j}}{\partial t} + \tau^{-1}\mathbf{j} = \sigma_0 \left(\frac{\partial \mathbf{e}}{\partial t} + \tau^{-1}\mathbf{e} \right), \quad (3.14)$$

This is equivalent to eq. (3.12) if we substitute $\mathbf{j} = \sigma_0\mathbf{e} - \mathbf{r}$, $\alpha = 1 - (1-\eta)^{-1}$, $\theta = \tau(1-\eta)$, and set $\beta = 1$. This shows that in the stretched exponential model, α plays the role of chargeability and θ that of a time constant. The parameter β plays a role analogous to that of c in the Cole-Cole model.

For $c, \beta \neq 1$ the models are not fully equivalent but they share many characteristics. These connections can be explored by comparing the Cole-Cole and




Figure 3.1: a) Stretched exponential impulse response for various values of β . b) Cole-Cole impulse responses for various values of c . For both plots $\eta = 0.1$, $\tau = 0.1$ and $\sigma_\infty = 0.1$.

stretched exponential impulse responses. The stretched exponential impulse response is

$$\hat{\sigma}_{\text{se}}(t) = -\alpha\sigma_0 \frac{\partial G}{\partial t} = -\alpha\sigma_0\theta^{-1}\beta t^{\beta-1} e^{-(t^\beta)/\theta}. \quad (3.15)$$

The α and σ_0 parameters control the amplitude of the response, τ controls the overall time scale of the response, and β the shape. The impulse response is very small for times much greater than τ . Recall that the residual current \mathbf{r} is the convolution of the impulse response with the electric field. This means that when gaining intuition about \mathbf{r} one should think of early times in the impulse response plots as affecting the behaviour of \mathbf{r} at late times. For example, a sharp drop in $\hat{\sigma}_{\text{se}}(t)$ at late times corresponds to a negligible IP current at early times after source shutoff and $\hat{\sigma}_{\text{se}}(t)$ increasing at early times as β decreases corresponds to slower decay of \mathbf{r} with decreasing β . This is illustrated in fig. 3.1a, which shows the stretched exponential impulse response $\hat{\sigma}_{\text{se}}(t)$ when $\eta = 0.1$, $\tau = 0.1$ and $\sigma_0 = 0.1(1 - \eta)$ plotted for several values of β . The Cole-Cole impulse responses for the same parameter values are shown in fig. 3.1b. The responses are very similar at early times but the stretched exponential curves decay much more quickly at late times than their Cole-Cole equivalents.

The Cole-Cole impulse responses were computed based on analytic series representations derived by Hilfer [44]. To my knowledge his paper is the only source

to give a general analytic expression for the Cole-Cole impulse response. He derives frequency and time domain representations of several models of relaxation in disordered systems (including the stretched exponential and Cole-Cole models) in terms of contour integrals of a specific form known as H-functions. Infinite series representations of these relaxation functions may be derived from the H-function representations. The series form of the Cole-Cole impulse response is

$$\frac{d}{dt}E_c [-(t/\tau)^c], \quad (3.16)$$

where E_c is the Mittag-Leffler function

$$E_c(x) = \sum_{k=0}^{\infty} \frac{x^k}{\Gamma(ck+1)}. \quad (3.17)$$

While this series representation is exact, it is very ill-behaved in certain parameter regimes and generally difficult to work with numerically. To my knowledge, more robust analytic formulae for the Cole-Cole conductivity impulse response only exist for the special cases $c = 1$ and $c = 0.5$.

Although the general form of the Cole-Cole impulse response is rather unwieldy, Marchant [62] was able to show that for $t \ll \tau$ and any $c \in (0, 1)$, it has the approximate form

$$\hat{\sigma}_{cc} \approx mt^{c-1} + d \quad (3.18)$$

for some constants m and d . Inspection of eq. (3.15) shows that

$$\hat{\sigma}_{se}(t) \approx -\alpha\sigma_0\theta^{-1}\beta t^{\beta-1} \quad (3.19)$$

for $t \ll \tau$, the same time dependence as the Cole-Cole response. This implies that at late times the Cole-Cole and stretched exponential models generate approximately equivalent current flow, again recalling that early time behaviour of the impulse response corresponds to late time current flow.

The relationship between the Cole-Cole and stretched exponential impulse responses can be further illustrated through the exercise of trying to fit the stretched exponential impulse response function to a given Cole-Cole response. I used non-linear least squares to find the stretched exponential parameters that best fit the




Figure 3.2: Cole-Cole impulse response with $c = 0.5$, $\tau = 0.1$, $\eta = 0.1$, $\sigma_\infty = 0.1$ plotted in blue alongside stretched exponential impulse responses with parameters chosen by non-linear least squares fitting. The green line attempted to fit the $0.001 < t < 100$ s time range and the magenta line the $0.1 < t < 100$ s range.

Cole-Cole impulse response with parameters $c = 0.5$, $\sigma_\infty = 0.1$ S/m, $\eta = 0.1$, and $\tau = 0.1$ s. Figure 3.2 shows the Cole-Cole response along with two fitted stretched exponential models, one to the time range $0.001 < t < 100$ s and one to the range $0.1 < t < 100$ s. As expected, when attempting to fit the wide time range it is easy to fit the early-time Cole-Cole response at the expense of capturing its late time persistence. The best fitting stretched exponential model had $\alpha = 0.088$, $\beta = 0.47$, $\theta = 0.23$. Using $1 - \alpha = 1/(1 - \eta)$ and $\theta = \tau(1 - \eta)$, α corresponds to $\eta = 0.081$ and θ corresponds to $\tau = 0.25$. Again, recall that this excellent fit at early times in the impulse response corresponds to equivalent residual current decays at late times.

It is not possible to fit the entire Cole-Cole model with a single stretched exponential. Trying to fit only the late time Cole-Cole impulse response gives a very rough approximation, as shown in fig. 3.2. This analysis also shows that the stretched exponential parameters do not correspond exactly to Cole-Cole parameters, even when the responses coincide. However, if one is interested in the time-scale where the models are equivalent it would be straightforward to find the

Cole-Cole parameters that match a certain stretched exponential electric or magnetic field decay. If one is concerned with matching the Cole-Cole behaviour very closely over a wide range of time-scales then a residual current model based on a sum of stretched exponential decays with different time constants could be used. I have not explored that possibility in this thesis but it would be a straightforward extension.

3.2 Discretization

I present two discretizations of the coupled system of differential equations composed of eqs. (3.1a), (3.1b) and (3.12). In both cases the spatial operators are discretized using the mimetic finite volume method on locally refined rectilinear meshes known as OcTrees, as discussed in chapter 2. In previous work of my Ph.D. supervisor’s research group (e.g. [41]), which this thesis builds upon, the backward Euler method has been used for discretization in time of the non-dispersive quasi-static Maxwell equations. Backward Euler was chosen for its simplicity and excellent stability properties. The quasi-static Maxwell equations in time are very stiff and implicit time-stepping was used to avoid having to take excessively small steps to preserve stability.

The backward Euler discretization described here extends that approach to include the stretched exponential auxiliary ODE formulation of Ohm’s law. A similar algorithm for the special case of Debye dispersion and for the Cole-Cole model described by Padé approximation was described by [63]. I will show how the systems of equations arising from this discretization can be solved efficiently using a combination of direct and iterative methods. In the second approach time-discretization is accomplished by the second order backward differentiation formula (BDF2). This more accurate time-stepping method may be used to parallelize the time-stepping process by the method developed in chapter 2.

3.2.1 Backward Euler Algorithm

I begin the first scheme by discretizing the time derivatives in Maxwell’s equations and Ohm’s law. I will then manipulate this semi-discrete system of equations to eliminate explicit dependence on the magnetic flux density \mathbf{b} and form an equation

for the electric field \mathbf{e} at time t in terms of quantities at earlier times. This equation will then be discretized in space to form a system of linear-algebraic equations to be solved for \mathbf{e} at each time-step.

Discretizing eqs. (3.1a), (3.1b) and (3.12) in time using the backward Euler method but keeping spatial variables continuous yields a system of semi-discrete equations

$$\frac{\mathbf{b}_{n+1} - \mathbf{b}_n}{\Delta t_n} = -\nabla \times \mathbf{e}_{n+1} \quad (3.20a)$$

$$\nabla \times \mu^{-1} \mathbf{b}_{n+1} - \sigma_0 \mathbf{e}_{n+1} = \mathbf{j}_{n+1}^s - \mathbf{r}_{n+1} \quad (3.20b)$$

$$\mathbf{r}_{n+1} - \mathbf{r}_n + \Delta t_n t_{n+1}^{\beta-1} \theta^{-1} \mathbf{r}_{n+1} = \alpha \sigma_0 (\mathbf{e}_{n+1} - \mathbf{e}_n), \quad (3.20c)$$

where Δt_n is the size of the n^{th} time-step and the $n + 1$ subscript on field variables refers to the field at time $t_{n+1} = \sum_{j=1}^n \Delta t_j$. These three equations can be combined to form a single equation expressing \mathbf{e}_{n+1} in terms of the residual current and electric field at earlier times, and the source current:

$$\begin{aligned} \nabla \times \mu^{-1} \nabla \times \mathbf{e}_{n+1} + \Delta t_n^{-1} \sigma_0 (1 - \gamma \alpha) \mathbf{e}_{n+1} = \\ - \Delta t_n^{-1} [\mathbf{j}_{n+1}^s - \mathbf{j}_n^s + (1 - \gamma) \mathbf{r}_n - (1 - \gamma \alpha) \sigma_0 \mathbf{e}_n], \end{aligned} \quad (3.21)$$

where I define $\gamma = (1 + \Delta t_n t^{\beta-1} \theta^{-1})^{-1}$ for brevity.

Next, following the methods described in section 2.1.1, I convert eq. (3.21) to weak form and discretize it. This results in a system of linear algebraic equations that can be solved for the approximate electric field discretized on cell edges

$$\begin{aligned} (\mathbf{C}^T \mathbf{M}_{\mu^{-1}} \mathbf{C} + \Delta t_n^{-1} \mathbf{M}_{\sigma_0(1-\gamma\alpha)}) \mathbf{e}_{n+1} = \\ - \Delta t_n^{-1} [\mathbf{s}_{n+1} - \mathbf{s}_n + \mathbf{M}_{(1-\gamma)} \mathbf{r}_n - \mathbf{M}_{\sigma_0(1-\gamma\alpha)} \mathbf{e}_n]. \end{aligned} \quad (3.22)$$

The matrix \mathbf{C} is the discrete curl operator. All \mathbf{M} matrices are mass matrices. Their subscripts indicate the PDE coefficients they discretize. The discrete source current \mathbf{s} is computed by discretizing the transmitter wire path onto mesh edges.

After solving for the electric field at a given time-step the residual current must also be updated, for use in the right hand side of the next electric field update.

To derive the residual current update I solve eq. (3.20c) for \mathbf{r}_{n+1} and discretize in space, giving the linear system

$$\mathbf{M}\mathbf{r}_{n+1} = \mathbf{M}_\gamma\mathbf{r}_n + \mathbf{M}_{\sigma_0\gamma\alpha}(\mathbf{e}_{n+1} - \mathbf{e}_n), \quad (3.23)$$

where \mathbf{M} is a unit coefficient mass matrix. \mathbf{M} is not diagonal in general but it is very sparse and has a structure amenable to efficient factorization by sparse direct methods. It may be factored once at the beginning of a simulation and eq. (3.23) can be solved quickly by triangular substitution thereafter.

Initial conditions are the last ingredient in the discretization. I assume steady state initial conditions, which have the convenient implication that the initial residual current is always zero. The initial electric field may either be zero or given by the solution to a DC resistivity problem, as occurs e.g. when modelling step-off galvanic sources. In this case I use a mimetic finite volume method to solve for the DC potential on mesh nodes. The initial electric field is then given by the negative discrete gradient of the potential. In this chapter I present numerical examples only for step off sources but the implementation is capable of handling arbitrary transmitter waveforms.

3.2.2 Solving the linear system

I now move to solving the electric field-update system eq. (3.22). For realistic parameter values and step-sizes the stretched exponential time stepping system can be considered a perturbation from the corresponding system arising from backward Euler discretization of the non-dispersive Maxwell equations. As such, experience with the non-dispersive equations has guided the approach here and I will frequently refer to the non-dispersive problem in the following discussion. In non-dispersive TEM simulation and inversion algorithms based on implicit time-stepping it is often preferable to solve the electric field update system using direct factorization methods rather than iterative techniques such as the conjugate gradient method. Stretched exponential modelling creates additional challenges but I can still make efficient use of direct methods and argue that they are preferable to iterative methods.

Over the last fifteen years, several excellent sparse matrix factorization soft-

ware packages have been developed [26] that make direct methods practical for very large matrices such as those encountered in 3D electromagnetic simulations. In this work I have chosen to use the MUMPS [3] software package because of its performance and numerical stability. Matrix factorization is expensive both in terms of memory and CPU time but once it has been completed, the factors can be reused to solve systems with the same matrix and different right hand sides extremely efficiently. In backward Euler discretization of the non-dispersive Maxwell equations (and in the dispersive TEM algorithm of [62]), the electric field update system depends on the step-size but not on the current time. Therefore, as long as the step size remains constant, one may factor the system matrix once and then advance the solution many time steps for relatively little computational cost. In an inversion using gradient based optimization, the factors may also be reused to accelerate the computation of cost function gradients.

In transient electromagnetic simulation, one is generally interested in a wide range of time-scales. Airborne systems typically collect data over several orders of magnitude in time. To maximize information about both inductive and IP effects in a grounded source survey one may need to simulate electric fields over four or five decades in time. In such scenarios, very small step-sizes must be used to accurately model early time behaviour and it is not feasible to use a single step-size for the entire simulation. However, using a small number of step-sizes, factoring the system matrix each time the step-size changes and advancing the solution several steps for each step size has still been more efficient than using iterative methods. This strategy was analyzed in detail by Oldenburg et al. [68].

Unfortunately, time appears explicitly in the stretched exponential electric field time-stepping system matrix in eq. (3.22), causing the matrix to change every step and destroying the main benefit of direct methods. Requiring a factorization at every time-step is too slow to be of practical use. For example, in one sample problem (the gradient array example discussed later in this paper with an additional source added) with two long grounded wire sources on a mesh of approximately 114000 cells with 80 time-steps and 6 unique step sizes, non-dispersive modelling, using 6 factorizations, took an average of 285 s on a high end laptop. Stretched exponential modelling took 2415 s factoring at each time-step.

Since it is now not possible to achieve the high reuse of matrix factors achieved

for non-dispersive simulations, I consider iterative Krylov subspace methods for solving the time-stepping system of equations. These methods do not manipulate the entries of the system matrix directly. They create sequences of vectors that hopefully converge to the solution of the system, requiring matrix-vector products at each iteration but treating the matrix as a black box.

The matrix in eq. (3.22) is symmetric positive-definite and so I turn to the preconditioned conjugate gradient method, the most effective iterative method for such systems [81]. Convergence of the conjugate gradient algorithm depends on the distribution of eigenvalues of the system matrix [81]. The algorithm will converge quickly if the eigenvalues are tightly clustered away from zero. Unfortunately the stretched exponential time-stepping matrix is highly ill-conditioned, making it difficult to use the preconditioned conjugate gradient algorithm with standard preconditioners to solve the system. The ill-conditioning is a consequence of the non-trivial the null-space of the curl operator and the fact that as a discretization of a differential operator the largest eigenvalue remains relatively constant, while the smallest eigenvalue shrinks with refinement of the mesh. Adding the mass term eliminates the zero eigenvalues associated with the null space of the curl but its values are generally not large enough to make the system well-conditioned. Thankfully, the explicit time dependence creates only a small perturbation in the system matrix away from the non-dispersive case. Furthermore, as the change in the mass term will be relatively uniform across the mesh, it seems likely that the change in the mass term at each time-step will cause a shift in the eigenvalues rather than increasing or decreasing their spread.

Motivated by the facts of the previous paragraph, I have developed a preconditioner for the stretched exponential time-stepping system matrix based on occasional matrix factorizations. Let us denote the system matrix at time t_0 as \mathbf{A}_0 and the matrix at time $t_n = t_0 + n\Delta t$ by \mathbf{A}_n . I emphasize again that these matrices depend on both the step size and the current time. While the step size Δt remains constant, the matrix \mathbf{A}_0 is a good approximation of the matrices \mathbf{A}_n at subsequent times $t_0 + n\Delta t$. I can write

$$\mathbf{A}_n = \mathbf{A}_0 + \delta\mathbf{A}_n. \quad (3.24)$$

Now, consider solving systems involving \mathbf{A}_n with the preconditioned conjugate

gradient method. If $\delta \mathbf{A}_n$ is not too large and \mathbf{A}_0^{-1} can be cheaply applied then \mathbf{A}_0 should be an excellent preconditioner for the conjugate gradient method applied to systems involving \mathbf{A}_n . If \mathbf{A}_0 is factored and the factors are stored then forward and back substitution can be used to cheaply and accurately apply \mathbf{A}_0^{-1} to arbitrary vectors, as is required when using \mathbf{A}_0 as a preconditioner.

This motivates the use of a hybrid direct-iterative approach to updating the stretched exponential electric field. At the first time-step, and subsequently when the step-size changes, the system matrix is factored. Otherwise the system is solved using the block preconditioned conjugate gradient method (PCG) [69] with the factored matrix as a preconditioner. I use the block variant rather than standard PCG to improve convergence in systems with multiple right hand sides. This gives the forward modelling algorithm shown in algorithm 1.

Algorithm 1 Electric field update algorithm

Input:

Number of time steps n
 List of step sizes Δt
 Initial electric field

Do:

Factor initial matrix $\mathbf{A}_0 = \mathbf{L}\mathbf{L}^T$
 Solve for \mathbf{e} by forward and back substitution
for $j = 2 \rightarrow n$ **do**
 Construct matrix \mathbf{A}_j
 if $\Delta t(n) = \Delta t(n-1)$ **then**
 Solve $\mathbf{A}_j \mathbf{X} = \mathbf{B}$ with block PCG
 else
 Factor \mathbf{A}_j
 Solve using forward and back substitution
 end if
end for

This approach yielded a significant decrease in forward modelling run-times. I observed a small degradation in accuracy, which was negligible except after a large number of time steps. I tested the algorithm by comparing run-times and computed electric fields when factoring at each time-step and when using the hybrid direct-




Figure 3.3: Error in PCG computed example problem electric fields for loose and tight PCG relative residual stopping tolerances.

iterative algorithm. I considered the fields computed by factoring at each step (\mathbf{e}_{dir}) to be correct and computed the error in the fields computed with the hybrid algorithm (\mathbf{e}_{iter}) as

$$\delta \mathbf{e} = \frac{\|\mathbf{e}_{\text{dir}} - \mathbf{e}_{\text{iter}}\|_2}{\|\mathbf{e}_{\text{dir}}\|_2}. \quad (3.25)$$

The highly ill-conditioned nature of the system makes it difficult to compute the electric field accurately, with very small PCG residuals not necessarily corresponding to accurate electric fields. These errors accumulate over the course of the time-stepping process since fields and residual currents computed at a given time step are used in the right hand side of future update linear systems. The relative error norm of the electric field for PCG residual stopping tolerances of 1×10^{-6} and 1×10^{-14} is shown in fig. 3.3. The fields were computed from the example model mentioned above, with 80 time steps, covering a time range 0.0025-3.1 s, using six unique step sizes. Hybrid forward modelling on that model using a PCG relative residual stopping tolerance of 1×10^{-14} took an average of 702 s, 2.5 times slower than non-dispersive modelling of the same example but only 29% of run-time when factoring at every step. The number of PCG iterations per time step varied between two and four. The time increase relative to non-dispersive modelling is mainly due to the PCG iterations requiring multiple applications of forward and back substitu-

tion during application of the preconditioner but also due to the cost of assembling mass matrices at each iteration and computation of the residual current, both of which are not required in non-dispersive modelling.

In my algorithm, \mathbf{A}_0 is a simple Neumann series preconditioner. For any square matrix \mathbf{A} for which the sequence $\{\mathbf{A}^k\}_{k=0}^{\infty}$ converges to zero, the following identity holds

$$(\mathbf{I} - \mathbf{A})^{-1} = \sum_{k=0}^{\infty} \mathbf{A}^k \quad (3.26)$$

and the series is called the Neumann series of \mathbf{A} . Defining $\mathbf{A} = \mathbf{A}_0 + \delta\mathbf{A}$, it follows easily from the Neumann series identity that

$$\mathbf{A}^{-1} = \sum_{k=0}^{\infty} (-\mathbf{A}_0^{-1} \delta\mathbf{A})^k \mathbf{A}_0^{-1} \quad (3.27)$$

as long as $\{(\mathbf{A}_0 \delta\mathbf{A})^k\}_{k=0}^{\infty}$ converges to zero. The eigenvalues of \mathbf{A}_0^{-1} are bounded so convergence will occur if $\delta\mathbf{A}$ is small in a suitable norm. Using a truncation or other approximation to this series as a preconditioner is called Neumann series preconditioning [81]. My algorithm uses zeroeth order Neumann preconditioning. If $\delta\mathbf{A}$ is too large then the Neumann series may converge slowly or not at all. In the sample problems we tested, zeroeth order Neumann series preconditioning worked extremely well when refactoring the system matrix each time the step size was changed. I believe that the method would continue to work for all time-scales, models and meshes of interest to applied geophysicists. However, in regimes with larger perturbations, a higher order Neumann series preconditioner could easily be implemented to accelerate convergence.

3.2.3 BDF2

As in the non-dispersive case of chapter 2, higher order time-stepping is desirable, both to improve the accuracy of sequential time-stepping and to allow the use of the parallel time-stepping algorithm of section 2.2. The second order time-stepping scheme for stretched exponential forward modelling is based on the second order backward derivative approximation. We follow the same sequence of steps as in the backward Euler discretization, discretizing the time derivatives in eqs. (3.1a),

(3.1b) and (3.12) with BDF2 approximations. Assuming a constant step-size, this gives the semi-discrete system

$$\frac{3}{2}\mathbf{b}_{n+1} - 2\mathbf{b}_n + \frac{1}{2}\mathbf{b}_{n-1} = -\Delta t \nabla \times \mathbf{e}_{n+1} \quad (3.28a)$$

$$\nabla \times \mu^{-1}\mathbf{b}_{n+1} - \sigma_0\mathbf{e}_{n+1} = \mathbf{j}_{n+1}^s - \mathbf{r}_{n+1} \quad (3.28b)$$

$$\begin{aligned} \frac{3}{2}\mathbf{r}_{n+1} - 2\mathbf{r}_n + \frac{1}{2}\mathbf{r}_{n-1} + t_{n+1}^{\beta-1}\Delta t\theta^{-1}\mathbf{r}_{n+1} = \\ \alpha\sigma_0\left(\frac{3}{2}\mathbf{e}_{n+1} - 2\mathbf{e}_n + \frac{1}{2}\mathbf{e}_{n-1}\right). \end{aligned} \quad (3.28c)$$

As with the backward Euler scheme, eliminating explicit dependence on \mathbf{b} and discretizing in space leads to a system of linear algebraic equations for \mathbf{e}_{n+1} at each time step in terms of fields at earlier times and the source current

$$\begin{aligned} (\mathbf{C}^T \mathbf{M}_{\mu^{-1}}^f \mathbf{C} + \frac{3}{2}\Delta t^{-1} \mathbf{M}_{\sigma_0(1-\gamma\alpha)}^e) \mathbf{e}_{n+1} = \\ \frac{3}{2}\Delta t^{-1} [\mathbf{s}_{n+1} - \frac{4}{3}\mathbf{s}_n + \frac{1}{3}\mathbf{s}_{n-1} + \mathbf{M}_{\sigma_0(1-\gamma\alpha)}^e (\frac{4}{3}\mathbf{e}_n - \frac{1}{3}\mathbf{e}_{n-1}) \\ - \mathbf{M}_{1-\gamma} (\frac{4}{3}\mathbf{r}_n - \frac{1}{3}\mathbf{r}_{n-1})], \end{aligned} \quad (3.29)$$

where Δt is now a constant step-size and $\gamma_2 = (1 + (2/3)t_{n+1}^{\beta-1}\Delta t\theta^{-1})^{-1}$.

3.2.4 Implementation

These algorithms were implemented in the Julia programming language. My code was designed to be compatible with jInv [80], a Julia software framework for partial differential equation (PDE) constrained parameter estimation problems, which was briefly described in section 2.1.1. In the context of this chapter, in which I am only discussing forward modelling, this is a minor point but it will make it simple to incorporate this simulation code into an inversion software package that will allow one to invert for intrinsic IP parameters.

I verified the correctness of the code using the method of manufactured solutions [82]. In this technique, one defines an electric field that meets the boundary conditions of the problem, then substitute it into Maxwell's equations and analytically calculates the magnetic flux density, residual current and source current that

match the manufactured electric field such that Maxwell's equations are satisfied. Then these exact formulae are used to compute discretized initial electric field, residual current and source currents. These values may be fed into the software to compute the approximate electric field and residual currents at later times. The rate of convergence of solutions in space and time can then be assessed by varying the mesh cell and time-step sizes.

3.3 Synthetic modelling examples

3.3.1 Grounded source example

I will now illustrate the capabilities of stretched exponential modelling of coupled electromagnetic induction and induced polarization (EMIP) effects using two synthetic examples where both EM induction and IP effects are significant—a gradient array survey and a coincident loop airborne time-domain electromagnetic (ATEM) sounding. The gradient array survey earth model was a homogeneous halfspace containing two anomalous bodies, one chargeable and the other conductive. The ATEM sounding model consisted of a buried conductive block in a resistive half-space with chargeable overburden.

First I consider the gradient array example. Gradient arrays allow for efficient collection of easily interpretable direct current induced polarization (DCIP) data over large areas but they have significant deficiencies when viewed from the DCIP perspective. The major deficiencies are a lack of depth resolution and significant EM coupling. The EMIP approach can significantly address these deficiencies. The survey layout and model for our example is shown in fig. 3.4. A 2 km long straight grounded wire is laid out over the survey area and an array of electric dipole receivers parallel to the transmitter dipole are laid out in a grid covering a 500×500 m rectangular area centred at the centre of the transmitter wire. The transmitter current waveform is modelled as a step-off. In the transmitter on-time, galvanic DC currents allow polarization to build up in the presence of chargeable bodies. When the transmitter current is shut off, the relaxation of these polarizations, as well as electromagnetic induction, will generate transient electric currents in the earth. It can be useful to think of the EM induction aspect of an EMIP gradient array




Figure 3.4: Gradient array example survey layout. a) Cartoon showing transmitter, receiver array bounding box, and sample cartoon receiver. b) Plan view layout of transmitter, receiver array and two block synthetic model.

survey as being analogous to a time-domain CSAMT survey for long grounded wire sources but near field effects can be important so the analogy is not exact. Away from the transmitter electrodes, in our area of interest near the centre of the wire, the primary magnetic field will be approximately solenoidal in the transmitter on-time. Transmitter shutoff will induce currents to flow parallel to the direction of the wire path.

The plan-view locations of the two anomalous blocks in the earth model are shown, along with the transmitter wire path, in fig. 3.4b. Both blocks are $150 \times 150 \times 75$ m rectangular prisms buried 50 m below the earth's surface. The left hand block had the same conductivity as the background halfspace ($\sigma_{\infty} = 0.01$), chargeability $\eta = 0.25$, IP time constant $\tau = 0.5$ s and stretched exponential decay exponent $\beta = 0.5$. Recall that the stretched exponential parameter α is related to η by the relation $1 - \alpha = 1/(1 - \eta)$ and that the parameter θ is related to τ by $\theta = \tau(1 - \eta)$. Recall also that chargeability is equivalent in the Cole-Cole and stretched exponential models. The time constants and exponents play roughly equivalent roles in both models but of course because the models are not fully equivalent we cannot in general choose a set of stretched exponential parameters that will lead to an electric field decay exactly equal to some Cole-Cole decay. With those reminders of the roles of the parameters I move to the right hand block which was one order of magnitude more conductive than the background ($\sigma_{\infty} = 0.1$) and was not chargeable. Air conductivity was set to 1×10^{-8} S/m.




Figure 3.5: Visualization of OcTree mesh discretization of the conductivity model.

The model was discretized on an OcTree mesh of 113772 cells, yielding a linear system of 319653 unknowns to be solved to compute the electric field on mesh edges at each time-step. Cubic cells with side length 25 m were used in the core region of the mesh, including the area of interest and the transmitter wire path, with cell size expanding away from the area of interest and transmitter. A vertical cross section of the mesh showing cell size and conductivity in the $y = 0$ plane is pictured in fig. 3.5. Starting from a DC initial electric field, the simulation ran to a final time of 3.11 s after current shutoff, using 80 time-steps with 6 unique step sizes ranging from 2.5×10^{-3} s to 0.1 s. Using Backward Euler time-stepping, solving linear systems using the iterative hybrid direct algorithm described above, this simulation had an average run-time of 620 s, including the time to solve the DC problem for the initial electric field.

The inline component of the DC electric field is shown in fig. 3.6. The plot shows a clear anomaly from the conductive block but only a small distortion of the field contours due to the chargeable block. The maximum discrepancy in the electric field between this model and an otherwise equal model with the chargeable block removed is approximately 5%. The difficulty of distinguishing between the IP and non-dispersive induction responses in the off-time electric field data depends




Figure 3.6: DC inline electric field at the earth’s surface for two block example model. The locations of the buried blocks are shown by black squares.

on several factors including the depth to conductive targets, depth to chargeable targets, and the time scale of IP responses. In scenarios with very good conductors and small IP time constants, traditional EM decoupling procedures may perform poorly. The ability to fully model the coupled EMIP response in 3D should allow for better isolation of the IP response. More generally, having such modelling capability allows inductively generated electric fields to be used as signal in grounded source IP, not just noise. Early-time transient electric fields provide additional information on ground conductivity, which can then be used to improve recovery of IP parameters [52].

The distinct characters of the decays of the electric fields near the conductive and chargeable blocks in this example model are illustrated in fig. 3.7. It is difficult to distinguish the IP response of the chargeable block at early times when induction dominates. The stretched exponential and non-dispersive fields are nearly identical before 1×10^{-2} s. After that the IP response appears clearly in the surface electric field near the chargeable block, as shown in figs. 3.7b and 3.7c. By contrast, fig. 3.7d shows that the stretched exponential surface electric field over the conductive block is almost identical to the non-dispersive field, as expected.

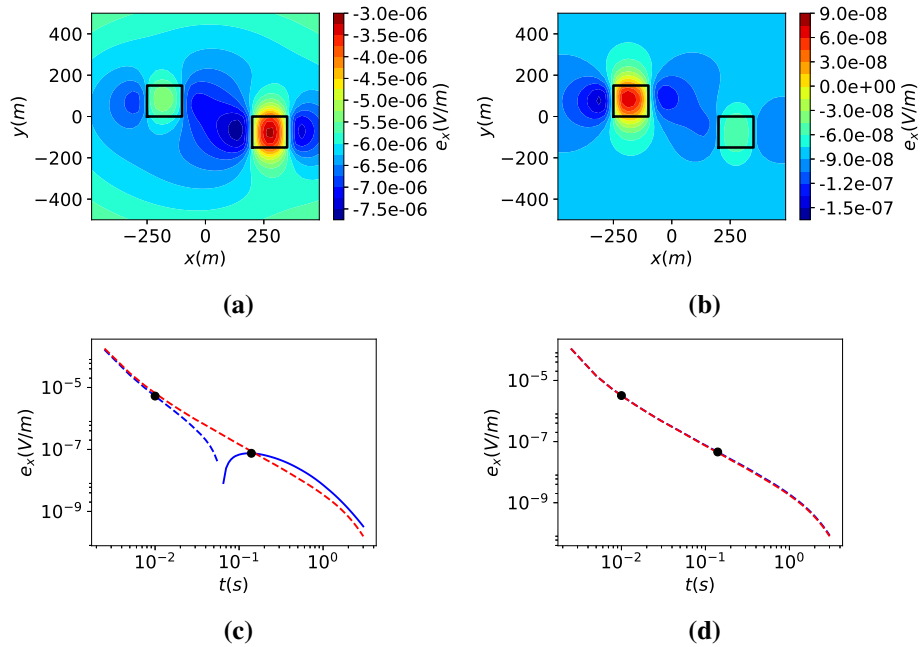


Figure 3.7: In-line component of surface electric field for gradient array example. a) Plan-view in-line field at 0.01 s after current shutoff. Left and right black squares show the boundaries of the chargeable and conductive blocks, respectively. b) Plan view field at 0.14 s. c) Electric field decays above the centre of the chargeable block, with stretched exponential field in blue and corresponding non-dispersive field ($\eta = 0$) in red. Dashed lines indicate negative values. d) Electric fields above the centre of the conductive block, showing that the stretched exponential and non-dispersive responses are almost identical.

3.3.2 Inductive source example

The following example illustrates the applicability of stretched exponential modelling to the simulation of airborne inductive source induced polarization. The phenomenon of negative transients in coincident and concentric loop EM systems has been known to exist for a long time. Weidelt [99] showed that these sign reversals could not occur with non-dispersive electrical conductivity and magnetic permeability, suggesting that IP effects could be responsible. There has been renewed interest in the phenomenon in recent years, as improvements in data acquisition




Figure 3.8: Section view of block in a half-space with chargeable overburden model. The overburden depth is 10 m and the block extends from 40-90 m depth with a plan view cross-section of 100×100 m. The background halfspace conductivity was 0.005 S/m and the conductivity of the block was 0.1 S/m. The overburden had $\sigma_\infty = 0.01$ S/m $\eta = 0.3$, $\tau = 1 \times 10^{-3}$ s and $\beta = 0.8$

have led to these signals appearing with increasing frequency in field data—see e.g. [61]. I illustrate how negative transients can be modelled with stretched exponential polarization with a simple model of a single sounding of a coincident horizontal loop ATEM transmitter with step off source over a resistive half-space containing a conductive rectangular prism buried 40 m beneath the earth’s surface and 10 m of chargeable overburden. The earth model is illustrated and the physical property values are given in fig. 3.8. The transmitter was a square loop with 10 m side length located 30 m above the surface, centred over the conductive block. This choice of IP parameters is representative of the type of fine-grained near surface material thought to be most commonly responsible for IP signals in ATEM data [60]. In addition to smaller time-constants, these materials tend to have sharper decays than typical ground IP targets, leading to β values closer to 1.

The effect of the chargeable overburden on the ATEM response of our example model is illustrated in fig. 3.9. It shows the total EMIP responses alongside the corresponding responses computed when neglecting the chargeable overburden. We used the stretched exponential algorithm with backward Euler time stepping to compute the transient electric field over the time interval $[1 \times 10^{-5}, 0.002]$ s. The vertical component of the magnetic flux density (db_z/dt) was then computed from




Figure 3.9: Effect of chargeable overburden on ATEM response from buried conductor. All plots show the vertical component of $d\mathbf{b}/dt$. The background halfspace response is shown in green in each plot. Dashed lines indicate negative values. a) Block with no overburden. The response with the conductive block at a depth of 40 m is shown in blue. The magenta line shows the response obtained when the top of the block is moved to a depth of 70 m. b) Stretched exponential response with 10 m of overburden ($\sigma_\infty = 0.005$, $\eta = 0.3$, $\tau = 1 \times 10^{-3}$, $\beta = 0.8$) is shown in black, along with the overburden free and background responses. c) Stretched exponential response with the top of the block at 70 m in black, along with the overburden free and background responses. d) Stretched exponential response with chargeable overburden and no conductive block in black, along with background EM response.

the electric field values. When the chargeable overburden is neglected (fig. 3.9a) the response from the block at 40 m depth is clearly distinguishable from the background response. The signal from the block remains but is much smaller when it is lowered to 70 m depth. The EM induction signal from the shallower block is strong enough and emerges early enough that both it and the IP response of the overburden are distinguishable in the EMIP data—fig. 3.9b. The EMIP response never changes sign due to the cancellation of the EM induction and IP signals but the decay is clearly indicative of an IP response. It is important to note here that a more subtle IP response could be difficult to distinguish from a resistive body with 3D geometry. I believe this is an important argument in favour of the use of 3D methods when considering the presence of IP effects in ATEM data.

When the top of the block is moved to a depth of 70 m (fig. 3.9c), the response of the conductor is completely masked by the IP response. However, as fig. 3.9d shows, the IP response generated by the overburden is much larger and is manifested in the data earlier in time without the presence of the block. Modelling the full physics of an ATEM system with IP effects included allows for quantitative simulations of the interactions of non-dispersive EM and IP responses. The ability to model these effects in 3D is very useful in its own right and the key step toward developing a 3D inversion algorithm for conductivity and IP parameters.

3.4 Conclusion

I have developed a novel computationally efficient forward modelling algorithm for 3D time-domain electromagnetic modelling in the presence of chargeable materials, based on stretched exponential relaxation. The Cole-Cole model is still by far the most common framework for modelling and classifying chargeable materials. Stretched exponential relaxation is not equivalent to Cole-Cole relaxation but it is capable of modelling a similar range of phenomena. It is a significant comfort to have shown that the two models do produce equivalent results at late times, showing that they can definitely model the same phenomena if the appropriate time interval is considered. This work also identified the Cole-Cole and stretched exponential responses within a more general context. When electric current may be represented as a convolution with the electric field, the two approaches simply

imply different convolution kernels.

I applied stretched exponential modelling to two synthetic examples possessing the key characteristics of two applied problems where both EM induction and IP effects are important. In the case of grounded source IP surveys with significant EM coupling, I believe that the ability to quantitatively simulate the transient EM and IP behaviour of the surveys opens up exciting possibilities for significantly increasing the amount of information about subsurface geology that they provide. In the case of ATEM surveying, the importance of treating IP effects in the analysis of field data is increasingly clear. I have developed an algorithm capable of efficiently modelling these effects. This algorithm forms the foundation of the three-dimensional (3D) inversion algorithm presented in the next chapter, which is to my knowledge the first algorithm to directly invert transient geophysical EM data for intrinsic IP parameters.

Chapter 4

Stretched exponential inversion

This chapter discusses my preliminary work into the inversion of transient electromagnetic data for stretched exponential parameters. I have developed routines for computing the sensitivity of electromagnetic data to these parameters and performed simple proof of concept synthetic inversions. These preliminary results are encouraging but much more work needs to be done in order to understand how to best extract induced polarization (IP) related information from transient electromagnetic (EM) data in the stretched exponential approach.

This is a multi-parameter inverse problem with a great deal of non-uniqueness. How to best handle that non-uniqueness is an open question. The best workflow is probably problem dependent. For example, should one invert for non-dispersive conductivity and the other stretched exponential parameters simultaneously, or is it necessary to have a good conductivity model before attempting to recover chargeability and the other parameters? Is it always possible to get a good conductivity model before at all considering IP effects?

In some problems it is possible a priori to separate a transient EM decay into three sections: an early time range where EM induction dominates, a late time range where IP effects dominate and a middle section where the two effects occur on similar scales and interact. In other problems this may not be possible and simultaneous inversion for conductivity and IP parameters might be the only way forward. In this chapter I focus on the former, simpler case. I demonstrate the capabilities of stretched exponential inversion on the gradient array model problem




Figure 4.1: Gradient array example survey layout. a) Cartoon showing transmitter, receiver array bounding box, and sample cartoon receiver. b) Plan view layout of transmitter, receiver array and two block synthetic model.

encountered in chapter 3.

4.1 The gradient array example

Recall the synthetic gradient array survey from section 3.3.1. Two perpendicular long grounded wire sources are positioned such that they both pass through the centre point of an array of 25 m electric dipole receivers oriented parallel to the transmitters that measure the direct current (DC) and transient voltages in the earth due to the transmitters. Thus, for the transmitter oriented parallel to the x coordinate axis only data from the x -oriented receivers were used. Similarly, only y receivers were used for the y -axis oriented transmitter. The survey layout for the x -oriented receiver is shown in fig. 4.1. The model consists of two $150 \times 150 \times 75$ m blocks buried 50 m below the surface of a homogeneous half-space with conductivity $\sigma_{\infty} = 0.01$ S/m. The true conductivity model is shown in fig. 4.2. The west block (left hand side of fig. 4.2a) has $\sigma_{\infty} = 0.018$ S/m, chargeability $\eta = 0.25$, time constant $\tau = 0.5$ s and exponent $\beta = 0.5$. The east block is conductive but not chargeable. It has $\sigma_{\infty} = 0.1$ S/m.

As in the forward modelling example in chapter 3, both transmitters used step-off waveforms to excite the earth. The initial DC voltages and transient voltages at 85 times ranging from 5×10^{-4} -1.128 s were computed. Standard deviations of 5% of the value plus a constant floor were assigned to each datum. The true noise-free




Figure 4.2: Selected slices of true conductivity model. Left hand block had $\sigma_{\infty} = 0.018$ S/m and right block $\sigma_{\infty} = 0.1$ S/m. Left block had $\eta = 0.25$, $\beta = 0.5$ and $\tau = 0.5$ s. Right block was not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.

synthetic data were used in the inversions. I also committed the further “inverse crime” of using the same mesh and time-stepping scheme in the inversions that was used to generate the synthetic data. As this work represents just an initial proof of concept for stretched exponential inversion, commission of the inverse crime was deemed acceptable.

4.2 Inversion methodology

Recall that the stretched exponential forward modelling algorithm computes the electric field on the cell edges of a computational mesh at a set of discrete times. The synthetic voltages are computed by approximating the line integrals of the electric fields over the receiver wire paths. In this example the receiver wire paths were always straight line segments coinciding with mesh edges.

I seek to recover the non dispersive conductivity σ_{∞} , the chargeability η , the time constant τ and the exponent β from the DC and voltage data. The parameter values are discretized onto the computational mesh, approximated as constant in each cell. I seek to recover values of the parameters in all subsurface mesh cells. I use the following four stage inversion procedure:

1. invert early time data, which is assumed to be negligibly effected by IP, for σ_{∞}

2. Hold σ_∞ fixed and invert for all three IP parameters simultaneously
3. Hold σ_∞ and η fixed, invert for IP parameters τ and β simultaneously.
4. Refine IP parameter estimates by inverting for all of them (η , τ , β) simultaneously

Stages 2-4 are initialized using the recovered models from the previous steps. Note that DC data are included in the stage 1 inversion despite the fact that they are affected by IP. However, even for DC, IP is a second order effect and it was judged that for this initial inversion it was acceptable to invert the DC data alongside the early off-time data for σ_∞ , ignoring chargeability.

At each stage the inverse problem is posed as a regularized least-squares parameter estimation problem, using Tikhonov regularization [93]. I seek the model \mathbf{m} that minimizes the following objective function subject to bound constraints on the parameter values

$$\begin{aligned} \underset{\mathbf{m}}{\text{minimize}} \quad & \Phi(\mathbf{m}) = \frac{1}{2} \|\mathbf{W}(\mathcal{F}(\mathbf{m}) - \mathbf{d})\|_2^2 + \gamma \mathcal{R}(\mathbf{m}) \\ \text{subject to} \quad & \mathbf{m}_l < \mathbf{m} < \mathbf{m}_h, \end{aligned} \quad (4.1)$$

where \mathbf{d} is the observed data, \mathbf{W} is a diagonal matrix holding the standard deviations of the data, \mathcal{F} is the forward modelling operator, \mathcal{R} is the regularization operator, and γ is the scalar regularization parameter that balances the weights of the misfit and regularization terms in the objective function. For single physical property inversions the regularization operator is the discretization of

$$\|W_i \nabla \mathbf{m}\|_2^2 \quad (4.2)$$

where W_i is a weighting operator that allows smoothness to be enforced more strongly in some regions of the model than others. The discrete regularization operator is

$$\mathcal{R}(\mathbf{m}) = \mathbf{m}^T \mathbf{G}^T \mathbf{W}_i \mathbf{G} \mathbf{m} \quad (4.3)$$

where \mathbf{G} is a cell-centered discretization of the gradient operator. It approximates the gradient of a cell-centre discretized scalar function on cell faces. \mathbf{W}_i is an

interface weighting matrix that discretizes W_i on cell faces. For multi-parameter inversions the model is simply a concatenation of the models for each parameter. The regularization operator is

$$\mathcal{R}_{\text{multi}}(\mathbf{m}) = \sum_j \mathbf{m}^T \mathbf{P}_j^T \mathbf{G}^T \mathbf{W}_{i(j)} \mathbf{G} \mathbf{P}_j \mathbf{m}, \quad (4.4)$$

where \mathbf{P}_j is a matrix that extracts the model for the j^{th} parameter from the full model and $\mathbf{W}_{i(j)}$ is the interface weighting matrix for the j^{th} parameter.

The optimization problem eq. (4.1) was solved using the Gauss-Newton algorithm with backtracking Armijo line search. I solved for the search direction using the projected preconditioned conjugate gradient algorithm with loose stopping tolerance.

There are several ways in which this procedure could be improved but they are beyond the scope of this thesis. For example, structural constraints such as cross-gradient regularization (see e.g. [33]) could be imposed to penalize differences in spatial structure between the different physical property models. Or in cases where there is reason to believe that the IP signal in a transient EM dataset comes from a geometrically simple subsurface body, the techniques of chapter 5 of this thesis could be used.

4.3 Inversion results

First I performed the conductivity only inversion, inverting the DC data and transient data from 8 times ranging from 0.002-0.0055 s. Those times correspond to the 5th to 12th steps of the forward modelling time-stepping. The first four steps could not be used because they contained artefacts from the approximation of the step-off waveform. In this synthetic example it was possible of course to know exactly when IP effects start to become distinguishable in the data. In this example I tried to choose a cutoff time just before IP effects could be clearly seen in the data without comparing it to the synthetic data generated from the true conductivity model without IP effects.

Previous work by my supervisor and I [10] has shown that the use of early time transient data can greatly improve the depth resolution of gradient array surveys.

However, in this work the focus is on studying IP and the interplay of EM and IP effects. In order to allow simulations to run to late enough times to see the IP effects overwhelm the EM effects without using a number of time steps that would render simulation too computationally demanding, I chose not to model early enough times to allow the conductivity inversion to recover the depths of the blocks correctly without the use of interface weighting. No sensitivity depth weighting was employed but the interface weights in the first subsurface layer of the mesh were set to 100 (they were 1 everywhere else) to penalize structures in the model being pushed to the surface. This worked very well. The recovered conductivity model is shown in fig. 4.3. The horizontal locations, depths, and




Figure 4.3: Selected slices of recovered conductivity model from DC and early off-time inversion. Thick black lines show true block boundaries. True conductivities are 0.01 S/m for the background, 0.018 S/m for the left block and 0.1 S/m for the right block. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.

relative conductivities of both blocks were recovered well. This inversion was run to completion, meaning it was able to fit the data, according to the χ^2 goodness of fit criterion, stopping the inversion when the root mean squared (RMS) misfit dropped below 1. The misfit convergence curve is shown in figure fig. 4.4

Next the inversion for η , τ , and β was performed with σ_∞ held fixed to the model recovered from the previous inversion. This inversion was halted after 5 Gauss Newton iterations. At that point the inversion began to struggle to decrease the data misfit and inspection of the models at each iteration showed that the inversion had significantly improved the chargeability model but made only very small




Figure 4.4: Misfit convergence curve for σ_∞ inversion.

changes to the time constant and exponent models. I have not performed a sensitivity analysis but I believe the data are much more sensitive to changes in η than to the other IP parameters, especially at early iterations when, indeed, the τ and β sensitivities depend on the chargeability model. The recovered chargeability model from this inversion is shown in fig. 4.5. The RMS misfit was approximately




Figure 4.5: Selected slices of recovered chargeability model from DC and mid to late off-time data inversion. Thick black lines show true block boundaries. Left hand block has $\eta = 0.25$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.

3.8 when the inversion was stopped. So far the inversion has captured the location

of the chargeable block well but underestimated its chargeability and depth extent.

Next I held both σ_∞ and η fixed and inverted for the time parameters τ and β . This inversion was able to fit the data, in the χ^2 goodness of fit sense, however, in doing so it created τ and β models with a great deal of spurious structure. The final recovered models from this stage 3 inversion are shown in figs. 4.6 and 4.7 the

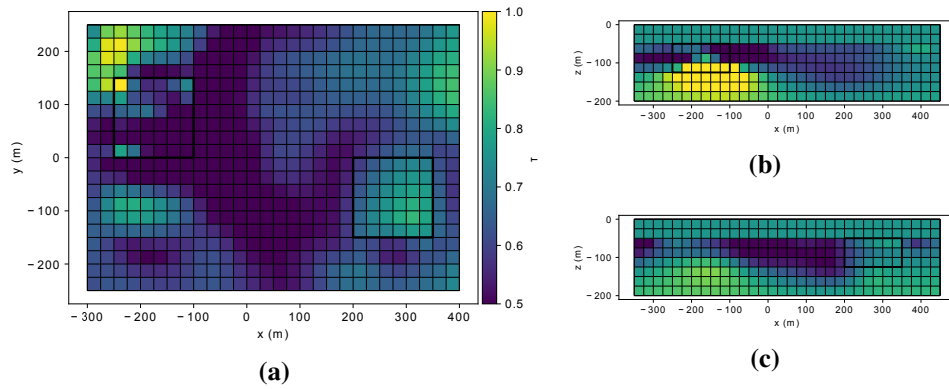


Figure 4.6: Selected slices of recovered time constant model from late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\tau = 0.5$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.

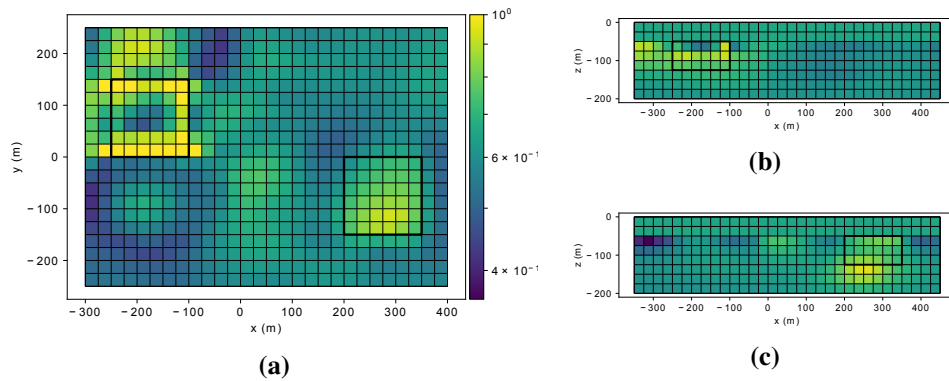


Figure 4.7: Selected slices of recovered exponent model from late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\beta = 0.5$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.

spurious structure is clearly seen. Instead of using this as the final model, the models from an earlier iteration of this inversion can be used to start another inversion for all three IP parameters η , τ , and β . The early iterates from the stage 3 τ and β inversion significantly improve on the results from the stage 2 inversion while remaining relatively smooth. Using these early stage 3 iterates and the recovered η model from the stage 2 inversion as starting models in a final inversion for all three IP parameters worked well. The recovered models from this final inversion are shown in figs. 4.8 to 4.10

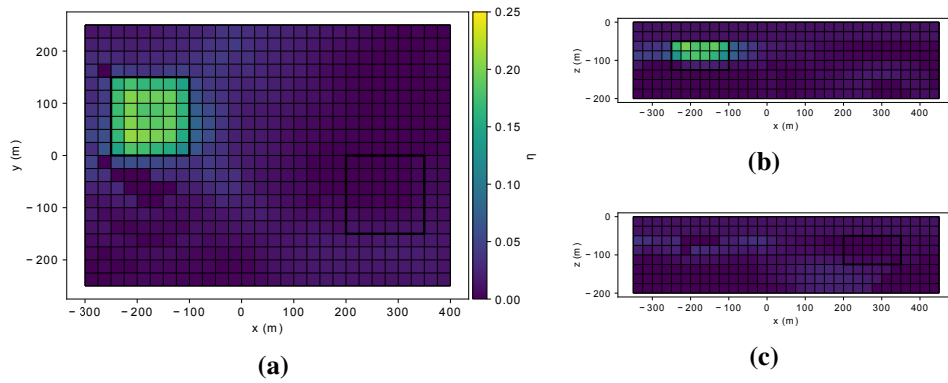


Figure 4.8: Selected slices of recovered chargeability model from stage 4 mid to late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\eta = 0.25$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.

The recovery of η is improved. The true depth extent is still not captured but the planview geometry and chargeability of the recovered block are improved. The time constant model is also improved. There are many artefacts outside the chargeable region but this is not a concern. The recovered τ and β models should only be interpreted in regions with non-negligible chargeability. These artefacts could likely be significantly reduced by penalizing spatial inconsistencies between the η , τ , and β models by a method such as cross-gradient. The β model contains artefacts on the boundary of the chargeable region but has good recovery in the centre of the block. I do not know the cause of these artefacts.

Images of the true and predicted data from the stage 4 inversion for the x-directed receivers due to the x-directed transmitter for times 0.02 s and 0.17 s after

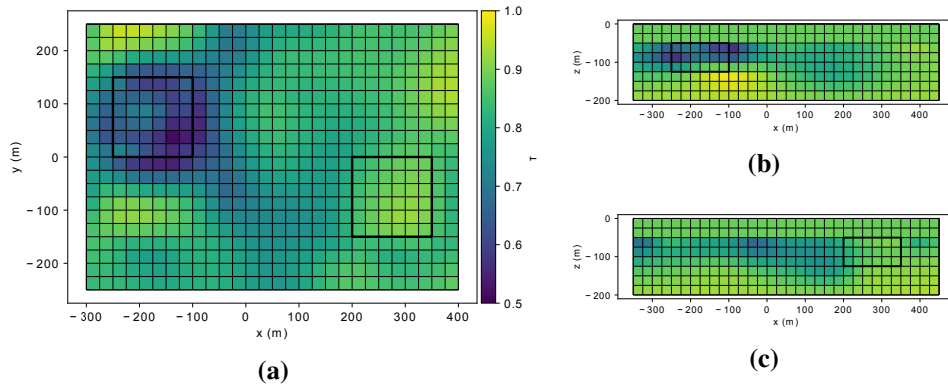


Figure 4.9: Selected slices of recovered time constant model from stage 4 mid to late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\tau = 0.5$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.




Figure 4.10: Selected slices of recovered exponent model from stage 4 mid to late off-time data inversion. Thick black lines show true block boundaries. Left hand block has true $\beta = 0.5$. Right hand block is not chargeable. a) $z = -85$ m b) $y = 75$ m c) $y = -75$ m.

transmitter shutoff are shown in figs. 4.11 and 4.12, respectively. The fit is quite good at early times, as can be seen in fig. 4.11. The fit is not as tight at late times, with the predicted data showing an anomaly that is geometrically equivalent to but stronger than that seen in the true data. The observed and predicted voltage decays at the centre of the chargeable block are shown in fig. 4.13.

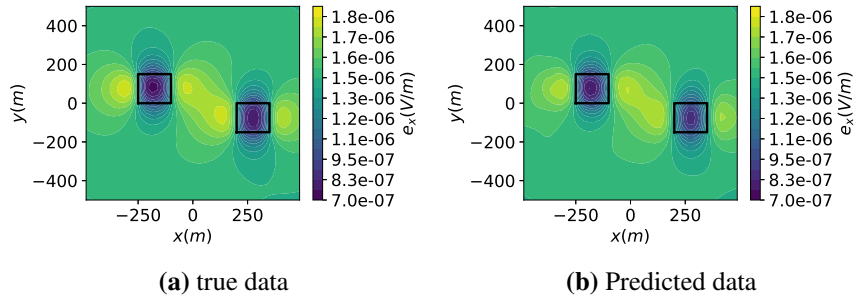


Figure 4.11: Voltage data for the x-directed receivers due to the x-directed transmitter at 0.02 s after source shutoff.

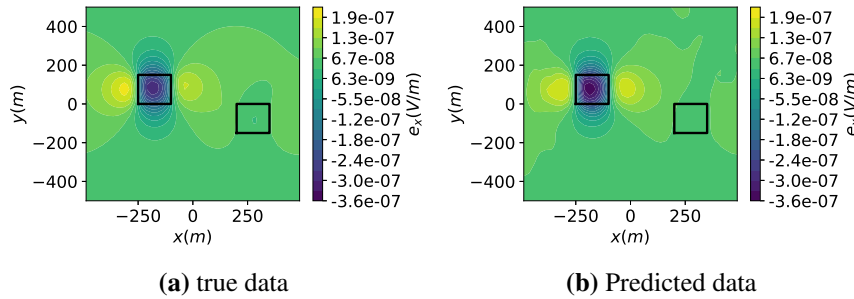


Figure 4.12: Voltage data for the x-directed receivers due to the x-directed transmitter at 0.02 s after source shutoff.

4.4 Conclusions

More work needs to be done in order to improve the stretched exponential inversion workflow but this chapter gives a simple proof of concept example showing the potential of the technique. The ease with which chargeability information was recovered from the transient EM data was encouraging, especially as this was the first work to recover the intrinsic chargeability directly from transient EM data in three-dimensional (3D). More work needs to be done to understand whether it is possible to achieve better recovery of the time parameters τ and β .

A potential approach to reducing the non-uniqueness of the problem is to use parametric methods. The next chapter develops a stochastic parametric inversion algorithm. Although IP is not addressed, the work of the chapter lays the groundwork for applying such an approach to stretched exponential inversion in the future.




Figure 4.13: true and predicted voltage data above the centre of the chargeable block from 0.02 s after source shutoff onward plotted on log scale. Blue line shows observed data and red line shows predicted data. Dashed lines indicate negative values.

Chapter 5

Randomized background stochastic inversion

5.1 Introduction

In many cases, the goal of inverse problems is to estimate the physical properties of relatively simple objects that can be described by small numbers of parameters. Unfortunately, to solve such problems it is often necessary to also estimate a potentially large number of parameters that consist of the background to this model. Such background can be viewed as nuisance parameters. The quality of these nuisance parameter estimates may then adversely affect the recovered estimates of the parameters of interest.

In order to discretize the simple object and estimate its properties, in this paper, we take the view of parametric level-set methods. These methods fall under the general setting of inverse shape reconstruction methods, which seek to estimate the position and shape of an object embedded in some background medium from indirect measurements such as geophysical or medical imaging data. Level set methods have been a popular approach to shape reconstruction inverse problems. See [19] for a survey of early work on the subject and, e.g., [50, 73] for examples of more recent work. In inverse level set methods, the boundary of the object of interest is represented by the zero level-set of some function. The level-set function may be represented and manipulated directly, in which case it must be

represented cell-wise over the computational domain [19]. This is a quite flexible approach but creates a high-dimensional parameter space to optimize over. This can lead to the inverse problem becoming highly ill-posed, requiring methods such as regularization of the level-set function to address the ill-posedness [28].

Alternatively, the level set function can be represented parametrically, using some sort of basis function expansion or geometric shape to represent the object of interest. For example, Aghasi et al. [1] used radial basis functions to reconstruct general 2D shapes. In our approach, we represent the object of interest as a skewed Gaussian ellipsoid. This is based on the method of McMillan et al. [64], who used Gaussian ellipsoids to represent high contrast subsurface bodies in electromagnetic geophysical inversions.

In the case of geophysical imaging, by which we are primarily motivated, the background medium in which the object of interest is hosted is usually heterogeneous, and for an accurate reconstruction of the object, good prior knowledge of the background is necessary. Typically, the object of interest exhibits a high contrast with the background but its response cannot be completely decoupled from the background response. Therefore, incorrect background estimation will introduce error into the estimate of the target object. A small amount of work has been done on quantification of uncertainties in shape reconstruction (see e.g. [21, 50]) but to our knowledge no work has specifically examined or tried to mitigate the effect of uncertainty in background estimation.

There are important practical applications in which it is reasonable to assume that one might have some reliable but imperfect background information based on data other than that used in the shape reconstruction problem. An example from geophysical imaging is time-lapse inversion of hydraulic fractures in the earth's subsurface. The physical properties of the background rocks before fracturing may be imaged by electromagnetic or seismic geophysical surveying. Then, after injection of conductive fluid into the earth during the hydraulic fracturing process it is possible to estimate the volume of rock effectively fractured by electromagnetic imaging, modelling the fractured zone as a thin conductive rectangular prism embedded in the previously estimated heterogeneous background medium [48]. In both these examples there will be uncertainty in the estimate of the background used in the shape reconstruction.

In shape reconstruction with uncertain background information, the parameters describing the background model are effectively nuisance parameters. They are not of direct interest but the accuracy of their estimation affects the accuracy of the shape reconstruction. To our knowledge, there has been very little work done in general on estimating large space nuisance parameters in inverse problems. The one notable example of which we are aware, which is very different in character from the present contribution, is [5]. They use variable projection to eliminate tuning parameters from regularized least squares inverse problems. Their technique only applies when the number of nuisance parameters is very small, relative to the number of parameters of interest. We are interested in the opposite regime, where the object of interest is described by a small number of parameters and the background is a cell by cell discretization of a continuous function onto some sort of computational mesh, making the number of nuisance parameters very large.

In this paper we present a stochastic optimization based method for improving the quality of shape recovery in parametric level-set shape reconstruction problems in the presence of uncertainty in the background parameters. In section 5.2 we will describe the method. In section 5.3 we present a simple linear inverse problem that we use as a first application. This example illustrates the properties of our algorithm and examine the relative strengths and weaknesses of the two optimization techniques we propose for its solution. Section 5.4 discusses an application to a non-linear three-dimensional inverse problem, before concluding and summarizing the paper.

5.2 Method

In general, we are concerned with the inverse problem of estimating the location, shape, and material properties of a homogeneous object embedded in a heterogeneous background medium in two-dimensional (2D) or three-dimensional (3D) space from a noisy set of indirect observations $\mathbf{d} \in \mathbb{R}^d$. We assume that the data are sensitive to some physical property θ of the material being probed. The data are related to θ through the forward problem

$$\mathbf{d} = \mathcal{F}(\theta) + \varepsilon, \tag{5.1}$$

where ε is additive noise and \mathcal{F} is the forward modelling operator, which is typically a functional of the solution of a partial differential equation (PDE).

In level-set shape reconstruction methods, the boundary of the object of interest is represented by the zero level-set of a function $\tau(\mathbf{x})$. The overall physical property model can then be written

$$\theta(\mathbf{x}) = \theta_0(\mathbf{x}) + H(\tau(\mathbf{x}))(\theta_1 - \theta_0(\mathbf{x})), \quad (5.2)$$

where H is the Heaviside step function θ_1 is the value of θ inside the object of interest and θ_0 is the heterogeneous background physical property model. In traditional level-set methods, the level-set function is discretized onto a grid and the inverse problem attempts to reconstruct it directly.

In parametric methods, the level-set function is represented with a low dimensional parametrization. For example, in the application problem in section 5.3 of this paper we model the object of interest as an ellipse, which can be described by its centre position, radii, and angle of orientation. The level set function for this model can be written

$$\tau(\mathbf{x}) = 1 - (\mathbf{x} - \mathbf{x}_0)^T \mathbf{M}(\mathbf{x} - \mathbf{x}_0), \quad (5.3)$$

where \mathbf{x}_0 is the the location of the centre of the ellipse and \mathbf{M} is a rotation and scaling matrix that depends on the radii of the ellipse and its angle of rotation about the horizontal axis of the mesh.

The goal of the parametric inverse problem is to estimate the shape and location parameters that describe the object of interest as well as the value of the diagnostic physical property inside the object. It should be noted here that in practice, in both traditional and parametric level-set methods, the Heaviside function in eq. (5.2) is replaced by a smooth approximation in order to facilitate optimization. This will be discussed further in the application sections.

We formulate the inverse problem as the least squares optimization problem

$$\begin{aligned} \underset{\mathbf{m}}{\text{minimize}} \quad & \Phi(\mathbf{m}) = \frac{1}{2} \|\mathbf{W}(\mathcal{F}(\mathbf{m}) - \mathbf{d})\|_2^2 \\ \text{subject to} \quad & \mathbf{m}_l < \mathbf{m} < \mathbf{m}_h. \end{aligned} \quad (5.4)$$

We seek the parametric model \mathbf{m} that minimizes the discrepancy between measured data and simulated data computed by solving the forward problem—subject to bound constraints on the parameter values. \mathbf{W} in eq. (5.4) is a diagonal matrix containing the inverse standard deviations of the data. We assume that our parametrization of the level-set function is sufficiently simple that it will have a regularizing effect and that therefore explicit regularization will not be required. This is a well known and well studied approach—see e.g. [1, 58, 64, 83]. The optimization problem eq. (5.4) can be solved by standard techniques. We have had success in solving such problems with the projected Gauss-Newton method, solving for the search direction using the preconditioned conjugate gradient (PCG) algorithm with large stopping tolerance.

Of course, perfect knowledge of the background is unlikely. To account for uncertainty in knowledge of the background, we model it as a random field. Then, since the objective function in eq. (5.4) is a function of the background model, we treat it as a random variable. This recasts the problem of estimating the model parameters of interest \mathbf{m} as a stochastic optimization problem. Rather than minimizing the objective function for a fixed background, we seek the \mathbf{m} that minimizes the expected value of the objective function with respect to the stochastic background model.

Let the background physical property model be a random field $\theta_0(\mathbf{x}, \omega)$, where ω is an element of the sample space Ω of a probability space (Ω, F, P) . For fixed ω (i.e. a realization of the random field), θ_0 is a deterministic function of space. We can write the expected value of our now random least squares objective function as a Lebesgue integral with respect to the probability measure P

$$\mathbb{E}[\Phi(\mathbf{m})] = \int \Phi(\mathbf{m}) dP(\theta_0). \quad (5.5)$$

We assume that P is continuous and that we can efficiently compute realizations of the background. Having defined the expectation we can define our stochastic shape-reconstruction problem

$$\begin{aligned} & \underset{\mathbf{m}}{\text{minimize}} && \mathbb{E}[\Phi(\mathbf{m})] \\ & \text{subject to} && \mathbf{m}_l < \mathbf{m} < \mathbf{m}_h. \end{aligned} \quad (5.6)$$

Stochastic optimization has been widely used in operations research for performing optimization under uncertainty in the formulation of or input to optimization problems—see e.g. [27, 95]. To our knowledge, it has not been used for the solution of inverse problems. A version of stochastic optimization has been used to reduce the computational complexity of solving large scale inverse problems involving many PDEs—e.g. [42, 78]. We are not aware of its use in improving the quality of the solution for an inverse problem.

The high-dimensional integral in the expectation eq. (5.5) cannot be evaluated analytically for most problems and so it must be approximated. We use the sample average approximation (SAA) to estimate this expectation and form a tractable optimization problem. The SAA approximates the expectation using a sample average [86],

$$\mathbb{E}[\Phi(\mathbf{m})] \approx \mathbb{E}_n[\Phi(\mathbf{m})] = \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{m}, \theta_0^{(i)}), \quad (5.7)$$

where $\Phi(\mathbf{m}, \theta_0^{(i)})$ denotes the misfit function evaluated with respect to a particular realization of the background model $\theta_0^{(i)}$. Applying this approximation to eq. (5.6) gives the optimization problem

$$\underset{\mathbf{m}}{\text{minimize}} \quad \mathbb{E}_n[\Phi(\mathbf{m})] = \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{m}, \theta_0^{(i)}). \quad (5.8)$$

Intuitively, the misfit function is now a sample average of the individual misfits computed with different realizations of the background. The variance in the estimate of the parameters of interest will decrease as the number of background samples increases.

Let \mathbf{m}_n^* be the minimizer of the SAA problem eq. (5.8) and \mathbf{m}^* the minimizer of the true problem. The theory of stochastic programming tells us [86] that \mathbf{m}_n^* converges in probability to \mathbf{m}^* such that

$$\|\mathbf{m}_n^* - \mathbf{m}^*\| = \mathcal{O}(1/\sqrt{n}). \quad (5.9)$$

Despite the slow theoretical convergence rate, we have found that in practice we can improve results in shape reconstruction problems using a number of samples that is computationally tractable.

Given a set of background samples $\{\theta_0^{(i)}\}$, eq. (5.8) may be solved using methods of deterministic optimization. We have used the Gauss-Newton (GN) and accelerated mini-batch stochastic gradient descent (SGD) algorithms. In the next section we will analyze the performance of these two optimization approaches on a simple linear inverse problem. We note that, computationally, each background sample requires separate forward problem, gradient, and possibly Hessian vector product evaluations at each iteration of the optimization procedure but that the computations for each background are trivially parallel. Looking ahead, we have found that given sufficient computing resources, the GN approach provides an efficient and robust solution while SGD can still give good results while using limited computational resources.

5.2.1 Computing Background Samples

Before turning to an analysis of the performance of the SAA approach on example problems, we consider the problem of choosing background model probability distributions and generating sample background models. Ideally the distributions would be determined empirically, based on real examples from past experiments. This is not typically possible in geophysical imaging. As an alternative, we model the background as a Gaussian random field (GRF) with mean and covariance functions that encode our prior knowledge of the background. The stochastic optimization methods used in this paper do not depend on the background distribution being Gaussian. We use GRFs in our examples because of their flexibility and the simplicity of sampling from them.

We use a truncated Karhunen-Loève expansion (KLE) to approximately sample from GRFs. The KLE provides a way to represent a random field via a spectral expansion of its covariance function [35]. Consider a Gaussian random field $m_0(\mathbf{x}, \omega)$ with mean \bar{m}_0 and covariance function $\kappa(\mathbf{x}, \mathbf{y})$, where $\mathbf{x} \in \mathbf{D} \subseteq \mathbb{R}^s$ is the spatial variable and ω is an element of the sample space in some probability space (Ω, F, P) . The KLE of m_0 is

$$m_0(\mathbf{x}, \vartheta) = \bar{m}_0 + \sum_{n=0}^{\infty} \xi_n(\omega) \sqrt{\lambda_n} \phi_n(\mathbf{x}), \quad (5.10)$$

where the λ_n and φ_n are the eigenvalues and eigenfunctions of the covariance function and the $\xi_n(\omega)$ are normally distributed random variables. For covariance function $\kappa(\mathbf{x}, \mathbf{y})$, the eigenvalues and eigenfunctions solve the eigenvalue problem

$$\int_{\Omega} \kappa(\mathbf{x}, \mathbf{y}) \varphi_n(\mathbf{y}) \, d\mathbf{y} = \lambda_n \varphi_n(\mathbf{x}). \quad (5.11)$$

In practice, except for special covariance functions, this must be solved approximately. We use the Nystrom method [14] to approximate the KLE eigenvalues and eigenfunctions. This method approximates the integral eq. (5.11) by a quadrature rule. This leads to an $n_\lambda \times n_\lambda$ matrix eigenvalue problem that gives the first n_λ eigenvalues and the values of the corresponding eigenfunctions at the quadrature points. The eigenfunctions can then be evaluated at points of interest by interpolation. See [14] for a detailed description. For smaller problems the quadrature points may be evenly distributed over the domain and the eigenfunctions evaluated for each cell in the computational domain. However, for larger 3D problems, distribution of quadrature points over the full domain becomes prohibitively expensive. We address this problem in the discussion of our non-linear 3D application problem in section 5.4.

5.2.2 Implementation

The software required to perform the numerical experiments in the remainder of this paper was developed in the Julia programming language [15], using the jInv framework [80]. jInv is a set of open source Julia software packages designed to simplify the process of building software for PDE constrained parameter estimation problems by providing modular building block routines commonly needed in these problems, including parallelized Gauss-Newton optimization and tools for finite volume PDE discretization. We used these building blocks to develop routines to solve our specific applications problems. To perform the Karhunen-Loève basis function computations we used the open source Julia package GaussianRandomFields [77], modified to our needs.




Figure 5.1: Cartoon schematic of straight ray tomography experiment showing straight ray paths from the transmitters to the receivers.

5.3 Application to a Simple Model Problem

We use two-dimensional (2D) straight ray cross-hole seismic tomography as an example problem that is computationally simple enough to permit extensive experimentation. Straight ray tomography is a linearization of seismic tomography in which seismic waves (pressure waves moving through the earth’s subsurface) are assumed to propagate in straight lines from the source of an excitation to a receiver point [18]. Transmitters are placed in a borehole in the earth and receivers in a separate borehole. The experiment then measures the times required for the signal to propagate from each transmitter to each receiver. The speed of wave propagation depends on the material properties of the propagation medium, allowing seismic tomography to be used as an imaging method. See fig. 5.1 for an illustration of a typical survey setup.

The straight-ray tomography forward problem is

$$\mathbf{t} = \mathbf{A}\mathbf{s}. \quad (5.12)$$

The data \mathbf{t} are the travel times between each source and receiver pair. \mathbf{s} is the inverse velocity, or slowness of each mesh cell, and each entry of \mathbf{A} , a_{ij} , gives the distance travelled by the straight ray through cell j along the path connecting transmitter-receiver pair i .

We attempt to recover an elliptical anomaly embedded in a heterogeneous background from travel time data. Our synthetic experiment is performed on a unit domain in \mathbb{R}^2 , with 50 transmitters spaced evenly along the left edge of the domain and 50 receivers spaced evenly along the right edge.

The background slowness model is described by a GRF with mean and covariance based on the geological scenario of a horizontally layered earth with seismic wave velocity increasing with depth—a scenario commonly encountered in seismic imaging applications [66]. The mean slowness model is composed of five homogeneous layers. It is shown in fig. 5.2a. The covariance function is

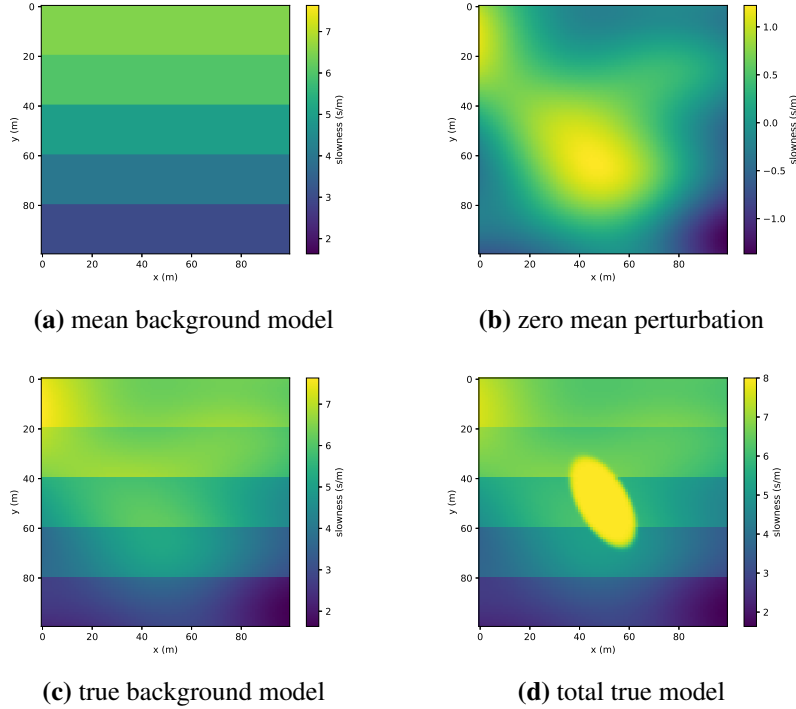


Figure 5.2: True model used for straight ray tomography inversions. The mean model a) consists of homogeneous horizontal layers. The random perturbation shown in b) is added to the mean model to form c), used as the background in the true model shown in d).

$$\kappa(\mathbf{x}, \mathbf{y}) = \sigma \exp(-(\mathbf{x} - \mathbf{y})^T A (\mathbf{x} - \mathbf{y})), \quad (5.13)$$

with

$$A = \begin{pmatrix} l_x^{-1} & 0 \\ 0 & l_y^{-1} \end{pmatrix}, \quad (5.14)$$

for characteristic length scales l_x and l_y . σ is the standard deviation, which is constant, giving the overall scale of fluctuations from the mean background. In this example we use $\sigma = 1$, $l_x = 0.3$ and $l_y = 0.1$, encoding our assumption that the slowness should vary more smoothly in the horizontal direction than the vertical.

The slowness model is completed by adding a homogeneous ellipse to the background model. The boundary of the ellipse is given by the parametric level-set function

$$\tau(\mathbf{x}) = 1 - (\mathbf{x} - \mathbf{x}_0)^T \mathbf{M}(\mathbf{x} - \mathbf{x}_0), \quad (5.15)$$

where \mathbf{x}_0 is the the location of the centre of the ellipse and \mathbf{M} is a rotation and scaling matrix that depends on the radii of the ellipse and its angle of rotation about the horizontal axis of the mesh. For more details on the level-set parametrization, see [64]. In order to use gradient based optimization techniques, the slowness model must be differentiable with respect ellipse parameters across the entire domain. To achieve this we use a smooth approximation to the Heaviside function to describe the transition from the ellipse to the background medium. The total slowness anywhere in the domain is given (for fixed background slowness model s_0) by

$$s(\mathbf{x}, s_0) = s_0 + \frac{1}{2}(1 + \tanh(a\tau(\mathbf{x})))(s_e - s_0(\mathbf{x})), \quad (5.16)$$

where s_e is the slowness inside the ellipse and a is a scaling factor that controls the sharpness of the transition between the ellipse and the background.

The model used to generate synthetic data for our inversion consisted of an ellipse in the centre of the domain embedded in a randomly chosen slowness model from the background GRF. We discretized the domain onto a 100×100 regular grid of square cells. A 176 term KLE was used to compute a sample from the zero mean GRF with covariance eq. (5.13). This was then added to the mean background to generate a sample from the background model GRF. The zero mean perturbation and resulting background sample are shown in figs. 5.2b and 5.2c, respectively. To this background model we added an ellipse with an anomalously

high slowness of 8 s/m. The ellipse was centred at the centre of the domain with an angle of 30° from vertical and radii of 0.1 m and 0.2 m. This total model is shown in fig. 5.2d. In the rest of this section this will be referred to as the true model. Using a random sample from the background GRF, rather than the mean background as the true model in our inversion study implies that the sample backgrounds used during inversion are being drawn from a distribution with the correct covariance function but wrong mean. In practice both the mean and covariance could only be approximately known. We have not attempted to assess the success of our method when the background covariance function is only approximately known.

Generally, the accuracy of the background model mean and covariance functions required for our SAA technique to be effective in practice will depend on specific characteristics of the application problem, and in particular, on how well separated the material properties of the object of interest are from the background. When the vast majority of the signal in the data comes from the object of interest, then adequate shape reconstruction might be possible without particularly precise characterization of the background. On the opposite end of the spectrum, if the background generates signals similar to that of the object of interest, then our approach will likely fail and shape reconstruction will only be possible with a very accurate description of the background. It is in the middle ground where we believe our method can be most useful.

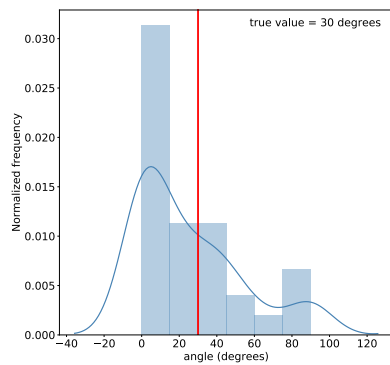
5.3.1 SAA inversion

We conducted several numerical experiments to study the behaviour of our method. All these experiments attempted to fit the same synthetic data, which was generated by applying the straight ray tomography forward modelling operator to the true model and then adding artificial independent Gaussian noise to each measurement. The added noise on each datum had standard deviation equal to 1% of its magnitude. We attempted to recover the orientation angle, centre position and slowness of the anomalous ellipse. For simplicity we did not attempt to recover the radii of the ellipse. They were fixed to their true values. The starting guesses and bounds on the active inversion parameters were the same in all inversions. They are listed in table 5.1.


parameter	starting guess	lower bound	upper bound
ϕ	1.5	0	90
x	0.2	0.1	0.9
y	0.4	0.1	0.9
s	4	0.01	15

Table 5.1: Starting guesses and parameter bounds for straight ray tomography inversion parameters. ϕ is the ellipse orientation angle in $^\circ$, x and y are the horizontal and vertical coordinates of the centre position, and s the slowness inside the ellipse.


Before studying the effectiveness of the SAA on the straight ray example problem, we wanted to characterize the variability of results in single background inversions caused by uncertainty in background estimation. Assuming our background model random field accurately captures the uncertainty in our knowledge of the background, this can be done by conducting many inversions, each using a single random background. We performed 100 such inversions. For each we solved the deterministic optimization problem eq. (5.4) using the projected Gauss-Newton algorithm with backtracking Armijo line-search. Inversions could end after attaining the desired data fit as determined by a χ^2 statistic, when the magnitude of the gradient was reduced below a threshold value relative to its initial magnitude, or when the line search failed to find a step that reduced the objective function. Results from all these categories were included in assessing the variance in the recovered parameters as a function of the background. The results are summarized in fig. 5.3, which shows histograms and kernel density estimates of the recovered values of each inversion parameter. There were a wide range of results. Excellent recovery was achieved for a few sample backgrounds but the variance in results was large and many inversions achieved very poor recovery of the anomalous ellipse. Orientation angle was particularly poorly resolved, with the angle being pushed to its lower bound being the mode result. This seems to show that the angle is poorly resolved by this particular straight ray tomography survey, rather than to give insight into the effect of uncertainty in the background model. Overall, these results show that the level of uncertainty in our knowledge of the background slowness model is too high for parametric inversion using a fixed background to be reliably effective




(a) rotation angle



(b) horizontal position



(c) vertical position



(d) slowness

Figure 5.3: Results of 100 independent single background inversions. The recovered values for each inversion parameter are plotted as histograms with accompanying kernel density estimates. The red vertical lines show the true parameter values. Note that the estimated densities may have support outside the parameter bounds. They are included in order to give general visual impressions of the spread of inversion results and these effects are not of concern.

n_s	1	10	100
ϕ	25.2	64.0	88.8
x	0.38	0.50	0.51
y	0.51	0.53	0.54
s	6.96	7.97	8.36

(a) sample mean

n_s	1	10	100
ϕ	28.2	31.2	8.29
x	0.17	0.051	0.0091
y	0.16	0.040	0.0073
s	3.03	1.11	0.286


(b) sample standard deviation

Table 5.2: Marginal statistics of SAA inversion results. The sample means of the recovered values of each parameter are shown for the 1, 10, and 100 background inversion experiments in a), with the corresponding sample standard deviations shown in b). Recall that the true parameter values are $\phi = 30^\circ$, $x = 0.5$, $y = 0.5$, and $s = 4$. See table 5.1 for parameter bounds and initial guesses.


for this particular problem.

To show the effect of the SAA approach in improving ellipse recovery in the presence of uncertainty in background estimation, we performed repeated multiple background SAA inversions and studied the variability of the inversion results. For each of two SAA sample sizes $n_s = 10, 100$ we performed 100 inversions, with n_s new independent background models sampled from the background GRF. Kernel density plots showing the range of recovered parameter values for the SAA and single background inversions are plotted together in fig. 5.4. It is clear from these results that the SAA approach did not help recovery of the orientation angle—in line with our expectations—but that it did a great deal to reduce the variance in the recovered values of the other parameters. The density plots also show that there are small but distinct biases in the distributions of recovered values. This stems from the fact that we don't know the true background model distribution and are thus not drawing our sample backgrounds from the correct distribution. Thus as the SAA sample size is increased we see the sample means of the recovered parameters converging to values slightly different than the true values.


The sample means and standard deviations of the recovered parameter values for each SAA sample size are shown in table 5.2. The values in the table show the $\mathcal{O}(1/\sqrt{n_s})$ convergence in standard deviation expected from the SAA. Overall these statistics show that a single SAA inversion with a large enough sample size can be expected to give much better parameter recovery than a single one back-




(a) rotation angle



(b) horizontal position



(c) vertical position



(d) slowness

Figure 5.4: Kernel density plots showing inversion results for single background, 10 background SAA, and 100 background SAA inversions. 100 background angle results are not shown because they took the upper bound value of 90° for all but two of these inversions rendering the corresponding kernel density plot meaningless. The vertical axes of the plots were clipped in order to show the single and 10 background results at reasonable scales.

ground inversion.

5.3.2 Optimization by stochastic gradient descent

A major downside of the SAA method described in this paper, when using traditional deterministic optimization methods, is that the computational cost grows linearly with the number of samples n_s , while the parameter estimates converge only at a rate of $\mathcal{O}(1/\sqrt{n_s})$. Each term in the SAA objective function requires separate forward modelling, gradient and hessian computations. When large scale parallel computing resources are available, a large number of terms can be handled in parallel, allowing the approach to remain efficient. But when such resources are not available, the computational cost per iteration becomes prohibitive as the number of samples increases.

In order to gain the statistical benefit of using a large number of samples in the SAA technique without incurring a prohibitively large computational cost, we turn to stochastic gradient methods. These algorithms converge much more slowly than the Gauss-Newton method but the per-iteration cost is drastically reduced and they can perform SAA inversions with large numbers of background samples when only modest computational resources are available since only a small number of background models are accessed at each iteration.

Stochastic gradient descent (SGD) and related algorithms such as Polyak’s momentum method [72], have been used in stochastic programming for a long time [86]. Recently, these methods, along with other variants, have gained great popularity in the machine learning community due to their ability to work with small subsets of large datasets and their surprising performance on difficult non-convex machine learning optimization problems—see e.g. [17, 57]. We are interested in SGD and related methods because of their ability to use information from a large number of sample background models without needing to process many of these models at each iteration.

We use the ADAM accelerated stochastic gradient descent algorithm [55]. The ADAM method combines ideas from momentum methods and gradient descent with adaptive step size control. The method is shown in algorithm 2. The model is updated by an exponentially weighted average of past gradients, with the indi-

Algorithm 2 The ADAM method, algorithm presentation adapted from algorithm 1 of [55]

Choose hyperparameters $\alpha, \beta_1, \beta_2, \varepsilon$.
Choose initial model \mathbf{m}_0 (Abuse of notation, do not confuse with background model)
Initialize $\mathbf{v}_0 = 0, \mathbf{w}_0 = 0, k = 0$
Draw a set $\Xi = \{\xi_i\}$ of n_s samples from the background model distribution
while not converged **do**
 Compute a random permutation of the samples and divide into mini-batches
 for j in mini-batches **do**
 $k \leftarrow k + 1$
 $\mathbf{g} \leftarrow \Phi[\mathbf{m}, \Xi(j)]$ (gradient of objective function w.r.t. samples in mini-batch j).
 $\mathbf{v}_k \leftarrow (1 - \beta_1^t)^{-1}(\beta_1 \mathbf{v}_{k-1} + (1 - \beta_1)\mathbf{g})$
 $\mathbf{w}_k \leftarrow (1 - \beta_2^t)^{-1}(\beta_2 \mathbf{w}_{k-1} + (1 - \beta_2)(\mathbf{g} \circ \mathbf{g}))$
 $\mathbf{m}_k \leftarrow \mathbf{m}_{k-1} - \alpha \mathbf{v}_k / (\sqrt{\mathbf{w}_k} + \varepsilon)$
 end for
end while
return \mathbf{m}_k

vidual components scaled by an exponentially weighted average of the magnitudes of past gradients. After reviewing several accelerated gradient descent algorithms, we chose ADAM because of its proven performance on many difficult large-scale problems and the fact that in our experience it seems to perform well with less hyperparameter tuning than other related methods. In order to achieve satisfactory performance it was still necessary to tune the step-size. After some experimentation, we settled on a value of $\alpha = 0.002$, which we used in the experiments discussed below. We set the other parameters to the default values recommended in [55]. These are $\beta_1 = 0.9$, $\beta_2 = 0.99$, and $\varepsilon = 1 \times 10^{-8}$.

At each iteration of the optimization, the approximate gradient of the SAA objective function is computed using a subset of the available background samples, called a mini-batch. We studied the performance of the ADAM method in solving the SAA optimization problem eq. (5.8) for the straight ray tomography problem by studying the variability of the recovery of ellipsoid parameters obtained with independent sets of background samples, as we did with GN. We studied the per-

n_s	100	500
ϕ	24.1	26.3
x	0.49	0.50
y	0.51	0.52
s	9.82	9.77

(a) sample mean

n_s	100	500
ϕ	1.10	0.515
x	0.0068	0.0038
y	0.0077	0.0038
s	0.288	0.139

(b) sample standard deviation

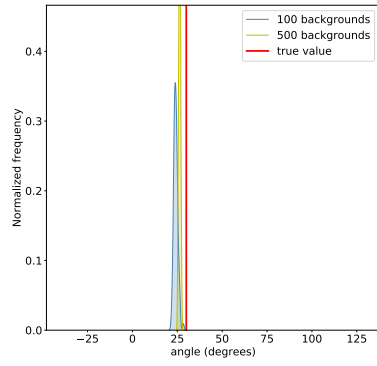
Table 5.3: Marginal statistics of SAA inversion results using the ADAM optimization algorithm. Recall that the true parameter values are $\phi = 30^\circ$, $x = 0.5$, $y = 0.5$, and $s = 4$. See table 5.1 for parameter bounds and initial guesses.

formance of ADAM using SAA sample sizes of $n_s = 100, 500$, with mini-batch sizes of 1 and 10. We performed 100 inversions in each of these configurations. The optimizations generally converged faster with the larger mini-batch size, as expected. The parameter recovery was equivalent in both cases. The summary statistics of the results (computed from the 1 mini-batch runs) are shown in table 5.3. Kernel density estimates for each parameter value, for both SAA sample sizes are plotted in fig. 5.5. The standard deviations of the ellipsoid position and slowness from the 100 background inversions was very close to what was observed in the 100 background GN experiments. Unexpectedly, the orientation angle was much better resolved by the ADAM inversions. The biases of the x and y position variables were similar to the GN results while the ADAM inversions tended to overestimate the slowness more than the GN inversions.


We cannot definitively say why the biases in the ADAM results were different from the GN biases. It may be because the ADAM algorithm is better able to escape from local optima and saddle surfaces, or that because of its much smaller step sizes and the scaling of the problem, ADAM is able to better explore the interior of the parameter space and find reasonable values for the position and slowness parameters without quickly pushing the angle to its upper bound value.

5.4 Application to non-linear 3D problem


We illustrate the application of our method to a larger scale non-linear problem using another example from geophysics, direct current resistivity (DCR) imaging.




(a) rotation angle



(b) horizontal position



(c) vertical position



(d) slowness

Figure 5.5: Kernel density plots of inversion results for 100 background and 500 background ADAM inversions with mini-batch size 1. Horizontal axis limits are the same as in fig. 5.4

In DCR imaging, a steady electric current is injected into the ground and resulting electrical potential differences due to electric currents flowing in the ground are measured at a small number of locations on the earth's surface. These voltages can be used to estimate the electrical conductivity (inverse of resistivity) of the earth's subsurface. Mathematically, DCR surveys are governed by the elliptic boundary value problem

$$\nabla \cdot (\theta \nabla \phi) = -\nabla \cdot \mathbf{j}_s \in (\Omega \subset \mathbb{R}^3) \quad (5.17)$$

$$\nabla \phi \cdot \hat{\mathbf{n}} \in \partial\Omega_1 \quad (5.18)$$

$$\phi = 0 \in \partial\Omega_2 \quad (5.19)$$

where ϕ is the electrical potential in the earth, θ the conductivity, here assumed to be a scalar function of space, and \mathbf{j}_s , the electric current driving the experiment, which is divergence free except at a finite number of isolated points where the source of current is connected to the ground. $\partial\Omega_1$ denotes the portion of the domain boundary that represents the surface of the earth and $\hat{\mathbf{n}}$ is the unit vector normal to $\partial\Omega$. $\partial\Omega_2$ is the rest of the domain boundary. The Dirichlet conditions model the fact that the potential tends toward zero far from sources of current. In order to approximately meet this condition in a finite domain, we must make the domain large enough that the non-surface boundaries are indeed far from any sources of current. We discretize eq. (5.19) on locally refined rectilinear meshes called OcTrees (see e.g. [45]), which makes it simple to enlarge the domain without significantly increasing the computational cost of solving the boundary value problem.

The potential difference measurements are linear functionals of the potential ϕ , a measurement being simply the difference in the potential between two locations on the earth's surface. The goal of this example problem is to use these measurements to characterize a plate-like highly conductive body embedded in a heterogeneous vertically layered background medium. We model the conductive body as a rectangular prism. The prism is incorporated with the background conductivity model using the parametric level-set approach. Similarly to the previous

example the overall conductivity model is

$$\theta(\mathbf{x}) = \theta_0(\mathbf{x}) + \tilde{H}(\mathbf{x})(\theta_p - \theta_0(\mathbf{x})), \quad (5.20)$$

where θ_0 is the heterogeneous background conductivity model, θ_p is the conductivity of the prism, and \tilde{H} is a smooth approximation to an indicator function that takes value 1 inside the prism and 0 outside. The indicator function is a product of one-dimensional (1D) indicator functions in the three Cartesian coordinate directions

$$\tilde{H}(\mathbf{x}) = \sigma_x(\mathbf{R}(\mathbf{x} - \mathbf{x}_0))\sigma_y(\mathbf{R}(\mathbf{x} - \mathbf{x}_0))\sigma_z(\mathbf{R}(\mathbf{x} - \mathbf{x}_0)), \quad (5.21)$$

where \mathbf{x}_0 is the position of the centre of the prism and \mathbf{R} is a rotation matrix that transforms $\mathbf{x} - \mathbf{x}_0$ to a coordinate frame whose axes are the principal axes of the prism. The 1D indicator function for coordinate i is

$$\sigma_i(\mathbf{r}) = \frac{1}{1 + \exp(-(r_i + h_i)/a_i)} - \frac{1}{1 + \exp(-(r_i - h_i)/a_i)} \quad (5.22)$$

where r_i is the i component of \mathbf{r} , h_i is the side length of the prism in the i direction and a_i the scaling factor that controls the sharpness of the prism boundary transition in the i direction.

The background conductivity θ_0 is modelled as a log-normal random field. The mean of the underlying GRF is constant in the y direction and piecewise constant in the x and z directions. It is illustrated in fig. 5.6 by a slice in the x - z plane showing the core region of the model. The covariance function of $\log(\theta_0)$ is

$$\kappa(\mathbf{x}, \mathbf{y}) = \sigma(\mathbf{x})\sigma(\mathbf{y}) \exp(-(\mathbf{x} - \mathbf{y})^T A(\mathbf{x} - \mathbf{y})), \quad (5.23)$$

where

$$A = \begin{pmatrix} l_x^{-1} & 0 & 0 \\ 0 & l_y^{-1} & 0 \\ 0 & 0 & l_z^{-1} \end{pmatrix}, \quad (5.24)$$

with $l_x = l_y = 2.5l_z$. The standard deviation $\sigma(\mathbf{x})$ is piecewise constant, being 50% of the mean value in each constant region of the mean model.

We take approximate samples of the $\log(\theta_0)$ GRF by using the Karhunen-





Figure 5.6: Mean background conductivity in the x - z plane at $y = 0$ in the core region of the mesh.

Loève expansion. In order to make the computation of the Karhunen-Loève eigenfunctions tractable and to make sampling efficient, we sample the background on a subset of the full computational domain. We use a large enough volume such that the lack of randomization outside it does not significantly effect the data. We then simply take the exponential of these GRF samples to get samples from the background random field θ_0 .


The 3D conductivity model used to generate synthetic data for our experiment (which we'll call the true model below) was taken to be a thin dipping rectangular prism embedded in a random sample from the background field. It is illustrated by orthogonal 2D slices in fig. 5.7.

With our computational resources, using the Gauss-Newton SAA inversion technique with many sample backgrounds was computationally intractable for this problem. However, we found that the stochastic gradient technique introduced in the previous section worked well. As in the last section we performed a set of inversions using single random samples from the background distribution. Due to the computational cost, we did not perform a series of SAA inversions. We performed a single SAA inversion using the ADAM stochastic optimization algorithm, with 400 background models and a mini-batch size of 4. This mini-batch size was chosen simply because this was the number of model realizations that we could process concurrently.


In this case the anomalous parametric body was better separated from the background than in the 2D example. Because of this, the single background DCR inver-



(a) Slice through model at $x = -25$ m



(b) Slice through model at $y = 0$ m



(c) Slice through model at $z = -60$ m

Figure 5.7: True model slices

sions gave better results than in the 2D tomography case. However, there was still a significant variability in the results. This is illustrated in fig. 5.8, which shows the recovered parameter values for the single background and SAA inversions. The single background inversions were able to somewhat reliably recover the position of the anomalous block but showed much larger variability in the recovery of the other parameters. None of the inversions were able to recover the thickness (z extent) well, suggesting that the data may not be very sensitive to this parameter. The single background inversions were able to reduce misfit by tweaking the extents of the block and varying its orientation angle but they struggled to capture its highly conductive nature. In most cases they increased the signal from the block only by increasing its thickness, and not by increasing the conductivity. The SAA inversion still pushed the thickness to its upper bound but, perhaps because it could not find descent directions by fitting small signals from background features, it was better able to capture the essential features of the anomalous block, leading to good recovery of its position, orientation, and conductivity.

5.5 Conclusions

In parametric shape reconstructions problems, estimating the properties of the background medium can be a major stumbling block to accurate reconstruction. We have shown that the effect of uncertainty in the background properties can be mitigated by the use of stochastic optimization methods. This can be thought of as a means to improve the shape recovery possible given a specified characterization of the background or as a justification to limit the effort one might expend on achieving a very accurate representation of the background.

The efficiency of this method could likely be improved in future. In this work we used ADAM, a first order stochastic optimization method, to make our approach tractable for large-scale problems. There may be an opportunity to use Stochastic Newton methods to improve the efficiency of optimization.

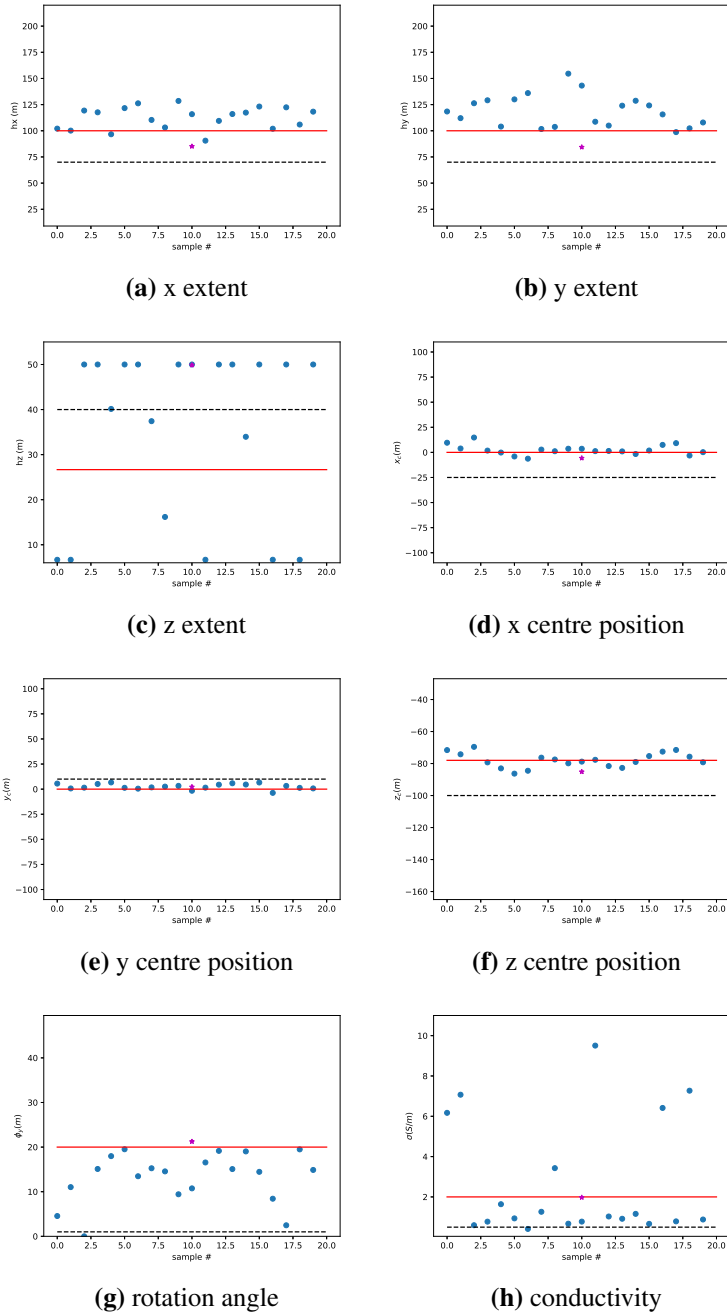


Figure 5.8: Results from single background and SAA ADAM inversions of DC data. Blue dots show the recovered values of each parameter from the 20 single background inversions and the stars show the parameter values recovered by the SAA inversion. In each plot, the vertical axis limits are the parameter bounds, the dashed black line shows the initial guess and the red line the true value of the parameter.

Chapter 6

Conclusions

This thesis has focused on the development of computational methods related to forward modelling and inversion of time-domain geophysical electromagnetic (EM) data in the presence of induced polarization effects. The work led to new, more efficient numerical methods for time discretization of the quasi-static Maxwell equations, a new forward modelling algorithm for simulating coupled EM induction and induced polarization (IP) effects, and a corresponding inversion algorithm for recovering three-dimensional (3D) conductivity and IP parameter earth models from transient EM data. The final main chapter concerned stochastic parametric inversion. It did not address IP directly but contributed a novel inversion algorithm that is well suited for application to IP inversion problems in the future.

This work was motivated by interest in induced polarization in the geophysical community and by a lack of appropriate computational tools for the investigation of coupled EM induction and IP effects. The first attempts to model these coupled effects in 3D directly in the time-domain have emerged only in the last five years, in the work of Marchant et al. [63] and Commer et al. [25]. This work sought to improve on the computational efficiency of these methods and to develop a purely time-domain EM and IP forward modelling algorithm scalable to large industrial problems that is amenable to use in inversion for IP parameters.

Efficient modelling of coupled EM induction and IP effects clearly requires efficient modelling of electromagnetic induction. Therefore, before addressing IP, Chapter 2 of the thesis examined the discretization in time of the quasi-static

Maxwell equations. Time-discretization in geophysical electromagnetics has traditionally used quite simple methods. As a parabolic problem, solutions tend to display simple smoothly decaying behaviour in time but short time-scale effects makes the problem very stiff. In order to maintain stability one is forced to use discretization methods that possess the property of stiff decay. In addition to stability, these methods have the advantage that they can use large time-steps to recover the late time behaviour of a solution without resolving the details of its early time behaviour. This key property allowed the development of the parallel in time forward modelling algorithm discussed in chapter 2. This simple modification eliminates the bottleneck in speeding up forward modelling caused by the inherently sequential nature of time-stepping. A key difference between the parallel time-stepping approach and lower level optimizations such as parallelization of the solution of the linear systems at each time-step is its scalability. The tests performed in chapter 2 showed that the approach could scale to multiple nodes of a computer cluster and achieve significantly faster forward modelling times than traditional methods.

Chapter 3 detailed the development of the coupled EM and IP forward modelling algorithm based on stretched exponential relaxation. The algorithm was implemented using finite volume spatial discretization on OcTree meshes. Backward Euler and second order backward differentiation formula (BDF2) time-stepping schemes were examined. Using the stretched exponential relaxation to model the IP effect was a departure from the Cole-Cole model that has become the de-facto standard model of IP in the geophysical community. These models are both phenomenological. The analysis of chapter 3 showed that they are not fully equivalent but can be used to model similar phenomena and are equivalent in some cases. The examples in chapter 3 showed that the stretched exponential model is capable of simulating electromagnetic fields qualitatively similar to those observed in the field by practicing geophysicists. However, the ability of the stretched exponential approach to model real data remains untested. This is an important area for future work. The stretched exponential forward modelling algorithm and the parallel time-stepping algorithm described in chapter 2 were published in [11].

Ultimately, the usefulness of the stretched exponential approach should be judged by its ability to generate useful earth models that fit experimental transient EM data through the process of geophysical inversion. Chapter 4 showed how

stretched exponential modelling could be used in an inversion algorithm to recover intrinsic IP parameters from grounded source transient EM data. The work was preliminary and tested only on a simple synthetic example. It leaves several open questions on how to best extract information on the IP properties of the earth from transient EM data. Some of these questions can be addressed by further synthetic inversion studies, such as the question of if the non-uniqueness of the inversion can be mitigated through regularization that couples the three IP parameters? The circumstances, if any, under which conductivity and chargeability can be recovered simultaneously, without a good starting conductivity model could also be investigated in synthetic studies? However, any such synthetic studies should be coupled with applications to field data.

A further avenue for mitigating the non-uniqueness in the IP inverse problem is to use parametric methods, which pose the inverse problem as that of recovering a small number of parameters that characterize a homogeneous body embedded in some background medium. This can be particularly useful in the common case where there is a priori reason to believe that the IP signal observed in a geophysical data-set is the result of a relatively homogeneous compact body. This scenario was explored in chapter 5 of the thesis, through the development of a stochastic parametric level-set inversion algorithm. In this class of methods, the domain of investigation into the homogeneous body of interest and a heterogeneous background medium that is held fixed in the inversion. The work in this thesis made two main contributions to this problem. It showed that better recovery of the homogeneous object of interest could be achieved by addressing the uncertainty in the model of the background medium. This was done by modelling the background as a random field and posing the inverse problem as a stochastic programming problem. The initial implementation of this approach was quite computationally intensive. It was shown that the algorithm could be adapted to make it scalable to large, industrial scale problems by employing stochastic gradient descent in the solution of the stochastic programming problem. The work in chapter 5 has been submitted for publication in the *SIAM/ASA Journal on Uncertainty Quantification*.

The stochastic parametric level-set inversion approach could be used to help the stretched exponential inverse problem in multiple ways that could be explored in future work. One potential use is to model the earth conductivity as part of the

background model when it is not possible to get a precise estimate of the conductivity before inverting for IP parameters. This would hopefully reduce the quality of the starting conductivity model needed to produce good models of the IP parameters. One could also use the stochastic approach to model a possibly polarizable background medium that has not been well characterized, when a compact conductive target is of interest.

Overall the problem of inverting transient EM data for IP parameters is still new and under explored. This thesis has developed computational methods that can now be used to investigate the recovery of IP information from real data. IP effects seem to be showing up more commonly in recent inductive source EM data and whether the target of interest in a particular survey is polarizable or if a conductive target is simply masked by IP signals from background media, it is important to understand how the presence of polarizable bodies is manifest in EM data and to incorporate this understanding into the modelling and inversion tools that are used to analyze the data. The study of IP in grounded source electrical data is more established but there is potential for much more information to be extracted from such data through the use of modelling and inversion tools, such as those developed in this thesis, that fully account for the behaviour of these data in time.

Bibliography

- [1] A. Aghasi, M. Kilmer, and E. L. Miller. Parametric Level Set Methods for Inverse Problems. *SIAM Journal on Imaging Sciences*, 4(2):618–650, 2011. doi:10.1137/1008002081.
- [2] F. Alvarez, A. Alegria, and J. Colmenero. Relationship between the time-domain Kohlrausch-Williams-Watts and frequency-domain Havriliak-Negami relaxation functions. *Physical Review B*, 44(14):7306–7312, 1991.
- [3] P. R. Amestoy, I. S. Duff, J. Koster, and J.-Y. L’Excellent. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal of Matrix Analysis and Applications*, 23(1):15–41, 2001.
- [4] S. Ansari and C. G. Farquharson. 3D finite-element forward modeling of electromagnetic data using vector and scalar potentials and unstructured grids. *Geophysics*, 79(4):E149–E165, 2014. ISSN 0016-8033. doi:10.1190/geo2013-0172.1.
- [5] A. Y. Aravkin and T. Van Leeuwen. Estimating nuisance parameters in inverse problems. *Inverse Problems*, 28, 2012. ISSN 02665611. doi:10.1088/0266-5611/28/11/115016.
- [6] U. Ascher and L. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, Philadelphia, PA, 1998.
- [7] U. M. Ascher. *Numerical Methods for Evolutionary Differential Equations*. SIAM, 2008. ISBN 9780898716528.
- [8] U. M. Ascher and C. Greif. *A First Course in Numerical Methods*. SIAM, Philadelphia, PA, 2011. ISBN 3717089528. doi:10.1137/1.9780898719987.

- [9] P. Belliveau. *Parallelizing the 2.5D Airborne Electromagnetic Inversion Program ArjunAir*. M.sc. thesis, Memorial University of Newfoundland, 2014.
- [10] P. Belliveau and E. Haber. Achieving depth resolution with gradient array survey data through transient electromagnetic inversion. In *SEG Technical Program Expanded Abstracts 2016*, number 1, pages 857–861, 2016.
- [11] P. Belliveau and E. Haber. Coupled simulation of electromagnetic induction and induced polarization effects using stretched exponential relaxation. *Geophysics*, 83(2):WB109–WB121, 2018. ISSN 0016-8033. doi:10.1190/geo2017-0494.1. URL <https://library.seg.org/doi/10.1190/geo2017-0494.1>.
- [12] K. M. Bennett, K. M. Schmainda, R. B. (Tong), D. B. Rowe, H. Lu, and J. S. Hyde. Characterization of continuously distributed cortical water diffusion rates with a stretched exponential model. *Magnetic Resonance in Medicine*, 50(4):727–734, 2003.
- [13] M. N. Berberan-Santos, E. N. Bodunov, and B. Valeur. Mathematical functions for the analysis of luminescence decays with underlying distributions: 1. Kohlsrausch decay function (stretched exponential). *Chemical Physics*, 315(1-2):171–182, 2005. ISSN 03010104. doi:10.1016/j.chemphys.2005.05.026.
- [14] W. Betz, I. Papaioannou, and D. Straub. Numerical methods for the discretization of random fields by means of the Karhunen-Loève expansion. *Computer Methods in Applied Mechanics and Engineering*, 271:109–129, 2014. ISSN 00457825. doi:10.1016/j.cma.2013.12.010. URL <http://dx.doi.org/10.1016/j.cma.2013.12.010>.
- [15] J. Bezanson, A. Edelman, S. Karpinski, and V. Shah. Julia: A Fresh Approach to Numerical Computing. *SIAM Review*, 59(1):65–98, 2017. URL <https://arxiv.org/abs/1411.1607>.
- [16] R.-U. Börner, O. G. Ernst, and S. Güttel. Three-dimensional transient electromagnetic modelling using Rational Krylov methods. *Geophysical Journal International*, 202(3):2025–2043, 2015. ISSN 0956-540X. doi:10.1093/gji/ggv224. URL <http://gji.oxfordjournals.org/lookup/doi/10.1093/gji/ggv224>.
- [17] L. Bottou. Large-Scale Machine Learning with Stochastic Gradient Descent. In *Proceedings of COMPSTAT'2010*. Physica-Verlag HD, 2010.

- [18] N. D. Bregman, R. C. Bailey, and C. Chapman. Crosshole seismic tomography. *Geophysics*, 54(2):200–215, 1989.
- [19] M. Burger and S. J. Osher. A survey on level set methods for inverse problems and optimal design. *European Journal of Applied Mathematics*, 16:263–301, 2005. ISSN 09567925. doi:10.1017/S0956792505006182.
- [20] H. Cai, X. Hu, B. Xiong, and M. S. Zhdanov. Finite-element time-domain modeling of electromagnetic data in general dispersive medium using adaptive Padé series. *Computers & Geosciences*, 109(January):194–205, 2017. ISSN 00983004. doi:10.1016/j.cageo.2017.08.017. URL <http://linkinghub.elsevier.com/retrieve/pii/S0098300417300079>.
- [21] M. Cardiff and P. K. Kitanidis. Bayesian inversion for facies detection: An extensible level set framework. *Water Resources Research*, 45(10):1–15, 2009. ISSN 00431397. doi:10.1029/2008WR007675.
- [22] K. S. Cole and R. H. Cole. Dispersion and absorption in dielectrics I. Alternating current characteristics. *The Journal of Chemical Physics*, 9(4): 341–351, 1941. ISSN 00219606. doi:10.1063/1.1750906.
- [23] M. Commer and G. Newman. A parallel finitedifference approach for 3D transient electromagnetic modeling with galvanic sources. *Geophysics*, 69(5):1192–1202, 2004. ISSN 0016-8033. doi:10.1190/1.1801936. URL <http://library.seg.org/doi/abs/10.1190/1.1801936>.
- [24] M. Commer, G. A. Newman, K. H. Williams, and S. S. Hubbard. 3D induced-polarization data inversion for complex resistivity. *Geophysics*, 76(3):F157–F171, 2011.
- [25] M. Commer, P. V. Petrov, and G. A. Newman. FDTD modelling of induced polarization phenomena in transient electromagnetics. *Geophysical Journal International*, 209:387–405, 2017. doi:10.1093/gji/ggx023.
- [26] T. A. Davis, S. Rajamanickam, and W. Sid-lakhdar. A survey of direct methods for sparse linear systems. Technical report, Texas A&M University, 2016.
- [27] U. Diwekar. Optimization Under Uncertainty. In *Introduction to Applied Optimization*. Boston, MA, 2008.
- [28] K. V. D. Doel and U. Ascher. On level set regularization for highly ill-posed distributed parameter estimation problems. *Journal of*

Computational Physics, 216:707–723, 2006.
doi:10.1016/j.jcp.2006.01.022.

- [29] V. Dolean, M. Gander, and L. Gerardo-Giorda. Optimized Schwarz Methods for Maxwell’s Equations. *SIAM Journal on Scientific Computing*, 31(3):2193–2213, 2009. ISSN 1064-8275. doi:10.1137/080728536. URL <http://epubs.siam.org/doi/abs/10.1137/080728536>.
- [30] V. Druskin and L. Knizhnerman. Spectral approach to solving three-dimensional Maxwell’s diffusion equations in the time and frequency domains. *Radio Science*, 29(4):937–953, 1994.
- [31] M. E. Everett. *Near-Surface Applied Geophysics*. Cambridge University Press, Cambridge, 2013.
- [32] G. Fiandaca, E. Auken, A. V. Christiansen, and A. Gazoty. Time-domain-induced polarization: Full-decay forward modeling and 1D laterally constrained inversion of Cole-Cole parameters. *Geophysics*, 77(3):E213–E225, 2012. ISSN 00168033. doi:10.1190/geo2011-0217.1.
- [33] L. A. Gallardo and M. A. Meju. Joint two-dimensional DC resistivity and seismic travel time inversion with cross-gradients constraints. *Journal of Geophysical Research*, 109(B3):B03311, 2004. ISSN 0148-0227. doi:10.1029/2003JB002716. URL <http://doi.wiley.com/10.1029/2003JB002716>.
- [34] M. J. Gander and S. Güttel. PARAEXP: A Parallel Integrator for Linear Initial-Value Problems. *SIAM Journal on Scientific Computing*, 35(2):C123–C142, 2013. ISSN 1064-8275. doi:10.1137/110856137. URL <http://epubs.siam.org/doi/abs/10.1137/110856137>
<http://dx.doi.org/10.1137/110856137>
<http://epubs.siam.org/doi/abs/10.1137/110856137>
<http://dx.doi.org/10.1137/110856137>.
- [35] R. G. Ghanem and P. D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer-Verlag, 1991. ISBN 9781461277958.
- [36] N. I. M. Gould, J. A. Scott, and Y. Hu. A numerical evaluation of sparse direct solvers for the solution of large sparse symmetric linear systems of equations. *ACM Transactions on Mathematical Software*, 33(2), 2007. doi:10.1145/1236463.1236465. URL <http://doi.acm.org/10.1145>.
- [37] E. Haber. *Computational Methods in Geophysical Electromagnetics*. SIAM, Philadelphia, PA, 2015. ISBN 9781611973792.

- [38] E. Haber and S. Heldmann. An octree multigrid method for quasi-static Maxwell's equations with highly discontinuous coefficients. *Journal of Computational Physics*, 223(2):783–796, 2007. ISSN 00219991. doi:10.1016/j.jcp.2006.10.012.
- [39] E. Haber and L. Ruthotto. A multiscale finite volume method for Maxwell's equations at low frequencies. *Geophysical Journal International*, 199:1268–1277, 2014. ISSN 0956-540X. doi:10.1093/gji/ggu268.
- [40] E. Haber, U. M. Ascher, and D. W. Oldenburg. Inversion of 3D electromagnetic data in frequency and time domain using an inexact allatonce approach. *Geophysics*, 69(5):1216–1228, 2004. ISSN 0016-8033. doi:10.1190/1.1801938. URL <http://library.seg.org/doi/abs/10.1190/1.1801938>.
- [41] E. Haber, D. W. Oldenburg, and R. Shekhtman. Inversion of time domain three-dimensional electromagnetic data. *Geophysical Journal International*, 171:550–564, 2007.
- [42] E. Haber, M. Chung, and F. Herrmann. An Effective Method for Parameter Estimation with PDE Constraints with Multiple Right-Hand Sides. *SIAM Journal on Optimization*, 22(3):739–757, 2012.
- [43] E. Hairer, S. Nørset, and G. Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*, volume 32. Springer, 1993. ISBN 354078862X. doi:10.1016/0378-4754(87)90083-8. URL <https://link.springer.com/content/pdf/10.1007/978-3-540-78862-1.pdf>.
- [44] R. Hilfer. H-function representations for stretched exponential relaxation and non-Debye susceptibilities in glassy systems. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, 65(6):9–13, 2002. ISSN 15393755. doi:10.1103/PhysRevE.65.061510.
- [45] L. Horesh and E. Haber. A Second Order Discretization of Maxwell's Equations in the Quasi-Static Regime on OcTree Grids. *SIAM Journal on Scientific Computing*, 33(5):2805–2822, 2011.
- [46] G. Horton and S. G. Vandewalle. A Space-Time Multigrid Method for Parabolic Partial Differential Equations. *SIAM Journal on Scientific Computing*, 16(4):848–864, 1995. ISSN 1064-8275. doi:10.1137/0916050. URL <http://epubs.siam.org/doi/abs/10.1137/0916050>.

- [47] G. Horton, S. G. Vandewalle, and P. Worley. An Algorithm with Polylog Parallel Complexity for Solving Parabolic Partial Differential Equations. *SIAM Journal on Scientific Computing*, 16(3):531–541, 1994.
- [48] G. M. Hoversten, C. Schwarzbach, P. Belliveau, E. Haber, and R. Shekhtman. Borehole to Surface Electromagnetic Monitoring of Hydraulic Fractures. In *79th EAGE Conference and Exhibition*, 2017. doi:10.3997/2214-4609.201700853.
- [49] J. M. Hyman and M. Shashkov. Mimetic Discretizations for Maxwell’s Equations. *Journal of Computational Physics*, 151:881–909, 1999.
- [50] M. A. Iglesias, Y. Lu, and A. M. Stuart. A Bayesian level set method for geometric inverse problems. *Interfaces and Free Boundaries*, 18(2): 181–217, 2016. ISSN 14639963. doi:10.4171/IFB/362.
- [51] S. G. Johnson. Wikimedia-YeeCube. URL <https://commons.wikimedia.org/wiki/File:Yee-cube.svg>.
- [52] S. Kang and D. Oldenburg. Revisiting the time domain induced polarization technique, from linearization to inversion. In *AGU Fall Meeting*, 2015.
- [53] S. Kang and D. W. Oldenburg. On recovering distributed IP information from inductive source time domain electromagnetic data. *Geophysical Journal International*, 207:174–196, 2016. ISSN 0956-540X. doi:10.1093/gji/ggw256.
- [54] A. Kemna, A. Binley, A. Ramirez, and W. Daily. Complex resistivity tomography for environmental applications. *Chemical Engineering Journal*, 77(1-2):11–18, 2000. ISSN 13858947. doi:10.1016/S1385-8947(99)00135-7.
- [55] D. P. Kingma and J. L. Ba. ADAM: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR 2015)*, 2015.
- [56] T. Kratzer and J. C. Macnae. Induced polarization in airborne EM. *Geophysics*, 77(5):E317–E327, 2012. ISSN 0016-8033. doi:10.1190/geo2011-0492.1.
- [57] Y. Lecun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553): 436–444, 2015. ISSN 14764687. doi:10.1038/nature14539.

- [58] Y. Li and D. W. Oldenburg. 3-D inversion of magnetic data. *Geophysics*, 61:394–408, 1996.
- [59] J.-I. Lions, Y. Maday, and G. Turinici. A parareal in time discretization of PDE's. *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 1(2):661–668, 2001.
- [60] J. Macnae. Quantifying Airborne Induced Polarization effects in helicopter time domain electromagnetics. *Journal of Applied Geophysics*, 2015. ISSN 09269851. doi:10.1016/j.jappgeo.2015.10.016.
- [61] J. Macnae. Fitting superparamagnetic and distributed Cole-Cole parameters to airborne electromagnetic data : A case history from Quebec. *Geophysics*, 81(6):B211–B220, 2016. doi:10.1190/geo2016-0119.1.
- [62] D. Marchant. *Induced Polarization Effects in Inductive Source Electromagnetic Data*. Ph.d. thesis, University of British Columbia, 2015.
- [63] D. Marchant, E. Haber, and D. W. Oldenburg. Three-dimensional modeling of IP effects in time-domain electromagnetic data. *Geophysics*, 79(6), 2014.
- [64] M. S. Mcmillan, C. Schwarzbach, E. Haber, and D. W. Oldenburg. Three dimensional parametric hybrid inversion of airborne time-domain electromagnetic data. *Geophysics*, 80(6):K25–K36, 2014. doi:10.1190/geo2015-0141.1.
- [65] P. Monk. *Finite Element Methods for Maxwell's Equations*. Oxford University Press, 2003.
- [66] O. Okko. Vertical increase in seismic velocity with depth in shallow crystalline bedrock. *Journal of Applied Geophysics*, 32(4):335–345, 1994. ISSN 09269851. doi:10.1016/0926-9851(94)90032-9.
- [67] D. W. Oldenburg and Y. Li. 3-D inversion of induced polarization data. *Geophysics*, 59(9):1327–1341, 1994. ISSN 00168033. doi:10.1190/1.1444877.
- [68] D. W. Oldenburg, E. Haber, and R. Shekhtman. Three dimensional inversion of multisource time domain electromagnetic data. *Geophysics*, 78(1):E47–E57, 2013. ISSN 0016-8033. doi:10.1190/geo2012-0131.1. URL http://apps.webofknowledge.com/full_{-}record.do?product=WOS{&}search_{-}mode=GeneralSearch{&}qid=4{&}SID=3D9g64O1pAi9CCJ9N1a{&}page=1{&}doc=3.

- [69] D. O’Leary. The block conjugate gradient algorithm and related methods, 1980. URL <http://www.sciencedirect.com/science/article/pii/0024379580902475>.
- [70] G. J. Palacky. Resistivity Characteristics of Geologic Targets. In *Electromagnetic Methods in Applied Geophysics, Volume 1, Theory*, pages 53–130. Society of Exploration Geophysicists, 1987.
- [71] W. Pelton, S. Ward, P. Hallof, W. Sill, and P. Nelson. Mineral Discrimination and Removal of Inductive Coupling with Multifrequency IP. *Geophysics*, 43(3):588–609, 1978.
- [72] B. T. Polyak and A. B. Juditsky. Acceleration of Stochastic Approximation by Averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992. ISSN 0363-0129. doi:10.1137/0330046. URL <http://epubs.siam.org/doi/10.1137/0330046>.
- [73] K. Prieto and O. Dorn. Sparsity and level set regularization for diffuse optical tomography using a transport model in 2D. *Inverse Problems*, 33, 2017.
- [74] I. T. Rekanos and T. G. Papadopoulos. FDTD Modeling of Wave Propagation in Cole-Cole Media With Multiple Relaxation Times. *Antennas and Wireless Propagation Letters, IEEE*, 9:67–69, 2010. ISSN 1536-1225. doi:10.1109/lawp.2010.2043410.
- [75] Z. Ren, T. Kalscheuer, S. Greenhalgh, and H. Maurer. A finite-element-based domain-decomposition approach for plane wave 3D electromagnetic modeling. *Geophysics*, 79(6):E255–E268, 2014. ISSN 0016-8033. doi:10.1190/geo2013-0376.1. URL <http://library.seg.org/doi/10.1190/geo2013-0376.1>.
- [76] A. Revil and N. Florsch. Determination of permeability from spectral induced polarization in granular media. *Geophysical Journal International*, 181:1480–1498, 2010. doi:10.1111/j.1365-246X.2010.04573.x.
- [77] P. Robbe. *GaussianRandomFields*, 2017.
- [78] F. Roosta-khorasani, K. V. A. N. D. E. N. Doel, and U. R. I. Ascher. Stochastic Algorithms for Inverse Problems Involving PDEs and Many Measurements. *SIAM Journal on Scientific Computing*, 36(5):S3–S22, 2014.

- [79] P. S. Routh and D. W. Oldenburg. Electromagnetic coupling in frequency-domain induced polarization data: a method for removal. *Geophysical Journal International*, 145:59–76, 2001. ISSN 0956-540X. doi:10.1046/j.1365-246x.2001.00384.x.
- [80] L. Ruthotto, E. Treister, and E. Haber. jInvA Flexible Julia Package for PDE Parameter Estimation. *SIAM Journal on Scientific Computing*, 39(5): 702–722, 2017.
- [81] Y. Saad. *Iterative Methods for Sparse Linear Systems*. SIAM, Philadelphia, PA, 2 edition, 2003. ISBN 9780898715347.
- [82] K. Salari and P. Knupp. Code Verification by the Method of Manufactured Solutions. Technical report, 2000.
- [83] F. Santosa. A Level-Set Approach for Inverse Problems Involving Obstacles. *ESAIM: Control, Optimisation and Calculus of Variations*, 1 (January):17–33, 1996. doi:10.1051/cocv:1996101.
- [84] O. Schenk and K. Gartner. Solving unsymmetric sparse systems of linear equations with PARDISO. *Journal of Future Generation Computer Systems*, 20(3):475–487, 2004.
- [85] H. O. Seigel. Mathematical Formulation and Type Curves for Induced Polarization. *Geophysics*, XXIV(3):547–565, 1959.
- [86] A. Shapiro, D. Dentcheva, and A. Ruszczycki. *Lectures On Stochastic Programming: Modeling and Theory*. SIAM, Philadelphia, PA, 2009. ISBN 089871687X. doi:http://dx.doi.org/10.1137/1.9780898718751.
- [87] R. S. Smith. Induced polarization effects in airborne electromagnetic data: estimating chargeability from shape reversals. In *SEG Technical Program Expanded Abstracts 2016*, pages 2211–2217, 2016.
- [88] K. Sørensen and E. Auken. SkyTEM a new high-resolution helicopter transient electromagnetic system. *Exploration Geophysics*, 35:191–199, 2004.
- [89] D. E. Stewart. *Dynamics with Inequalities: Impacts and Hard Constraints*. SIAM, Philadelphia, PA, 2011. ISBN 978-1-61197-070-8.
- [90] D. M. Sullivan. Z-Transform Theory and the FDTD Method. *IEEE Transactions on Antennas and Propagation*, 44(1):28–34, 1996.

- [91] Y. Takayama and W. Klaus. Reinterpretation of the auxiliary differential equation method for FDTD. *IEEE Microwave and Wireless Components Letters*, 12(3):102–104, 2002. ISSN 15311309. doi:10.1109/7260.989865.
- [92] W. M. Telford, L. P. Geldart, and R. E. Sheriff. *Applied Geophysics*. Cambridge University Press, second edition, 1990.
- [93] A. Tikhonov and Arsenin. *Solutions of Ill-posed Problems*. Winston & Sons, Washington, D.C., 1977.
- [94] G. VanVoorhis, P. Nelson, and T. Drake. Complex resistivity spectra of porphyry copper mineralization. *Geophysics*, 38:49–60, 1973.
- [95] S. W. Wallace and W. T. Ziemba, editors. *Applications of Stochastic Programming*. SIAM, Philadelphia, PA, 2005.
- [96] S. H. Ward. Resistivity and Induced Polarization Methods. In *Geotechnical and Environmental Geophysics: Volume 1: Review and Tutorial*, pages 147–190. Society of Exploration Geophysicists, 1990. doi:10.1190/1.9781560802785.ch6.
- [97] S. H. Ward and G. W. Hohmann. Electromagnetic theory for geophysical applications. In M. N. Nabighian, editor, *Electromagnetic Methods in Applied Geophysics, Volume 1, Theory*, pages 131–311. Society of Exploration Geophysicists, 1987.
- [98] W. Weedon and C. Rappaport. A general method for FDTD modeling of wave propagation in arbitrary frequency-dispersive media. *IEEE Transactions on Antennas and Propagation*, 45(3):401–410, 1997.
- [99] P. Weidelt. Response characteristics of coincident loop transient electromagnetic systems. *Geophysics*, 47(9):1325–1330, 1982. ISSN 1070485X. doi:10.1190/1.1441393. URL <http://geophysics.geoscienceworld.org/content/47/9/1325.abstract>.
- [100] G. A. Wilson, A. P. Raiche, and F. Sugeng. 2.5D inversion of airborne electromagnetic data. *Exploration Geophysics*, 37:363–371, 2006.
- [101] M. G. Wismer and R. Ludwig. An Explicit Numerical Time Domain Formulation to Simulate Pulsed Pressure Waves in Viscous Fluids Exhibiting Arbitrary Frequency Power Law Attenuation. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 42(6): 1040–1049, 1995. ISSN 08853010. doi:10.1109/58.476548.

- [102] Z. Xu and M. S. Zhdanov. Three-Dimensional Cole-Cole Model Inversion of Induced Polarization Data Based on Regularized Conjugate Gradient Method. *IEEE Geoscience and Remote Sensing Letters*, 12(6):1180–1184, 2015. ISSN 1545598X. doi:10.1109/LGRS.2014.2387197.
- [103] K. Yee. Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media. *IEEE Transactions on Antennas and Propagation*, 14(3):302–307, 1966.
- [104] M. Zaslavsky and V. Druskin. Solution of time-convolutionary Maxwell’s equations using parameter-dependent Krylov subspace reduction. *Journal of Computational Physics*, 229(12):4831–4839, 2010. ISSN 0021-9991. doi:10.1016/j.jcp.2010.03.019. URL <http://dx.doi.org/10.1016/j.jcp.2010.03.019>.
- [105] M. S. Zhdanov. Generalized effective-medium theory of induced polarization. *Geophysics*, 73(5):F197–F211, 2008. ISSN 00168033. doi:10.1190/1.2973462.
- [106] K. Zonge, J. Wynn, and S. Urquart. Resistivity, Induced Polarization, and Complex Resistivity. In *Investigations in Geophysics No 13*, page 265. 2005.
- [107] K. L. Zonge and J. C. Wynn. Recent advances and applications in complex resistivity measurements. *Geophysics*, 40(5):851–864, 1975.