Geophysical Survey Decomposition and Efficient 3D Inversion of Time-domain Electromagnetic Data

by

Dikun Yang

B. Sc., China University of Geosciences, 2005
M. A. Sc., China University of Geosciences, 2008

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)

September 2014

Abstract

Rigorous three-dimensional (3D) forward and inverse modeling of geophysical electromagnetic (EM) data can be time-consuming and may require a large amount of memory on expensive computers. In this thesis, a novel framework, called survey decomposition, is proposed to make the 3D EM modeling more efficient.

Recognizing the multi-scale nature of the EM modeling problems, the fundamental idea is to break down an EM survey, which consists of many transmitters, receivers and times/frequencies, into a number of subproblems, each of which is only concerned about data modeled by a localized source, receiver and time/frequency. The modeling is then carried out on the subproblems at different scales, instead of the original problem as a whole. Such a decomposition is able to speed up the numerical modeling, because: (1) A subproblem can have highly efficient discretizations in space and time customized to its localized source, receiver, time/frequency and the specific scale of investigation, for example, it uses a local mesh that is much smaller than the one used in the original global problem; (2) A subproblem is a self-contained EM modeling problem that does not depend on other subproblems, so it is suitable for massive parallelization; (3) Upon decomposition, no modeling is carried out on the global mesh and the amount of computation is proportional to the number of subproblems, so the scalability improves significantly.

After decomposition, the large number of subproblems is further reduced by adaptive, random and dynamic subsampling of the data. The adaptive scheme matches the number of samples to the scale of investigation so that only the data necessary for the model reconstruction are selected.

The framework of survey decomposition is applied to two types of time-domain

EM (TEM) surveys: airborne TEM and ground large loop TEM. Both synthetic and field data are inverted using this new approach. I show that survey decomposition is capable of producing modeling and inversion results similar to those from the conventional methods with greatly reduced time and memory usage. Further speed-up by massive parallelization and generalization to other types of EM surveys is straightforward.

Preface

This thesis presents my original researches completed at the Department of Earth, Ocean and Atmospheric Sciences and the Geophysical Inversion Facility (GIF) at the University of British Columbia (UBC), Vancouver, Canada. Some chapters contain the materials from the papers published on peer-reviewed scientific journals and other sources.

Chapter 2 is a revised version of a published paper (Yang, D. and D. Oldenburg, 2012, Three-dimensional inversion of airborne time-domain electromagnetic data with applications to a porphyry deposit: *Geophysics*, 77, B23-B34). I carried out the numerical experiments and prepared the maniscript. Oldenburg designed and guided the research and edited the manuscript.

Chapter 3 contains the text and figures excerpted from a published paper (Yang, D., D. Oldenburg, E. Haber, 2014, 3-D inversion of airborne electromagnetic data parallelized and accelerated by local mesh and adaptive soundings: *Geophysical Journal International*, 196, 1942-1507) and unpublished materials. I initiated the research project, carried out all the experiments and prepared the manuscript. Oldenburg advised the formation of the basic ideas and revised the manuscript.

Chapter 4 is a revised version of a published paper (Yang, D., D. Oldenburg, E. Haber, 2014, 3-D inversion of airborne electromagnetic data parallelized and accelerated by local mesh and adaptive soundings: *Geophysical Journal International*, 196, 1942-1507). I initiated the research project, carried out all the experiments and prepared the manuscript. Oldenburg supervised the project and revised the manuscript. Haber provided computational basis and technical advice for the research.

The contents in Chapter 5 have not been published. I initiated the research project, carried out all the experiments and prepared the manuscript. Oldenburg supervised the project and revised the manuscript. This work will be presented at the Society of Exploration Geophysicists Annual Meeting in 2014 and a journal paper is in preparation.

Appendix A and B contain the materials derived from the course notes of MATH 522 and EOSC 555B lectured by Eldad Haber.

I prepared the manuscript of the entire thesis with advice and help from Oldenburg on the organization, proof-reading and editing.

Table of Contents

List of Tables

List of Figures

Acknowledgments

It is UBC-GIF's tradition to thank Roman Shekhtman, the man who answers the endless requests from all GIF members, at the very beginning of the acknowledgments. I am particularly grateful to his help and assistance on programming and computers. My reseaches could not take place if Roman did not walk to the computer room and reboot the computers by hand after my codes crashed the whole cluster on holidays.

My supervisor, Dr. Douglas Oldenburg, is amazing. He is always academically, emotionally and financially supportive. I have learned not only skills and knowledges but also visions and attitudes from him. Dr. Eldad Haber's expertise in computational geophysics has enlightened my minds; my research is built on the numerical foundations that he and Dr. Douglas Oldenburg had developed. They are the giants whose shoulders I can stand on. I also thank Dr. Catherine Johnson and Dr. Christian Schoof, who kindly served on my supervisory committee and candidacy committee.

Major funding for my research is from MITEM Consortium at UBC-GIF and NSERC CRD and IRC programs. I received tuition awards from the Faculty of Graduate and Postdoctoral Studies and the Faculty of Science, UBC. Geoscience BC has generously offered me the financial support for two consecutive years (2009∼2010) with Geoscience BC Scholarship program. I am also grateful to Egil Harold Lorntzsen Scholarship and Thomas and Marguerite MacKay Memorial Scholarship.

Some organizations and individuals have kindly provided help on data and other technical information. In particular, Geoscience BC made the VTEM field data at Mt. Milligan available; HudBay Minerals Inc. and Dennis Woods at Discovery International Geophysics Inc. provided the SQUID data and background information at the Lalor Mine deposit; The discussions with Peter Kowalczyk inspired the investigation on the airborne EM systems' footprint in Chapter 2; Terrane Metals Corp. (now Thompson Creek Metals Company) and Dr. Dianne Mitchinson provided information about the geology at Mt. Milligan.

My studies would be less productive and boring if there was no invaluable discussion and curling with David Marchant, Dr. Kristopher Davis, Dr. Laurens Beran and many other colleages in UBC-GIF.

As an international student, I should not forget the people who had made my academic pursuit in Canada possible. Dr. Xiangyun Hu and Prof. Jiaying Wang in China University of Geoscienecs (Wuhan) encouraged and recommended me to study in UBC. The St. John's College in UBC and the lovely residents there had offered great comfort to me when I just came to Vancouver.

I personally trace the origin of this thesis back to the two books my parents gave to me as the birthday gifts when I was a first grader, one folk tale that unleashed my power of imagination and the other *One Hundred Thousand Whys* that triggered my curiosity in science. Thank you, Mom and Dad.

Maggie, who gives me unconditional understanding, patience, support and accompany through the years, deserves the highlighted spot at the end of my acknowledgements.

Dedication

To whoever is impressed by the following statement.

“ Imagination is more important than knowledge. For knowledge is limited, whereas imagination embraces the entire world, stimulating progress, giving birth to evolution. It is, strictly speaking, a real factor in scientific research.”

– Albert Einstein

Chapter 1

Introduction

The physical properties of the earth are the keys to the understanding of the process, evolution, and status of the planet, for examples, sedimentation, tectonic events, mineralization, groundwater flow, rock deformation, etc. However, direct measurement is usually very difficult, especially for the deep interior and vast volume of the earth. Geophysicists have to rely on observation of a number of different types of physical field to infer the earth's physical structures. Modeling of those physical processes is non-trivial given the complexity of the earth, the underlying governing equations and the variety of survey methods. In this thesis, I develop a novel framework for advanced geophysical modelings, and use it to improve the efficiency of three-dimensional inversion of time-domain electromagnetic induction data.

This chapter introduces the background of my research, including motivation, literature reviews, fundamental methodologies, objectives, along with the basic equations and assumptions used in this thesis.

1.1 Electromagnetic Surveys for Conductivity

1.1.1 Conductivity: a diagnostic physical property

Among the many physical properties that the earth possesses, electrical conductivity is one of the most diagnostic (Palacky, 1988). Conductivity reflects how easily the medium conducts electric current when an electric field is present.

Previous studies in petrophysics have established the links between the conductivity and the geologic targets and activities.

- Mineralization. Many minerals are characterized by high conductivities; for example, different types of natural pyrite exhibit considerable variation in conductivity (Abraitis et al., 2004); some massive sulphide deposits can be many orders of magnitude more conductive than the background rock (Stolz, 2000); and the conductive graphites are indicative in finding unconformity uranium (Mwenifumbo et al., 2004).

- Hydrocarbon reservoir. The existence of hydrocarbons can reduce the conductivity of sandstone, making a resistive anomaly an interesting target in oil and gas exploration (Waxman et al., 1974; Constable, 2010).

- Tectonic activity. It is known that there is partial melting and/or dehydration of rock, a cause of high conductivity, in the crust and upper mantle, so the distribution of the conductive zones can be informative in the study of the earth's interior (Hyndman & Hyndman, 1968; Waff, 1974). The deep fault zones can also be imaged in terms of their electrical conductivity (Ritter et al., 2005).

- Groundwater flow. When water moves underground, whether it is a geothermal source or sea water intrusion, the impacted zone can show very different conductivity signatures copmpared to the surrounding (Schwarz et al., 1985; Fitterman & Stewart, 1986).

- Environmental problems. Both natural processes and human activities can alter the earth's conductivity near the surface, for example, the development of sinkholes (van Schoor, 2002), the contamination by landfill leachate (Chambers et al., 2006), the unexploded ordnance (Huang & Won, 2003), and so on.

Therefore, there is great interest of detecting the earth's conductivity structures at different scales.

1.1.2 Maxwell's equations

Direct measurement of the earth's conductivity using well logging or rock samples is very expensive or even impossible, so exploration geophysics uses the observations of electromagnetic fields to infer the conductivity.

Conductivity is connected to the EM fields through Maxwell's equations that usually consist of four partial differential equations (PDEs). If stated in the differential form, they are

$$\nabla \cdot \mathbf{D} = \rho_f, \tag{1.1a}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{1.1b}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{1.1c}$$

$$\nabla \times \mathbf{H} = \mathbf{J}_f + \frac{\partial \mathbf{D}}{\partial t} + \mathbf{J}_s, \tag{1.1d}$$

where t is time, $\mathbf{E}$ is electric field, $\mathbf{H}$ is magnetic field intensity, $\mathbf{B}$ is magnetic flux density, $\mathbf{D}$ is electric displacement field, ρ_f is free charge density, and $\mathbf{J}_f$ is free electric current density and $\mathbf{J}_s$ is the artificial external source. Equation 1.1d shows the possibility of driving the EM fields using artificial sources and observing the responses of the fields. This motivates the use of the controlled source EM (CSEM) survey methods in the exploration and is a topic of this thesis introduced in the next subsection.

Fields and fluxes are associated through the constitutive relations

$$\mathbf{D} = \varepsilon \mathbf{E}, \tag{1.2a}$$

$$\mathbf{B} = \mu \mathbf{H}, \tag{1.2b}$$

$$\mathbf{J} = \sigma \mathbf{E}, \tag{1.2c}$$

where ε is the electrical permittivity of the medium, μ is the magnetic permeability, and σ is the electrical conductivity; they are the (bulk) physical properties of the EM medium and geophysicists see the earth as a composition of materials with different ε, μ and σ. In order to fully describe the earth's materials, these properties must be expressed as tensors to model anisotropy, and/or as functions of frequency for frequency-dependence. As those complications are out of the scope of this thesis, I assume the earth is made of linear, isotropic and frequency-independent materials.

More insights can be gained by viewing Maxwell's equations in frequency domain. Taking the Fourier transform of equation 1.1c and 1.1d, with equation 1.2 substituted in, yields

$$\nabla \times \mathbf{E}(\omega) = -i\omega\mu\mathbf{H}(\omega), \tag{1.3a}$$

$$\nabla \times \mathbf{H}(\omega) = \sigma\mathbf{E}(\omega) + \mathbf{J}_s(\omega) + i\omega\varepsilon\mathbf{E}(\omega), \tag{1.3b}$$

where $\omega = 2\pi f$ is the angular frequency. For the geophysical surveys of interest in this thesis, the quasi-static approximation $\sigma \gg \omega\varepsilon$ holds for most media encountered in geosciences. The physical implication is that the magnetic field emanates dominantly from the current of moving charges in the manner of diffusion; the quasi-static approximation allows the (wave) term $i\omega\varepsilon\mathbf{E}$ representing the displacement current to be safely neglected.

The impact of small variation of μ in most geoscientific objects is considered secondary in comparison with the fluctuation of EM fields due to σ, which can vary in a range of $10^2 \sim 10^{-6}$ S/m. Therefore, it is a reasonable to assume the earth has free space values $\varepsilon_0 = 8.85418 \times 10^{-12}$ farad/m and $\mu_0 = 4\pi \times 10^{-7}$ H/m, and conductivity σ is the sole physical property varied in modeling and inversion.

I note the assumptions made in this thesis are meant to facilitate the development of the methodology and do not restrict generalization of my results to other more sophisticated cases after necessary modifications.

The process of EM induction can be viewed as a linear system, for which the source $\mathbf{J}_s$ is the input and the measured fields are output. Computing the system's responses at an infinite number of frequencies for the entire spectrum provides complete information. That spectrum can be inverse transformed to generate the impulse response. The system's responses, the electric and magnetic fields, are then the convolution of the impulse response and the waveform of $\mathbf{J}_s$ in time. In principle, working at infinite times from zero to infinity (seconds) is equivalent to working at infinite frequencies from zero to infinity (Hz). This means the measurements and modeling undertaken in the time domain and their counterparts in the frequency domain are interchangeable.

1.1.3 Controlled source EM surveys

The advantages of controlled source EM (CSEM) is that by changing the survey parameters, for example, the transmitter/receiver geometries and locations, and the current waveform, the earth can be illuminated in different ways. Also, if the source current is known, the conductivity information can be obtained by collecting either the electric field or the magnetic field data.

Although there is no fundamental difference between understanding EM induction in frequency domain and in time domain, significant distinction exists in terms of field surveys, because the instruments are made specific to work with either specific frequencies or delay times at a specific base frequency:

- **Frequency-domain EM (FEM)**. The transmitter, either grounded electrodes (galvanic source) or a closed loop (inductive source), carries a time-harmonic current at some discrete frequencies. The receiver measures the EM fields at the same frequencies. The measured signals can be decomposed into two components: one is in-phase (real) with the source and the other is out-of-phase (imaginary or quadrature). So in a FEM survey, there are real and imaginary data at each frequency for each source-receiver pair.

- **Time-domain EM (TEM)**. The transmitter carries a certain amount of current that varies as a function of time with a specific base frequency, usually involving a pulse or a turn-on and turn-off of current. This time-varying current induces the secondary fields in the earth, which decay over time. The receivers measure the fields as functions of time at a sequence of delay times (time channels).

A variety of CSEM survey configurations in time and frequency domain have been developed for many different applications and they have provided invaluable information about the earth's conductivity. For example:

- Airborne EM method can be used to quickly map the geology (Palacky, 1981, 1993), target the mineralization zones (Fraser, 1978), and investigate the groundwater (Fitterman & Deszcz-Pan, 1998) and other near-surface problems (Pfaffhuber et al., 2010).

- Ground EM using a large and fixed loop is useful to find highly conductive mineralizations buried under cover (Zang, 1993).
- Marine CSEM is able to locate the resistive region in the sea bed that is potentially associated with hydrocarbon reservoirs (Edwards, 1997; MacGregor & Sinha, 2000; Constable, 2010).
- DC resistivity is considered a special case of CSEM, in which the source waveform is constant. Depending on the array configurations, DC resistivity can be used to look for minerals (Spitzer & Chouteau, 2003; Legault et al., 2008), groundwater (Benkabbour et al., 2004), geo-hazard (Johnson, 2003) and other targets (Bernstone et al., 2000).
- A small portable loop-loop system can be used in the environmental and engineering applications to investigate the near surface (De Jong et al., 1979; Jardani et al., 2007).
- LOTEM (long-offset TEM) is often used in large-scale investigation of the earth's crust (Strack et al., 1990) and hydrocarbon prospecting (Strack et al., 1989).

In this thesis, I focus upon results in the time domain EM, particularly two most commonly-used TEM configurations in mining explorations: airborne and large loop ground TEM, but there is continued reference to frequencies since my methodology is equally applicable to both and some concepts are easier to explain or understand in the frequency domain.

In practice, a horizontal transmitter loop is usually used due to its easy deployment and good coupling with the ground. A trade-off between resolution and penetration depth exists when choosing the size of loop: a small loop captures great detail of the variation in the conductivity structure near surface and enjoys good mobility, but is more likely to suffer from low signal-noise ratio (SNR) at late delay times that have information from deep in the earth; a large loop, on the other hand, creates a large primary field illuminating a large volume of the earth, but may not be able to differentiate the variations of conductivity at the small scale and is cumbersome to move around. Therefore, both small and large loop source are widely used in TEM surveys for different purposes and they are specifically studied in this thesis.

Small loop source: airborne TEM. When a transmitter loop is much smaller than the scale of investigation, it can be treated as a magnetic dipole source. Such a local source induces EM fields in its close proximity, so it is important to move the loop around to sample a large area and take measurement close to the source for strong anomalous signals. Although surface operations of moving a small loop are possible, an airborne survey is more appealing because of its ability to sweep a large area at a much lower cost. An airborne TEM system usually consists of a transmitter loop mounted on, or towed by, an aircraft and a towed receiver coil at the center of the transmitter or slightly offset. As the aircraft flies along flight lines, the transmitter emits an EM excitation in a particular waveform, then the induced fields (usually dB/dt in volts) are measured at a sequence of delay times by the receiver; this emit-receive cycle completes an airborne TEM sounding at that location. Over a measurement period of a fraction of a second, while the aricraft has moved a few meters, a sounding is acquired. This distance moved is assumed small compared with the scale of investigation of the EM system (see the sketch of a helicopter-borne system in Figure 1.1). By choosing different parameters of the survey configuration, for example, the base frequency, waveform, pulse width, flight height, loop radius and sounding spacing, airborne TEM systems can be tuned to geo-electrical explorations at different scales. The applications of surveys include regional geologic mapping, mineral deposit targeting, hydrological and environmental studies. Most commercial systems for mineral explorations deliver data at spacings from a few meters to about 10 meters. The spacing between parallel flight lines can vary from about kilometers for regional reconnaissance to tens of meters for near-surface investigations.

Large loop source: ground TEM. In contrast to moving a small loop in air, one can energize a large loop fixed on the surface to illuminate deep targets at depth. In order to maximize the depth of investigation, such loops are generally a few square kilometers in area. In such a survey, the transmitter induces EM fields in the ground using a particular current waveform; then the receiver measures the

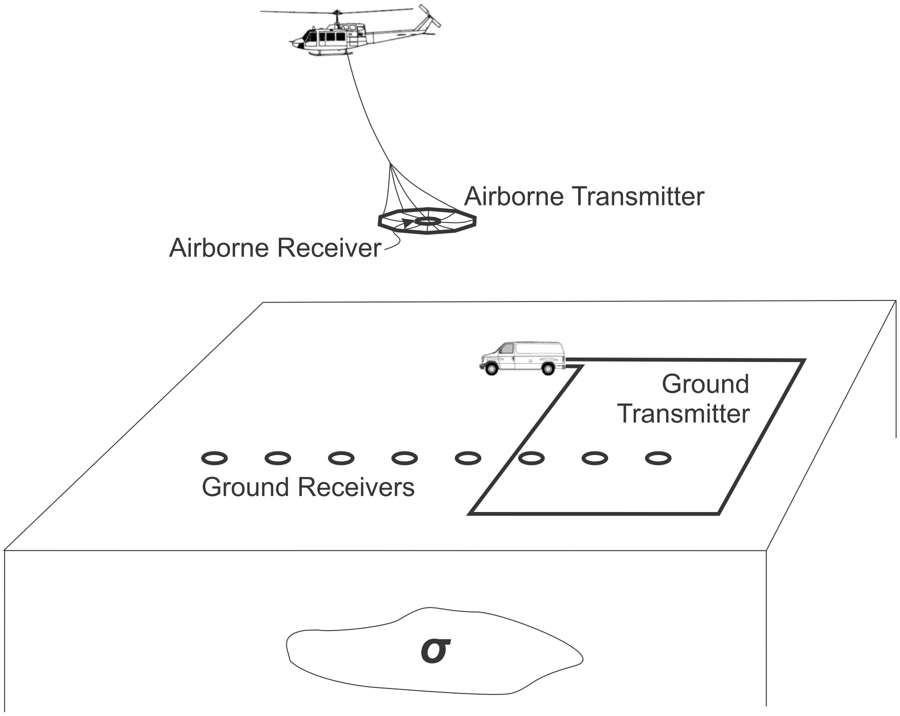

Figure 1.1: Airborne and large loop ground TEM survey.

fields of B or dB/dt at some delay times at a particular sounding location; after each sounding, the receiver moves to the next location to repeat the same procedure of measurement while the transmitter loop is immobile. Due to the cost of moving receivers on the surface, the common practice of ground TEM so far is following a few receiver lines, on which soundings are located a certain distance apart (see the sketch of a ground loop system in Figure 1.1). This drawback is compensated by its stability allowing much later times to be measured. Therefore, ground TEM is primarily used in applications like the delineation of geologic basement, buried mineral deposits and others for which depth penetration is important.

In addition to measuring the magnetic field (B or dB/dt) on the surface, it is possible to have receivers down boreholes or flying up in the air. However, there is no fundamental difference between measuring data in borehole/air and on the surface from the modeling's point of view.

1.2 Interpretations of TEM Data

Interpretation is the process of extracting information from the data for the use of decision making. This process is critical as the measured data can only be useful if they are interpreted properly. Historically, the interpretation methods have evolved over time and there are qualitative, semi-quantitative and quantitative methods of interpretation.

1.2.1 Qualitative and semi-quantitative methods

Because of the complexity of electromagnetic induction, TEM data have been historically interpreted by using qualitative methods. For example, one can differentiate regional rock units by looking at time channel grids of an airborne TEM survey; or compare the field data on a profile or at a sounding with some well-known canonical responses and infer the geometry of the target. These methods are easy to implement and have been working well for some simple targets or in the field for the data quality control.

When more specific information is requested, for example, the conductivity (conductance) and the geometry of the target, some semi-quantitative methods can be used by assuming simplified earth models:

- Time constant analysis (Palacky & West, 1973; Nabighian & Macnae, 1991; Macnae, 1998). This method approximates the decay of the field with an exponential function; the time constant derived from the function is related to the geometry and conductivity of the target, for example, soundings over large and/or good conductor have slow decays and thus large time constants. Calculation of time constant is very fast but it is difficult to separate the size and the conductivity based on the time constant alone.

- Apparent conductivity (Fraser, 1978; Palacky & West, 1991; Palacky, 1993; Annan et al., 1996). This method finds a representative uniform half-space that would produce the same responses as measured data at a particular delay time; it converts data to quantities in conductivity, eliminating the data complications contributed by various survey configurations, but the depth information is still in time.

- Conductivity depth transform/imaging (CDT/CDI) (Wolfgram & Karlik, 1995; Macnae, 1998; Eaton, 1998; Reid & Fullagar, 1998; Macnae et al., 2010). CDT relates time to depth through the concept of diffusion distance, so it can do the imaging by putting the apparent conductivities at estimated depths. Because a diffusion distance is defined for a uniform half-space, the depth estimation can be unreliable if the subsurface contains complicated structures.

- Parametric inversion. Under certain circumstance, an earth model can be effectively represented by some simple bodies, like plate (Keating & Crossley, 1990; Raiche, 2004) or sphere (Smith & Lee, 2002); and the geometric parameters and conductivity of the bodies are recovered as model parameters in the inversion. Parametric modeling can be as good as rigorous methods if the ground truth is similar to the assumed bodies, but, in reality, the simple bodies can hardly explain the entire data set due to the complexity of the actual geology.

Using the interpretation methods above, we can quickly get a good handle on the first-order structure of the target. However, when the earth is more complicated or advanced information is required from the data, more rigorous and quantitative methods have to be used.

1.2.2 Rigorous modeling

Rigorous modeling, sometimes referred as pixel-based modeling, rigorously solves Maxwell's equations. This method represents the 3D earth by a number of pixels, each of which has a constant physical property. Pixel-based models can have more flexibility to characterize the earth's structures, but the cost of computation is increased because of the increase in the number of model parameters. Rigorous modeling methods are often characterized by the dimensionality of the discretization in space.

The simplest discretization is one-dimensional (1D) layered model, in which the earth consist of a stack of layers each with a different value of a physical property and thickness. This assumes that the physical property is not a function of horizontal location and only changes with depth; the surface also has to be flat. This assumption is compatible with geological sedimentation and thus deemed valid in many cases. 1D inversion provides great insights about depth resolution and works very well when the strata are nearly 1D. As many 1D algorithms have been published (Raiche et al., 1985; Farquharson & Oldenburg, 1993; Lane et al., 2000; Wolfgram et al., 2003; Sattel, 2005; Brodie & Sambridge, 2006; Vallee & Smith, 2009; Fullagar et al., 2010), it is a well-established approach in real applications. However, if severe topography is present or the lateral variation of conductivity is significant, the credibility of 1D inversion results can be questioned.

If the geology varies with depth and one horizontal direction, but is constant in the other horizontally orthogonal direction, one can construct a two-dimensional (2D) model consisting of many "cells" on a 2D plane, which works best with geophysical profiles along lines perpendicular to the strike direction. The extra dimension means there are more parameters in 2D models. As a transition from 1D to 3D, some 2D forward modeling and inversion algorithms have been proposed with applications (Wilson et al., 2006; Yu, 2012).

Eliminating all the restrictions on dimensional homogeneity, a three-dimensional (3D) model can image arbitrarily complicated models, including 1D and 2D pixel-based models. Therefore, a 3D model is a very flexible tool for geophysical interpretation. A pixel in a 3D model is called a "cell" and it resides in a "mesh". The number of cells increases geometrically as one goes from 1D to 2D, then to 3D. Despite the computational costs that increase nonlinearly with the increase in the number of cells, the flexibility of 3D modeling is very appealing in real applications.

In the following section, I briefly overview the two essential components involved in the rigorous modeling:

- **Forward modeling**: computing the theoretical responses of a model in a survey;

- **Inverse modeling (inversion)**: finding a model that reproduces the field observations.

1.3 3D Numerical Modeling

1.3.1 Forward modeling

Given a parameterized model $\mathbf{m}$, a forward modeling simulates field observation for data $\mathbf{d} = \{d_1, \ldots, d_N\}$ in data space

$$\mathbf{d} = F(\mathbf{m}), \tag{1.4}$$

where F is a symbolic forward operator determined by the physics of the problem and applied sources. Real data always contain noise, but the noise can be complicated, and is generally treated as a zero-mean random variable not explicitly modeled due to its complications.

If the model is in 3D, the forward problems have to be solved with numerical techniques. The spatial domain of modeling is first discretized into many cells based on a mesh grid. The physical property of every cell then is considered a model parameter. By placing the fields at nodes, edges, centers or on faces of cells, the differential or integral operators can be converted to vectors and matrices using the finite volume (Haber et al., 2000; Haber & Ascher, 2001), finite difference (Wang & Hohmann, 1993; Commer & Newman, 2004), finite element (Pridmore et al., 1981; Badea et al., 2001; Schwarzbach & Haber, 2013) or integral equation methods (Hohmann, 1975; Newman et al., 1986). Finally the fields are obtained by solving some matrix equations.

Regardless of the specific discretizing and formulation techniques, eventually one has to solve a linear system of equations at each time or frequency in the form of

$$\mathbf{A}\mathbf{u} = \mathbf{q}, \tag{1.5}$$

where the Maxwell matrix $\mathbf{A}$ is due to the discretization and conductivity, $\mathbf{u}$ is the field or potential, and the right-hand side $\mathbf{q}$ represents sources, boundary conditions and/or the primary fields. Then the TEM forward modeling algorithms differ in how equation 1.5 is solved using existing numerical solvers.

Rapid development of computers has made 1D and 2D modeling relatively easy, but the numerous model parameters from a 3D discretization can lead to a linear system (equation 1.5) of unmanageable size. Geophysicists continue to struggle to find better ways of solving 3D problems with good accuracy and speed. The improvement of efficiency in 3D modeling has been an active research topic for many decades and is still one of the top items on the geophysicists' wish list, since being able to forward model geophysical responses efficiently is critical to many advanced analysis of data.

1.3.2 Inversion

A geophysical inversion infers the physical property of the earth from the observed data in field. Fundamentally an inversion can be viewed deterministically or stochastically. The former finds one, or some, models that reproduce the data and hopes the models, to some extent, are representative to the ground truth. It leaves the uncertainty untreated. Stochastic inversion instead regards the model parameters as random variables and converts data to probability density distributions of model parameters (Tarantola, 2005). Stochastic inversion is more philosophically sound, but severely suffers from the "curse of dimensionality" and is useful only when the model has a limited number of parameters, for example in 1D layered inversion. My thesis will focus on the deterministic approach and "finding a model" will implicitly refer to that methodology in the rest of my thesis. Nevertheless, because stochastic inversion shares the same foundation of forward modeling with deterministic inversion, some of my research results are also valid in stochastic inversions whenever applicable.

Regardless of the specific survey and type of data, inversions all have the similar formulations and methodologies: suppose the current model is $\mathbf{m}$ and the forward responses based on $\mathbf{m}$ is $\mathbf{d} = F(\mathbf{m})$, the inversion finds a model update $\delta\mathbf{m}$ so that the misfit between the observed data $\mathbf{d}^{obs}$ and the predicted data $\mathbf{d}^* = F(\mathbf{m} + \delta\mathbf{m})$ is sufficiently reduced. This process requires two pieces of information: (1) how current predicted data differ from the observed data, i.e. data misfit; (2) how the pertubations of model parameters affect the predicted data, i.e. sensitivity. Provided these two pieces of information are available, variants of inversion schemes can be formulated with different manipulations.

Due to the ill-posedness of the inverse problem, the information in data is not enough to uniquely determine a model update. A sensible way of tackling this problem is to impose some prior information on the relationship among model parameters and formulate the inversion as an optimization problem with Tikhonov regularization (Tikhonov & Arsenin, 1977)

$$\underset{\mathbf{m},\beta}{\text{minimize}} \quad \phi_d(\mathbf{m}) + \beta\phi_m(\mathbf{m}), \tag{1.6}$$

where ϕ_d is a functional of data misfit, ϕ_m is a regularization term measuring the model norm, and β is a trade-off (regularization or Tikhonov) parameter. Equation 1.6 states that in addition to reducing the data misfit, the model must also comply with the structural constraints imposed by ϕ_m, which generally needs the model to be "simple" as measured by certain criteria. A "simple" model has large data misfit, while a model fitting the data better is more "complex". An inversion also finds a β that ensures the data are reasonably fit while the model has reasonable amount of structures.

A variety of formulations for ϕ_d and ϕ_m have been proposed in previous publications (Portniaguine & Zhdanov, 1999; Guitton & Symes, 2003; Haber, 2004; Farquharson, 2007; Sun & Li, 2014). Since the choice of ϕ_d and ϕ_m is not essential to my research, I assume the most commonly-used L-2 norm for ϕ_d and ϕ_m. The regularization term ϕ_m contains the smallest component and the flattest component. Specifically, they are defined as (Li & Oldenburg, 1996)

$$\phi_d = \frac{1}{2}\sum_{i=1}^{N}\left(\frac{F_i(\mathbf{m}) - \mathbf{d}_i^{obs}}{\varepsilon_i}\right)^2, \tag{1.7}$$

$$\phi_m = \frac{1}{2}\alpha_s\int_{\Omega}\{\mathbf{w}_s(\mathbf{m}-\mathbf{m}^{ref})\}^2 dv + \frac{1}{2}\sum_{i=x,y,z}\alpha_i\int_{\Omega}\left\{\mathbf{w}_i\frac{\partial(\mathbf{m}-\mathbf{m}^{ref})}{\partial_i}\right\}^2 dv. \tag{1.8}$$

In the data misfit term (equation 1.7), N is the total number of data, F is the forward modeling from model to data, $\mathbf{d}^{obs}$ is the observed data, and ε is the estimated uncertainty. In the model norm term equation 1.8, Ω is the modeling domain, $\mathbf{m}^{ref}$ is a reference model, $\alpha_s, \alpha_x, \alpha_y, \alpha_z$ are scalar weighting parameters adjusting the relative importance of different components in the model norm, and $\mathbf{w}_s, \mathbf{w}_x, \mathbf{w}_y, \mathbf{w}_z$ are vector weighting parameters fine tuning the penalty to the local structures cell by cell.

Expressing equation 1.7 and 1.8 in discrete form, the objective functional in equation 1.6 can be written as

$$\phi = \frac{1}{2} \left\| \mathbf{W}_d \left[F(\mathbf{m}) - \mathbf{d}^{obs} \right] \right\|_2^2 + \beta \frac{1}{2} \left\| \mathbf{W}_m \left(\mathbf{m} - \mathbf{m}^{ref} \right) \right\|_2^2, \tag{1.9}$$

where $\mathbf{W}_d$ can be a diagonal weighting matrix containing the information about the data uncertainty, $F(\mathbf{m})$ and $\mathbf{d}^{obs}$ are vectors of data, and $\mathbf{W}_m$ can be an assembly of one diagonal matrix and three directional first-order differential matrices with corresponding weights coded.

Equation 1.9 can be minimized using different optimization techniques that utilize gradient direction, an approximated second-order derivative, or even a full second-order derivative Hessian matrix. Generally gradient methods take less effort to compute but converge more slowly than those making use of curvature information. I use the Gauss-Newton (GN) method as the basis of my inversion algorithms.

Differentiating equation 1.9 with respect to m and setting the gradient to zero yields a system

$$g(\mathbf{m}) = F'(\mathbf{m})^\top \mathbf{W}_d^\top \mathbf{W}_d \left[F(\mathbf{m}) - \mathbf{d}^{obs} \right] + \beta \mathbf{W}_m^\top \mathbf{W}_m (\mathbf{m} - \mathbf{m}^{ref}) = 0, \tag{1.10}$$

where $F'(\mathbf{m})$ is the derivative of the forward modelling operator with respect to $\mathbf{m}$. $F'(\mathbf{m})$ is often called the sensitivity or Jacobian matrix and denoted by $\mathbf{J}$. Suppose the current model is $\mathbf{m}^k$, a model update $\delta\mathbf{m}$ is sought so that $\mathbf{m}^{k+1} = \mathbf{m}^k + \delta\mathbf{m}$. For non-linear problem $\mathbf{J}$ depends on $\mathbf{m}$ and local linearization gives

$$F(\mathbf{m}^{k+1}) = F(\mathbf{m}^k + \delta\mathbf{m}^{k+1}) \approx F(\mathbf{m}^k) + \mathbf{J}(\mathbf{m}^k)\delta\mathbf{m}^{k+1}. \tag{1.11}$$

Equation 1.11 is exact for a linear problem.

Substituting equation (1.11) into equation (1.10) and rearranging yields an equation for the model update $\delta\mathbf{m}^{k+1}$

$$\left[\mathbf{J}(\mathbf{m}^k)^\top \mathbf{W}_d^\top \mathbf{W}_d \mathbf{J}(\mathbf{m}^k) + \beta \mathbf{W}_m^\top \mathbf{W}_m\right] \delta\mathbf{m}^{k+1} = -g(\mathbf{m}^k). \qquad (1.12)$$

where $\mathbf{J}(\mathbf{m}^k)$ is the Jacobian matrix of the model $\mathbf{m}^k$. The updated model $\mathbf{m}^{k+1} = \mathbf{m}^k + \alpha\delta\mathbf{m}^{k+1}$ where α is a step length parameter obtained by a line search. Since the optimization is an iterative process and the model update $\delta\mathbf{m}^{k+1}$ does not need to be exact, equation 1.12 is often solved by iterative solvers, for example, the conjugate gradient (CG) methods (Hestenes & Stiefel, 1952). In a CG solver, the coefficient matrix in equation 1.12 multiplies a vector $\mathbf{v}$ at every CG iteration. Essentially this means that the efficiency of optimization depends on the operation of dense matrix $\mathbf{J}$ or $\mathbf{J}^\top$ times a vector $\mathbf{v}$ ($\mathbf{J}\mathbf{v}$ or $\mathbf{J}^\top\mathbf{v}$). There are basically two ways of doing this. The explicit method first computes the elements in $\mathbf{J}$ by carrying out many forward-equivalent computations and storing $\mathbf{J}$ in the memory, then does $\mathbf{J}\mathbf{v}$ or $\mathbf{J}^\top\mathbf{v}$ upon request during CG iterations. The implicit method expresses $\mathbf{J}$ as a multiplication of some other matrices, including the forward modeling operation $\mathbf{A}^{-1}$ in equation 1.5; when $\mathbf{v}$ is applied to $\mathbf{J}$, the forward modeling actually takes place at every CG iteration. The choice of explicit or implicit method should be problem-dependent and is discussed in later chapters when specific problems are concerned. (See details about the explicit and implicit methods in Appendix B.)

The minimizer $\mathbf{m}$ of equation 1.6 obtained by solving equation 1.12 may not be satisfactory if β is too large and the data are under-fit. Some approaches have been developed to find an appropriate β (Constable et al., 1987; Hansen & O'Leary, 1993; Haber & Oldenburg, 2000). I use an approach called the "cooling method" to find the largest β that offers an acceptable data misfit (Haber, 1997). The procedure can be described as following. The inversion starts with an over-estimated β implying the simplicity of the model is more important than the reduction of data misfit and the data are only roughly fit; for a particular β, it takes a few GN updates δm to achieve the optimality; then β is reduced for another round of GN updates with the new β and the starting model from the previous β; the "cooling" of β terminates once the model contains enough structure to reasonably reproduce the observed data. The two levels of iteration involved here are called "outer iteration" for the β reduction and "inner iteration" for every GN update. The cooling approach is a process of progressively transferring information from the data to the model.

One single solve of equation 1.12 requires computations of forward responses $F(\mathbf{m})$ and the sensitivity matrix $\mathbf{J}(\mathbf{m})$, which counts for many forward-equivalent computations. On top of that, the outer and inner iterations involved in the optimization problem add even more forward modeling computations. One should also note the inversion is non-unique and more inversions must be carried out for robust interpretation. All these together make the inversion process very expensive as hundreds or even thousands of forward modeling computations are required.

1.3.3 State of the art

Since the pioneering work of 3D numerical modeling in 1970s and 1980s, primarily based on the integral equation method (Hohmann, 1975; SanFilipo & Hohmann, 1985; Newman et al., 1986; Wang et al., 1994), a variety of 3D algorithms using different numerical techniques have been proposed. Despite the success of the laboratory development, people have found the technology very difficult to use in practice because the 3D modeling is usually very time-consuming and requires substantial amount of computing resources. Therefore, in the last decade we have seen much research focusing on the improvement of the efficiency using different approaches. In particular, I have found the following work relevant and inspiring:

- **Domain decomposition**. The difficulty of 3D modeling is mostly associated with the large number of cells within the modeling domain. The most straightforward solution is to decompose the large domain and solve the subdomain problems using parallel computing (Alumbaugh et al., 1996; Newman & Alumbaugh, 1997; Xie et al., 2000; Commer & Newman, 2004; Holtham & Oldenburg, 2012), and as many as 32768 processors have been reported in the parallelization (Commer et al., 2008). All of the efforts have made the large-scale problems more scalable, but problems still exist. For example, a fine-grained domain decomposition has many small subproblems at the cell level and in theory is able to utilize a large number of processors, but the communication overhead may be significant enough to prohibit the use of additional processors. Also, balancing load across processors needs sophisticated task scheduling and assignment (Commer & Newman, 2004, 2008).

- **Direct solvers**. In the past few years, there has been a trend of migrating from iterative solvers to direct solvers when solving the forward modeling in equation 1.5. For a single forward modeling of a few sources, equation 1.5 can be inexpensively solved by iterative solvers such as Krylov subspace methods (Haber et al., 2007; Commer & Newman, 2008). However, when solving the inverse problem, which involves many $\mathbf{A}^{-1}$ operations for the forward modeling and sensitivity, a direct solver is preferable, because if $\mathbf{A}^{-1}$ is expressed as factorized triangular matrices many different sources and sensitivities of different data can be rapidly solved. With the help of the newly developed parallel direct solvers, like MUMPS (Amestoy et al., 2006), the large Maxwell matrix $\mathbf{A}$ now can be efficiently factorized so that direct solvers can speed up the modeling of realistically-sized problems (Oldenburg et al., 2013; Grayver et al., 2013). This thesis develops new strategies based on the algorithm outlined in Oldenburg et al. (2008, 2013), and refers to that as the "standard algorithm" since it was state-of-the-art when proposed, but it is a benchmark that I want to surpass. Development of the basic EM algorithm is not original work of this thesis. Rather, I begin with those basic developments and generate a new methodology to solve problems that were too large for the standard algorithm and to greatly speed up those that could be solved. Some necessary details on the derivations and implementations of the algorithm are provided in Appendix A and B.

- **Multiple meshes**. Besides the above, another recent trend of modeling is emerging in different areas. That is the use of multiple meshes in one modeling: Commer & Newman (2006) used a fine mesh for the early time modeling and switched to a coarse mesh for the late times during time stepping; as a further development, Commer & Newman (2008) decoupled the model mesh and the simulation mesh, so that the simulation mesh can be optimized to the frequency and source-receiver orientation after the data decomposition of marine CSEM; Wang et al. (2009) inverted borehole EM data by using multiple meshes with moving refinements at different small intervals along the well and modeling the regions outside of the focused sections by coarse background cells; Cox et al. (2010, 2012) used the concept of footprint in the integral equation formulation to invert the airborne EM data. Thus every airborne sounding only uses a small part of the entire 3D mesh and the sensitivity matrix is dramatically sparsified leading to a more efficient **Jv** computation. The speed-up due to the use of multiple meshes has been shown to be significant.

While the previous works each improve efficiency, none is a panacea and a technique with only a particular improvement soon finds itself in a situation of "too big to solve", if the size of the survey continues to expand. This imposes serious challenges, but offers great opportunities for my research. My thesis gains inspirations from the ideas scattered in the above seemingly separate improvements and consolidates them into one generalized theory of survey decomposition. The new theory does not only reveal why the previous improvements work, but also points out a new technical route that uses the theory to the ultimate extend for even better efficiency.

1.4 Research Topics

1.4.1 3D modeling: a pending paradigm shift

Now we are in a particular time, when the basic 3D EM modeling algorithms have been developed in the laboratory, high performance computing devices have become easily available and more quantitative information in 3D or even 4D is expected from geophysics. An era of interpretation based on full 3D modeling may soon become a reality. However, while it is true that more TEM data sets have been processed by 3D inversion than ever before, the fact remains that the majority of TEM data are still processed by non-3D approaches. In particular, practitioners are facing some serious challenges in 3D EM modeling:

- **Expense**. In spite of progress made, 3D modeling is still too expensive in that it usually takes a long time and substantial amount of memory on high-performance computers.

- **Big data**. Advanced acquisition systems now deliver very large data sets. Modeling the entire large-scale survey can be extremely expensive or even impractical with current technologies.
- **Emerging computing resources**. 3D EM modeling is traditionally carried out on dedicated work stations or computer clusters. Although new computing devices/platforms, for examples, the general purpose GPU (graphics processing unit) and the cloud computing, are invented to handle big data, it is still not clear how 3D EM modeling can effectively benefit from the large number of distributed processors.
- **Necessary**? The high costs associated with the 3D modeling make it a hard decision when choosing to work in 3D. Without solid evidence of why 3D modeling must be used, people tend to lean towards other more economical non-3D solutions.

I realize there are fundamentally two questions that must be properly answered before the paradigm shift can take place: (1) Is 3D modeling really necessary? (2) If yes, with the help from massive parallelization, what can be done to further improve its efficiency for the large-scale problems? In this thesis, the first question motivates my research and the second question is the ultimate problem I solve.

1.4.2 Thesis outline

Before diving into full 3D modeling, I investigate the rationale behind 3D inversion: that is, why and under what circumstance, 3D inversion is necessary. In Chapter 2, I demonstrate the superior results obtained by 3D inversion, and also show how the restrictions of dimensionality in model parameterization can harm an inversion. I further critically examine the conditions of using 3D inversions. This sets the fundamental motivation of working in 3D.

Chapter 3 introduces the concept of survey decomposition: why geophysical EM data should be modeled in this way, and how it is implemented for different survey configurations. This forms the proposed theoretical framework upon which the remaining chapters are built.

The framework of survey decomposition is applied to airborne and ground loop

EM surveys in Chapter 4 and 5, respectively. These two parallel chapters include the details about the implementations of the proposed methods and provides both synthetic and field data examples as proof of usefulness.

Chapter 6 concludes the research and makes statements on the overall importance of my work within the context of the paradigm shift. Further improvements and future work are also discussed.

Chapter 2

Why 3D Modeling

While 3D inversion in principle was always thought to be desirable, geophysicists have long been convincing themselves that lower dimensional inversions could at least yield a reasonable approximate model to the ground truth for interpretation. This chapter shows how 1D airborne TEM inversions can fail when the assumption of dimensionality is violated. It is based on a real example from a porphyry mineral deposit. Analysis from another airborne EM data set from the same site however shows 1D inversion may be acceptable sometimes. Generally, except in certain circumstances, 3D inversions should be carried out to prevent misinterpretation.

Insights developed from this chapter, regarding the sensing volume of particular survey equipment, form part of the foundation for a later analysis in this thesis.

2.1 Field Example: Airborne TEM Data at Mt. Milligan

2.1.1 Geologic background

Mt. Milligan is an alkalic porphyry Cu-Au deposit situated approximately 155 km northwest of Prince George in central British Columbia, Canada. The deposit was formed within the Early Mesozoic Quesnel Terrane, a Late Triassic to Early Jurassic magmatic arc complex that lies along the western North American continental margin where many similar porphyry deposits have been discovered in the past (Welhener et al., 2007).

Extensive exploration programs employing drilling, geologic, geochemical and geophysical surveys have been carried out since the 1980's. Mt. Milligan consists of several mineralization zones of copper and gold (MBX, DWBX, 66 Zone, Southern Star shown in pink in Figure 2.1). The mineralized zones are associated with monzonite stocks that have intruded into basaltic volcaniclastic rocks. Three types of alteration: potassic, albitic (sodic-calcic) and propylitic, can develop in a spatial geometry represented in Figure 2.2a. The alteration model, obtained by drilling around the MBX stock (Figure 2.2b), has a pattern similar to the theoretical model. The whole deposit is overlain by a layer of Tertiary sedimentary overburden that becomes thicker on the east side.

Despite extensive drilling, little logging for the electrical conductivity has been done, so the true conductivity model at the deposit scale is not available. However some general information about the conductivity is thought to be known. There are basically four different conductivity targets at Mt. Milligan. The most conductive unit is expected to be the overburden that is saturated or semi-saturated by ground water. Faults and fractured rocks are also likely to be conductive, especially when they are largely connected and serve as the pathway for groundwater. On the other hand, solid and unaltered igneous rocks, whether intrusive or host, should be resistive. The alteration zone is the most complicated unit and its conductivity may vary over a broad range of values because of different types and degrees of alteration. Some laboratory conductivity measurements of hand samples funded by Geoscience BC (Mitchinson & Enkin, 2011) show that the propylitically and albite-altered basalt can be relatively more conductive than other types of alteration due to high abundances of sulphides (pyrite and chalcopyrite). The monzonite, even altered, is still very resistive. Since airborne TEM is most sensitive to conductors, I expect that the airborne TEM survey at Mt. Milligan can delineate the overburden, water-saturated fractures and some conductive alteration that contrast with resistive rock units, including potassically-altered basalt, unaltered basalt and all types of monzonite.

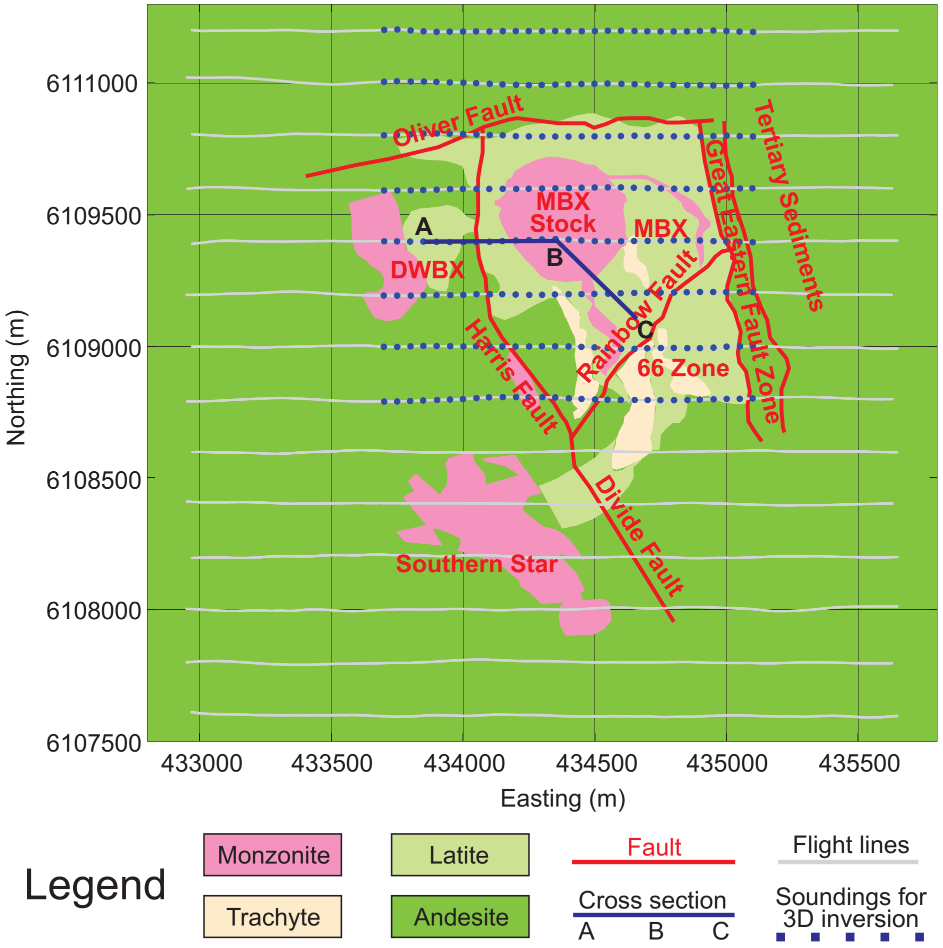

Figure 2.1: General geologic structure and airborne TEM flight lines at Mt. Milligan.

2.1.2 VTEM survey

In 2007, an airborne TEM survey (versatile time-domain electromagnetic, VTEM, a system operated by Geotech Ltd.) was carried out over the Mt. Milligan deposit. The survey consisted of 14 lines, each 2.7 km long and 200 m apart (grey lines in Figure 2.1). The VTEM survey covered all the mineralized zones associated with the Mt. Milligan porphyry complex: the MBX, DWBX, 66 Zone and Southern Star.

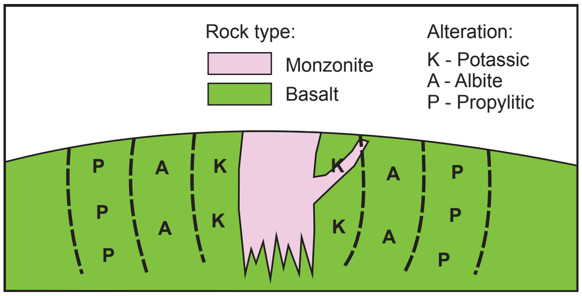

(a) Theoretical model with intrusive monzonite and surrounding alteration shells (Mitchinson & Enkin, 2011).

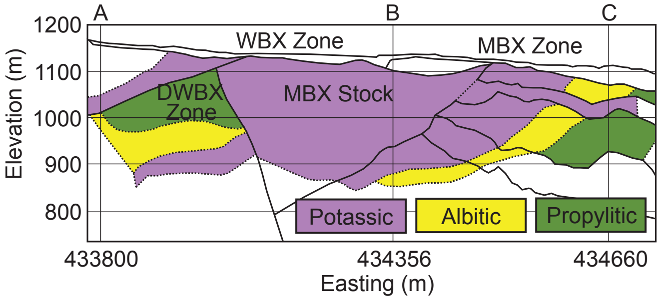

(b) Alteration model around the MBX stock at Mt. Milligan by Jago (2008); location of this cross section (A-B-C) is indicated in Figure 2.1.

Figure 2.2: Cross section of porphyry system at Mt. Milligan.

The VTEM system (2007) has a transmitter loop of 13 m radius and dipole moment of 503000 Am^2. The bird, including coincident transmitter and receiver, is towed about 42 m below the helicopter. The helicopter is positioned by onboard GPS and a radar altimeter. The system measures the vertical component of dB/dt data at 27 off-time time channels from 99 to 9245 μs. These time channels are marked on the half-cycle of the transmitter current waveform in Figure 2.3. The VTEM system at Mt. Milligan has a base frequency of 30 Hz. In the final deliverables, the 30 waveforms per second are stacked (averaged) and then the data are filtered to give a reading every 3 to 5 m along each traverse.

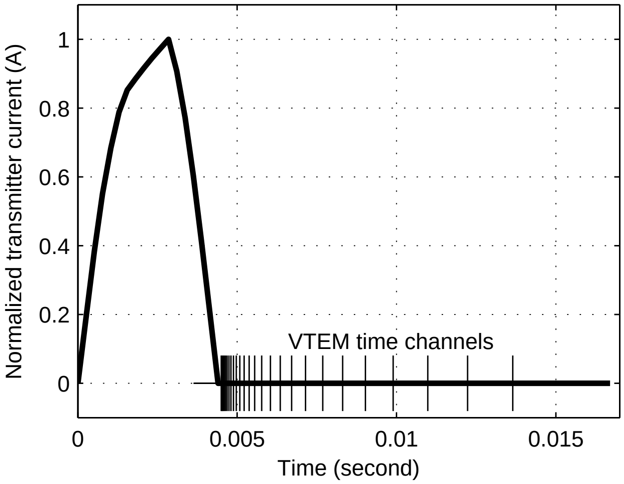

Figure 2.3: Transmitter current waveform and time channels of VTEM survey at Mt. Milligan.

2.2 Failure of 1D Interpretation

Plate modeling can be used to interpret airborne TEM data, but it works best with isolated good conductors and is not likely to generate a useful model for a porphyry deposit. There are also some other methods that assume a 1D conductivity structure routinely used in the industry: apparent conductivity, time constant analysis and 1D layered model inversion.

2.2.1 Apparent conductivity

Apparent conductivity is the conductivity of a uniform half-space that produces the same response as measured in a single time channel, or a best fit to a number of time channels. In this calculation, the flight height is taken into account. Here I use a look-up table similar to Sattel (2005) to rapidly find the half-space conductivity of a particular time channel at a certain height. If a solution is double-valued, the one minimizing the lateral conductivity variation is chosen. Figure 2.4 shows an apparent conductivity map of the entire survey at 0.4 ms (VTEM time channel 9).

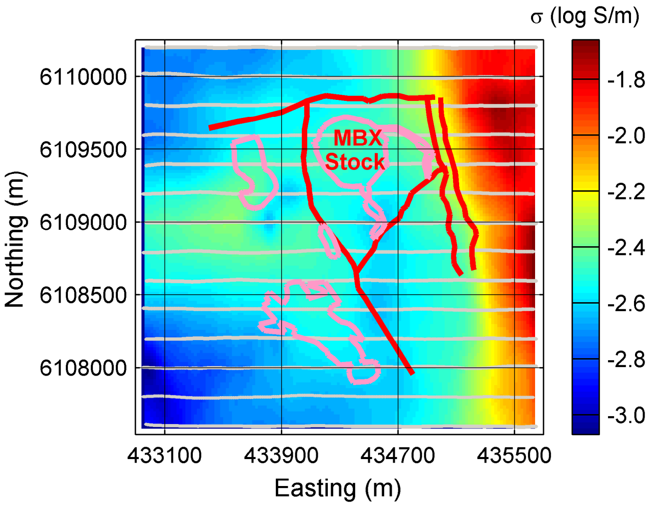

Figure 2.4: Apparent conductivity map of VTEM data at Mt. Milligan.

2.2.2 Time constant analysis

The time constant (or decay constant, τ) can provide a measure of the decay rate and is commonly used to rank conductors (Nabighian & Macnae, 1991). It is estimated by fitting the transient signals with an exponential decay curve. For each sounding I compute a τ that best fits VTEM time channels $6 \sim 12$. These time channels correspond to the expected depth of the deposit in this conductivity environment. The time constant map is shown in Figure 2.5.

2.2.3 1D layered model inversion

Finally I invert the VTEM data in 1D. While there are many 1D TEM inversion algorithms available, I use that outlined by Farquharson & Oldenburg (1993). For each sounding I seek a 1D layered model that reasonably reproduces the observed data and has minimum complexity of structures. In the 1D inversion, the top 600m is subdivided into 40 logarithmically spaced layers. An uncertainty of 10% of the datum value plus a floor value of 10^{-13} V/A is assigned to each datum. The starting model for the 1D inversion at each sounding is the best-fitting half-space.

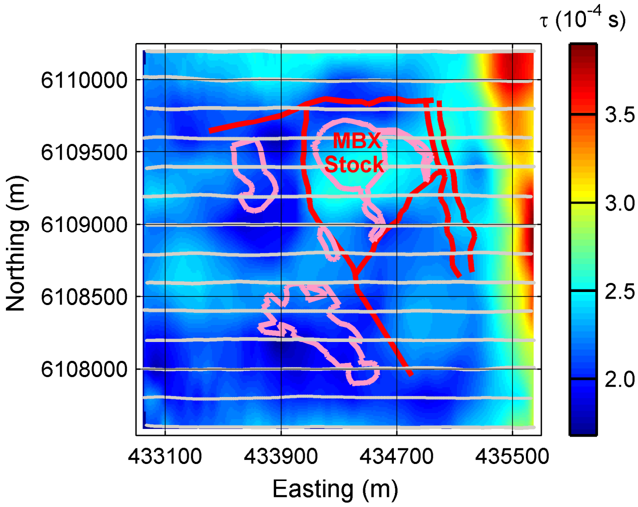

Figure 2.5: Time constant map of VTEM data at Mt. Milligan.

The 1D models recovered from the inversion are stitched and interpolated to form a pseudo-3D volume of conductivity. Figure 2.6 shows the depth slice at an elevation of 950 m.

These three conventional data interpretation methods reveal very similar patterns of conductivity. At the large scale there is good correspondence with regional geology, particularly the contrast between the conductive sediments on the east side and more resistive igneous rocks dominating the rest of survey area. There are similarities regarding smaller features, and in particular they all suggest that the conductivity associated with the MBX stock is higher than the host. This stock however is well-known to be monzonite and thus electrically resistive. In the following section, I will demonstrate that the misinterpretation arises from the assumption that the earth is 1D.

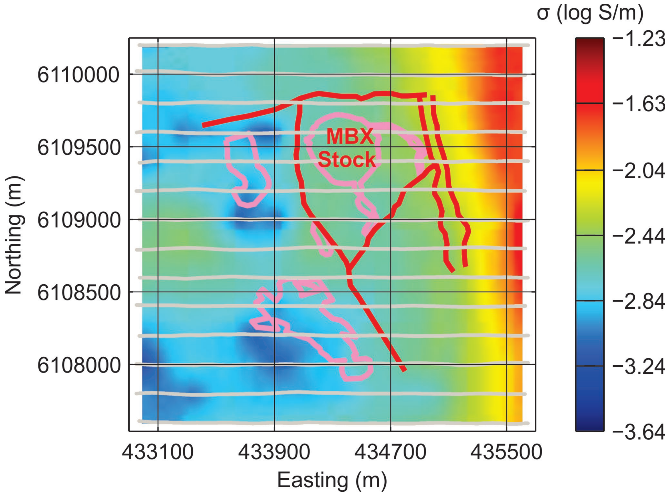

Figure 2.6: Depth slice of stitched 1D inversion models of VTEM data at Mt. Milligan at an elevation of 950 m.

2.3 Studies on a Synthetic Model

2.3.1 Synthetic 3D model

To investigate whether the high conductivity that coincides spatially with the location of the monzonite stock is a 1D inversion artifact because of 3D geometry, I design a synthetic porphyry-type model that consists of a thin conductive overburden overlying a resistive intrusive stock surrounded by a conductive alteration shell. The host rock is more conductive than the stock but less conductive than the alteration shell. The synthetic survey consists of 91 soundings at 50m above the surface on a 7 × 13 data grid (Figure 2.7). The transmitter current waveform is a step-off. Airborne TEM data are forward modeled at 20 VTEM time channels from 0.099 to 2.745 ms using the standard algorithm (Oldenburg et al., 2013). The data are noise-free but a 10% standard deviation is assigned to each datum for inversion.

2.3.2 1D inversion

To emulate the previous analysis on the field data, I attempt to generate an approximate conductivity model using 1D inversion. The inversion parameters are the

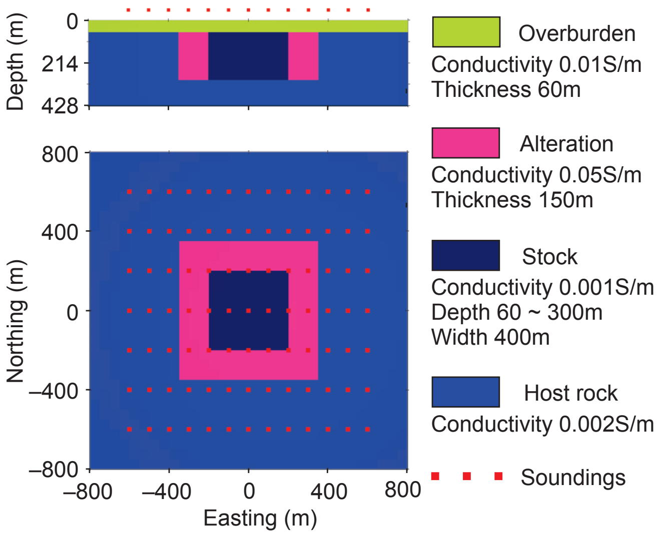

Figure 2.7: Synthetic conductivity model of a porphyry system.

same as for the 1D inversion of field data in the previous section. Figure 2.8 shows the central cross section and depth slice (160 m below surface) of the pseudo-3D volume. The black lines outline the true model. The 1D inversion recovers the overburden but the remaining features are incorrect. There is a large conductor sitting at where the resistive stock is supposed to be and the conductive shell is missing. The deep extension of the conductor on the cross section of Figure 2.8 also conveys very misleading information that the conductor is rooted at the depth of bedrock.

In an attempt to understand the above results, I model the currents and secondary magnetic fields in the earth at various times due to a transmitter located directly above the resistive stock. At early times the induced currents are near the surface and localized beneath the transmitter (Figure 2.9a). Hence a 1D inversion of the early time data should yield a reasonable estimate of the conductivity beneath the transmitter. However, at middle and late times, the peak currents propagate outward and downward and are concentrated in the conductive shell around the stock (Figure 2.9b). Once the currents reach the conductive material they decay slowly in accordance with the local conductivity. These trapped currents produce

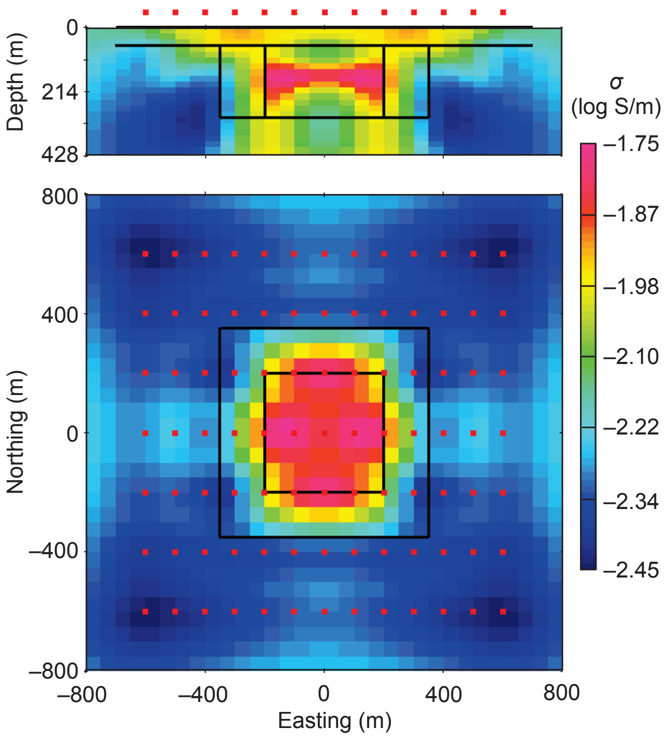

Figure 2.8: The recovered model from 1D inversion of the synthetic data.

significant dB/dt signals measured at the receiver and that signal can be translated into high conductivity beneath the receiver when the data are inverted in 1D.

The above phenomenon is well known and in fact forms the basic concept of why TEM surveys are useful. What was not appreciated before this analysis is how dramatic the effect can be in 3D. The circular geometry of the porphyry deposit and the location of the sounding maximize the potential for a significant artifact but other simulations, where the transmitter is displaced well away from the center of the deposit, produce a similar artifact. Once any appreciable amount of current reaches the conductive material, a significant signal will be recorded at the receiver location.

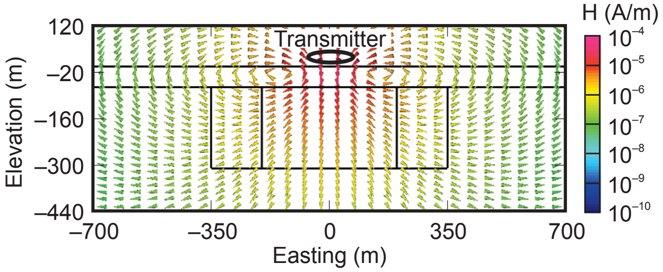

(a) At early time the currents are localized in the overburden close to the transmitter.

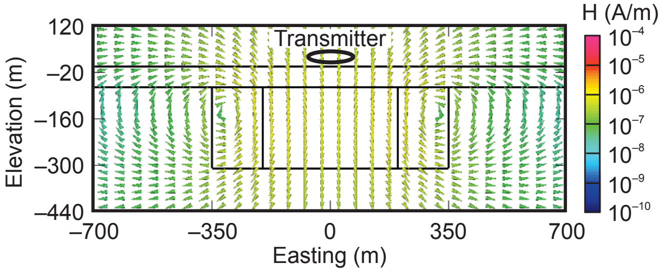

(b) At late time the currents are mostly trapped in the conductive shell.

Figure 2.9: Snapshots of the secondary magnetic field excited by a transmitter loop above the center of the resistive stock.

2.3.3 3D inversion

Having now understood that the correct assumption of dimensionality is crucial, we proceed with 3D inversions of the same synthetic data using Oldenburg et al. (2013). For numerical accuracy, the smallest cell, $50 \times 50 \times 20$m, is about one third of the diffusion distance of the earliest time for a 0.0067 S/m half-space. Padding cells are also needed so that appropriate boundary conditions are satisfied. The final mesh, $48 \times 48 \times 44$, consists of 101376 cells. The uncertainty assignment is the same as for the synthetic 1D inversion shown in the previous section. Initial and reference models are uniform half-spaces of 0.005 S/m. In addition to obtain a 3D inversion model, I would also like to examine how the number of soundings affects the images of inversion. Therefore, three individual inversions using 16, 49 and 91 soundings are carried out. The CPU times are summarized in Table 2.1. All inversions are carried out on 3 nodes of computer cluster, with 6 Intel Xeon E5410

Table 2.1: Summary of three inversions using different numbers of soundings for the synthetic data.

	Synthetic I	Synthetic II	Synthetic III
Number of transmitter	16 (4 × 4)	49 (7 × 7)	91 (7 × 13)
In-line sounding spacing	400 m	200 m	100 m
Cross-line sounding spacing	400 m	200 m	200 m
Number of iterations	4	5	5
Time for one forward modeling	230 s	780 s	1320 s
Total time for inversion	290 min	1108 min	1951 min

CPUs and 96GB RAM available.

Cross sections and depth slices of the three synthetic inversion models, and corresponding data misfit curves, are summarized in Figure 2.10. Even with only 16 soundings 400 m apart (Synthetic I in Figure 2.10a), the 3D inversion shows the conductivity contrast between the central resistive stock, the surrounding conductive shell and the host. The conductivities and geometry of anomalies are distorted, but the concept of a conductive shell around a resistive stock is evident. This contrasts with the incorrect image shown in Figure 2.8. Increasing the number of soundings to 49 in Synthetic II (Figure 2.10b) greatly improves the image. All three geologic units, overburden, conductive shell and resistive stock, are correctly delineated and have more realistic conductivities. A further increase to 91 soundings in Synthetic III produces only marginal improvement by sharpening up the boundaries (Figure 2.10c). This improvement may not be justified by the increased computational time required for the inversion (see CPU time in Table 2.1).

To summarize, by carrying out synthetic 1D and 3D inversions I have learned that: (1) 3D inversion is necessary for geologic units with 3D geometry because 1D inversion can produce serious artifacts; (2) 3D inversion, even with a few soundings can still recover a reasonable large picture that out-performs a 1D inversion result; (3) there is a redundancy of information in the airborne TEM data so an efficient inversion should strive to reduce the number of soundings involved while maximizing the amount of information from the data that can be incorporated into the model.

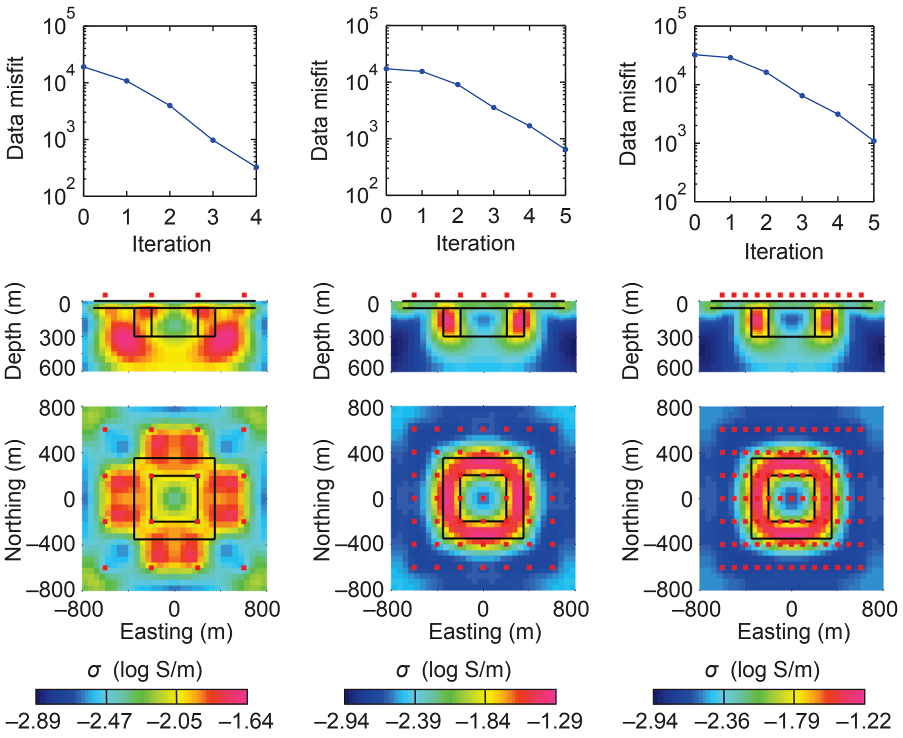

(a) Synthetic I **(b)** Synthetic II **(c)** Synthetic III

Figure 2.10: Data misfit convergences and models of three synthetic 3D inversions. The red dots indicate the sounding locations and the black lines outline the true model.

2.4 3D Inversion of Mt. Milligan Data

2.4.1 Inversion set-up

With the knowledge learned in the synthetic inversions, I invert the VTEM data set at Mt. Milligan in 3D to see how 3D inversion can improve interpretation. As a proof of concept, I focus upon a 1.4 × 1.4 km region surrounding the MBX stock (433700∼435100E, 6108800∼6111200N). The earliest time channels in an airborne TEM data set are usually difficult to model accurately because the exact transmitter current close to the turn-off time may vary from sounding to sounding, and thus I invert the data from more stable time channels 6 ∼ 20. The diffusion distance of the first time channel used in inversion, determined by the earliest time and the overall conductivity obtained from 1D inversion, suggests the 3D inversion would not have resolution finer than about 50 m. Therefore, the airborne soundings are down-sampled to 50 m spacing, giving a total of 232 soundings on a 8×29 data grid (blue dots on grey lines in Figure 2.1).

The smallest cell size in horizontal directions is chosen to be 50 m, so the finest information in the data can be adequately modeled. In the vertical direction, the smallest cell size is 20 m. This is needed to capture the topography and accommodate the transmitter flight heights varying from 20 m to 90 m above surface. The final mesh is $50 \times 50 \times 64$ and has a total of 160000 cells.

The assigned uncertainty for each datum is 10% plus a noise floor of 10^{-13} V/A; this is the same as the 1D inversion of Mt. Milligan VTEM field data. Discretization of the transmitter waveform is important since this affects the numerical accuracy, time and memory usage when solving the forward problem. I discretize the waveform shown in Figure 2.3 to have a minimum number of partitions, each of which is discretized in a uniform step length, and still retain desired accuracy. This results in two on-time and four off-time Maxwell matrix factorizations and a total of 50 time steps in one forward modeling with the algorithm in Oldenburg et al. (2013).

2.4.2 A multi-level approach

With my current computer the time for carrying out 50 time steps for 232 transmitters on a mesh with 160000 cells requires 33 minutes. Unfortunately an inversion typically requires hundreds of equivalent forward modelings and this makes 3D inversion extremely time-consuming. For practical purposes, it is always desirable to reduce the computing time. The reason why 3D forward modeling is slow is that there are a large number of cells in the mesh and many soundings (sources) need to be modeled. One way to speed up the computation is to adopt a multi-level concept: use approximations at the early stage of inversion when the model is far away from the solution; then gradually refine the modeling to resolve finer-scale features. I note that in airborne TEM inversion, a coarser mesh that moderately violates the general rules of spatial discretization may not hurt the model update at

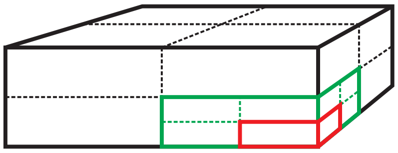

Figure 2.11: Dimensions of the smallest cells in the extra-coarse (black), coarse (green) and fine (red) meshes.

the early stage of 3D inversion, because the data misfit is usually orders of magnitude greater than the discretization error. Likewise, including all the soundings in an early stage of inversion is not necessary and only a down-sampled subset is actually needed. I therefore propose a multi-level strategy to increase the efficiency of the large inversion problem. That is, first invert a small number of soundings on a coarse mesh, then pass the recovered model to a finer mesh and refine the model by inverting additional soundings; repeat until the last inversion that is carried out on the finest mesh and includes all the soundings.

For this particular VTEM data inversion around the MBX stock at Mt. Milligan, I design the dimensions of the smallest cells in the multi-level meshes to have an octree-like relationship as portrayed in Figure 2.11. For the current problem I use three levels of mesh which are referred to as extra-coarse (X), coarse (C) and fine (F). The nested octree design reduces the size of the problem by factors of 8 and 64 on the coarse and extra-coarse meshes and guarantees exact model conversion between meshes. All three meshes have the same padding cells extending to 4200 m, which is calculated based on the latest delay time and a representative resistive background conductivity 0.001 S/m. Specifications of the extra-coarse, coarse and fine mesh inversions are listed in Table 2.2.

Numerical accuracy of the three meshes can be monitored by testing 3D forward modeled data against a 1D algorithm using conductive (0.025 S/m) and resistive (0.001 S/m) half-spaces. CPU times are calibrated for one forward modeling on an Intel i7 960 Quad-Core desktop computer with 16 GB RAM. As shown in Table 2.3, all three meshes do well at late delay times over resistive ground because the same padding distance is used. However, modeling errors at early times

Table 2.2: Specifications of the extra-coarse, coarse and fine mesh inversions.

	Extra-coarse (X)	Coarse (C)	Fine (F)
Smallest cell size (vertical)	200 m	100 m	50 m
Smallest cell size (horizontal)	80 m	40 m	20 m
Number of cells	18750	48552	160000
In-line sounding spacing	200 m	100 m	50 m
Cross-line sounding spacing	200 m	200 m	200 m
Number of soundings	64	120	232

Table 2.3: Numerical performances of the extra-coarse, coarse and fine meshes used in the multi-level inversion.

	Extra-coarse (X)	Coarse (C)	Fine (F)
Number of soundings	64	120	232
Early time channel modeling error	35%	18%	6%
Late time channel modeling error	3%	3%	3%
Total time for factorizations	17 s	75 s	734 s
Total time for time steps	67 s	331 s	3722 s

increase as the mesh discretization becomes coarser. Poor modeling often creates artifacts in the inversion model, but fortunately those artifacts can be corrected in a subsequent finer mesh inversion. Comparison of CPU time required by one forward modeling on the three meshes shows significant potential of speed-up by using coarser meshes in the early stages of the inversion.

2.4.3 3D inversion

I first implement the extra-coarse mesh inversion. The inversion starts with a 0.002 S/m half-space with topography and achieves a normalized data misfit 2.55 (blue curve in Figure 2.12) in about 1.73 hours (X-iterations in Figure 2.13). The model (X5) is converted to the coarse mesh and 100-meter-spacing soundings are used in the coarse inversion. Because of the information contained in those newly-added soundings, the initial normalized data misfit of the coarse mesh inversion rises up to 5.72 (C0 in Figure 2.12). The coarse mesh inversion takes another 5.74 hours (C-iterations in Figure 2.13) to reduce the normalized data misfit to 2.32 (red curve

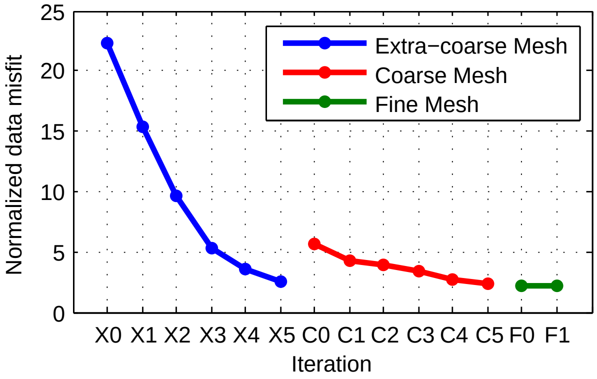

Figure 2.12: Normalized data misfit as a function of iterations for the inversions using three multi-level meshes.

in Figure 2.12). Finally the inversion is switched to the fine mesh and uses all of the soundings. However, these 50-meter-spacing soundings do not provide significantly more information than the 100-meter-spacing soundings. The initial normalized data misfit of the fine mesh inversion is 2.23 (F0 in Figure 2.12), which is even less than the final misfit of the coarse mesh inversion. The fine mesh inversion, which takes 10.65 hours for one iteration (F-iterations in Figure 2.13), is terminated when there is a reasonably good data fit; 80% of the 3445 data have normalized data misfit less than unity and there are only a few outliers that contribute large misfit (Figure 2.14). Note the extra-coarse mesh inversion (X) is carried out on a Intel i7 960 Quad-Core desktop computer with 16GB RAM available; the coarse mesh inversion (C) on one node of a cluster with 2 Intel Xeon X5660 CPUs and 64GB RAM available; the fine mesh inversion (F) on two nodes of the same cluster with 4 Intel Xeon X5660 CPUs and 128GB RAM available.

Depth slices at an elevation of 970 m of the three mesh inversion models (X5, C5 and F1) are shown in Figure 2.15. In spite of the substantial modeling error (Table 2.3), the extra-coarse mesh does a good job of delineating the large scale conductivity features over the entire area (Figure 2.15(a)). This inversion is carried out quickly on a desktop computer. Increased resolution is evident on the subsequent meshes. However, little difference is found between the coarse (Figure 2.15(b)) and fine (Figure 2.15(c)) mesh inversion models. This suggests that the 50-meter-

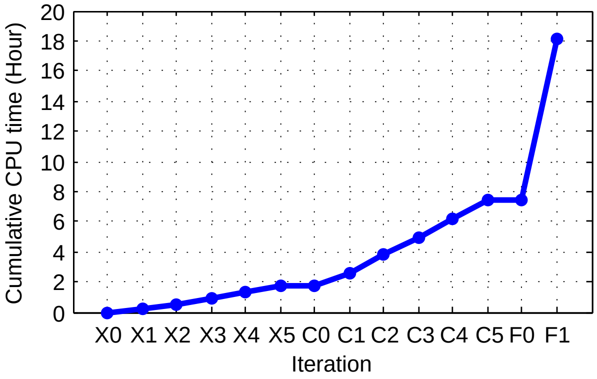

Figure 2.13: Cumulative CPU time of iterations in the extra-coarse, coarse and fine mesh inversions.

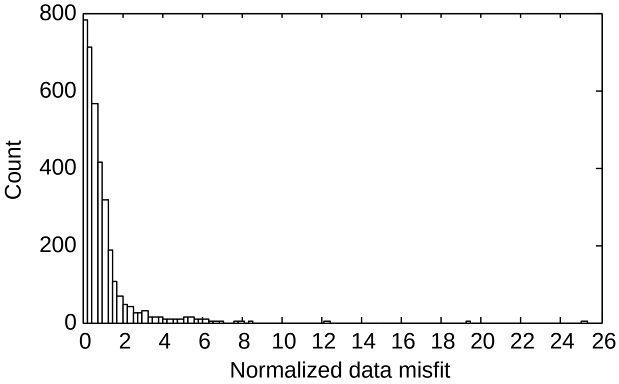

Figure 2.14: Histogram of normalized data misfit after the fine mesh inversion.

spacing data do not provide much more information than 100-meter-spacing data and that the 100-meter-resolution mesh (the coarse mesh) would probably be fine enough for this data set. The total time to complete the inversion at the coarse mesh scale is 7.46 hours and the total time, with a fine scale mesh update, is 18.11 hours. These are both significantly smaller than solving the problem on the fine scale from the beginning, which took 82.43 hours to achieve the same data misfit in a follow-up benchmark inversion.

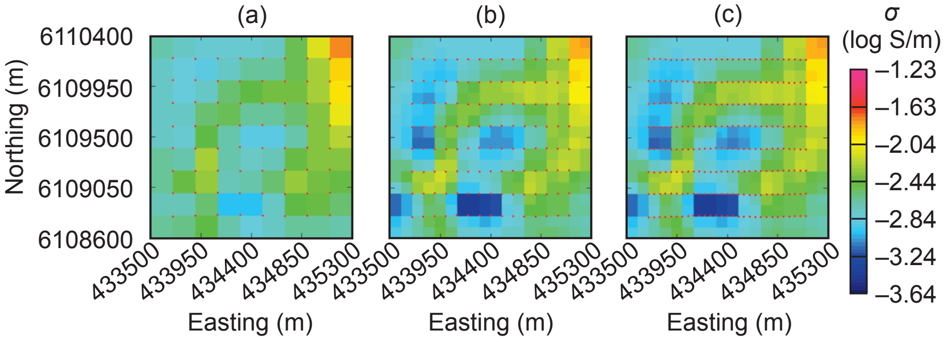

Figure 2.15: Depth slices of the inversion models at Mt. Milligan at an elevation of 970 m after the (a) extra-coarse, (b) coarse and (c) fine mesh inversions.

2.4.4 Model interpretation

Depth slices of the smoothed final model are shown in Figure 2.16. The most conductive feature in this model is a near-surface conductor that cuts the northeastern corner, which is interpreted as the Tertiary sediments. Our target, the porphyry system, is well recovered as a deeply rooted massive resistor in the center of study area and a relatively conductive ring around it with an open towards the south. This feature can only been seen in the 3D inversion model and not in other methods or maps assuming a 1D earth model.

Figure 2.17(a) is a depth slice of the inversion model at an elevation of 1030 m (about 100 m below the surface) with the geology overlaid. The general features are a resistive core, a surrounding conductive halo, followed by a resistive outer region. On the east, this sequence is modified by the presence of the highly conducting sediments. This conductivity model is also supported by a previous 3D inversion of DC resistivity data (Oldenburg et al., 1997). In Figure 2.17(b) I overlay contours of the conductive anomaly from the DC resistivity inversion at the 1030 m elevation. There is general correspondence, especially for the west-wing of the conductive ring, and the MBX stock is found resistive by both 3D airborne TEM inversion and DC resistivity inversion. Note that, in addition to the resistive stock which I have been focusing upon, there is a prominent dark-blue-colored resistor on the southern part of the image. This feature is required by the data which are characterized by

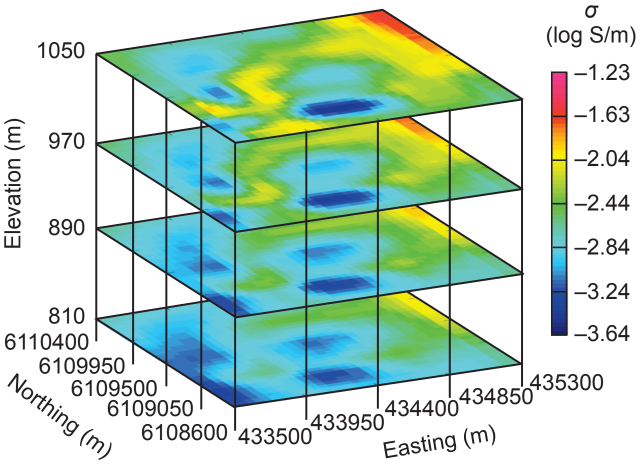

Figure 2.16: Depth slices of the final conductivity model at elevations of 1050 m, 970 m, 890 m and 810 m.

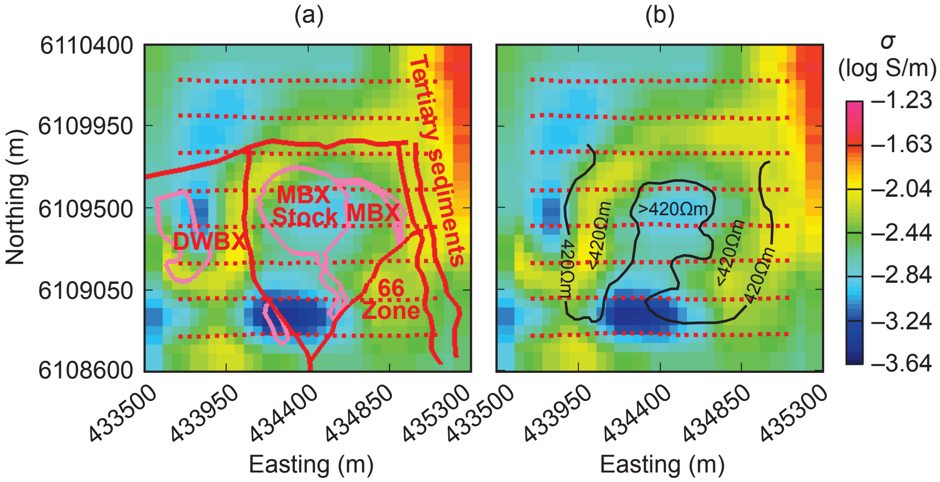

Figure 2.17: Depth slice of the final interpretation model at an elevation of 1030 m overlain by (a) geology and (b) 420 Ωm contour of the DC resistivity model.

unusual decay constants and negative transients in that region. Interpretation of those data may require modeling of complex conductivity for a dispersive medium (Marchant et al., 2013), and will not be discussed in this thesis.

The application of 3D inversion to VTEM data at Mt. Milligan has yielded a conductivity model that is geologically and geophysically more plausible than the one obtained in 1D inversion. This study shows that the correct dimensionality is crucial and by interpreting TEM data in 3D, misleading information, like that obtained from 1D modeling, can be prevented.

2.5 Critical Rethink: When is 3D Inversion Necessary?

I have shown, with a synthetic example and field data inversion, that a 3D inversion of airborne TEM data is necessary if the geologic setting is 3D similar to the Mt. Milligan porphyry deposit. The 1D inversions fail because of the lateral variation of the conductivity. However, in a previous study of Mt. Milligan (Oldenburg et al., 1997), 1D inversion of another airborne EM data set (frequency-domain DIGHEM, a system operated by CGG) successfully revealed the resistive nature of the MBX stock. This paradox leads to a critical question pertaining to circumstances about when a 3D inversion is necessary. I attribute this to the fact that the VTEM system, for the time channels used here, has a much larger footprint than the DIGHEM system. In this section I quantify the concept of footprint of an airborne EM system by using the distribution of electric current density and then discuss the conditions for undertaking 1D and 3D inversions.

2.5.1 Quantification of the EM footprint

The concept of footprint has been used to quantify the size of the spatial domain within which an airborne sounding can effectively sense the conductivity. Different approaches were proposed to calculate the size of footprint (Liu & Becker, 1990; Kovacs et al., 1995; Beamish, 2003; Reid & Vrbancich, 2004; Reid et al., 2006; Yin et al., 2014) and other closely related measurements (Smith & Wasylechko, 2012), but some limitations exist. I therefore use the similar idea in Liu & Becker (1990), but extend it to realistic 3D conductivity model and make it easier for the comparison between the time-domain and frequency-domain systems.

My approach is based on a quantity called "pseudo-sensitivity" that is calculated using the distribution of the induced current and the distance between the currents and the location of the receiver. Using the finite volume algorithm in Haber et al. (2000) and Oldenburg et al. (2013), the total strength of the current density vector at every cell for a particular frequency/delay time can be computed for any complicated 3D model and source; then the pseudo-sensitivity of a cell to the datum at the receiver for that frequency/delay time is calculated as the current density strength of that cell divided by the squared distance between the receiver and that cell. The pseudo-sensitivity honors the current density and the geometric decay of signal in the Biot-Savart law that describes how the induced currents can give rise to the magnetic field measured at the receiver.

Next, using superposition, the data at the receiver can be treated as the contribution made by many subsurface current density vectors. By summing up the pseudo-sensitivities of the cells within a progressively expanding radius R centered at the sounding, we can tell whether a system at a particular frequency/delay time is able to probe a relatively broad or small region. For the convenience of comparison, the summed pseudo-sensitivity within a particular R is further normalized by the summed pseudo-sensitivity when $R = \infty$, giving a relative measurement p $(0 < p < 1)$ called the "cut-off index". If p is given, the corresponding R can be searched by stepping at small increment. Then it is possible to compare different systems' R at the same p. Because the conductivity can have arbitrary variation in 3D, it may be necessary to examinate the entire p-R curve at different cut-off indices.

2.5.2 Comparison between VTEM and DIGHEM

This approach is applied to VTEM and DIGHEM surveys at Mt. Milligan. In my investigation, I take the 3D conductivity model of the VTEM inversion at Mt. Milligan, and place an airborne sounding above the center of the MBX stock (UTM: 434500E, 6109500N, elevation 1180 m). The earliest and latest times of VTEM, 234 μs and 2745 μs, and the highest and lowest frequencies of DIGHEM, 56000 Hz and 900 Hz, are considered here. The estimated footprints R versus different cut-off indices p are plotted in Figure 2.18.

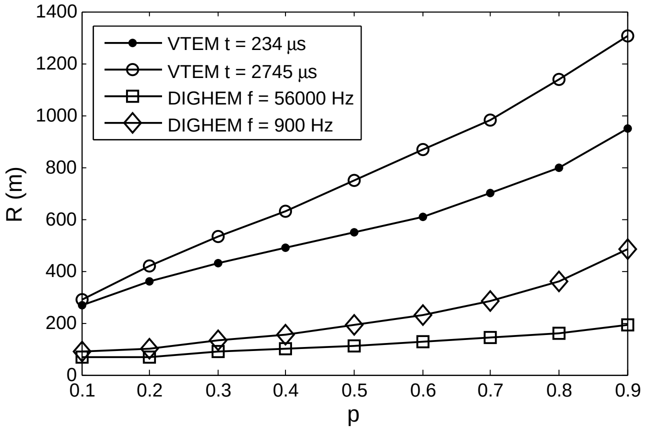

Figure 2.18: The footprints *R* of the VTEM and DIGHEM systems at different cut-off indices *p* at Mt. Milligan.

Unsuprisingly, later delay times, or lower frequencies, have larger footprint compared to earlier delay times, or higher frequencies within one system (Figure 2.18). But a big gap in the footprint between VTEM and DIGHEM is obvious: even the lowest frequency of DIGHEM has much smaller footprint than the earliest time of VTEM. At Mt. Milligan, the size of our target, the MBX complex, is comparable to the footprints of most VTEM delay times and is larger than the footprints of most DIGHEM frequencies. Therefore, the ground conductivity can be considered to be one-dimensional at the scale of DIGHEM data and a 1D inversion is able to correctly recover the local conductivity. At the scale of the VTEM data, the same conductivity structure becomes complex in three dimensions so that a 3D inversion with multiple soundings is needed.

When it comes to the applicability of 1D and 3D inversion, one needs to consider the measurement times/frequencies of the system and the conductivity. If the time channels are early enough, or if the frequency is high enough so that the footprint is smaller than the scale of conductivity variation, then the local conductivity information can be picked up by an individual sounding and a 1D inversion can produce an interpretable result. However, for later time channels, or lower frequencies, the EM footprints can match the scale of conductivity variation and a 3D inversion is necessary. If the footprint continues to expand and significantly exceeds the scale of the inhomogeneity, the 3D strucutres can be effectively represented by a uniform value by up-scaling.

2.6 Summary

My studies on the airborne TEM data at Mt. Milligan shows that under some conditions 3D inversion is more essential than previously thought. Although the circular geometry of the conductor at Mt. Milligan seems the most vulnerable case to artifacts in an 1D inversion, numerical experiments in other less extreme 3D and even 2D environments present similar issues that result in significant misinterpretations when the footprint of the EM system becomes comparable with the geologic targets in size. When using a large ground transmitter loop, the footprint could reach as far as tens of kilometers. For this reason it is desirable to have 3D inversion as a standard procedure for most TEM data sets, especially for applications, like mineral exploration, in which 3D structures are common.

However, improved interpretation comes at high costs. Even with advanced computers and sophisticated numerical solvers, 3D inversion of a data set as small as the focus area at Mt. Milligan MBX complex still takes dozens of hours on large computer clusters. This is far more expensive than 1D inversion and parametric 3D inversion, which can be carried out within minutes on regular desktop computers.

To speed up the Mt. Milligan airborne TEM data inversion, I used a multi-level method to save unnecessary fine-scale computing at early inversion iterations. The idea of “inaccurate computing” has proven to work well, but some problems exist:

- The scalability of 3D modeling is poor. As the survey gets large, the computation grows geometrically and eventually one needs to solve a huge problem on a massive mesh at later iterations; this could be difficult or impractical.

- Massive parallelization is difficult. If a problem is large, information used in computing has to be stored on different nodes across a computer clusters. So we will soon reach a point that the overhead of communication dominates, and adding more processors makes the inversion even slower.

- Soundings (data) are down-sampled on a regular grid based on the size of the cell in the mesh. The down-sampling rate, made to match the resolution

> of the mesh, was somewhat arbitrary. It was hard to know whether the rate was too high (loss of information) or too low (over-computing).

While the usefulness of 3D inversion is justified with a field example, its effectiveness in practice is heavily restricted by the efficiency of 3D modeling. The rest of this thesis will concentrate on a fundamentally novel framework of modeling that accelerates the existing algorithms at a generic level and makes 3D inversion much more affordable.

Chapter 3

Survey Decomposition

The benefits and necessity of 3D modeling demonstrated in Chapter 2 show that it is desirable to have EM data routinely inverted in 3D. However, such a shift of paradigm to full 3D is not easy because of the time and significant amount of computing resources required by realistic 3D modeling. This is exemplified by the 3D inversion of a small subset of the Mt. Milligan data in Section 2.4: a survey of 1.5 × 1.5 km took dozens of hours on two nodes of a computer cluster, even with a multi-level approach for speed-up. So, it becomes obvious that without a breakthrough in the efficiency, 3D modeling can hardly make an impact in real-world applications of exploration. Motivated by this issue of efficiency, the rest of my thesis research tries to make improvement by fundamentally changing the architecture of 3D modeling using an unconventional framework, called *survey decomposition*.

This chapter first discusses the rationale behind this novel framework by analyzing the computational complexities that have slowed the modeling. Then survey decomposition, along with other related concepts like subproblems, local discretizations, is explained in detail. I also discuss the performance and scalability of survey decomposition in parallel computing. Finally, an approach of data subsampling using cross validation, particularly useful in dealing with a large number of subproblems after decomposition, is introduced with a demonstrative example.

3.1 Computational Complexities

The idea of decomposing a large EM problem comes from a critical examination of the complexities involved in 3D modeling and the attempt to eliminate over-computing. Because of the multi-scale nature of EM modeling problems, the computational complexity arises from three aspects: space complexity, time complexity and optimization complexity (for inverse modeling only). The philosophy of "adapt to the scale" is used to fight against the over-computing in all three of these aspects.

3.1.1 Space complexity

The space complexity is often the primary concern in 3D modeling due to the large number of discretized cells (model parameters) and possibly the large number of sources. There are several typical scenarios that can lead to highly expensive computations:

- Large surveys. A survey that spatially covers an extensive area can result in numerous cells and thus a large system of equations, regardless of the type of mesh. In airborne EM, the surveys are usually flown over tens or hundreds of square kilometers demanding millions of model parameters; even for ground loop TEM, in which a few sources are fixed on the surface, the size of the ground loop can still require a considerably large volume in the modeling domain.

- Numerous EM sources. This is common in airborne EM surveys that consist of hundreds of thousands of sources (transmitters). Without modification this may require the million-parameter problem to be solved hundreds of thousands of times for just one complete forward modeling.

- Complicated model structures. If complicated structures, like topography, dykes and faults, are to be modeled, the mesh must be finely discretized; this adds more cells in addition to those needed to work with a large survey.

Space complexity is a prohibitive issue because of the notoriously poor scalability of 3D modeling: the increase of the costs (time and memory) is not proportional to the increase of the number of model parameters. I demonstrate the

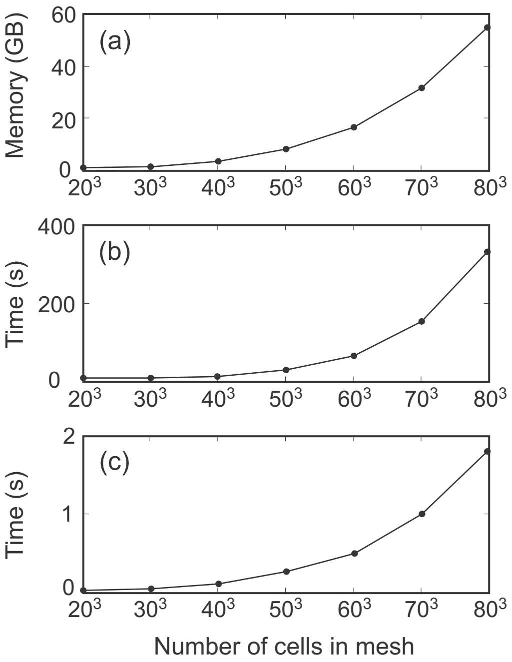

Figure 3.1: Computational costs of a single airborne TEM sounding as poorly scaled functions of the number of mesh cells: (a) memory usage, (b) time for one Maxwell matrix factorization, and (c) time for one time step.

demand of memory and time by a sequence of single-sounding airborne TEM forward modelings on meshes with a variable number of cells using the algorithm described in Appendix A. The forward modeling experiments are measured by the memory required by storing one factorized matrix (Figure 3.1a), the time required to carry out one Maxwell matrix factorization (Figure 3.1b), and time required to solve for the fields at one time step (Figure 3.1c). The poor scalability makes the forward solution computationally expensive when solving realistically-sized problems.

The first obvious solution is to reduce the number of model parameters as much as possible. Parametric 3D modeling can be seen as an extreme example using strong and explicit information for the model simplification. Some sophisticated meshes, like tetrahedral mesh (Schwarzbach & Haber, 2013), are used to do local refinement, so that the additional cells due to topography, or other complications of the model, can be minimized. The number of model parameters can also be reduced by representing the model in a transformed domain with lower dimension of model space (Oldenburg et al., 1993; Oldenburg & Li, 1994a,b).

If many sources are present, a direct solver of Cholesky or LU decomposition is usually used to save time when dealing with multiple right-hand sides in equation 1.5. Direct solvers are generally expensive in matrix factorization, but once the factorization is finished, many different sources can be quickly solved. This technique has now gained great popularity, thanks to the recent developments of matrix factorization algorithms (Amestoy et al., 2006; Oldenburg et al., 2013; Grayver et al., 2013).

Another approach which has been very actively pursued is to carry out a domain decomposition and solve the problem in parallel. Typically, the entire modeling domain is broken down into subdomains that can be solved in parallel. The solution on one subdomain usually depends on the information from other subdomains, so there are frequent intercommunications. If the original problem is decomposed into very fine grains, for example at a cell level (Newman & Alumbaugh, 1997), massive parallelization can be implemented but the benefit of adding more processors may be quickly marginalized due to the overhead of intercommunication. With fewer subdomains, a coarse-grained domain decomposition does not have excessive overhead on communication, but it is harder to utilize a large number of processors and maintain a good balance of workload among parallel processors.

All of the previous improvements have made 3D modeling more tractable, but when challenged by large data sets from large surveys, the size of the mesh and/or the number of sources are simply beyond the capability of the existing methodologies. This can be understood in terms of the scale of modeling: a mesh must have cells small enough close to every source/receiver, and the domain size must be large enough to encompass the entire data area and handle the boundary conditions properly. This scale contrast usually generates meshes with a very large number of cells.

Realizing the difficulty of using a global mesh for modeling, survey decomposition seeks a logical decomposition based on the sources and receivers that are spatially confined so that the meshes for the subproblems are localized and can have fewer cells than the global mesh. This leads to a reduced scale of contrast in a local mesh. Unlike straightforward subdivision of the space in a standard domain decomposition approach, one subproblem involving a physical source and a physical receiver is a self-contained EM problem and can be modeled independently with minimal communications with other subproblems. Collectively, the space complexity from a large mesh is conquered by modeling multiple small EM problems and putting them together at the end, rather than working on the original problem as a whole.

3.1.2 Time complexity

The time derivatives in equation 1.1c and 1.1d need the modeling of the time dynamics of the EM fields, which is usually implemented by either stepping in time or solving Maxwell's equations at selected frequencies and then transforming to the time domain. Time complexity arises because the phenomenon of EM diffusion span a range of time from 10^{-6} to 1 s, or a range of equivalent frequencies, for most earth models in exploration geophysics. If the equations are solved in the frequency domain, there is still complexity, since the EM problem must be solved at multiple frequencies.

If Maxwell's equations are solved by a time stepping method, the step length must be carefully chosen so that the step is neither too large, and thus cause excessive numerical errors, nor too small, and thus lead to many steps. One widely-adopted solution to this multi-scale problem in time is to adapt the step length to the scale of delay times: at the beginning of stepping, small steps are used to model the rapid change of fields over time; as the fields diffuse, longer steps are employed for efficient modeling without significant loss of accuracy. The choice of step length in fact also depends on the type of solver. Iterative methods often have logarithmic step lengths (Haber et al., 2007), but the solution time is too long. Direct methods are efficient if the time step is constant, but a single step length for the whole time range yields too many steps; therefore a trade-off is usually made so that the time discretization uses a few different step lengths from small to large, each of which is constantly used in a certain range of time at different scales (Oldenburg et al., 2013; Um et al., 2010).

Although in the time-stepping schemes mentioned above the step length is already adjusted to the delay times, there can still be significant contrast of scale within a global time discretization designed for all of the times, as the small step lengths at the beginning are only useful for early times and are unnecessarily fine for late times. This over-computing in time is the result of scale contrast between early and late time channels, similar to the modeling on a global mesh as discussed in the previous subsection. This complexity can be minimized by designing the time discretization for each delay time and modeling each time individually after the decomposition in time.

The similar minimum scale contrast can also be achieved by doing frequency-domain modeling at some discrete frequencies since each frequency represents a certain scale of modeling and many frequencies can be modeled separately. Then the question becomes choosing the frequencies. It has been shown that the optimal set of frequencies can be found by the reduced Krylov subspace (RKS) method (Börner et al., 2008; Zaslavsky et al., 2011, 2013). However, if the data are in time domain, the time complexity taxes the frequency-domain approach in other ways: (1) a complex system of equations has to be solved for every frequency and (2) a wide frequency spectrum may be required for an accurate frequency-time transform. Because of (1), a single frequency modeling in frequency domain is likely to be more expensive than a single delay time modeling in time domain.

The analyses of space and time complexities have revealed that the overriding idea of having an one-size-fits-all global discretization in either space or time may be problematic. The modeling can become cumbersome due to the large number of cells and/or time steps required by the different scales in the global discretizations. Such scale contrasts and associated over-computing can be greatly reduced if data are modeled separately with customized meshes and time discretizations via survey decomposition.

3.1.3 Optimization complexity

Another complexity is the amount of computation required to solve the inverse problem. This exists in all types of inversion and is related to the time needed to complete a forward modeling because: (1) an inversion needs the forward responses and the sensitivities to make a model update, and computing sensitivity is equivalent to forward modeling; (2) the forward responses and sensitivities must be re-computed every iteration because of the nonlinearity of the problem.

The sensitivity, or Jacobian, is an essential part in the minimization of the objective functional used in an inversion. Previous studies have concentrated on the development of optimization algorithms using full or approximate Hessians or the gradient with the information from the sensitivity. For example, Newton, Gauss-Newton, quasi-Newton and nonlinear conjugate gradient methods have been applied to geophysical inversions. Generally, an optimization method that spends more time on computing the quadratic information may converge in fewer iterations, while one that concentrates on gradients finishes an iteration quickly but requires a large number of iterations to achieve the stationary point. So, no algorithm is absolutely superior to others.

This thesis uses the Gauss-Newton method as the primary optimization method to demonstrate the framework of survey decomposition. In particular, the Gauss-Newton method solves a normal equation for a model update $\delta\mathbf{m}^{k+1}$ at the $k+1$th iteration. The normal equation could look like

$$\left[\mathbf{J}(\mathbf{m}^k)^\top\mathbf{J}(\mathbf{m}^k)+\beta\mathbf{W}^\top\mathbf{W}\right]\delta\mathbf{m}^{k+1} = -\mathbf{J}(\mathbf{m}^k)^\top\left[F(\mathbf{m}^k)-d\right]-\beta\mathbf{W}^\top\mathbf{W}(\mathbf{m}^k-\mathbf{m}^{ref}), \tag{3.1}$$

where $F(\mathbf{m})$ represents forward response of model $\mathbf{m}$, $\mathbf{J}(\mathbf{m})$ is the Jacobian matrix, $\mathbf{W}$ is the regularization matrix, and $\mathbf{m}^{ref}$ is the reference model. The matrix $\mathbf{J}$ is large, dense and then equation 3.1 is generally solved by iterative methods, like conjugate gradient (CG), that only need to compute $\mathbf{J}$ and $\mathbf{J}^\top$ times a vector $\mathbf{v}$. This sensitivity-vector multiplication ($\mathbf{J}\mathbf{v}$ or $\mathbf{J}^\top\mathbf{v}$) is a common operation found in other optimization algorithms. The optimization complexity also depends upon the two options of making the sensitivity available.

The first option is to implicitly represent $\mathbf{J}$ or $\mathbf{J}^\top$ in terms of the multiplication of some matrices (equation B.2 and B.3). When doing $\mathbf{Jv}$ or $\mathbf{J}^\top\mathbf{v}$, the vector $\mathbf{v}$ is sequentially multiplied by those matrices, and it takes the same amount of computation as one complete forward modeling. This approach implicitly computes the sensitivity on the fly upon request, so the amount of computing is more dependent on the number of CG iterations, instead of the number of data. Implicit sensitivity also needs substantial amount of memory to store the matrices that build $\mathbf{J}$ or $\mathbf{J}^\top$.

The second option is to explicitly compute $\mathbf{J}$ and store it in the memory for the multiplication. If $\mathbf{J}$ is computed row by row (more efficient when the number of data is smaller than the number of model parameters), the forward modeling must take place for N (the number of data) times to form the full $\mathbf{J}$. Once $\mathbf{J}$ is computed, $\mathbf{Jv}$ and $\mathbf{J}^\top\mathbf{v}$ can be quickly carried out for many CG iterations. When computing $\mathbf{J}$, the explicit method requires the same amount of memory as the implicit method does, but afterwards, all of the matrices in equation B.2 and B.3 can be deleted, leaving only $\mathbf{J}$, relatively small, in the memory.

The reduction of computational complexity in sensitivity should use different strategies for the implicit and explicit sensitivities. The implicit method is efficient if there is a small number of CG iterations, but this may potentially harm the search for optimality. The efficiency of the explicit method is less sensitive to the number of CG iterations, but the number of data should be as small as possible. Here is why survey decomposition can make a difference: (1) if the data are modeled separately and concurrently, the number of data will not hinder the efficiency of the explicit method; (2) if the data are modeled on the localized mesh, the explicit method can store $\mathbf{J}$ locally with less memory. For these reasons, the explicit sensitivity method is preferred under the framework of survey decomposition.

Now that the optimization complexity is associated with the number of data, it is desirable to reduce the number of data as much as possible. There are two types of over-sampling that possibly lead to more data than are necessarily needed in the inversion. Firstly, the number of data in a survey is usually determined by the field acquisition, which tends to be redundant in both space and time. Secondly, at the early iterations of an inversion, only a small portion of the information in the data is needed, and the constructed model is obtained by minimizing an objective functional that is dominated by the regularization term, so computing every single datum to high accuracy is unnecessary.

The survey decomposition itself, although not able to intrinsically reduce the number of data and the number of iterations in an inversion, does offer more flexibility to allow the development of better strategies. Once the subproblems are separated, we can easily control the selection of data for computation. If the relationship between the selection of data and the desired amount of information can be established, more savings can be achieved by only modeling the very necessary data (subproblems). The data subsampling problem, tackled by an adaptive and random sampling scheme, is discussed in Section 3.4.

3.2 Identification of Subproblems

The insight extracted from Section 3.1 is that modeling data in highly contrasted scales (space and/or time) together with global discretizations is very inefficient. Savings can be obtained by separating the data into subproblems and allowing each subproblem to have its own localized discretizations that match the scale of the subproblem modeling.

Bearing this rationale in mind, this section identifies the subproblems in the controlled source EM (CSEM) surveys with the concept of an "atomic" problem as a building block.

3.2.1 Atomic problem

While a CSEM survey can be decomposed in many different ways, there are two keys to decompose a survey in an efficient manner:

- Fine-grained decomposition. Every subproblem must be small enough to allow customized high-efficiency discretization and many small subproblems for massive parallelization and platform flexibility.
- Self-contained problem. Every subproblem itself is a complete EM modeling entity with actual sources, receivers working at some time or frequency. The self-containedness that enables the subproblems to be solved independently is important in maintaining good scalability for a large-scale modeling problem.

Because every subproblem only models a subset of the entire data set, it is easier to consider the survey decomposition as a data grouping problem. We can imagine the data are organized by a spreadsheet, in which each row corresponds to each measurement datum and the columns are the attributes of data, such as source index (S), receiver index (R), delay time or frequency (T/F), field or component (F/C) and so on. Then the data are grouped based on some conditions that can be applied to one or more attributes; the attributes that are not used in the conditions are ignored in the grouping process. For example, if a grouping process requires the data measured before 0.01 s to be put into a subproblem and the data after 0.01 s to be put into another subproblem, the grouping condition only uses the attribute of delay time, and other attributes are not considered; so within in each of the two subproblems, the data from different transmitters and receivers can be found.

The less attributes considered in the grouping conditions, the more inclusive a subproblem can be. The conventional methods, modeling the entire data set in one problem, have the largest subproblem possible using no grouping condition at all. If plotted on the spectrum of subproblem grouping, the conventional global discretization modeling falls at the most inclusive end (Figure 3.2). An inclusive subproblem is more likely to see high contrast of modeling scales, as the data are more indiscriminately grouped together with fewer conditions. In the move towards the more exclusive subproblem, we can find other previously attempted strategies that effectively achieved better efficiency by data grouping. For example, the straightforward domain decomposition or tiling approach can be seen as a grouping using the condition that requires the sources or receivers (S/R) to be located in some certain subdomains. The multi-level grid approach, modeling different times/frequencies on different meshes, is logically related to the grouping using the condition on the attribute of time/frequency (T/F).

Unsurprisingly, the lowest scale contrast and the highest efficiency are found all the way at the exclusive end of the spectrum, where the grouping conditions involve all of the attributes. In this case, every subproblem only models a single datum with its own distinct source, receiver, time/frequency and field/component (S-R-T/F-F/C). It appears that now we are at the physical limit of decomposing a survey and no smaller subproblems can be found. However, the subproblem can be made even more exclusive by demanding a subproblem to model the data from a

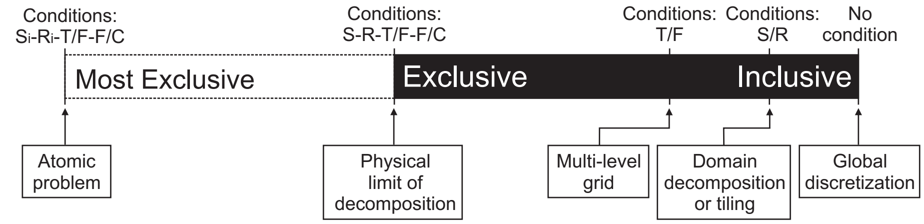

Figure 3.2: The spectrum of subproblem data grouping.

particular part of the source or receiver if the loop or wire can be decomposed. By making that "particular part" as small as possible, we arrive at an ultimate limit of survey decomposition: an *atomic problem* that models a point source and a point receiver at an particular time/frequency for a datum of a particular field/component (S_i-R_i-T/F-F/C). Now the conditions of fine-grainedness and self-containedness requested at the beginning of this subsection are both satisfied.

The point source in an atomic problem can be either an electric dipole representing a short straight wire, colloquially called "linelet", or a magnetic dipole representing a small loop, colloquially called "looplet". Both electric and magnetic dipoles can carry a certain amount of current as a function of time. The point receiver measures the real or complex datum of the electric or magnetic field, or other quantities of interest at a certain orientation. The atomic problems with different S_i-R_i-T/F-F/C parameters are able to build any type of CSEM survey through simple combination or superposition.

3.2.2 Practical recipes

In theory, an EM modeling that is directly using the atomic building block as the subproblems would have the best efficiency from the single datum simulation's point of view. However, solving the entire problem with such fine-grained decomposition may be impractical because of the large number of subproblems and associated complications of management. So practically, we need to find a trade-off point on the data grouping spectrum (Figure 3.2) that is as close to the atomic problem as possible while yielding a manageable number of subproblems to be solved. I therefore provide some recipes as a guideline for identifying the subproblems in any CSEM survey.

The first recipe is to combine atomic problems, where possible, into one subproblem that is computationally the same as one atomic problem. For example, if the EM modeler compute the three-component of magnetic field at the receiver at the same delay time, then the three atomic problems working for $\mathrm{H_x}$, $\mathrm{H_y}$, $\mathrm{H_z}$ can be grouped in one subproblem.

The second recipe is to eliminate the atomic problems that show negligible differences. In many surveys, the sources and receivers are small enough to be approximated by point sources and receivers, like the transmitter used in the airbrone surveys, so subdivision of such source/receiver is not necessary.

The third recipe is to consider encapsulating the atomic problems with small or moderate contrast of modeling scales. For example, if a survey has frequencies/times that spans a relatively small range, the atomic problems with different frequency/time can be merged. Unlike the first two principles, this simplification of decomposition always comes with extra computational costs. The users need to be aware of this trade-off as it is shifting towards the more inclusive subproblem data grouping in Figure 3.2.

The first recipe is often automatically applied to combine the subproblems with the same source, receiver and time/frequency but different fields/components in most modern EM simulators. The applications of the second and third recipes are largely dependent on the geometries of the survey – the “distances” between the subproblems in space or time. In realistic field surveys, this boils down to the geometry of the transmitter and the span of time/frequency range. In the next two subsections, I examine the decompositions of some CSEM surveys categorized by the geometry of source.

3.2.3 Decomposition of surveys with localized sources

Many CSEM surveys use the simplest electric or magnetic dipole source, which is effectively a point source requiring no subdivision according to the second recipe in Subsection 3.2.2.

The first type of geometry I consider is a dipole source and a receiver that are close to each other and move together; the source-receiver pair makes measure-

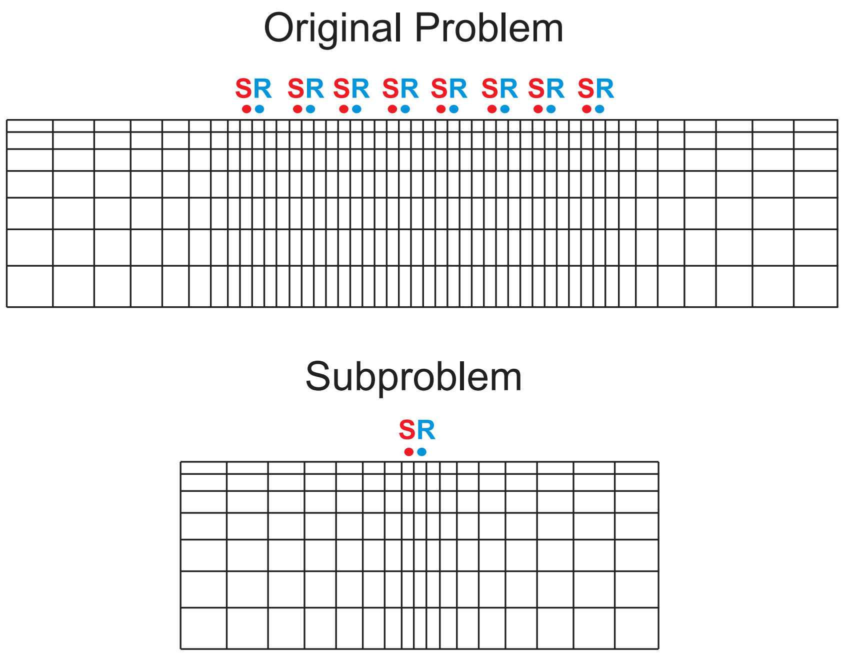

Figure 3.3: Decomposition of dipole source survey with small source-receiver offset.

ments at every new location, so the data set consists of many individual soundings at different locations over the survey area (see a cross section diagram in Figure 3.3). This geometry is commonly used to quickly sweep a large area. Examples include the airborne EM using magnetic dipoles and the DC resistivity profiling using electric dipoles (dipole-dipole array). The global mesh designed for the original problem must be large and fine enough to encompass the entire survey area; upon decomposition, each subproblem only has one dipole source and receiver and the local mesh designed for that subproblem can be much smaller (Figure 3.3). A DC problem has no time complexity, so there is no decomposition in time. An airborne survey usually operates within a relatively narrow band of times or frequencies; as a result, the decomposition of time/frequency is optional. If all time channels/frequencies are modeled together, the subproblems can be obtained by a common source-receiver (S-R) grouping.

The second type of geometry is similar to the first type in mobility, but the

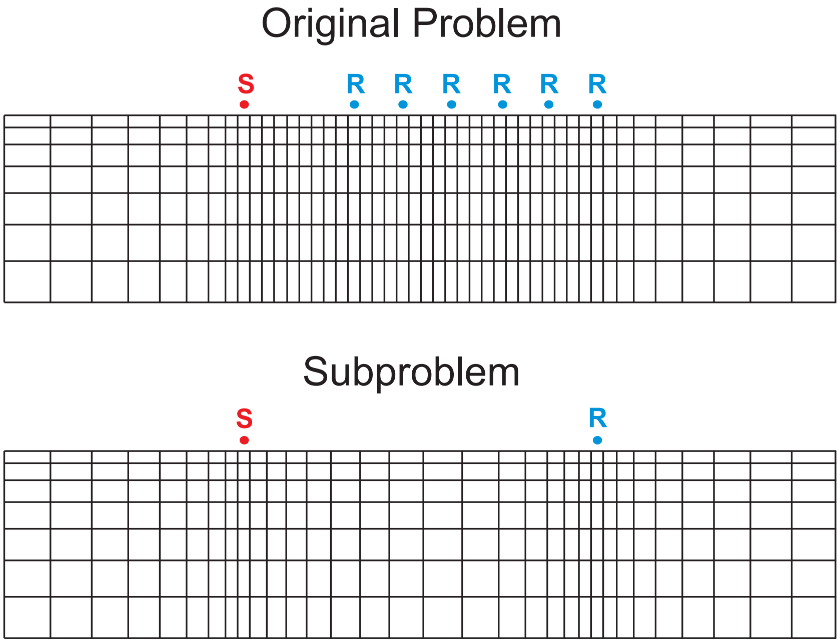

Figure 3.4: Decomposition of dipole source survey with large source-receiver offset.

separation between the source and receiver is large. This geometry can be found in the DC resistivity and marine CSEM using towed streamers. In this geometry, it is common to have multiple receivers for one dipole source to better illuminating the underground (see a snapshot of this geometry in Figure 3.4). Modeling such problems using the conventional approach requires the global mesh to be finely discretized everywhere in a domain large enough to contain the spread of the system and also the entire survey area. If decomposed, every subproblem only has one dipole source and one dipole receiver at a certain offset; then the local mesh for a subproblem can be designed in the manner shown in Figure 3.4; this is still much smaller than the mesh that a global discretization would require. Marine CSEM can measure data at times/frequencies with a few orders of magnitude contrast. Therefore the decomposition in time/frequency is recommended, leading to a common source-receiver-time/frequency (S-R-T/F) grouping.

3.2.4 Decomposition of surveys with distributed sources

Some CSEM surveys use a source as large as the entire survey area. This kind of distributed source can be considered as superposition of many dipole sources. Then the modeling results are obtained by summing up the subproblems' results.

The first example is the ground loop survey using a large and closed loop source fixed on the surface and measuring the fields in and outside of the loop (Figure 3.5). The global mesh has to contain the entire loop as well as all receiver locations with fine cells everywhere. If the large loop is broken down into many small looplets, each of which is essentially a magnetic dipole, then the looplet subproblem can have much smaller local mesh (see one of the looplet subproblems in Figure 3.5). Given the capability of measuring wide-band data in the ground survey, I choose to do time/frequency decomposition. A subproblem in a ground loop survey is therefore a result of common looplet-receiver-time/frequency (S_i-R-T/F) grouping.

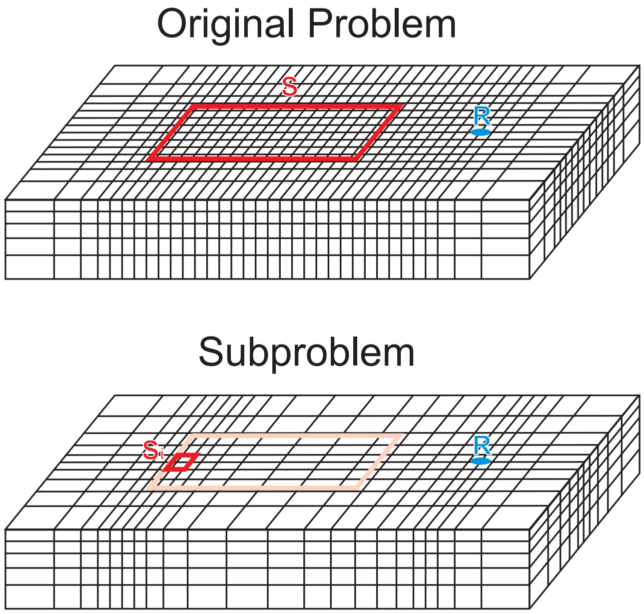

Figure 3.5: Decomposition of survey with large loop source.

The distributed source can also be a open wire galvanically grounded by the electrodes, as the one used in LOTEM (Strack et al., 1990), GREATEM (Mogi et al., 1998), MMR (Edwards, 1974) and MIP (Seigel, 1974). The decomposition and mesh design are similar to those for the closed loop, except many electric dipole linelets are used, instead of looplets (Figure 3.6). This type of survey uses common linelet-receiver-time/frequency (S_i-R-T/F) grouping.

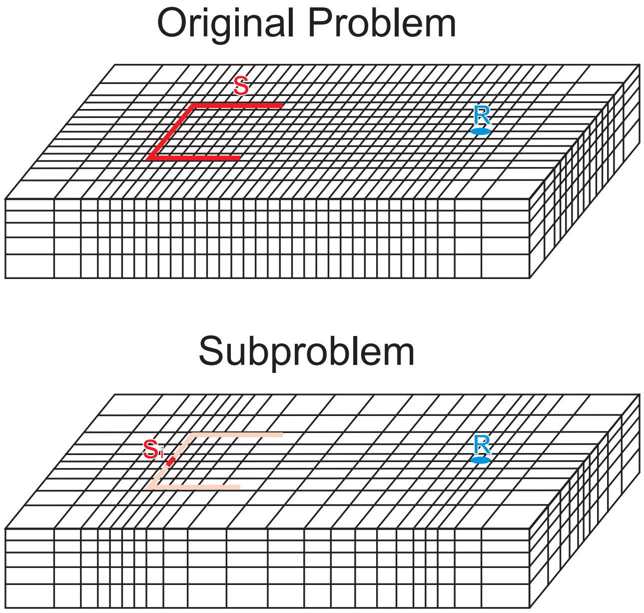

Figure 3.6: Decomposition of survey with long wire source.

3.3 Localized Discretization

The fundamental motivation of decomposing a survey is to allow more efficient discretization customized to the individual subproblems. In contrast to the standard discretizing design that seeks global capability for all sources, receivers and delay times at one attempt, my survey decomposition framework promotes "localized discretizations" to avoid the unnecessary contrast of modeling scale in space and time. The global-local design of the framework involves interactions between the two levels, which also takes into account how the subproblems can be solved using massive parallelization.

3.3.1 Local mesh

A mesh needed for a subproblem is called a "local mesh", and is designed to only accommodate the source, receiver and time/frequency (S-R-T/F) it concerns. It can be highly efficient compared to a global mesh because: (1) the cell sizes can be adapted to the scale of the time/frequency; (2) the domain size of a local mesh may be much smaller than the original problem; (3) the EM diffusion allows geometric coarsening of the mesh cells away from the source and receiver locations.

A generic example of a local mesh in a structured rectilinear grid (regular mesh) is demonstrated in Figure 3.7. The regions near the source (S) and receiver (R) are locally refined because EM data have higher resolution in the domain near the source and receiver. The separation of the local refinements can be adjusted according to the source-receiver offset in specific subproblems. Other type of mesh structures, unstructured or semi-structured, can have a similar appearance. If there is uneven topography, the topography close to the source/receiver is finely discretized, whereas the mountains and valleys far away are approximated by bulky cells.

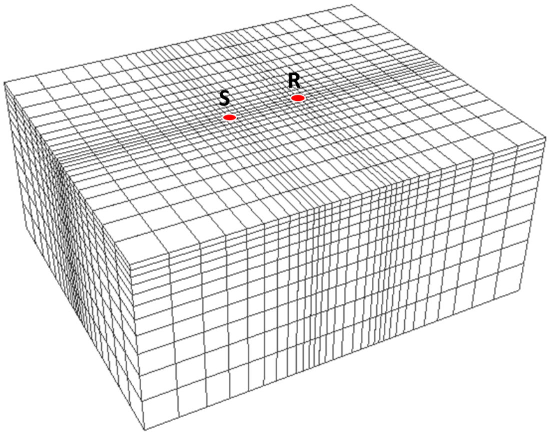

Figure 3.7: A local mesh designed for a subproblem modeling the source (S) and receiver (R) marked by the red dots.

A rectilinear local mesh in Figure 3.7 can be characterized by some design parameters:

- **Smallest cell size**. This parameter specifies the cell size at the center of local refinements and it is usually determined by the ambient conductivity and the earliest time (or highest frequency) it models. In practice, I use an iterative approach to adaptively find the appropriate smallest cell size: the program starts with a cheap initial guess, and then gradually reduces the smallest cell size until the difference made by the discretization is negligible.

- **Domain size or boundary condition distance**. This parameter specifies the distance from the refinement center to the outmost cells in the modeling domain. In order to iteratively determine the size of the modeling domain, I begin with a small distance, carry out the forward modeling, and then keep expanding the boundary until the difference between two successive forward modelings falls beneath a prescribed tolerance. This adaptive search procedure is carried out for every subproblem and for every updated model during an inversion, so subproblems at different locations/times/frequencies and at different iterations of inversion have their own customized domain size.

- **Cell expansion rate**. This parameter specifies how fast the cells are coarsened when moving away from the source/receiver location. It is defined as the ratio of size of the larger cell to the smaller cell next to it along a certain direction. Choice of this parameter depends on the variation of conductivity. Empirically I have found a value between 1.2 and 1.5 can produce acceptable results for a good balance between the accuracy and efficiency for most of the earth models.

The local meshes designed for the subproblems with different S-R-T/F can look very different. Here I provide two illustrative examples of local meshes and compare with the global mesh. The first case is for the subproblems having zero-offset source-receiver pair (Figure 3.8), which is common in airborne surveys. The source and receiver, marked by the red dots, are located at the center of each local mesh. Because of the scale variations controlled by t/σ (delay time over conductivity) or ρ/f (resistivity over frequency), it is possible to have local meshes in different

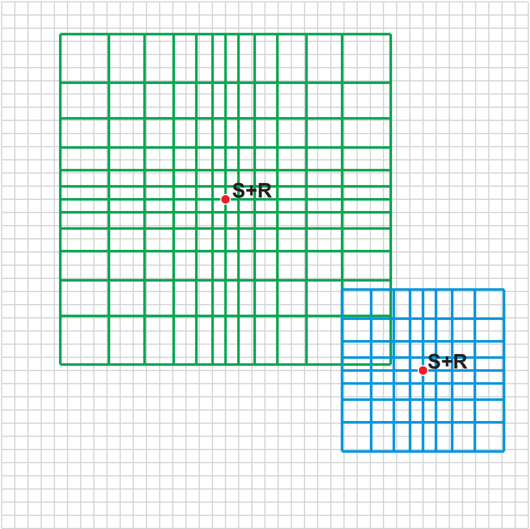

Figure 3.8: Examples of local mesh for zero-offset source-receiver pair overlaying the global mesh in plan view.

sizes within one survey, as exemplified by the green and blue local meshes in Figure 3.8. The model parameters outside of the green mesh are deemed irrelevant to the data modeled by the green-colored subproblem; in other words, the domain size is determined by the domain of sensitivity. For a survey as extensive as the global mesh (grey grids), there may be many local meshes whose local refinements collectively cover the entire global mesh. The difference of space complexity between the local meshes and the global mesh, in terms of the number of cells, is evident and can be even more exaggerated in 3D.

The second case is when the source and receiver in a subproblem are relatively far apart, which happens in a variety of EM surveys, like ground loop EM and marine CSEM. This requires two local refinements in a local mesh. In the two examples shown in Figure 3.9, the blue and green local meshes are oriented in accordance with the source (S) and receiver (R) locations of each subproblem. The source and receiver with longer offset usually demands a larger local mesh, but the reduction of space complexity is still significant compared to the global mesh.

Another important aspect of using local meshes is to represent the global model

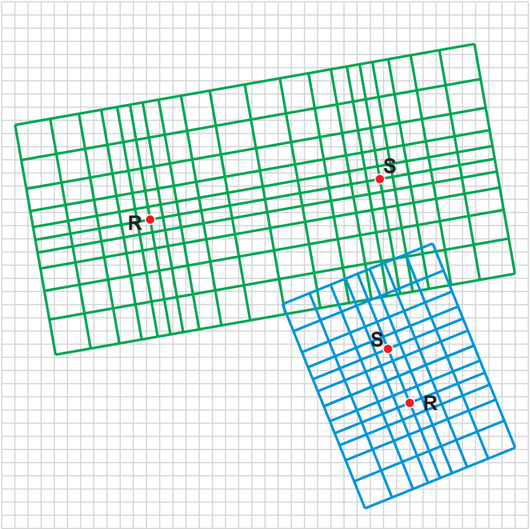

Figure 3.9: Examples of local mesh for long-offset source-receiver pair overlaying the global mesh in plan view.

on a local mesh. A technique called "up-scaling" or "up-gridding" or "optimal gridding" may be involved when converting the parameters of many small cells in the global mesh to one single parameter of a large cell in the local mesh (Li et al., 2001; Fincham et al., 2004). Since it is reasonable to assume a 3D isotropic conductivity model in mining geophysics, I adopt the averaging approach in Commer & Newman (2008) and make modifications and extensions so that the mesh conversion is a linear operation based on the interception of global-local cells and the local coordinate system can rotate with respect to the global system (Figure 3.9). The specific details of this operation can be found in Subsection 3.3.3.

3.3.2 Local time discretization

The time complexity can be reduced by using local time discretization. Unlike a global time discretization that models all of the delay times, a local time discretization only works for one particular time in a subproblem (S-R-T/F).

An illustrative example is presented in Figure 3.10, in which a global problem modeling two delay times, t_1 and t_2, is decomposed into two subproblems with

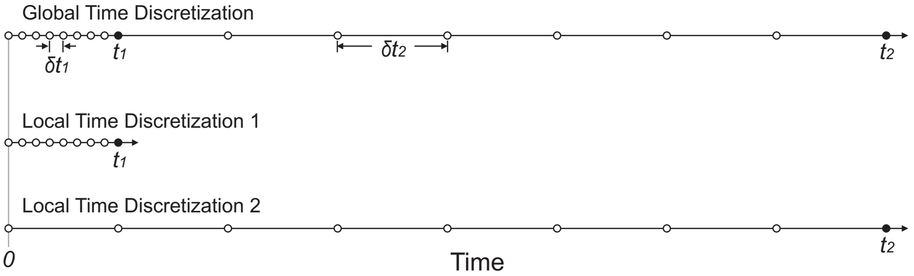

Figure 3.10: Global and local time discretizations.

their own local time discretizations 1 and 2. Suppose the modeling at time t requires a step length of one eighth of t, using the step-adaptive idea in Oldenburg et al. (2013), the global time discretization has two different step lengths δt_1 and δt_2 to offer adequate modeling accuracy to t_1 and t_2. It takes the global problem 15 steps to finish the modeling. Upon decomposition, t_1 and t_2 are modeled independently using the local time discretizations 1 and 2, so each only needs 8 steps with the step length of δt_1 and δt_2 respectively. If the subproblems are computed in parallel, the required time is for 8 steps, which is half of what would be needed in the global problem. The saving of time can be further exaggerated if more delay times over a large range are modeled. It is important to note that this reduction of time complexity after decomposition does not compromise the modeling accuracy because, for example, t_2 is still modeled by the step lengths δt_2 and using the excessively small step lengths δt_1 would have contributed to the over-computing for t_2.

The saving by using local time dicretization can be even more significant when explicitly computing the sensitivity matrix. If the sensitivities for t_1 and t_2 are computed using the method in Appendix B, the total number of time steps would be 30 (15 steps × 2 data) with the global discretization in Figure 3.10. Separate modeling with the two local discretization results in 8 steps for each subproblem and 16 steps in total. The example in Figure 3.10 shows two delay times in moderate contrast; for a realistic survey with time scale contrasts up to three or four orders of magnitude, the overall efficiency can be improved by an order of magnitude or more.

There are additional benefits of using local time discretizations in comparison with the global one:

- The size of the mesh can be adapted to the scale of the time channel, leading to more efficient local meshes.

- The change of step lengths requires the Maxwell matrix to be re-factorized and stored if using a direct solver. A constant time step used in a local time discretization avoids this, making the algorithm more memory-flexible.

- The design of a global time discretization has to trade-off between the number of different step lengths (number of factorizations) and the number of total time steps. This can result in a step length that reduces the accuracy of the simulation immediate after the change of step length. A local discretization is likely to have better modeling accuracy by choosing a step length that is tuned to a specific delay time. When using the implicit time stepping method, there are usually 10$\sim$20 steps from $t = 0$ to the desired delay time.

3.3.3 Global-local interactions

Under the framework of survey decomposition, the information about the data and model must be managed at, and exchanged between, the global and local levels. The basic flows of information include: (1) the global process broadcasts the model/data and survey information to many parallel local processes; (2) after the actual computations are carried out at the local level, the local processes return the results, usually a few scalars or a vector, to the global process; (3) the global process summarizes the results and makes decisions about the model update. These steps constitute one iteration of using survey decomposition for numerical modeling. A forward-only modeling needs one such iteration, while an inversion requires many iterations.

The conductivity models, as well as the model updates produced during an inversion, are all presented on a pre-defined global mesh. The survey decomposition approach transfers the global model to the local meshes on which the subproblems are solved. A key step in this procedure is to be able to convert models from one mesh to the other, so that the forward responses and sensitivities for the global mesh are correctly evaluated by solving local problems. No such operation can be exact and perfect. While I acknowledge mesh conversion is a complicated topic that deserves sophisticated treatments, I use a simple formula that makes use of the linear operation of intersecting volumes. The essential procedure is to find the mutual intersections between cells in the global mesh and local mesh. This can be represented by a sparse matrix **R**, in which r_{ij} is the intersecting volume of the ith cell in the local mesh and the jth cell in the global mesh. Every subproblem has its own **R** matrix reflecting the volumetric relation between the local mesh and the global mesh. The material averaging operation for the ith local mesh is expressed as

$$\mathbf{m}_{li} = \mathbf{V}_{li}^{-1}\mathbf{R}_i\mathbf{m}_g, \tag{3.2}$$

where $\mathbf{m}_l$ and $\mathbf{m}_g$ are the local and global models respectively. $\mathbf{V}_l$ is a diagonal matrix of local cell volumes and its inverse normalizes each row of **R** by the volume of corresponding cell in the local mesh. Matrix **R** can be computed using either the rigorous computational geometry or fast approximate methods, for example, the point-cloud algorithm documented in Appendix C.

After computations on a local mesh, the modeled data, for example, the magnetic fields in x, y, z directions, are transformed back to the global coordinate system through

$$\begin{pmatrix} d_{xgi} \\ d_{ygi} \\ d_{zgi} \end{pmatrix} = \mathbf{M}_i^{-1} \begin{pmatrix} d_{xli} \\ d_{yli} \\ d_{zli} \end{pmatrix}, \tag{3.3}$$

where the subscript g and l indicate "global" and "local" coordinate system respectively, the subscript i means the ith subproblem, and **M** is the 3D rotation matrix mapping a vector from the global to the local coordinate system of the ith subproblem.

Equation 3.3 brings the simulated data from local meshes with different rotations and translations to a universal system consistent with the global mesh. There may be one more operations of linear combination, denoted by another sparse matrix **S**, before the transformed data can be useful. Matrix **S** bears the information about how the modeling results from the subproblems are summed to render the desired data. In an airborne survey, **S** is the identity; the structure of **S** for a ground loop survey is discussed in Chapter 5. The final data vector is expressed as

$$\mathbf{d}_g = \mathbf{S} \begin{pmatrix} \vdots \\ \mathbf{M}_i^{-1} \mathbf{d}_{li} \\ \vdots \end{pmatrix}, \tag{3.4}$$

where

$$\mathbf{d}_{li} = \begin{pmatrix} d_{xli} \\ d_{yli} \\ d_{zli} \end{pmatrix}. \tag{3.5}$$

The simulated data from the ith subproblem in equation 3.5 can be any kind of EM field. If more than one delay time is modeled in the ith subproblem, $\mathbf{d}_{li}$ can have multiple columns corresponding to multiple delay times.

The computation of sensitivity, which has the same amount of computing as forward modeling, is also carried out on the local meshes in parallel subproblems, where the coarsening of cells matches the decay of sensitivity. Entries of the sensitivity matrix have dimension of “data per unit model parameter”. Column-wise, they can be considered as a special kind of data; every column of the sensitivity matrix is a vector pertaining to one model parameter. Row-wise, they can be considered as a special kind of model; every row is a vector pertaining to one datum. Therefore, a local sensitivity matrix is transformed to its global counterpart for use in inversion in two steps. Firstly, the sensitivity, as a model on the local mesh is mapped to the global mesh. This local-to-global operation uses the same cell-intersecting information needed for the material averaging, but here the process is additive. The sensitivity matrix of the data in the ith subproblem on the global mesh in the local coordinate system is

$$\mathbf{J}_{li} = \begin{pmatrix} \cdots \mathbf{j}_{xli} \cdots \\ \cdots \mathbf{j}_{yli} \cdots \\ \cdots \mathbf{j}_{zli} \cdots \end{pmatrix} \mathbf{V}_{li}^{-1} \mathbf{R}_i, \tag{3.6}$$

where matrices $\mathbf{V}_{li}^{-1}$ and $\mathbf{R}_i$ together make the sensitivity available on the global mesh. In the case where the local mesh is smaller than the global mesh, the global cells not intersecting with any local cells are set to zero sensitivity. Secondly, the sensitivities, just as the data were, must be rotated back to the global system. Similar to equation 3.4, the sensitivity matrix in the global coordinate system and on the global mesh is

$$\mathbf{J}_g = \mathbf{S} \begin{pmatrix} \vdots \\ \mathbf{M}_i^{-1}\mathbf{J}_{li} \\ \vdots \end{pmatrix} . \tag{3.7}$$

The full sensitivity $\mathbf{J}_g$ is rarely formed in practice. Instead, $\mathbf{J}_g$ or $\mathbf{J}_g^\top$ is more commonly multiplied by a vector $\mathbf{v}$ in a variety of optimization algorithms. In the survey decomposition framework, when using an explicit method and direct solver, the sensitivity matrices along with all the local auxiliary matrices are stored locally. When a sensitivity-vector multiplication is requested, the operation can be carried out concurrently within the subproblems.

To implement $\mathbf{J}_g$ times $\mathbf{v}$, the vector $\mathbf{v}$ is broadcast to all of the subproblems and applied to equation 3.6 in parallel; then every local process returns a short vector, which is equivalently $\mathbf{J}_{li}\mathbf{v}$, to the global process; the global process assembles those short vectors into a long vector; finally this long vector is applied to $\mathbf{S}$ for the final result. This procedure can be expressed as

$$\mathbf{J}_g\mathbf{v} = \mathbf{S}\mathbf{J}_l\mathbf{v} = \mathbf{S} \begin{pmatrix} \vdots \\ \mathbf{M}_i^{-1}\mathbf{J}_{li}\mathbf{v} \\ \vdots \end{pmatrix} . \tag{3.8}$$

$\mathbf{J}_g^\top$ times $\mathbf{v}$ requires the transpose of a dense matrix $\mathbf{J}_{li}$ so it can be slow. This is worked around by computing $\mathbf{v}^\top\mathbf{J}_g$ then taking the transpose. $\mathbf{v}^\top$ first left-multiplies $\mathbf{S}$, yielding another row vector with the length equal to the number of data; then this vector is divided into many short row vectors based on data ordering and segmentation, and the short row vectors are sent to the corresponding subproblems; within each subproblem, the short row vector left-multiplies $\mathbf{M}_i^{-1}$ and $\mathbf{J}_{li}$, returning another row vector to the global process; the global process sums up those row vectors for the final result of $\mathbf{v}^\top \mathbf{J}_g$. This procedure can be expressed as

$$\mathbf{v}^\top \mathbf{J}_g = \sum_i \mathbf{v}^\top \mathbf{S} \begin{pmatrix} \vdots & \vdots & \vdots \\ \mathbf{e}_{xi} & \mathbf{e}_{yi} & \mathbf{e}_{zi} \\ \vdots & \vdots & \vdots \end{pmatrix} \mathbf{M}_i^{-1} \mathbf{J}_{li}, \tag{3.9}$$

where $\mathbf{e}_{xi}$, $\mathbf{e}_{yi}$ and $\mathbf{e}_{zi}$ have all zero entries except for the unit value in the row that corresponds to the datum from the ith subproblem.

Sometimes it is particularly useful to compute the diagonal of the square matrix $\mathbf{J}_g^\top \mathbf{J}_g$. This piece of information can be used to assess the well-posedness of the inverse problem, determine the initial trade-off parameter β in equation 1.6, and precondition the CG solver when solving equation 1.12. In this calculation, every row of $\mathbf{S}$, denoted by $\mathbf{s}_j^\top$ for the jth row, left-multiplies $\mathbf{J}_g$ at the local level in the same way of equation 3.9 and then all of the elements in the resulting row vector are squared; finally those row vectors across different rows of $\mathbf{S}$ are summed up and transposed to obtain the diagonal of $\mathbf{J}_g^\top \mathbf{J}_g$. This procedure can be expressed as

$$\mathbf{diag}(\mathbf{J}_g^\top \mathbf{J}_g)^\top = \sum_j \big(\sum_i \mathbf{s}_j^\top \mathbf{M}_i^{-1} \mathbf{J}_{li}\big)^{\circ 2}, \tag{3.10}$$

where $\circ 2$ denotes the operation of element-wise (Hadamard) square. Because $\mathbf{s}_j^\top$ is sparse, a short-cut can be made by sending only non-zero segments of $\mathbf{s}_j^\top$ to the active subproblems if the data are properly ordered.

Equations discussed in this subsection are generic, and can be modified according to the specific needs in different surveys for improved performance.

3.3.4 Massive parallization

Solving the modeling problem after survey-parameter (S-R-T/F) based decomposition is appealing because the resultant subproblems can be independently modeled, which makes the modeling much more scalable in the environment of massive parallelization compared to the domain-based decomposition with heavy inter-subproblem dependency. In this subsection I examine a wall-clock time model of using survey decomposition and discuss its scalability in massive parallelization. For simplicity, the computing environment is assumed to be a uniform array of processors, each of which has its allocated random-access memory.

In massive parallelization, the computing time consists of two parts, the solving time spent on the actual modeling and the communication time spent on passing data through the network. The solving time, assuming every subproblem is computationally identical, can be calculated as

$$T_s = t_s \cdot \frac{N_{sp}}{N_p}, \tag{3.11}$$

where t_s is the time required by solving one subproblem on one processor, N_{sp} is the number of subproblems, and N_p is the number of processors.

The communication time is taxed because a vector, either model or data, must be first sent from the global process (host) to the local processes (workers) and then retrieved from the local processes to the global process. This global-local-global communication is shown by the solid and dashed arrows in Figure 3.11 and its time can be approximated by

$$T_c = 2 \cdot t_c \cdot \log_2 N_p, \tag{3.12}$$

where t_c is the time of passing a vector between two processors. The model of parallelization in Figure 3.11 assumes there are enough workers (processors) for every subproblem; otherwise every worker needs to be assigned more than one subproblems. The high efficiency of $O(\log n)$ in equation 3.12 is due to the fact that the parallel workers barely talk to each other, except at the beginning and end of subproblem computing when vectors are broadcast through the networks.

Considering the sum of T_s and T_c as a function of N_p, the total time for a given N_{sp} has a minimum at

$$N_p^* = \frac{\ln 2}{2} \cdot \frac{t_s \cdot N_{sp}}{t_c} \approx 0.35 \cdot R_{s/c} \cdot N_{sp}, \tag{3.13}$$

in which $R_{s/c}$ is the ratio of t_s to t_c. If the optimal number of processors is greater than the number of subproblems ($N_p^* > N_{sp}$), the computation can always gain further speed-up by adding extra processors (up to N_{sp} processors). This condition,

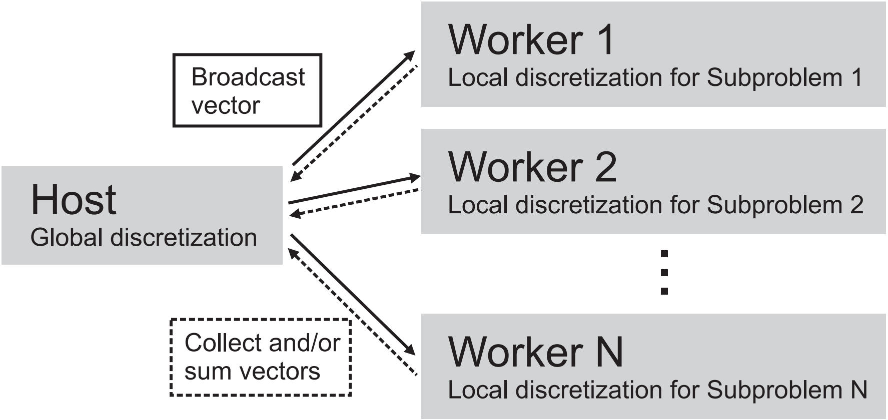

Figure 3.11: Model of parallel computing for survey decomposition.

referred to as continual speed-up, occurs if

$$R_{s/c} > \frac{2}{\ln 2} \approx 2.88. \tag{3.14}$$

Assuming $N_{sp} = 10000$, $t_s = 100$, Figure 3.12a shows the decay of the solving time T_s and logarithmic growth of the communication time T_c for three different example scenarios $t_c = 1, 5, 10$. The minimizer of the total time function, calculated by equation (3.13) for three t_c values, are 350000, 70000 and 35000 (processors), all greater than N_{sp}, so continual speed-up is possible by using up to N_{sp} processors and the minimum time required for $N_p = N_{sp}$ is 127, 233 and 366 for $t_c = 1, 5, 10$ respectively. If t_c is relatively large, the communication time begins to dominate the total time more quickly at a smaller number of processors. Figure 3.12b shows that the linearity of speed-up depends on $R_{s/c}$.

In reality, t_c is a hardware-dependent parameter and may vary significantly but it is generally true that, after survey decomposition, most computationally intensive jobs, including solving PDEs and dense matrix-vector multiplication, are carried out within the subproblems at a local level so $R_{s/c}$ is large enough that the end users are more bound by the constraint of computational resources than by the algorithm's ability of taking advantage of more processors.

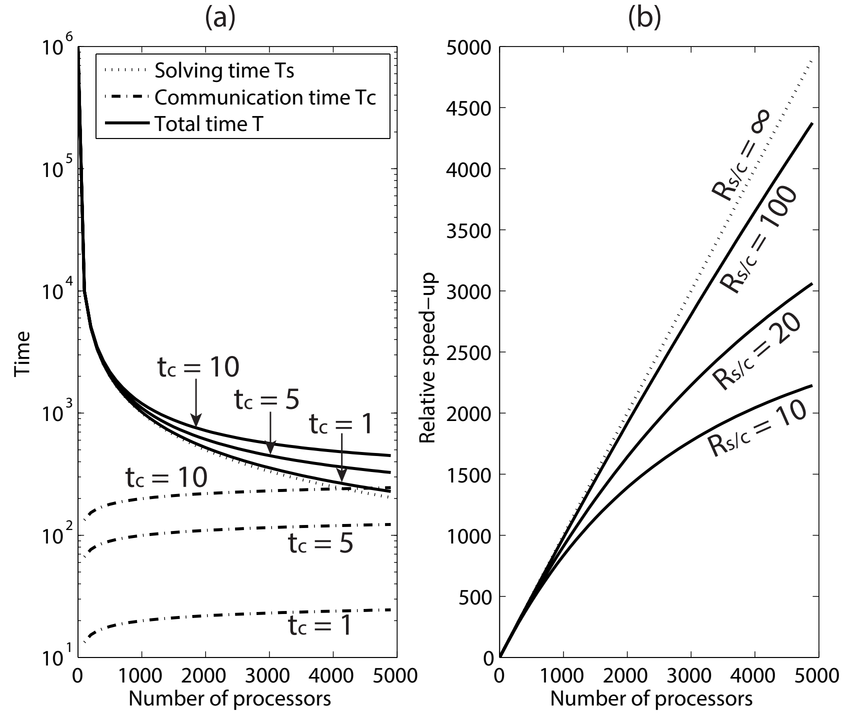

Figure 3.12: Hypothetical performance of massive parallelization using survey decomposition.

3.4 Data Subsampling

Survey decomposition breaks the entire EM survey down to many physically self-contained subproblems in small size. Although massive parallelization is expected to tackle the large number of subproblems, it is always desirable to solve as few subproblems as possible. I have noticed the collected data are usually oversampled in space and time and are more than actually needed in inversion. This section discusses the feasibility of modeling only a small portion of available data and introduces the general idea of adaptively subsampling the data using cross validation.

3.4.1 Oversampling in EM modeling

Given an area of exploration, the amount of data from an EM survey depends on the sampling rate of the equipment in space and time/frequency. The data acquisition systems usually samples at the highest possible rate allowed by the budget. The redundancy of data can improve SNR in the presence of random noise from the data processing point of view. However, this does not necessitate the modeling of every single datum. Recognizing why there is oversampling provides a basis for effectively reducing the number of subproblems after survey decomposition.

Oversampling happens if the information density of the signals is relatively low compared to the sampling rate, or in other words, the sampled data are providing almost redundant information. In EM modeling, this can be semi-quantified by comparing the measurement spacing with the scale of survey, which is usually calculated by “diffusion distance” or “skin depth”. Larger scale means lower information density and possibly more oversampling.

Due to the diffusive nature of the EM induction, the EM responses are generally considered “smooth” in both time and space at the scale of most surveys. For example, for a half-space model of 0.002 S/m and TEM data at 10^{-4} s, the diffusion distance is about 300 m, which is significantly larger than the in-line sounding spacing in most airborne surveys. Information density falls off rapidly from early to late time channels. For a ground survey that measures data as late as 1 s, the scale of the survey can be so large that meaningful change of the EM responses may not be seen even if the receiver moves across the entire survey area. Nevertheless, data at that time were still measured on a dense grid. The similar smoothness is also evident in time/frequency domain. Therefore, the first type of oversampling is due to the mismatch between the information density of EM data and the sampling rate in realistic surveys. Some regular or smart subsampling schemes have been proposed to work on this (Siripunvaraporn & Egbert, 2000; Foks et al., 2014).

The second type of oversampling, caused by the regularization in inversion, can also be explained using the concept of information density and the scale of survey. A regularized inversion often starts with strong emphasis on the simplicity of the model, constructing only large-scale features in the model and using little information from the data. At this stage, not all the information in the data is needed; so even a sampling rate tuned to properly match the information density of data yields oversampling. When the inversion proceeds and the regularization parameter decreases, more model structures at smaller scale will be allowed and thus demand a higher sampling rate.

To conclude, in reducing the oversampling, effort can be made by matching the sampling rate to (1) the static scale exhibited in the data and/or (2) the dynamic scale concerned in the process of inversion.

3.4.2 Adaptive subsampling using cross validation

A common practice of choosing a subset of data is to subsample the original data using a coarser grid with uniform spacings. However, such a static grid may have a higher risk of biasing the inversion result. For example, an insufficient subsampling rate on a fixed grid can prevent the inversion from having complete information about the noise; in a worst case scenario, when the subsampling happens to be in-phase with a particular noise source, like the swing of the bird in an airborne survey, or any kind of periodic structure of the earth, significant bias may be created in the model. One solution is to sample the continuous signals dynamically and randomly, which, unlike a uniform grid that only passes low-frequency signals, is not frequency selective (Herrmann, 2010). Using the scheme of random sampling with a uniform probability distribution (constant probability density function) also simplifies the process of subsampling to finding the optimal number of samples over the sampling domain.

The number of samples required in inversion depends on the scale of investigation. Cross validation can be used in the adaptive search for the practically optimal number of samples that matches the scale. This procedure usually involves two subsets of the data: a training subset and a test subset, which are independently chosen. Ideally, a certain type of operation or investigation is supposed to apply to the entire data set; however, due to the cost of accessing all the data, one can choose to only operate on the training subset, from which a decision or result is then obtained. To prevent the decision/result from being misled by the subsampled training subset, the test subset is used as an independent check to verify whether the decision/result made based on the training subset also satisfies the test subset.

If the validation fails, the number of samples in the training subset is deemed not representative enough, so more samples should be added. This adaptive procedure can keep adding new samples until an extreme situation that the training subset is the entire data set, and the decision/result is guaranteed to pass the test. The test subset can have the same number of samples as the training subset or some other reasonable number.

Here I use a demonstrative example to show how cross validation can be indicative in matching the required number of samples to the scale of investigation. The operation in this example is to simply calculate the mean value of a data vector. For a given vector of length 100, different numbers of random samples per subset (N, the size of subset) from 1 to 40 are used to estimate the mean value in both training and test subsets. Then the consistency error, evaluating the subsampling performance, is defined as the absolute difference between the means from the training subset and the test subset for a given N.

The first experiment is operated on a vector having high-frequency signals (Figure 3.13a). The mean values estimated by the training and test subsets at different Ns are indicated by the solid dots and circles in Figure 3.13b. There are local fluctuations in the consistency error, but the overall trend shows that larger N can improve the accuracy of operation. For this particular data vector, about $N = 30$ are required to maintain a performance measured by the consistency error of 1 (Figure 3.13c).

In the second experiment, the high-frequency component in the original vector is removed, leaving a smooth signal representing a broader scale of investigation (Figure 3.13d). The mean values from the two subsets converge quickly at a smaller N (Figure 3.13e), so that the consistency error of 1 is achieved for an N less than 10 (Figure 3.13f). Therefore, by checking the training subset against another independent test subset, it is possible to tell whether the current number of samples is enough to match the current scale of investigation.

In the context of inversion, the training and test subsets are samples from the full data set. Every sample of data is a subproblem characterized by source/receiver location, time, or their combination. The operation on the training subset is computing a model update at a certain scale determined by the trade-off parameter. The test subset validates the proposed model update by comparing the data misfits

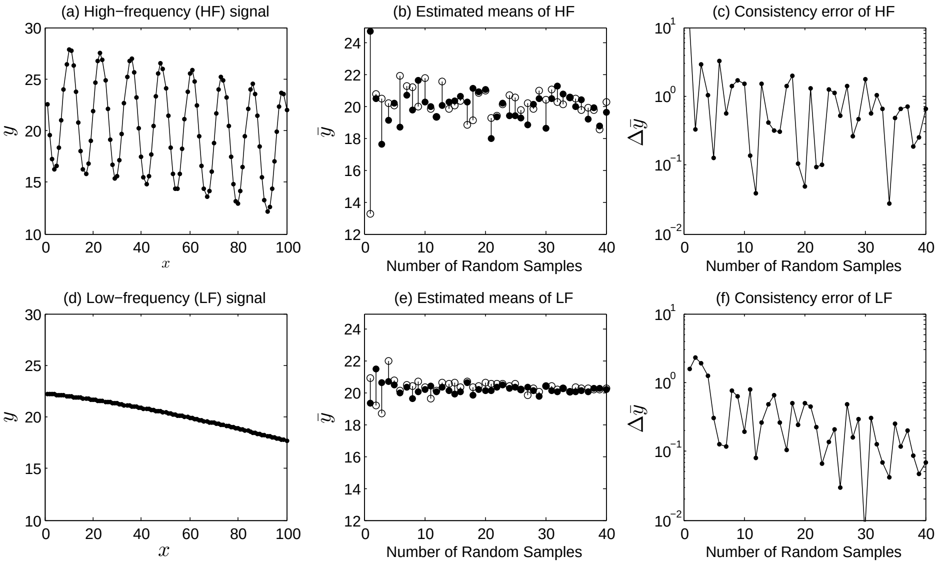

Figure 3.13: Demonstrative example: adapting the number of samples to the scale of investigation using cross validation.

of the test subset before and after the update. If the model update satisfactorily improves the misfit of the test subset, the model update can be accepted and the inversion can proceed to the next iteration with reduced trade-off parameter; otherwise the model update is rejected. In the latter case, the current number of samples is deemed insufficient for the current scale of investigation, so more samples need to be added into the training subset to generate a new proposed model update. This adaptive approach of data subsampling can be implemented in different ways depending on how the data are organized; the specific workflows for airborne and ground loop surveys are discussed in Chapter 4 and 5 respectively.

I also note that this “adapt-to-the-scale” approach exploits the multi-scale nature of EM inversion. The requirement for the modeling accuracy can be relaxed when the data and model at a broader scale is considered, leading to significant saving on the number of data. The implementation of such selective modeling is difficult in the standard algorithm, but becomes a natural choice if the survey is decomposed as proposed here.

3.5 Summary

This chapter lays the theoretical foundation of the framework of survey decomposition in geophysical EM modeling.

Firstly, three computational complexities associated with EM modeling, namely space complexity, time complexity and optimization complexity, are critically examined to reveal the over-computing hidden in the standard algorithm. A significant amount of unnecessary over-computing exists if the conventional global discretization is used. The one-size-fits-all discretization, exemplified by a global mesh that serves all sources, receivers and times/frequencies, is not economical because of the great contrast of modeling scales in one discretization scheme. Any discretization with reduced scale contrast has reduced computational complexity and thus can be more efficient. This idea prompts the decomposition of the survey so that the source, receiver and time/frequency are localized and the subproblems can have customized discretizations.

Secondly, an atomic building block is established as the smallest component for assembling the subproblems of any type of controlled source EM survey in time and frequency domain. An atomic problem only has one dipole source, one point receiver and one time channel/frequency. With the help of superposition, I was able to decompose most of the commonly used CSEM surveys, including airborne EM, ground large loop EM, marine CSEM, DC resistivity, MMR/MIP, etc., into subproblems built by atomic problems.

Once the subproblems are identified, I show how the local mesh and local time discretization are designed for each subproblem. Unlike a global mesh that has fine cells everywhere over the entire survey area, a local mesh is locally refined at the dipole source and receiver locations and usually only covers the region that the subproblem is sensitive to. A local time discretization only models one time channel with a constant step length adapted to the right scale of investigation. The local discretization, being different from the global one, mandates a two-level (global and local) architecture of computation. The actual computations and storage of forward responses, sensitivity and sensitivity-vector multiplication ($\mathbf{Jv}$ or $\mathbf{J}\top\mathbf{v}$) are all carried out within the subproblems at the local level, while the global process is only responsible for summarizing the computing results from the local level. As no PDE solving is run directly on the global mesh, the size of the global mesh is no longer a prohibitive issue when working with large-scale problems.

The logical decomposition based on survey parameters makes the subproblems fine-grained and also independent to each other. This is ideal for massive parallelization. Each subproblem, with its own source, receiver and time/frequency, is a self-contained EM problem. No communication is required after the parallel workers receive the initializing broadcast from the host (global) process and before the computing results are sent back to the host. Thanks to this independence, each datum from the survey can be assigned to a subproblem running on one processor at the maximum scale of parallelization. A simplified run time model is discussed to show the scalability of the new framework.

After decomposition, the original problem becomes a number of small subproblems. I notice that due to oversampling of data, it is not necessary to model all of them in inversion. An adaptive, random and dynamic data subsampling approach using cross validation is therefore proposed to minimize the number of data (subproblems) used in inversion. Again, the fundamental idea is still to match the subsampling rate to the scale of modeling.

In the next two chapters, the theory of survey decomposition is applied to two surveys important in mineral explorations, airborne EM and ground loop EM, with both synthetic and field data examples.

Chapter 4

Airborne TEM

Given the theoretical background and the general idea of survey decomposition in Chapter 3, this chapter uses the framework of survey decomposition to solve the practical modeling problems in airborne TEM. Because of its unique way of collecting data, the entire airborne survey is decomposed into many soundings. Every sounding is a subproblem with its source, receiver and measurements at a sequence of time channels. This chapter demonstrates the benefits of carrying out forward modeling and sensitivity computation for airborne TEM based on local meshes. A synthetic example is provided as a proof-of-concept inversion. The VTEM data at Mt. Milligan, partly inverted in Chapter 2, is more efficiently inverted as a whole using survey decomposition.

4.1 Local Mesh for Airborne Sounding

The generic formula in Chapter 3 are still valid for airborne soundings. Treating the airborne survey as a special case, some simplifications can be made:

- Every sounding has its own local mesh, on which a subproblem concerning all of the delay times is solved.
- The separation of transmitter and receiver is small enough so there is only one center of refinement. A local mesh for airborne looks similar to Figure 3.7 but without the cells between the transmitter and receiver.

- Such a local mesh does not need rotation, so **M** in equation 3.3 is the identity.
- The non-rotating local mesh has grids parallel to the global mesh grids; **R** in equation 3.2 can be more quickly computed using the rigorous computational geometry.
- Every sounding has its own magnetic dipole transmitter, so there is no superposition and **S** in equation 3.4 is also an identity.
- If **M** and **S** are ignored, the computation of $\mathbf{J}_g^\top \mathbf{J}_g \mathbf{v}$ and $\mathbf{diag}(\mathbf{J}_g^\top \mathbf{J}_g)$ only need one global-local-global cycle of communication, cutting the communication time in half.

4.1.1 Forward modeling on local mesh

To show that airborne soundings can be effectively modeled on local meshes, I use a synthetic model with a random conductivity structure. The original model is defined on a global mesh finely discretized to 50 m resolution for a survey area of 4×4 km (Figure 4.1a). I choose the sounding at the center of the model (indicated by a red dot). The first local mesh (Local Mesh 1) has an expansion rate of 1.2 and a boundary distance 1500 m. This mesh performs well at early times compared to the global mesh but has up to one order of magnitude error at late times due to an insufficient distance to the boundary. The local mesh is then expanded to a boundary distance of 3000 m (Local Mesh 2, Figure 4.1b), which is still much smaller than the global mesh. With the additional three cells in every direction, Local Mesh 2 is able to provide forward modeled data within 5% of the global mesh results at late times (Figure 4.2).

Although fine structures of the model away from the sounding location are represented by bulky cells, a local mesh is still capable of producing a good simulation at a much reduced cost. Table 4.1 summarizes the three forward modelings on the global and local meshes. The global mesh, being large for the entire survey, is expensive in both time and memory. Because a large Maxwell matrix needs to be factorized, 12 processors are required. The local meshes are much smaller and thus more efficient, even though only one processor is used.

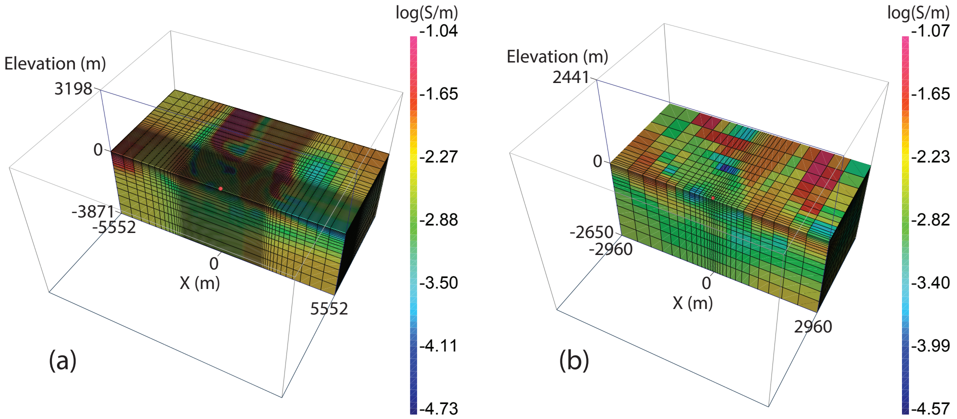

Figure 4.1: The synthetic model on the global mesh and a local mesh.

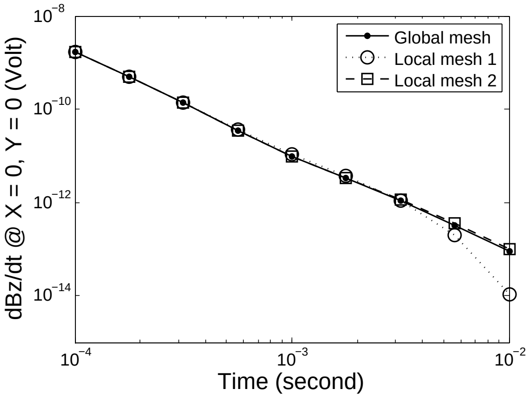

Figure 4.2: Forward modeled data on the global mesh and the two successive local meshes in the test.

Table 4.1: Forward modeling of an airborne sounding on the global mesh and two local meshes.

	Global Mesh	Local Mesh 1	Local Mesh 2
Number of cells	$108 \times 108 \times 33$	$22 \times 22 \times 27$	$28 \times 28 \times 30$
Time for one factorization	123 s	1.5 s	4.1 s
Time for one time step	1 s	0.02 s	0.06 s
Number of processors	12	1	1
Memory (one $\mathbf{A}^{-1}$)	40.5 GB	174 MB	395 MB

Multiplying the cost by the number of soundings further signifies the benefits of using a local mesh approach in airborne inversion. If the entire survey contained 1080 soundings and one forward modeling has 4 factorizations and 50 time steps, the global mesh would require at least 54492 s for one complete forward modeling on 12 processors and further speed-up by adding more processors is difficult. However Local Mesh 2 needs 19.4 s for one sounding (subproblem) and a total of 20952 s if all computations were carried out on a single processor. Distributing these jobs over the 12 processors reduces this to 1746 s. Increasing the number of processors continues to achieve this linear benefit.

4.1.2 Sensitivity on a local mesh

The same synthetic model in the forward modeling example (Figure 4.1a) is used to show the global sensitivity can be reasonably reconstructed using a local mesh. At the same sounding, I choose a row of the sensitivity matrix corresponding to the dB_z/dt datum at time $t = 0.001$ s and present the sensitivity on the mesh. The sensitivity is computed directly on the global mesh (Figure 4.3a) and on Local Mesh 2 then using equation 3.6 for interpolation (Figure 4.3b). The local mesh result matches well with the sensitivity computed on the global mesh. The computational costs are similar to those of forward modeling.

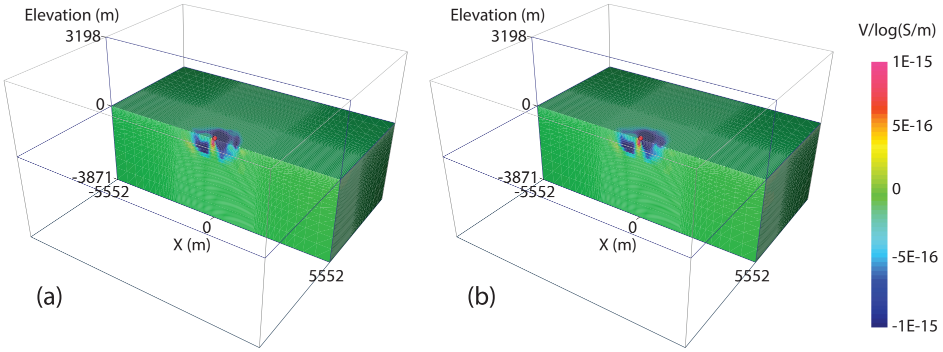

Figure 4.3: Sensitivities of dB_z/dt datum at $t = 0.001$ s for the synthetic model: (a) the global mesh result and (b) the local mesh result with interpolation.

4.1.3 Synthetic inversion of a two-block model

A synthetic example is designed to test the inversion using survey decomposition. The true model consists of two conductive prisms, 0.1 and 0.05 S/m, buried in a 0.01 S/m uniform half-space (see depth slice at 150 m in Figure 4.4a). A synthetic airborne TEM data set is created on a 13×37 data grid over a 1.2×3.6 km area marked by the red dots. Seven time channels of dB_z/dt data from 10^{-4} to 10^{-2} s are simulated at 481 sounding locations 100 m apart. The synthetic data are noise-free, but I require the inversion to fit the data with an assigned uncertainty of 5%. The mesh used for the creation of the synthetic data, and which served as the global mesh, has 155820 cells (53×98×30). A complete forward modeling consists of four factorizations and 48 time steps, and modeling all 481 soundings on this mesh takes about one hour.

As a benchmark, I first invert the synthetic data set directly on the global mesh with a 0.01 S/m half-space as the initial and reference models using the standard algorithm. The inversion, using 128 GB memory on 24 processors, takes 7 Gauss-Newton iterations and achieves the target misfit in 156 hours (Figure 4.4b).

For survey decomposition, I use the identical inversion parameters but with local meshes to compute the forward modeling and sensitivities. The first inversion test is carried out by equally distributing 481 sub-problems (soundings/local meshes) on 12 processors. The target misfit is achieved within 15 hours after 8 iterations. Figure 4.4c shows the conductivity model recovered using the local meshes; the two prisms are delineated with correct geometries and conductivities; this result is similar to the model recovered by the standard algorithm directly on the global mesh (Figure 4.4b).

The second run of the same inversion on 24 processors is carried out to test the scalability. In this case the workload on each processor is cut in half to 20 sub-problems. The second inversion produces an identical inversion result, convergence curve, and data misfit, but the total CPU time is reduced to about 7.5 hours. Even greater potential for speed-up exists if more processors are available.

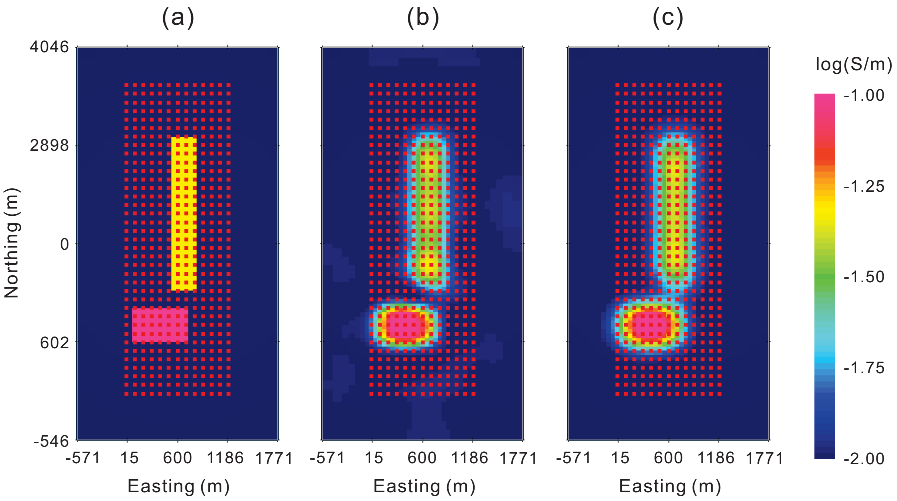

Figure 4.4: Synthetic airborne TEM inversion usign survey decomposition. Depth slices at 150 m of (a) the true model, (b) global discretization inversion model and (c) survey decomposition inversion model.

4.2 Adaptive Subsampling of Soundings

4.2.1 Subsampling workflow

Although the framework of survey decomposition greatly reduces the time of 3D inversions, and further acceleration is achievable by adding more processors, the total number of soundings in an airborne EM survey typically numbers hundreds of thousands and this can be formidable. Therefore, data subsampling is very necessary in airborne data inversion.

Using the concept discussed in Section 3.4, I develop an adaptive subsampling workflow for airborne soundings (Algorithm 1). This workflow features random and dynamic subsampling of soundings and the number of soundings N at every iteration is adapted to the scale of investigation in inversion, which is essentially controlled by the trade-off parameter β.

Adaptive subsampling of soundings requires frequent solves of the inverse problem on the training subset for the proposal of model update $\delta\mathbf{m}$, and the forward problems need to be solved for the test subset to cross validate the model update. In practice, I want the computation of the local mesh design, matrix factorizations and forward problems done for the test subsets to be recycled for the training subset (inversion) in the next iteration, so N is always doubled when increased. Each Gauss-Newton iteration involves two steps: (1) selecting the training subset and computing $\delta\mathbf{m}$ using the training subset; (2) selecting the test subset and evaluating whether the proposed model update is acceptable by comparing the misfits of the data in the test subset before and after the update. There are three possible consequent actions:

- End of inversion. This happens when the data misfit is below a tolerance.
- Enough number of soundings (N is large enough). If the data misfit of the test subset is sufficiently reduced (controlled by a factor μ) after model update, $\delta\mathbf{m}$ is accepted and the test subset, already having the forward solutions $F(\mathbf{m}+\delta\mathbf{m})$ for the updated model, can be used as the new training subset in the next iteration after additional computations of the sensitivity $\mathbf{J}(\mathbf{m}+\delta\mathbf{m})$. There is no need to increase N; β is cooled to a smaller value to allow more model structures in the next iteration.
- Not enough number of soundings (N is not large enough). If the model update does not improve the data misfit of the test subset, $\delta\mathbf{m}$ is declined, and the test subset, already having the forward solutions $F(\mathbf{m})$ for the original model, is appended to the existing training subset after additional computations of the sensitivity $\mathbf{J}(\mathbf{m})$; this double-sized training subset is used in another attempt to make a model update with the unchanged β. N is doubled.

The solve for $\delta\mathbf{m}$ in Algorithm 1 can be attained by using methods other than the Gauss-Newton method implemented here. Both explicit and implicit sensitivity can be used at the end users' discretion.

4.2.2 Synthetic inversion with adaptive soundings

The combination of survey decomposition and adaptive soundings can greatly reduce the CPU time for the inversion. I illustrate this with the same synthetic data set from the two-prism model in Figure 4.4 and invert with the same parameters

Algorithm 1 Inversion with adaptive subsampling of airborne soundings

```
Initialization:
  Select $N$ soundings for training subset $\mathbf{S}^{train}$
  Compute forward responses $F(\mathbf{m})^{train}$ at $\mathbf{S}^{train}$
  Compute sensitivity $\mathbf{J}(\mathbf{m})^{train}$ at $\mathbf{S}^{train}$
  Empty test subset $\mathbf{S}^{test}$
repeat
    Solve for $\delta\mathbf{m}$ using $F(\mathbf{m})^{train}$, $\mathbf{J}(\mathbf{m})^{train}$, $\beta$ and $\mathbf{m}$
    Select $N$ soundings for test subset $\mathbf{S}^{test}$
    Compute forward responses $F(\mathbf{m})^{test}$ at $\mathbf{S}^{test}$
    Compute data misfit $\phi(\mathbf{m})^{test}$ at $\mathbf{S}^{test}$
    Compute forward responses $F(\mathbf{m}+\delta\mathbf{m})^{test}$ at $\mathbf{S}^{test}$
    Compute data misfit $\phi(\mathbf{m}+\delta\mathbf{m})^{test}$ at $\mathbf{S}^{test}$

    if $\phi(\mathbf{m}+\delta\mathbf{m})^{test} \leq tol$ or $\phi(\mathbf{m}+\delta\mathbf{m})^{test} < \mu \cdot \phi(\mathbf{m})^{test}$ at $\mathbf{S}^{test}$ then
        Compute sensitivity $\mathbf{J}(\mathbf{m}+\delta\mathbf{m})^{test}$ at $\mathbf{S}^{test}$
        $\mathbf{S}^{train} \leftarrow \mathbf{S}^{test}$
        $F(\mathbf{m})^{train} \leftarrow F(\mathbf{m})^{test}$
        $\mathbf{J}(\mathbf{m})^{train} \leftarrow \mathbf{J}(\mathbf{m})^{test}$
        Empty $\mathbf{S}^{test}$, $F(\mathbf{m})^{test}$, $\mathbf{J}(\mathbf{m})^{test}$
        Reduce $\beta$
        $\mathbf{m} \leftarrow \mathbf{m}+\delta\mathbf{m}$
    else
        Compute sensitivity $\mathbf{J}(\mathbf{m})^{test}$ at $\mathbf{S}^{test}$
        $N = N \times 2$
        $\mathbf{S}^{train} \leftarrow \mathbf{S}^{train} \cup \mathbf{S}^{test}$
        $F(\mathbf{m})^{train} \leftarrow F(\mathbf{m})^{train} \cup F(\mathbf{m})^{test}$
        $\mathbf{J}(\mathbf{m})^{train} \leftarrow \mathbf{J}(\mathbf{m})^{train} \cup \mathbf{J}(\mathbf{m})^{test}$
        Empty $\mathbf{S}^{test}$, $F(\mathbf{m})^{test}$, $\mathbf{J}(\mathbf{m})^{test}$
    end if
until $\phi(\mathbf{m})^{test} < tol$
```

using 24 processors. The initial and reference models are still 0.01 S/m half-space. The initial number of soundings is set to 48 at the beginning of adaptive soundings. The number of soundings gradually increases to 384 until the target misfit is achieved within 2.5 hours after 6 iterations (Figure 4.5).

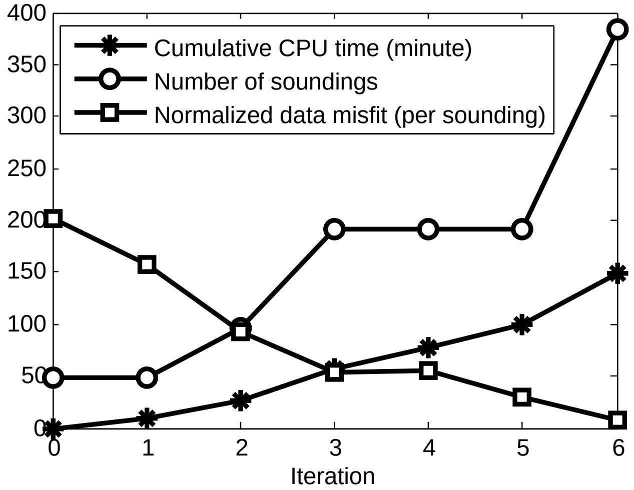

Figure 4.5: Summary of synthetic airborne inversion using both survey decomposition and adaptive soundings.

The final inversion model (Figure 4.6d) is almost identical to the recovered model in the inversion which only uses the local mesh method (Figure 4.4c). The intermediate models created by adaptive soundings are somewhat different (Figure 4.6). The first iteration only has 48 soundings, so the inversion recovers the global trend of the conductivity plus significant artifacts near the surface (Figure 4.6a). As the number of soundings increases, the desired targets become clearer (Figure 4.6b and c). The final model in Figure 4.6d is visually the same as the inversion model in Figure 4.4 if they are rendered in the same color scale. By using adaptive soundings, about 60% of the CPU time is saved compared to modeling all 481 soundings for every iteration.

Although random subsets of soundings have been used, the overall coverage of data, in terms of how many times a sounding is used during the inversion, is adequately balanced in space (Figure 4.7). On average, each sounding is selected

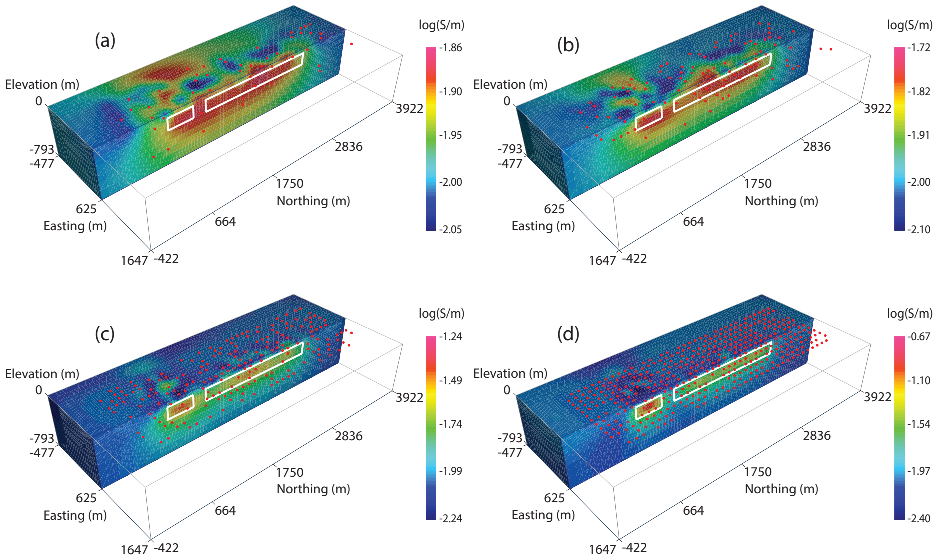

Figure 4.6: Conductivity models recovered at different iterations of adaptive soundings: (a) Iteration 1 with 48 soundings, (b) Iteration 2 with 96 soundings, (c) Iteration 4 with 192 soundings, (d) Iteration 6 with 384 soundings. The red dots indicate the sounding locations, and the white boxes outline the exact locations of the two prisms.

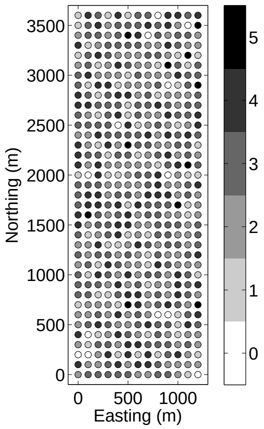

Figure 4.7: Counts of selection for every sounding throughout the entire inversion.

for approximately 2.4 iterations. At the end of the procedure there were only 10 soundings that were not used at all. We also note that most commercial airborne EM systems have a much denser spatial sampling rate (5∼10 m) than our synthetic survey (100 m). Therefore, using adaptive sounding approach to find the necessary number of soundings for an inversion is more economical than inverting every sounding in the overly redundant data set. This is shown in the field example in the next section.

4.3 Inversion of the Entire Mt. Milligan Data Set

4.3.1 3D inversion

When solving EM problems on local meshes, the forward modeling is never carried out on the global mesh, so the number of cells in the global mesh becomes less crucial. For the VTEM data at Mt. Milligan I design a global mesh with 443520 ($88\times84\times60$) cells to hold the entire survey area at 50 m horizontal resolution and 20 m vertical resolution for the topography (Figure 4.8. This mesh, along with the large number of soundings, is too large for the standard algorithm in Oldenburg et al. (2013) to be practically carried out.

Figure 4.9 summarizes three key parameters of the inversion: the cumulative run time, the number of soundings used at each iteration and the normalized data misfit. For early iterations only 24 soundings were needed to build up the large-scale conductivity distribution. Smaller-scale features were built up by adding soundings. The number of soundings used in the final iteration is only 192, which implies that the necessary number of soundings for 3D VTEM inversion at Mt. Milligan is about 200, only 1.4% of the total number of soundings acquired in the survey. During the entire inversion procedure, information from 744 soundings (5% of the total number of soundings) was incorporated into the final model and thus 744 subproblems (local meshes) have been solved. The total time was 4.3 hours and I anticipate more speed-up if additional processors were available.

A depth slice of the 3D inversion model over the whole survey area at an elevation of 950 m and a cross section A-B at 6109500N cutting the major stock MBX are shown in Figure 4.10. White lines outline the major monzonite stocks

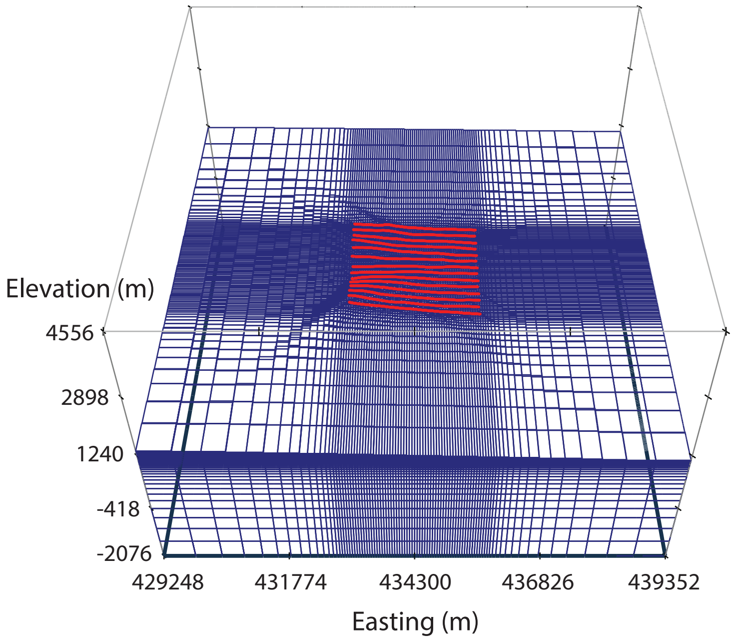

Figure 4.8: Global mesh for the 3D inversion of the VTEM data at Mt. Milligan. The VTEM sounding locations are indicated by the red dots.

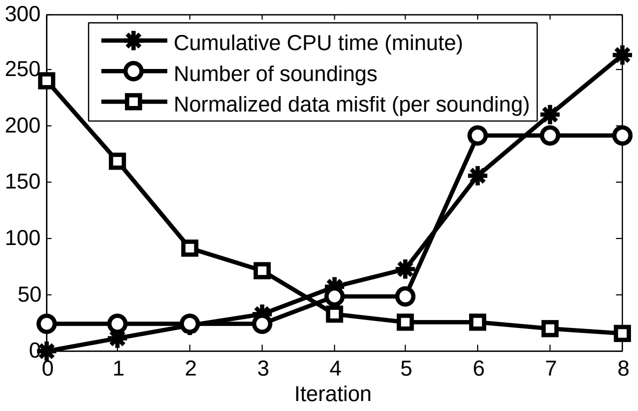

Figure 4.9: Summary of Mt. Milligan VTEM 3D inversion.

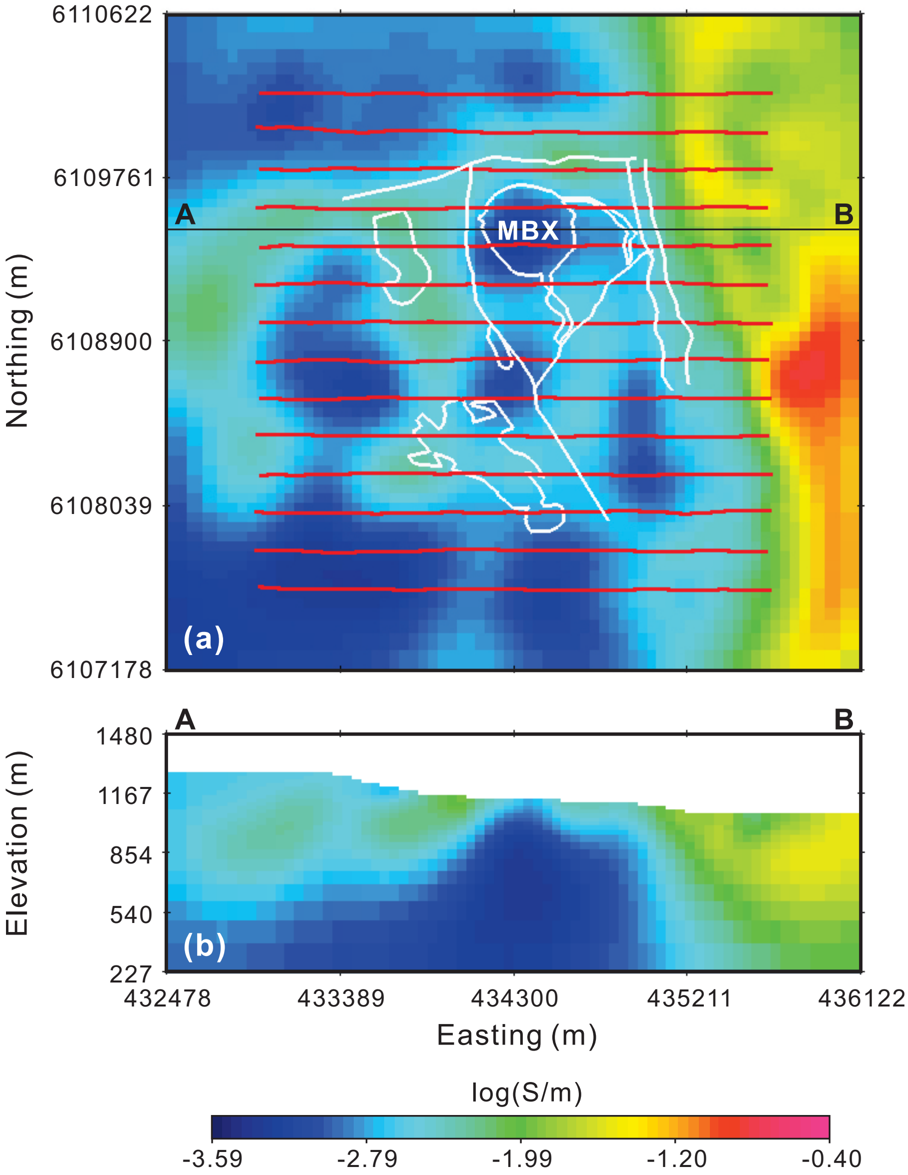

Figure 4.10: Conductivity model of the Mt. Milligan VTEM 3D inversion: (a) a depth slice at 950 m elevation, (b) a cross section A-B at 6109500N.

and the faults. VTEM flight lines are indicated by the red lines. Like the small-scale inversion in Section 2.4, the resistive MBX stock is clearly delineated by the 3D inversion.

4.3.2 Final model validation

Lastly, I carry out a complete forward modeling using the recovered model in Figure 4.10 at all 14362 soundings.

In this data set there are 14362 soundings and 8 data (time channels) at each sounding. The total number of data is 114896. The misfit of this validation forward modeling, if normalized by the total number of data, is $\phi_d/N = 2$, which compares satisfactorily with $\phi_d/N = 1.87$ estimated by the 192 random soundings at the last iteration and $\phi_d/N = 2.1$, the data misfit achieved in Section 2.4 for a smaller area around the MBX.

The time channel grids of the observed and predicted data at 0.68 ms (a delay time corresponding to the expected depth of the deposit) in Figure 4.11 show the 3D model recovered by using local mesh and adaptive soundings can reasonably reproduce the observed data. The white lines indicate 14362 soundings approximately 3 m apart from each other. The scattered black and red dots show the 744 and 192 soundings used in all iterations and the last iteration of inversion respectively.

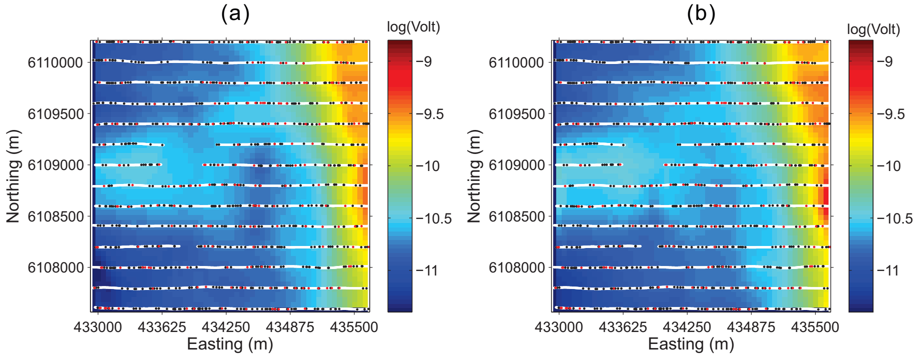

Figure 4.11: Data grid of time channel at 0.68 ms for the observed and predicted dB_z/dt data at Mt. Milligan.

4.4 Summary

Airborne EM data are difficult to invert in 3D because of the large computations required to handle multiple transmitters and large meshes needed to represent the volume being modeled. The computational difficulties are tackled by using the framework of survey decomposition.

Modeling an airborne EM survey has large space complexity and moderate time complexity, so such a survey can be conveniently decomposed so that every airborne sounding is a subproblem modeling all delay times. A local mesh is optimally designed to compute the forward responses and sensitivities at every sounding and can yield good accuracy with a modest number of cells. As such an entire forward modelling or sensitivity calculation is easily distributed among an array of processors in parallel.

The second area of progress pertains to the number of soundings used at each iteration in the inversion. It is well known that modeling every sounding in an airborne data set is not necessary. I propose a random and dynamic subsampling method, called adaptive sounding. It is essentially a random sampler with an adaptive number of soundings selected for each iteration. Fewer soundings are selected in early iterations to build up large-scale structure and more soundings are added later as the regularization is relaxed and additional structure is needed to fit the data. The procedure has an additional optional check in that the user may carry out a complete forward modeling at the end, using all the soundings, to validate the constructed model.

In a test using synthetic data from a two-prism model I reduced the inversion time, compared to the standard approach, by more than a factor of 60. That example was small and further disparity between the standard and new method will increase as the size of the problem and the number of processors increase.

In a field example, which was too large to be practically solved in Section 2.4, the inversion using the survey decomposition takes only 4.3 hours on 24 processors. This is a very substantial speed-up and represents a major step toward making routine 3D inversion of airborne EM data a reality.

Chapter 5

Ground Loop TEM

While the decomposition of an airborne EM survey is straightforward, a ground survey using large loops can be more complicated because:

- Ground loop surveys are less standardized. The layout of a survey, including the geometry of the loop and locations of receiver, varies greatly from one site to another.
- The transmitter loop is usually too large to be local, so it must be properly decomposed to a number of small dipole sources using superposition.
- Data at large dynamic range of times are measured. The time complexity associated with ground loop EM necessitates decomposition for individual delay times.

This chapter first introduces a field data example, whose specifications are used for my investigation of ground loop EM survey. The large loop source is efficiently decomposed into looplets by using an adaptive looplet approach. Similar to the airborne case, data subsampling is also an integrated part of the workflow; this minimizes the over-computing and the number of subproblems to be solved. The framework is tested with a synthetic example and finally applied to the field data to show its efficiency and effectiveness.

5.1 Field Example: SQUID Data at the Lalor Mine

The Lalor Mine deposit is located in the Chisel Basin portion of the Flin Flon Greenstone Belt, and about 8 km west to Snow Lake in central Manitoba, Canada. The deposit, discovered in 2007 within the meta-volcanics, meta-sediments and granitoids of the Churchill province near Chisel Lake and Snow Lake, is believed to be the largest volcanogenic massive sulfide (VMS) ever found in this area (Figure 5.1).

VMS deposits in this region have geophysical signatures of good conductors, usually discrete bodies, buried in a relatively resistive host. Due to VMS's strong EM responses, a variety of EM surveys have been conducted to delineate the complex of conductors, which are most likely to be directly associated with economic Zn-Au-Cu mineralizations. One of the surveys, featuring an advanced JESSY HTS SQUID (superconducting quantum interference device) magnetometer, is of particular interest. Compared to the conventional dB/dt (time derivative of the B-field) data, SQUID data have two advantages: (1) a SQUID magnetometer measures three-component data in unprecedented quality (Chwala et al., 2001; Osmond et al., 2002; Leslie et al., 2008), while the dB/dt data measured by induction coils tends to be sensitive to noise at late times; (2) B-field data can better detect the slow decays in a good conductor (Smith & Annan, 1998), while the dB/dt data may confuse the slow decays with low conductivity. For these reasons, the SQUID is particularly useful in exploration of VMS at the Lalor Mine as SQUID data enhances the signals from a good conductor and is capable of measuring the weak B-field signals from deep targets at very late times. As a relatively new technique, B-field SQUID data have not previously been interpreted using rigorous 3D modeling; therefore, work on 3D inversion of this data set has additional value of novelty.

The original SQUID survey at the Lalor Mine consists of multiple transmitter loops and receiver lines. In this thesis, I concentrate on a 900 × 1500 m rectangular transmitter loop, called the "East Loop", and data from two receiver lines perpendicular to each other; this provides some degree of 3D coverage. The locations of the transmitter and receivers are shown in survey coordinate system in Figure 5.2. The transmitter, powered by a Phoenix TXU-30, transmits a step-off waveform at a base frequency of 0.5 Hz. The magnetic field data measured by the SQUID were

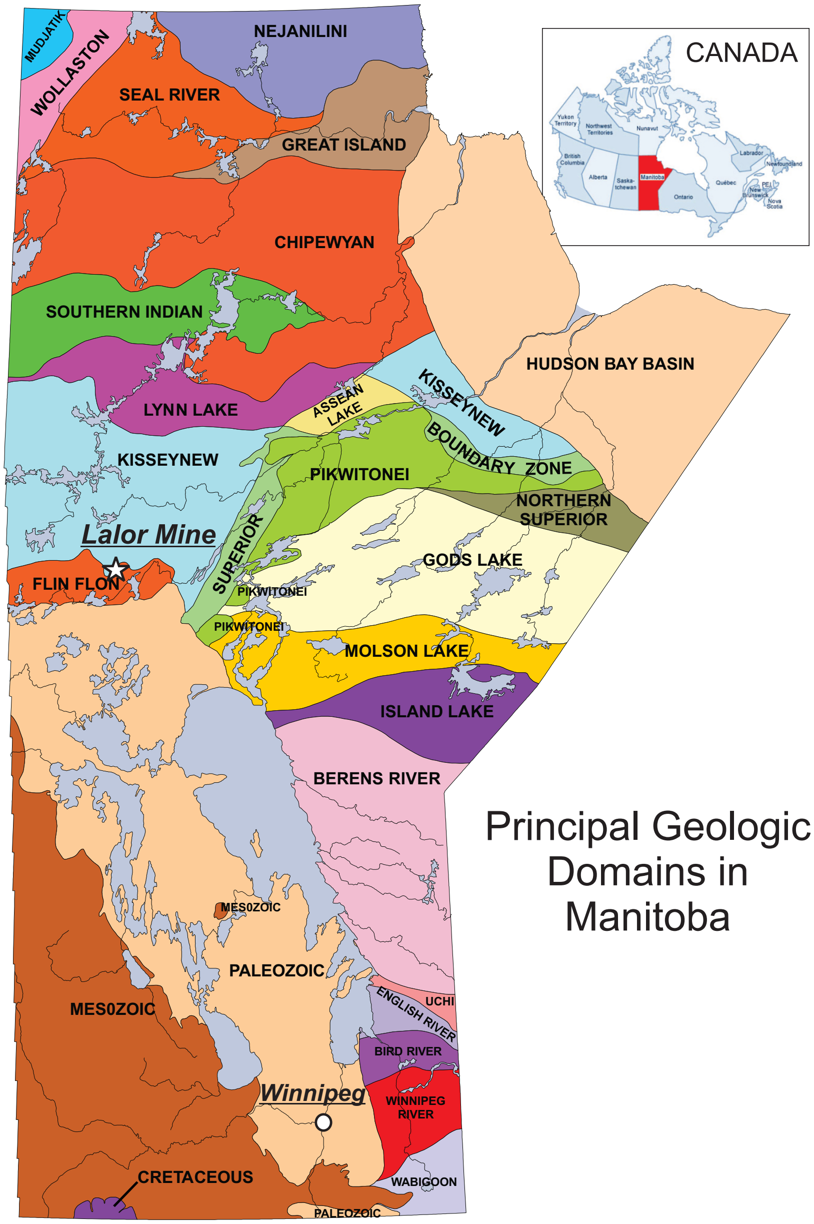

Figure 5.1: Location of the Lalor Mine deposit on the map of geologic domains in Manitoba (adopted from the government of Manitoba's website).

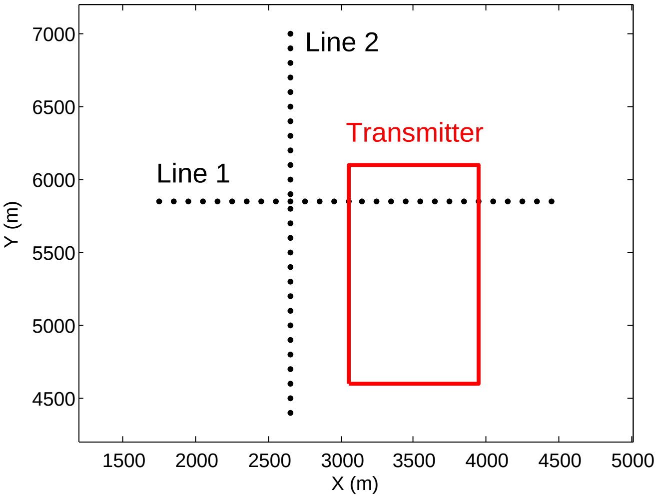

Figure 5.2: Layout of the SQUID ground loop TEM survey at the Lalor Mine.

recorded using an EMIT SMARTem24 at 39 time channels from 0.1 to 371 ms. The data before 0.87 ms are highly contaminated by correlated noise; as a result, only the data from 0.87 to 371 ms are used in the inversion. The survey layout in Figure 5.2 is to be used as a representative example for both synthetic and field data studies throughout this chapter.

5.2 Decomposition of Large Transmitter

One of the complications of decomposing a ground TEM survey is that the transmitter is too large to be treated as a localized source. In this section I represent the large transmitter loop by using the linear combination of some magnetic dipole sources (looplets), and show how they perform in computing the forward responses and sensitivity.

5.2.1 Atomic problem with long-offset source-receiver pair

In most cases, the source and receiver in the subproblems of a ground survey can have offset from zero to a few kilometers. As a result, the local meshes used in the subproblems of a ground survey need two local refinements. This subsection tests a sequence of atomic problems with an offsetting source-receiver pair on a

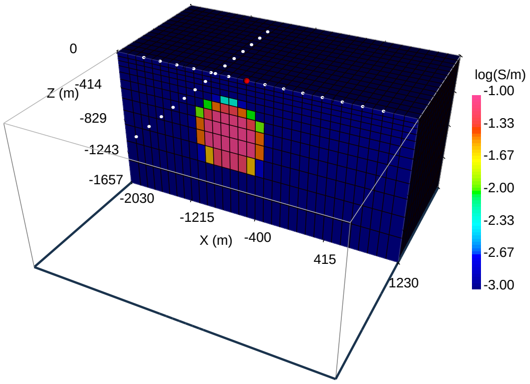

Figure 5.3: Conductivity model and locations of sources and receivers for the atomic problem modeling test. The magnetic dipole transmitter is marked by the red dot and the receivers by the white dots.

synthetic model.

The synthetic model is a 10^{-1} S/m conductive sphere buried in a 10^{-3} S/m uniform half-space (Figure 5.3); the magnetic dipole source is very close to the target on the surface (red dot in Figure 5.3), giving a challenging scenario for numerical modeling. The receivers move along two lines (white dots in Figure 5.3 Line 1 along the x-direction and Line 2 along the y-direction). This yields many different types of orientations with respect to the source, receiver and target.

Four delay times, 10^{-4}, 10^{-3}, 10^{-2} and 10^{-1} s, are modeled in this test. Each atomic problem has its own time discretization requiring 15 steps from $t = 0$ to its modeled delay time, and its own local mesh based on the location of the source and receiver and also the time. The adaptive strategy described in Section 3.3 is employed. The initial smallest cell size and domain size of a local mesh are first inferred using the diffusion distance. The domain size then expands at a rate of 1.5; the cell size refines at a rate of 0.6.

The results at the four time channels obtained with local discretizations (plotted as dots) are compared to the results from the global discretization modeling (solid lines) in Figure 5.4. A high degree of agreement is observed for most of the data, especially for the two early time channels 10^{-4} and 10^{-3} s. The $\mathrm{H_y}$ data of Line 1 are supposed to be zero due to the symmetry, and the atomic modeling results are too small to be on the chart. Some moderate degree of discrepancy are seen for $\mathrm{H_x}$

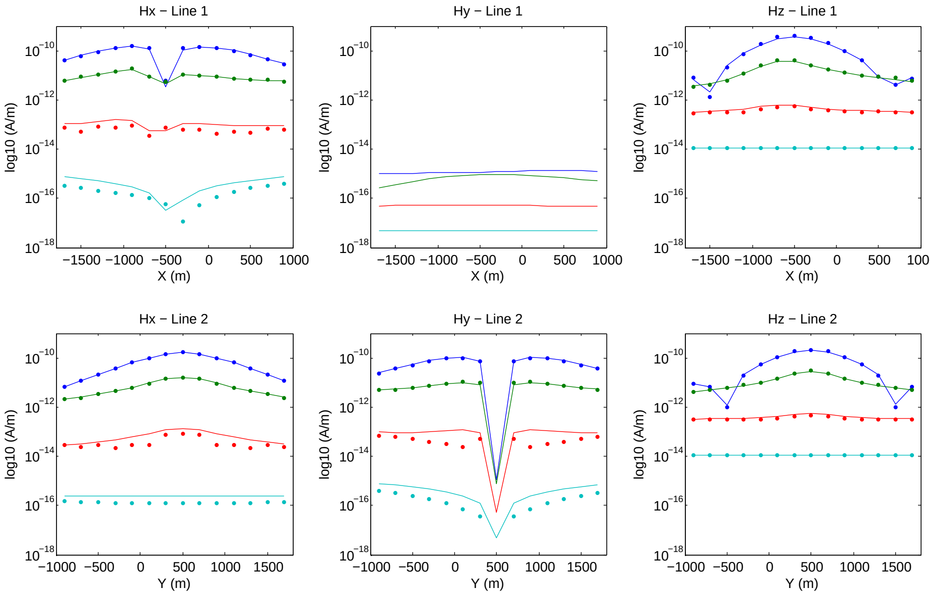

Figure 5.4: Atomic problem modeling results obtained with local and global discretizations along the two perpendicular lines in Figure 5.3. The local and global discretization modelings are plotted as dots and solid lines respectively. The time channels are distinguished by colors.

and H_y at 10^{-2} and 10^{-1} s when the local meshes rotate about 45° with respect to the global coordinate system. This can be improved by specifically fine-tuning the parameters of the local mesh design at the cost of increased computation; however, I have noticed that the H_x and H_y components are mostly much weaker than the H_z component due to the coupling with the horizontal transmitter loop; H_z is modeled reasonably well using local discretization everywhere at all time channels. The slightly poorer modeling results for H_x and H_y becomes less problematic in the inversion when the data misfits are normalized by the total field strengths. Then the null-coupled component data are down-weighted.

5.2.2 Adaptive looplet

A typical ground TEM survey uses a large transmitter loop on the ground and measures the magnetic-field data at many receiver locations around the transmitter at a sequence of time channels. The survey is first decomposed to many receiver-time pairs (R-T) sharing the same large transmitter. Within each R-T, the problem of modeling is solved by using many looplets inside the transmitter loop and then summing the results with defined weights; every subproblem has a looplet as the source and an offsetting receiver measuring data at a particular delay time.

The first task is to find the locations of the looplets and the weights associated with each of them. This problem is mathematically identical to the numerical integration of a 2D function, which in the context of this EM modeling is the EM response (or sensitivity) measured at a fixed R-T when moving the looplet source to different locations. The sampling rate that accurately computes the integration varies depending on the complexity of the 2D function. As discussed in Section 3.4, R-T at late time channels may require far fewer looplets than at early times due to the multi-scale nature of EM diffusion. The ideas of "adapt-to-the-scale" and "refine-until-no-change" are used to develop an adaptive scheme of looplet selection featuring random sampling.

The adaptive looplet procedure starts with only one looplet at a random location in the area bounded by the large transmitter wires; then another randomly located looplet is added to the existing collection of looplets and the sum of forward modeling responses from all the looplets is updated; the addition stops if there is no meaningful change of the sum when more looplets are added. For a collection of N looplets, the forward responses from them are summed with the weights, physically equivalent to the magnetic dipole moments or effective looplet area calculated from the Voronoi diagram (Aurenhammer, 1991). Figure 5.5 shows an instance of adaptive looplets refinement in a time-lapse manner; six Voronoi diagrams are generated as the number of looplets gradually increases from 1 to 6; the weights (moments or effective area) assigned to the looplets change if the tessellation changes; no additional looplets are needed if the summed responses from $N = 6$ and $N = 5$ are sufficiently close.

Every R-T has its own number of looplets, as well as its own summing matrix **S** described in equation 3.4 and 3.7 to combine the results from looplet subproblems together. If the data/sensitivity from a looplet is ordered in x-y-z (equation 3.5 and 3.6), **S** can be assembled using weighted identity matrices. For example, for the ith

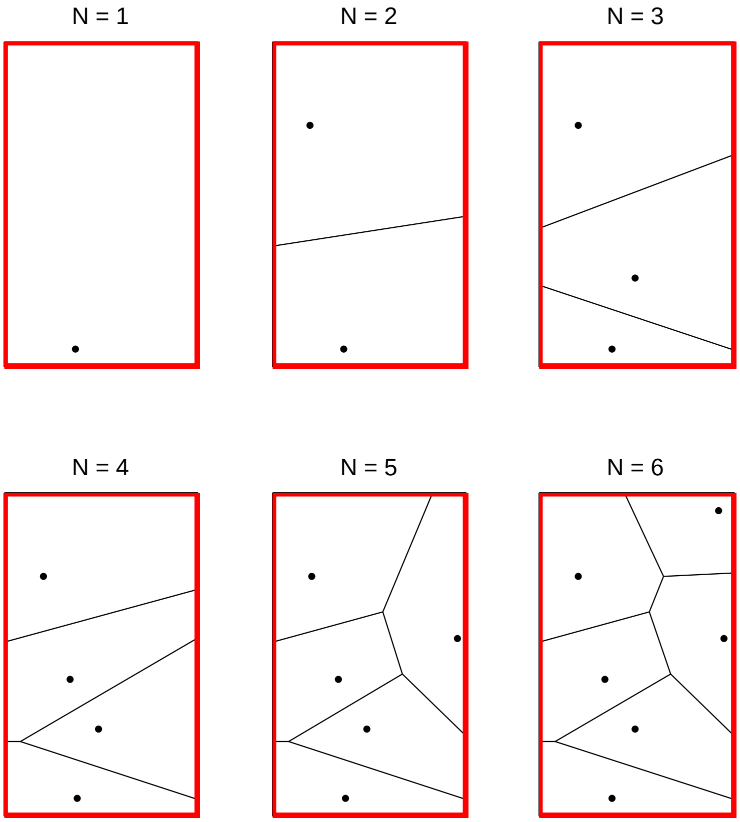

Figure 5.5: Decomposition of a transmitter loop using adaptive looplets refinement based on the Voronoi tessellation. The tessellation and the looplets' moments are updated after every addition of looplets (black dots).

R-T pair that has n looplets, the summing matrix can be

$$\mathbf{S}_i = \begin{pmatrix} w_1 & & & w_2 & & & & w_n & & \\ & w_1 & & & w_2 & & \cdots & & w_n & \\ & & w_1 & & & w_2 & & & & w_n \end{pmatrix}, \tag{5.1}$$

where w_n is the weight for the nth looplet derived from the Voronoi diagram. The final $\mathbf{S}$ for the entire survey is therefore obtained by putting many R-T summing matrices (equation 5.1) together in a block-diagonal pattern

$$\mathbf{S} = \begin{pmatrix} \mathbf{S}_1 & & & \\ & \mathbf{S}_2 & & \\ & & \mathbf{S}_3 & \\ & & & \ddots \end{pmatrix}. \tag{5.2}$$

5.2.3 Forward modeling using looplets

The adaptive looplet approach is tested on the same synthetic model used in Subsection 5.2.1, but the source is the actual large loop and the receiver spacing is now refined to 100 m. Figure 5.6 shows a cross section of the model (padding cells removed), as well as the transmitter loop (red lines) and receivers (white dots). The layout of the synthetic survey emulates the one used in the Lalor Mine SQUID survey, but in a synthetic coordinate system. The data collected at the receiver location (-900, 500), marked by a yellow dot in Figure 5.6, at 10^{-4}, 10^{-3}, 10^{-2} and 10^{-1} s corresponding to four R-T pairs (one for each delay time), are taken as examples for a close-up examination.

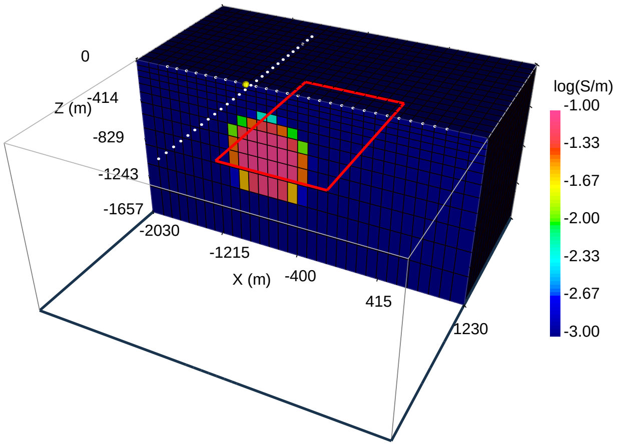

Figure 5.6: Conductivity model and survey layout for the adaptive looplets modeling test. The transmitter loop is indicated by the red lines and the data at one of the receivers marked by a yellow dot are tested for comparison.

For the purpose of benchmarking, the desired data are first modeled directly on the global mesh with 169136 (62×62×44) cells. A complete forward modeling of this survey for time from 10^{-4} to 10^{-1} requires 4 different step lengths and 15 time steps per step length (60 steps in total). The CPU time is 234 s on a 6-core computer in parallel mode. The peak memory usage is 11.4 GB.

The adaptive looplet algorithm determines that 10, 7, 4 and 3 looplets are required for the four R-T pairs at 10^{-4}, 10^{-3}, 10^{-2} and 10^{-1} s, given a tolerance of 2%. The latest delay time may only need one or two looplets, but a minimum of 3 looplets per R-T is required by the adaptive algorithm to detect the stop criterion. The final looplet weights of every R-T are obtained from the Voronoi diagrams in Figure 5.7; it is evident that the number of looplets for a R-T pair has a strong dependence on the delay time of the data and oversampling at later times can be minimized by using adaptive looplets.

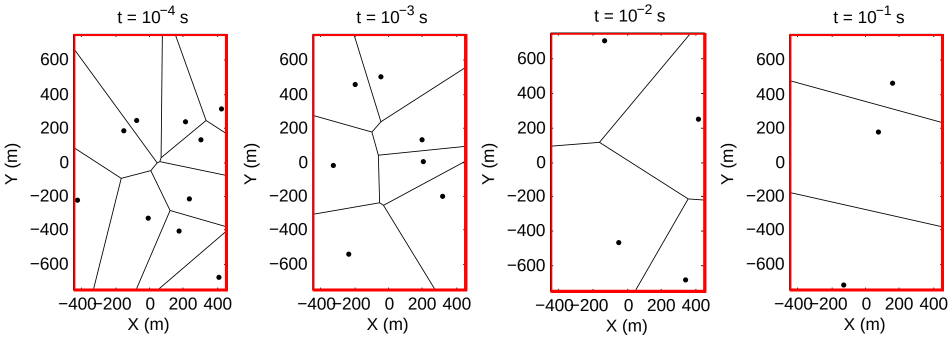

Figure 5.7: Locations of looplets and the associated Voronoi tessellations of the transmitter loop for the four delay times at the example receiver.

The local mesh and model from one of the 10 looplet subproblems used for the modeling at 10^{-4} s are shown in Figure 5.8 (coordinates in local system after rotation and translation). This mesh serves one looplet-receiver pair marked by the red and yellow dots and has only 12000 (30×20×20) cells, considerably fewer than the cells in the global mesh. Since only one time channel is modeled, there are only 15 time steps at a constant step length from 0 s to the time of measurement needed. Modeling on such a mesh takes 2 s on the same computer and uses one

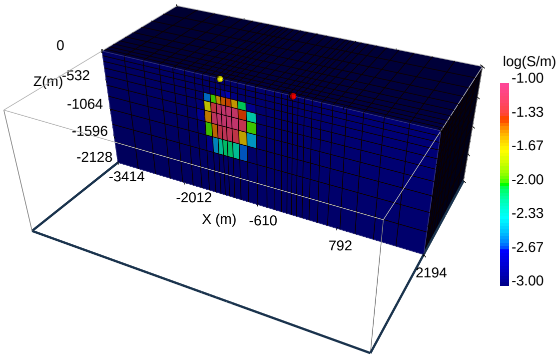

Figure 5.8: Local mesh and model used in one of the looplet subproblems modeling at 10^{-4} s.

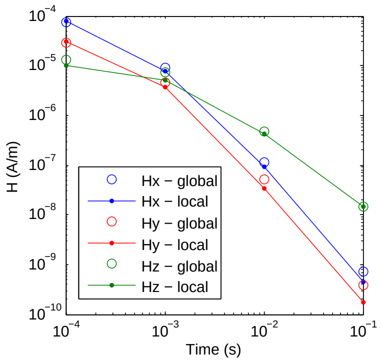

Figure 5.9: Forward modeled data using the global discretization and local discretization (adaptive looplets).

core only and the memory usage is hardly noticeable. The forward modeling CPU time of other local meshes/models ranges from about less than 1 s to 5 s depending on the looplet-receiver offset and delay time. The total CPU time required by modeling the entire survey will depend on the number of processors available and this is where massive parallelization can play a role.

The modeled H-field data using the global discretization and local discretization (looplets) are compared in Figure 5.9. Good overall agreement is observed, and the strong components always have the best accuracy.

5.2.4 Sensitivity using looplets

Once the looplets are chosen, the weights are calculated and the associated local mesh and model are made. The sensitivities can then be computed just as done in the airborne case. The sensitivities can be stored on the local meshes of the looplet subproblems, if an explicit sensitivity is to be used. The sensitivities for an R-T pair is then obtained by summing up looplet sensitivities in a way similar to the forward modeling. The summing matrix **S** in the previous subsection (equation 5.1 or 5.2) is also valid for sensitivity.

In order to show that the survey decomposition, after decomposing the transmitter into looplets, is capable of computing the correct sensitivity using the local discretization, the sensitivity of a particular datum measured at a receiver location (-900, 500) at 10^{-3} s for H_z is computed using equation 3.7. This is presented as a 3D model and compared to the results obtained directly from the global discretization modeling in Figure 5.10. The sensitivities on the global mesh are computed without decomposing the transmitter. The sensitivity computed using decomposition is very similar to the global mesh result. The positive and negative sensitivities caused by the conductive sphere and transmitter/receiver orientations are seen on both cross sections at the same locations. The difference at the highest positive sensitivity is noticeable; and this is likely due to the coarser local meshes used in the subproblems. For sensitivity, the correctness of the sign and having the correct character is important in the inversion, so the sensitivity in Figure 5.10b should be effective in the inversion.

5.3 Adaptive Subsampling of Receivers

After an R-T pair can be efficiently modeled, there is another oversampling problem: the large numbers of receivers and delay times in a ground survey yield many R-T pairs, most of which are probably redundant. Therefore, special treatment is desirable to reduce the number of R-T pairs from a field survey, exactly like reducing the number of the soundings in airborne case.

In addition to adapting the number of data to the scale of inversion, like the strategy used in the airborne case (Algorithm 1), I also want to sample unevenly in time to take into account the large contrast of time-scales in a ground survey,

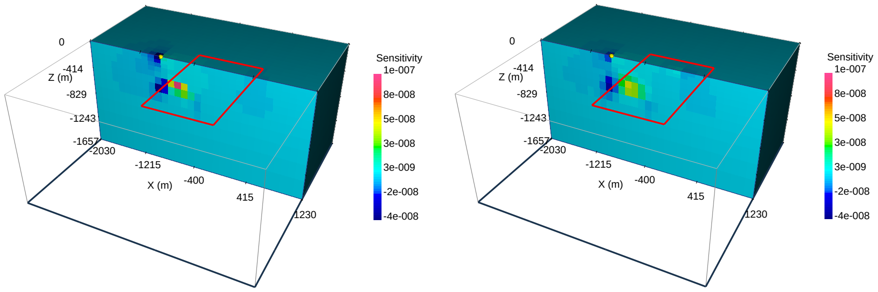

(a) Global discretization result

(b) Local discretization result

Figure 5.10: Example sensitivity computed using the global discretization and local discretization (adaptive looplets). The transmitter loop is indicated by the red lines and the receiver location by the yellow dot.

as the late delay times, with smoother EM fields, may require fewer samples than the early times. Simultaneous subsampling with a non-uniform distribution in the space spanned by the receiver locations and times is complicated and there are likely many viable strategies. I choose the following. First I restrict the sampling to some select time channels that are logarithmically spaced in time; then I do uniform random sampling of receivers within each time channel. The data subsampling problem now reduces to find the optimal numbers of receivers at different time channels at a particular iteration of inversion.

The same adaptive approach used in the airborne survey (Algorithm 1) can also be adopted here, because the optimization part is identical. However, some modifications have to be made to allow time-channel-wise cross validation and receiver refinement. In the adaptive receiver algorithm designed for the ground survey (Algorithm 2), there are still the training and test subsets, but every delay time has its own number of samples (receivers) N_i (for the ith time), its own subset $\mathbf{S}_i$ and its own normalized data misfit ϕ_i. N_i is the same in the training and test subsets. Each adaptive receiver iteration computes the following: (1) a model update using the forward responses and sensitivities from the R-T pairs in the training subset; (2) the data misfits of the pre-update model from the R-T pairs in the test subset; (3) the data misfits of the post-update model from the R-T pairs in the same test subset.

Then the two misfits from the same test subset, taken before and after the model update, are compared by delay time to decide whether the model update can be accepted. The misfit check at a particular delay time is acceptable if at least one of the following two statements is true: (1) the post-update misfit is below a prescribed tolerance; (2) the post-update misfit has made reasonable improvement from the pre-update misfit. This check is done for all delay times. There are three possible outcomes:

- **End of the inversion**. This happens when the pre-update misfit at every delay time is below a prescribed tolerance.
- **Accept the model update**. This happens when all delay times pass the check. In this situation the model update is accepted and the number of receivers at each time remains unchanged. The trade-off parameter is reduced for the next iteration.
- **Reject the model update**. This happens when one or more delay times cannot pass the check. The delay times that fail the check will increase N_i for both the training and test subsets, while other delay times do not need to change anything. Then the inversion will start over again without changing to the model and trade-off parameter.

Gradually, this adaptive scheme will add more receivers across all delay times, but emphasis is given to the early times or the times of high information density. The notations in Algorithm 2 are the same as in Algorithm 1: N_t is the number of delay times; N_i is the number of receivers at the ith delay time. I also use the idea of recycling the computing results of the test subset from the previous iteration for the training subset at the next iteration, so the computations done for the test subset have additional use. It is important to note that there are two levels of dynamic selection: (1) the random choice of R-T pairs changes at every iteration in Algorithm 2; (2) the numbers and locations of looplets within the changed R-T pairs are also dynamically sought at every iteration. This mechanism ensures the inversion model is not biased by any fixed parameter. The effectiveness of this adaptive receiver method is shown by a synthetic inversion in the next section.

Algorithm 2 Inversion with adaptive subsampling of receivers in ground TEM

```
Initialization:
  Select $N_i$ receivers at $i$th time channel for training subset $\mathbf{S}_i^{train}$ $(i = 1, \ldots, N_t)$
  Compute forward responses $F(\mathbf{m})_i^{train}$ at $\mathbf{S}_i^{train}$ $(i = 1, \ldots, N_t)$
  Compute sensitivity $\mathbf{J}(\mathbf{m})_i^{train}$ at $\mathbf{S}_i^{train}$ $(i = 1, \ldots, N_t)$
  Empty test subset $\mathbf{S}_i^{test}$ $(i = 1, \ldots, N_t)$
repeat
  Solve for $\delta\mathbf{m}$ using $\beta, \mathbf{m}, F(\mathbf{m})_i^{train}$ and $\mathbf{J}(\mathbf{m})_i^{train}$ at $\mathbf{S}_i^{train}$ $(i = 1, \ldots, N_t)$
  for $i = 1, \ldots, N_t$ do
    if $\mathbf{S}_i^{test}$ is empty then
      Select $N_i$ receivers at $i$th delay time for test subset $\mathbf{S}_i^{test}$
      Compute forward responses $F(\mathbf{m})_i^{test}$ at $\mathbf{S}_i^{test}$
      Compute data misfit $\phi(\mathbf{m})_i^{test}$ at $\mathbf{S}_i^{test}$
    end if
  end for
  Compute forward responses $F(\mathbf{m} + \delta\mathbf{m})_i^{test}$ at $\mathbf{S}_i^{test}$ $(i = 1, \ldots, N_t)$
  Compute data misfit $\phi(\mathbf{m} + \delta\mathbf{m})_i^{test}$ at $\mathbf{S}_i^{test}$ $(i = 1, \ldots, N_t)$

  if $\phi(\mathbf{m} + \delta\mathbf{m})_i^{test} \leq tol$ or $\phi(\mathbf{m} + \delta\mathbf{m})_i^{test} < \mu \cdot \phi(\mathbf{m})_i^{test}$ at $\mathbf{S}_i^{test}$ $(i = 1, \ldots, N_t)$
then
    Compute sensitivity $\mathbf{J}(\mathbf{m} + \delta\mathbf{m})_i^{test}$ at $\mathbf{S}_i^{test}$ $(i = 1, \ldots, N_t)$
    $\mathbf{S}_i^{train} \leftarrow \mathbf{S}_i^{test}$ $(i = 1, \ldots, N_t)$
    $F(\mathbf{m})_i^{train} \leftarrow F(\mathbf{m})_i^{test}$ $(i = 1, \ldots, N_t)$
    $\mathbf{J}(\mathbf{m})_i^{train} \leftarrow \mathbf{J}(\mathbf{m})_i^{test}$ $(i = 1, \ldots, N_t)$
    Empty $\mathbf{S}_i^{test}$, $F(\mathbf{m})_i^{test}$, $\mathbf{J}(\mathbf{m})_i^{test}$ $(i = 1, \ldots, N_t)$
    Reduce $\beta$
    $\mathbf{m} \leftarrow \mathbf{m} + \delta\mathbf{m}$
  else
    for $i = 1, \ldots, N_t$ do
      if $\phi(\mathbf{m} + \delta\mathbf{m})_i^{test} > tol$ and $\phi(\mathbf{m} + \delta\mathbf{m})_i^{test} \geq \mu \cdot \phi(\mathbf{m})_i^{test}$ at $\mathbf{S}_i^{test}$ then
        Compute sensitivity $\mathbf{J}(\mathbf{m})_i^{test}$ at $\mathbf{S}_i^{test}$
        $N_i = N_i \times 2$
        $\mathbf{S}_i^{train} \leftarrow \mathbf{S}_i^{train} \cup \mathbf{S}_i^{test}$
        $F(\mathbf{m})_i^{train} \leftarrow F(\mathbf{m})_i^{train} \cup F(\mathbf{m})_i^{test}$
        $\mathbf{J}(\mathbf{m})_i^{train} \leftarrow \mathbf{J}(\mathbf{m})_i^{train} \cup \mathbf{J}(\mathbf{m})_i^{test}$
        Empty $\mathbf{S}_i^{test}$, $F(\mathbf{m})_i^{test}$, $\mathbf{J}(\mathbf{m})_i^{test}$
      end if
    end for
  end if
until $\phi(\mathbf{m})_i^{test} \leq tol$ $(i = 1, \ldots, N_t)$
```

5.4 Synthetic Inversion of a Sphere-in-halfspace Model

This section tests the inversion under the framework of survey decomposition using adaptive looplets and adaptive receivers with the synthetic sphere model and the "East Loop" layout previously described in Subsection 5.2.1 (Figure 5.14a). The synthetic data set is created by modeling the soundings along the two perpendicular receiver lines at seven delay times from 10^{-4} to 10^{-1} s. The global mesh used here is the same as the one used in the adaptive looplets modeling test (Subsection 5.2.3) and the global time discretization has five different step lengths and 76 steps in total. The data are noise-free but 5% of the total field strength is assigned as the uncertainty.

5.4.1 Global discretization inversion

This synthetic data set is first inverted using the standard algorithm with the global discretizations and an initial/reference model of 0.001 S/m half-space. The global discretization inversion takes about 1 hour for the entire inversion: 40 s per factorization, 25 s per complete time stepping and about 13 minutes per model update in parallel mode on 12 cores of CPU. The inversion has to constantly use about 64 GB of memory. The large memory consumption is due to the fact that 5 factorizations, required by the time stepping from 0 to 10^{-1} s, and the simulated fields on the entire mesh at all 76 time steps are stored in the memory. The trade-off parameter β, total normalized data misfit and model norm as functions of iteration in the inversion are plotted in Figure 5.11. The recovered model resembles the true model with a conductive body in a resistive host (see cross section in Figure 5.14b).

5.4.2 Survey decomposition inversion

The same data set is then inverted using survey decomposition with adaptive looplets and adaptive receivers. All seven delay times are used and the inversion parameters are the same as those in the global discretization inversion. Table 5.1 records the number of receivers used at different delay times and iterations. For a particular trade-off parameter β, the adaptive approach may take a few iterations to add more receivers until satisfying the cross validation check. Interestingly, the last two time channels have not seen any addition of receivers indicating there is very

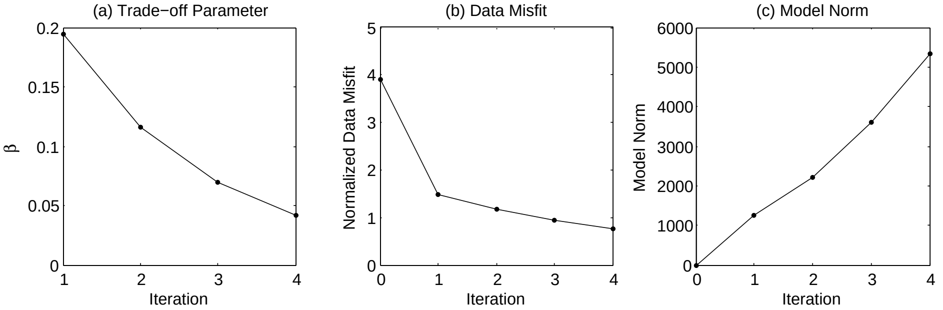

Figure 5.11: Summary of the synthetic inversion using the standard algorithm with global discretizations.

Table 5.1: Synthetic inversion using adaptive looplets and adaptive receivers.

Iteration	Number of receivers (time channel 1∼7)							CPU time (s)	Result
1	2	2	2	2	2	2	2	182	add
1	4	2	2	2	2	2	2	83	add
1	8	4	2	2	2	2	2	133	pass
2	8	4	2	2	2	2	2	205	add
2	8	4	4	4	4	2	2	158	pass
3	8	4	4	4	4	2	2	284	add
3	16	4	8	8	4	2	2	266	pass
4	16	4	8	8	4	2	2	370	add
4	16	8	16	16	8	2	2	481	add
4	32	16	16	16	8	2	2	444	pass

little information in the data at that late times; on the other hand, as the constraint on the model norm relaxes, more receivers at early times are demanded to construct fine-scale structures in the model. The randomly sampled R-T pairs used in the test subset at the last iteration are plotted in Figure 5.12. This synthetic example demonstrates that the adaptive receivers approach, together with the adaptive looplets, is able to substantially reduce the number of subproblems needed in an inversion. The trade-off parameter β, total normalized data misfit as estimated by the test subset, and the model norm at every iteration are provided in Figure 5.13. The recovered model from this survey decomposition inversion is shown in Figure 5.14c.

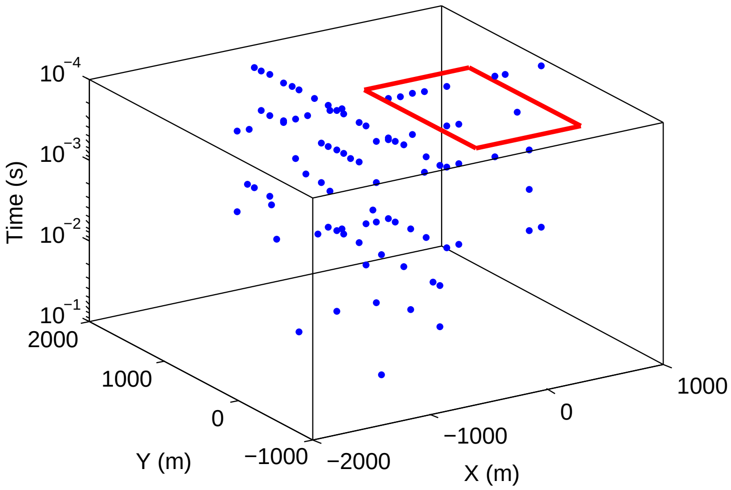

Figure 5.12: Location of random samples of receiver-time (R-T) pair used at the last iteration of the synthetic inversion. The dimension of time is depicted as the depth and the transmitter loop is indicated by the red lines.

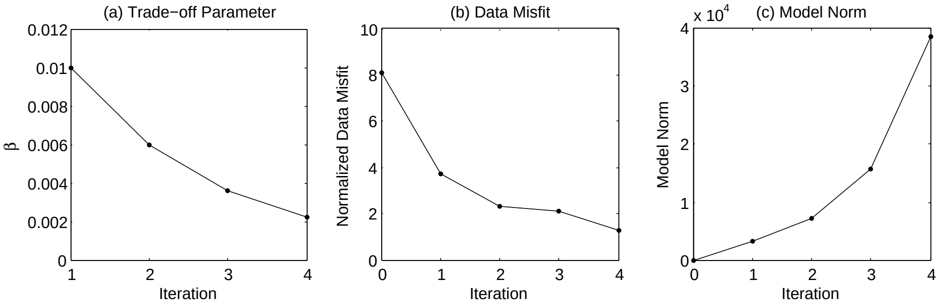

Figure 5.13: Summary of the synthetic inversion using the survey decomposition with local discretizations.

The total CPU time of survey decomposition, about 43 minutes on the same computer (Table 5.1), is only marginally less than the time (1 hour) required using the global discretization inversion. However, there are strong reasons for using survey decomposition even though this synthetic example does not show an astonishing speed-up:

- The two synthetic inversions use a different number of CG iterations when solving equation 3.1. The standard algorithm has maximum 10 CG iterations for one Gauss-Newton step, whereas the survey decomposition approach has 60 CG iterations. If the standard algorithm also uses 60 CG iterations, its total CPU time would be about six times longer (6 hours). This clearly shows the computational advantages of using the survey decomposition framework and working with explicit sensitivity on the local meshes.
- The global mesh of this example, having 169136 cells and only one transmitter, does not have a space complexity as large as that used in the airborne examples. If the mesh becomes larger, the poor scalability will slow down the modeling significantly as shown in Figure 3.1. But the survey decomposition scales much better to the number of subproblems since it does not do any forward or sensitivity modeling directly on the global mesh.
- Adding more processors does not help to speed up the global discretization inversion and sometimes it even slows down the inversion because of excessive communication overhead. In contrast, survey decomposition is able to benefit from a large number of processors. At the last iteration in Table 5.1, there are 418 looplet subproblems in 92 R-T pairs, which takes 208 s for a complete forward modeling using 12 cores. If every subproblem has its own processor, the modeling time could be reduced to a few seconds.
- The memory consumption of survey decomposition, which is about 5 GB for the storage of the sensitivity and other matrices at the last iteration, is very low compared to that needed for the global mesh inversion. This results because every subproblem only has 15 steps at a constant step length and most importantly after the forward response and sensitivity are computed, all of the dense matrices (factors and fields) can be erased from the memory,

leaving only the sensitivities stored on the local meshes. In fact, the memory required by the survey decomposition modeling is scaled by the number of concurrently running subproblems, and has nothing to do with the size of global discretizations. Such flexibility of memory allows 3D inversion to be inexpensively carried out on low-end computers.

5.4.3 Model interpretation and validation

The model recovered using survey decomposition is compared to the true model and global discretization inversion result in Figure 5.14. Both inversions recover the conductive sphere at the correct location and depth; the weakened conductivity of the sphere in the inversion models is compensated by the smeared “tails” around and below the recovered conductors, which is what we should expect when the target is compact but the regularization measures the model norm based on L2. The initial/reference model, a resistive 0.001 S/m half-space, also contributes to the underestimation of the conductivity of the sphere. The survey, which is not ideal for 3D inversion because there is only a single large transmitter loop and the receiver locations are sparse, also permits further non-uniqueness to the recovery of the model. Regardless, the consistent images from the global discretization inversion and survey decomposition inversion show that the forward responses and sensitivities are reasonably computed in the subproblems, and the entire framework works as well as the standard algorithm does for the modeling of ground loop survey.

Finally, a quantitative model validation is carried out by forward modeling the survey decomposition inversion model with the global mesh and global time discretization. The true data misfit can then be calculated without any assumption, approximation and reduction used in the survey decomposition inversion. The normalized data misfits of the data are plotted as a histogram in Figure 5.15. For each datum, the normalized misfit is the absolute difference between the observed and predicted data divided by the uncertainty, which is 5% of the observed total field strength. For the 1134 data (54 receivers $\times$ 7 time channels $\times$ 3 components) used here, there were 780 data fit to less than the prescribed uncertainty. The overall normalized data misfit is 1.08, a value close to the misfit estimated by the test sub-

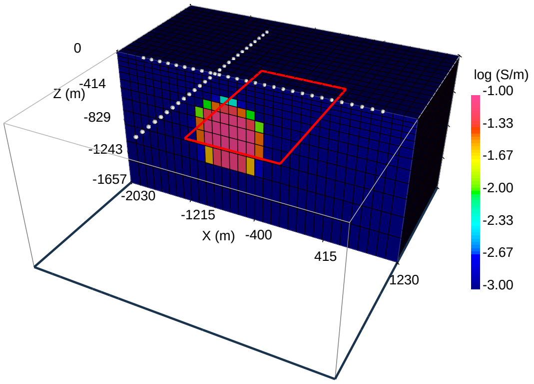

(a) True Model

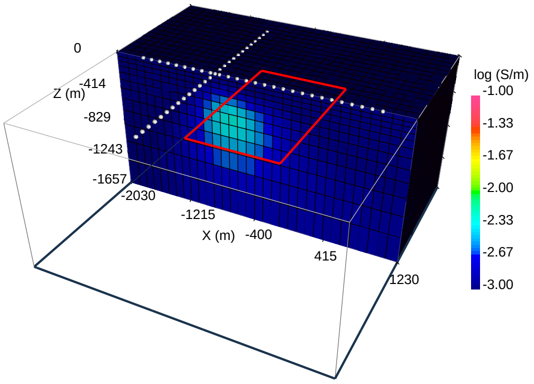

(b) Global mesh inversion

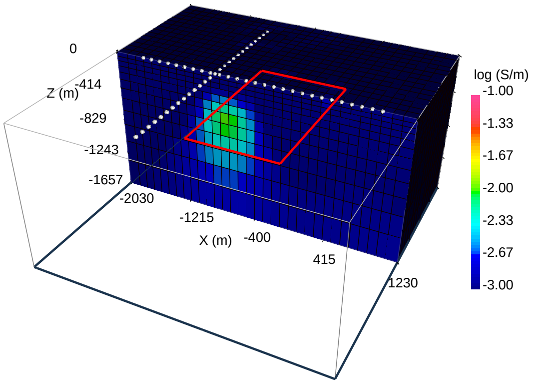

(c) Survey decomposition inversion

Figure 5.14: Synthetic inversion models for the ground loop TEM survey.

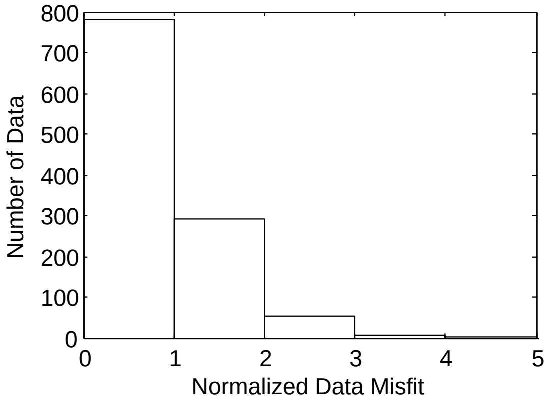

Figure 5.15: True data misfit of the final survey decomposition inversion model reassessed using the global discretizations.

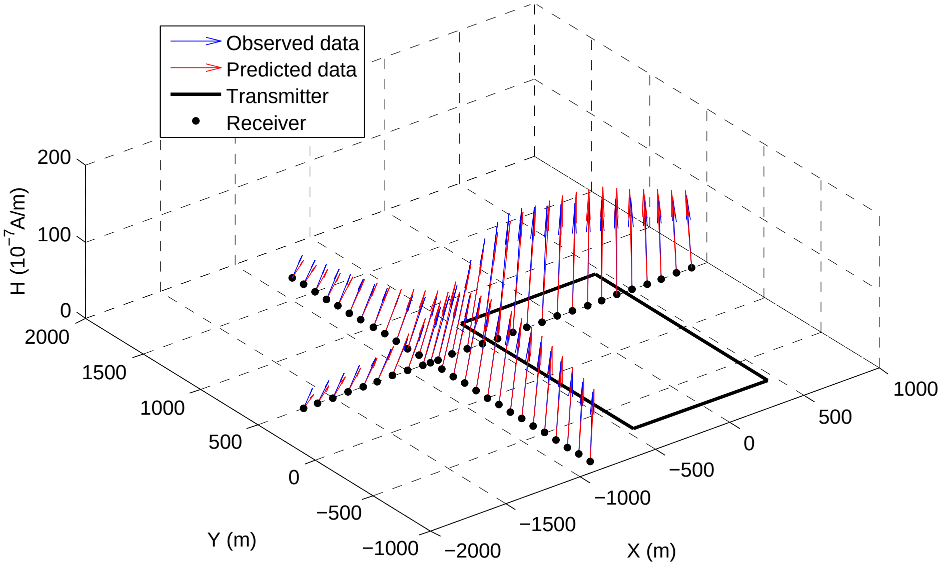

Figure 5.16: Data fit of the final survey decomposition inversion model (showing delay time at 10^{-3} s) for the synthetic data.

set (Figure 5.13b) and to that achieved in the global inversion. The predicted data of the final survey decomposition inversion model computed in the validation are plotted as 3D H-field vectors measured on the surface along with the observed data vectors in Figure 5.16. The distortion of the magnetic field, caused by the conductor near the cross point of the two receiver lines, is recognized and reproduced by the survey decomposition inversion.

5.5 Inversion of SQUID Data at the Lalor Mine

In this final section, I test the survey decomposition inversion with the SQUID field data at the Lalor Mine previously mentioned in Section 5.1 and compare it with the result obtained from the standard method. The data set is assigned an uncertainty of 10% of total field strength plus a noise floor of 10^{-7} A/m.

5.5.1 Global discretization inversion

The global mesh used here is geometrically identical with the one used in the synthetic study in Section 5.4, but in the realistic coordinate system. The global time discretization models 29 time channels from 1.47 to 71 ms using four different time step lengths and 74 time steps in total. In terms of the numbers of cells and time steps, the computational complexities of this inversion are almost the same as those of the synthetic inversion in the previous section, so the time and memory required by the factorization and time stepping for one model update (Gauss-Newton step) are similar (13 minutes).

Starting from a 0.001 S/m half-space model, the inversion ran for 12 Gauss-Newton steps. However, the target misfit (normalized data misfit = 1) is achieved after about 6 model updates in about 80 minutes. The trade-off parameter β, total normalized data misfit and model norm as a function of iteration in the inversion are plotted in Figure 5.17. The recovered model from Iteration 6 is shown on two cross sections in Figure 5.19a and 5.19b.

5.5.2 Survey decomposition inversion

When doing the survey decomposition inversion, the subsampling of receiver is restricted to 8 delay times from 1.47 to 371 ms. As a result, there are 432 R-T pairs

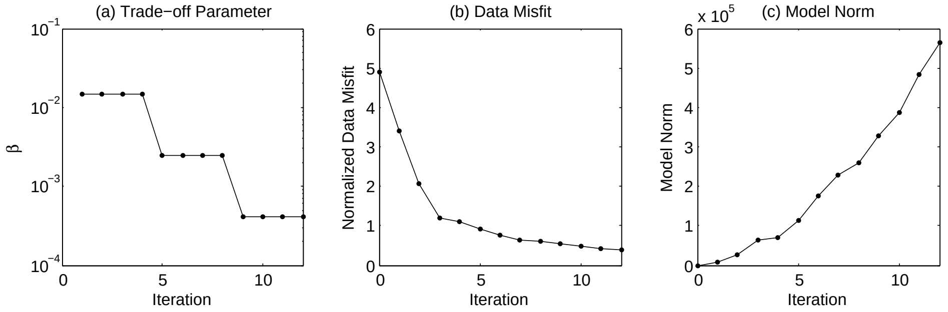

Figure 5.17: Summary of the Lalor Mine field data inversion using the standard algorithm with global discretizations.

(54 receivers × 8 delay times) available in the database for selection. Table 5.2 records the addition of receivers at different delay times as the inversion proceeds. Similar to the synthetic example, the field inversion also finds that the most R-T pairs are needed at the early times and at those times encoded with rich information about the anomaly. Throughout the inversion, a total number of 402 R-T pairs have been used in the inversion, with no more than 44 in any single iteration. The initial/reference model is a 0.001 S/m half-space.

The target data misfit is achieved after 11 Gauss-Newton steps in 42 minutes, about two times faster than the global discretization inversion, and can be further speeded up with massive parallelization. About 2.8 GB memory is required to store the sensitivities and other matrices at the last iteration.

The trade-off parameter β, total normalized data misfit estimated by the test subset and model norm at every iteration are provided in Figure 5.18. I note that the model norm of Iteration 11 in this survey decomposition inversion is close to the model norm of Iteration 6 in the global discretization inversion. The recovered model (Iteration 11) of the survey decomposition inversion is shown on two cross sections in Figure 5.19c and 5.19d.

More insights can be gained by further comparing the synthetic and field data inversions. The two data sets, even though they are collected using the same survey layout, are actually in different scales: the synthetic model has a shallower target conductor and earlier delay times, while the field example has a deeper target and uses later times. It is difficult for the global discretization inversion, working on

Table 5.2: Field data inversion using adaptive looplets and adaptive receivers.

Iteration	Number of receivers (time channel 1$\sim$8)								CPU time (s)	Result
1	2	2	2	2	2	2	2	2	207	pass
2	2	2	2	2	2	2	2	2	166	pass
3	2	2	2	2	2	2	2	2	172	pass
4	2	2	2	2	2	2	2	2	159	add
4	4	2	2	2	2	2	2	2	100	pass
5	4	2	2	2	2	2	2	2	159	pass
6	4	2	2	2	2	2	2	2	173	add
6	8	2	2	2	2	2	2	2	147	pass
7	8	2	2	2	2	2	2	2	222	pass
8	8	2	2	2	2	2	2	2	204	pass
9	8	2	2	2	2	2	2	2	206	add
9	8	4	2	2	2	2	2	2	134	pass
10	8	4	2	2	2	2	2	2	234	pass
11	8	4	2	2	2	2	2	2	231	pass
12	8	4	2	2	2	2	2	2	246	pass
13	8	4	2	2	2	2	2	2	236	add
13	8	4	4	4	4	4	2	2	224	add
13	16	8	4	4	4	4	2	2	255	pass

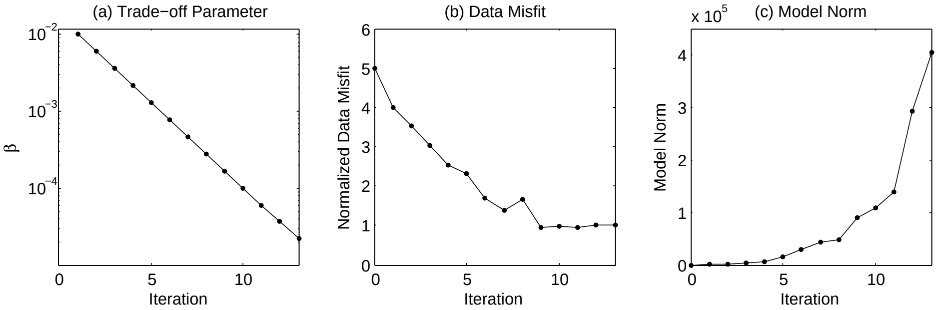

Figure 5.18: Summary of the Lalor Mine field data inversion using the survey decomposition with local discretizations.

the scales set by the global mesh and global time discretization, to adapt to the appropriate scales of investigation. However, the new approach is able to recognize the difference and only uses 194 looplets in 44 R-T pairs in the last iteration of the field data inversion. This is about half the number used in the synthetic inversion and it is a clear indication about the over-computing carried out in the global discretization inversion.

5.5.3 Model interpretation and validation

The final conductivity models, recovered using the standard algorithm (Iteration 6 of the global discretization inversion) and the survey decomposition (Iteration 11 of the survey decomposition inversion), are compared on two perpendicular cross sections in x-direction and y-direction in Figure 5.19. Both inversions provide very similar images of the deep major conductor, in terms of location, depth, size and conductivity. They also have similar model norms if measured quantitatively.

There are noticeable differences of the small-scale structures near the surface. This is likely because the near-surface features are minor anomaly compared to the large conductor at depth. The survey decomposition inversion captures the primary anomaly first by using fewer samples, and gradually builds more small-scale features as the adaptive receiver procedure adds more receivers to the early delay times. The model of Iteration 13 in the survey decomposition inversion has the near-surface structures similar to the global discretization inversion result. This shows that the quality of survey decomposition is not compromised after significant computational savings.

Although the recovered conductor is a good candidate for VMS mineralization, further investigation involving other data sets and other geological information is required to verify the inversion model for the geological interpretation, which is out of the scope of this thesis. Therefore, the emphasis of model validation is given to the similarity between the global discretization inversion and survey decomposition inversion.

Finally, I again validate the survey decomposition inversion model using a forward modeling with the global discretizations. Note, although the survey decomposition is restricted to sample at 8 out of 29 available delay times in the inversion,

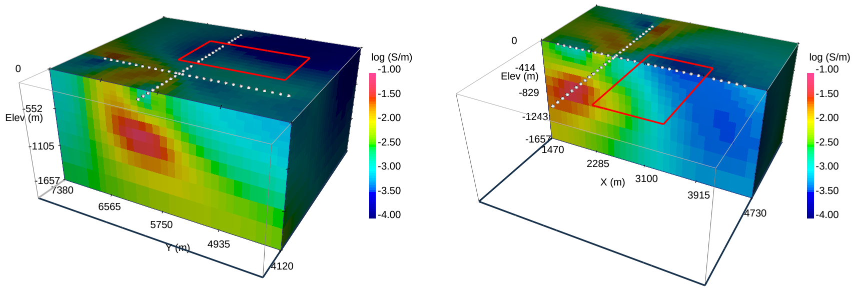

(a) Global discretization inversion (y-direction cross section)

(b) Global discretization inversion (x-direction cross section)

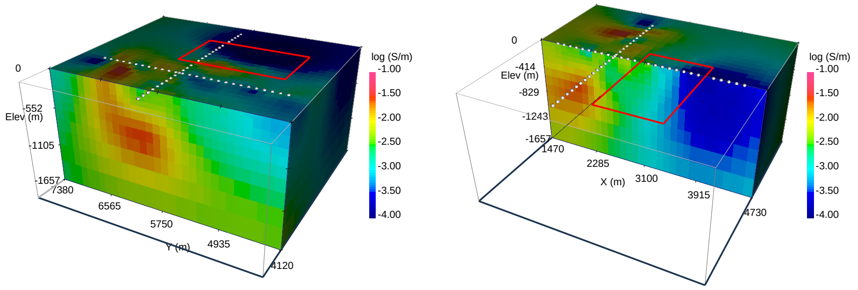

(c) Survey decomposition inversion (y-direction cross section)

(d) Survey decomposition inversion (x-direction cross section)

Figure 5.19: Inversion models for the SQUID field data at the Lalor Mine VMS deposit. The two models are cut along x-direction and y-direction for cross sections.

the validation uses all 29 time channels to calculate the data misfit. The normalized data misfits of the predicted data are plotted as a histogram in Figure 5.20; the overall normalized data misfit is 0.86 for the final model.

The field SQUID data and the predicted data of the final survey decomposition inversion model computed in the validation are plotted as 3D H-field vectors measured on the surface to show that the most important anomalies due to the major conductor in the data are reasonably reproduced by the survey decomposition inversion (Figure 5.21).

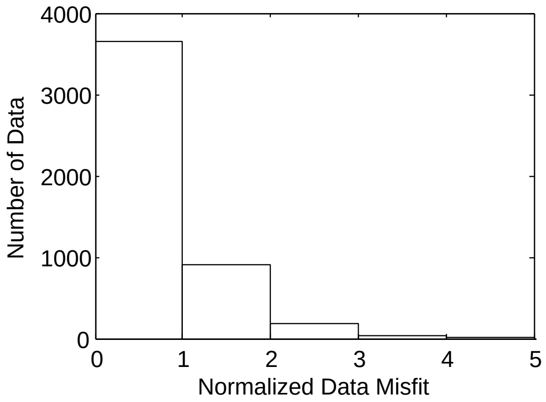

Figure 5.20: True data misfits of the final survey decomposition inversion model at the Lalor Mine reassessed using the global discretizations.

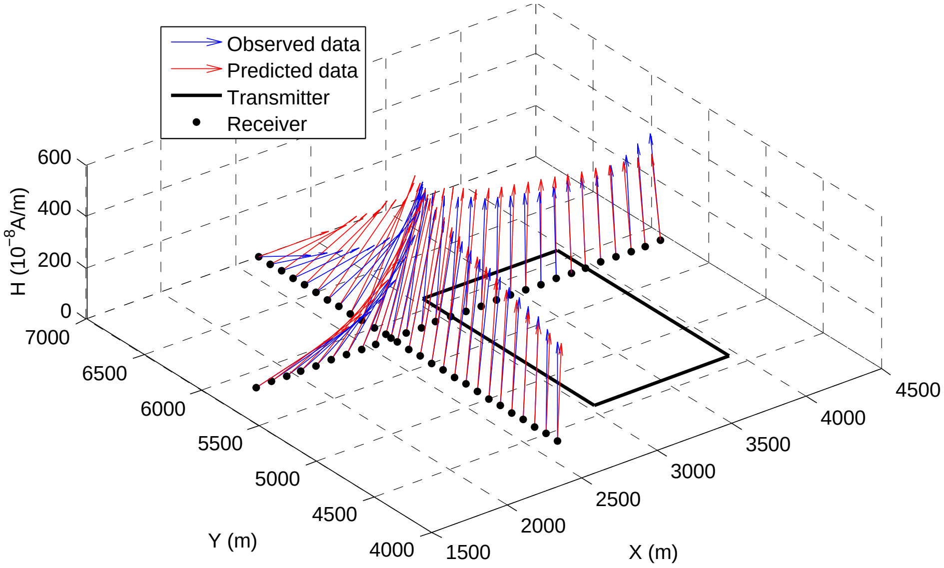

Figure 5.21: Data fit of the final survey decomposition inversion model (showing delay time at 3.78×10^{-3} s) for the Lalor Mine SQUID data.

5.6 Summary

This chapter implements the framework of survey decomposition in ground loop TEM data inversion. A ground large loop survey has two important computational complexities: the large transmitter and a large contrast of modeling scales between early and late delay times. Survey decomposition reduces these computational complexities by modeling many subproblems, each of which is in fact an atomic problem that only concerns a dipole source (looplet), a receiver and a particular time.

The implementation consists of two parts. The first part is to compute the data or sensitivity at a particular receiver and delay time (R-T), given a large transmitter loop. This is solved by a procedure called adaptive looplets that finds a minimum collection of randomly located looplets within the transmitter loop, and the principle of superposition, to effectively represent the original loop. The number of looplets needed adapts to the scale of EM fields, so later times may need fewer looplets than early times. A new round of search for the looplets is initiated when the selection of R-T pair is changed in the second part.

The second part, once an R-T can be modeled, is to choose a minimum subset of all possible R-T for the inversion since the data set is always redundant. I propose another procedure called adaptive receivers to find the minimum collection of receivers at each delay time by using a time-by-time cross validation. This procedure also has the ability of adapting the number of receivers to the scale of investigation, both in time and in the regularization of inversion.

Together, the adaptive looplets and adaptive receivers can be used to ensure that only the very necessary subproblems are modeled. Both procedures are tested with a synthetic conductive sphere model. I have shown the forward responses and sensitivities can be reasonably computed in the subproblems on local discretizations at much reduced costs. A synthetic data set based on the sphere model is inverted using survey decomposition with adaptive looplets and adaptive receivers, yielding an inversion model comparable with the global discretization inversion results. The new methods show that savings in time and memory can be realized and further speed-up by massive parallelization is very promising.

Finally a field SQUID data set at the Lalor Mine VMS deposit is inverted using the newly-developed framework and procedures. The survey decomposition inversion obtains results similar to that from the standard inversion, but in a much shorter time even using a small-scale parallelization. The adaptive procedures have shown the ability of adapting the number of subproblems to the scale of the EM problem.

In both the synthetic and field data inversions, the final survey decomposition inversion models are validated by forward modeling using the standard algorithm with the global discretizations. The true data misfits of the final models, calculated without any assumption, approximation and reduction used in the new approach, have been shown to be reduced to an acceptable level.

Chapter 6

Conclusions

It is widely accepted that inverting electromagnetic data acquired from geophysical surveys will yield a distribution of electrical conductivity that will be useful in geoscience problems. Despite a handful of modeling algorithms developed in the laboratory and greatly improved computers, practitioners are still facing prohibitive challenges in using rigorous 3D modeling as an integrated tool in their exploration decision makings. Much of the problem is traced to the computational costs of 3D modeling and can be summarized in two critical questions: (1) Is 3D modeling/inversion really necessary? (2) If yes, what can we do to make it more efficient? This thesis attempts to answer these questions and the findings are summarized in the following sections.

6.1 Conditions of Using 3D Inversion

In the first part of this thesis, I explore in considerable detail a specific example of why, and under what circumstance a 3D interpretation is needed.

With a field example of airborne TEM survey (VTEM) from Mt. Milligan, I found there was a significant discrepancy between the model obtained by 1D interpretation and 3D geology. This contradiction was studied by a synthetic model that emulates the Mt. Milligan deposit: a resistive rock stock surrounded by a conductive halo. The 1D inversion of the synthetic data showed a misleading conductive anomaly at the location of the stock. This is explained by the fact that an airborne sounding has a stronger sensitivity away from the sounding location than beneath it, so a sounding above the resistive stock can have high response because of the conductive halo around it. However, a 3D inversion, capable of modeling 3D sensitivity and constructing a 3D model, correctly recovered the true model and the geometry of the resistive stock and the conductive halo were reasonably delineated. Then a small region of the field data around the main stock was inverted in 3D, producing a much more geologically plausible model. A multi-level strategy involving multi-resolution mesh and data grid was developed to speed up the 3D inversion.

Following the success of 3D inversion of airborne TEM data, I revisited a frequency-domain airborne EM data set (DIGHEM) at Mt. Milligan. It posed a paradox. A 1D inversion of the DIGHEM data had previously recovered the resistive nature of the stock, which seemingly discredited the necessity of 3D inversion. A critical analysis using a semi-quantitative measure of the footprint of an airborne system revealed that VTEM had a much larger footprint than DIGHEM, and the scale of the conductivity variation at Mt. Milligan just happened to be comparable to the VTEM footprint and larger than the DIGHEM footprint. This discovery leads to the condition of using 3D inversion: when the scale of the target is larger than the scale of the survey, a non-3D modeling could be justified; but if the scale of the target is comparable to the scale of the survey, which is common in practice, a 3D modeling must be used.

6.2 Over-computing in 3D EM Modeling

The technical route I choose to improve the speed of 3D modeling is somewhat unconventional. After reviewing the existing technology, I categorize the computational difficulties into three complexities:

- Space complexity. This is primarily caused by the large size of the modeling domain and the large number of sources. A single forward modeling and inversion mesh covering a large area results in a large system to be solved for each source.
- Time complexity. This becomes a concern when times/frequencies with a large contrast of scale are modeled with a single discretization because: (1)

The mesh must have a large number of cells to model both the fine-scale features for the early times and the large-scale background for the late times; (2) If a time stepping scheme is used, modeling all time channels with one global time discretization can cause over-computing.

- Optimization complexity. An inversion requires many iterations of forward modelings and solutions of an optimization problem on a 3D mesh. Over-computing is especially prevalent at the early stage of inversion, when only large-scale information is needed but the computations micromanage the details at an unnecessarily fine level.

All three computational complexities share the same fundamental issue that can be found in the conventional methods: they attempt to work simultaneously or indiscriminately with problems involving multiple scales. This leads to large matrix systems to be solved for the forward problem and results in a slower and less flexible algorithm, and also unnecessarily models the redundant field data set from a survey.

6.3 Survey Decomposition: A New Path

The difficulties associated with over-computing can be overcome by adopting a survey decomposition framework. This new framework isolates the source, receiver and time/frequency in a survey so that the decomposed subproblems have more efficient localized discretizations: (1) local meshes in space, and (2) local time discretizations. A local mesh is much smaller than the global mesh and is only refined around source/receiver locations. A local time discretization only serves one time channel, so it only needs one constant step length in an implicit time stepping approach.

If a transmitter or receiver is spatially large, then the modeling can be carried out at an even finer level involving a dipole source and a dipole receiver working with single time or frequency and concerned with a particular component of the observed field. I refer to this as an atomic building block problem. Any CSEM survey can be decomposed into these atomic problems. One or more atomic problems can be combined into a subproblem that is treated as a self-contained computing unit in a parallel environment. In my implementation, an airborne survey is straightforwardly decomposed to many sounding subproblems, so every airborne subproblem is assembled by the atomic building blocks sharing the same source and receiver; all delay times from a particular sounding are modeled together in one subproblem. A ground survey, assuming one large transmitter loop, is first decomposed to many receiver-time (R-T) pairs; then within each R-T, the transmitter is further decomposed to many dipole looplet sources using superposition, yielding many looplet-receiver-time units (i.e. atomic building blocks) as subproblems.

A general workflow of survey decomposition can be described as follows:

- **Global mesh.** A global mesh, the same as the one used in a standard method inversion, must be designed beforehand. The global mesh must be large enough to accommodate the entire survey. The number of cells in the global mesh is not important because it is only used as a container for the model and no forward modeling is ever carried out on it. Decoupling the inversion and forward modeling meshes is a critical item in the overall workflow.

- **Subproblem assignment.** The global model on the global mesh, along with the specifications of the subproblems, is sent to many parallel workers. The local mesh and time discretization are specifically made within each individual subproblem. The global model is projected onto the local meshes.

- **Computing.** The parallel workers solve the subproblems with the local discretizations. Both the forward responses and the sensitivities are computed and stored on the local meshes. The modeled data and sensitivity are made available to the global mesh through simple linear transformations, including rotation, interpolation and possibly summation if the transmitter is decomposed. During the inversion, the operation of sensitivity-times-vector can be done within the subproblems. In this way, the entire process of forward and inverse modeling is parallelizable.

- **Data subsampling.** A survey decomposition usually yields many subproblems. This is tackled by adaptively choosing a subset of the entire data set for the inversion. The adaptive subsampling, assisted by random sampling and cross validation, ensures that the number of data in a particular subproblem is appropriate for the current scales of modeling in time and space, regardless of how many data are collected in the field.

The workflows designed for airborne and ground surveys are both tested using synthetic and field data. I have shown that the new framework of survey decomposition is able to produce satisfactory results with significantly reduced time and memory usage. Further improvement of efficiency via massive parallelization is easily achievable since the subproblems are both fine-grained and mutually independent.

6.4 Further Development

In this thesis, I have developed a novel framework to improve the efficiency of 3D EM modeling and shown that it is viable. There are certain items in my procedures that can be improved in the future. There are also other research opportunities based upon my work.

6.4.1 Possible improvement

This thesis lays out the basic elements and workflows required by the survey decomposition. I provide some simple ways to implement some of the steps. However, these deserve more careful investigation and potentially yield better accuracy and greater efficiency.

- **Material averaging.** The linear operation of converting the model from the global mesh to the local mesh works well so far, but it can be made to better deal with more complicated situations, for example, high conductivity contrast.
- **Un-structured or semi-structured mesh.** This thesis uses regular mesh as the example to show the effectiveness of survey decomposition. The overall performance may be further improved if other types of mesh, like tetrahedral mesh, curvilinear mesh and octree mesh, are implemented.
- **Full space-time subsampling.** In my examples, I restricted the subsampling of the receivers at some preselected delay times. Eventually, the procedure

should be able to do random sampling simultaneously in space and time. The ultimate case is that a dipole source and a dipole receiver move freely in 3D space and measure data at an arbitrary time/frequency, and the data in this 7D space (S_x,S_y,S_z,R_x,R_y,R_z,T/F) are subsampled for the inversion.

- **Smarter sampler.** As the simplest strategy, I used a random sampler with a uniform probability distribution. This can be made more efficient and accurate if the sampling concentrates on the area with complex signals.
- **Frequency-domain approach.** My local meshes are designed for solving time-domain problems. If the problem is modeled in the frequency domain the mesh may be changed accordingly.
- **Galvanic source/receiver.** The concept of "linelet" for the galvanic source has been proposed, and there should be no difficulty in implementing it. However, I have thus far focused attention on the inductive sources and receivers.

6.4.2 Other research opportunities

The framework of survey decomposition based on the atomic building blocks opens up a host of new research applications.

Many ideas used in the survey decomposition, for example, "localized discretization", "adapt-to-the-scale", and "refine-until-no-change", are all generic. Exact the same approach can be applied to other types of surveys with little modification:

- **Marine CSEM**. The modeling of marine CSEM also suffers from the similar problems as the airborne survey. It usually covers a large area with many source and receiver locations. Survey decomposition can certainly speed it up.
- **Magnetotelluric (MT)**. MT is a good example of large (natural) source, large area and large contrast of scales. Because MT measures the electric field, MT modeling may require very small cells everywhere in the survey

area. This can be potentially improved by modeling MT soundings separately on local meshes.

- **DC resistivity**. A DC survey generally uses a large number of sources, either localized "dipole" arrays, or distributed "pole" arrays. Therefore, it is also a good candidate for a survey decomposition but with using linelet.

Because of the blindness of the global thread to the computational details at local level, it is possible to use different source waveforms/frequencies, different meshes, different solvers, different types of data in one inversion. There is certainly great opportunities to jointly invert multiple data sets under the framework of survey decomposition. Such a "hybrid inversion" can be highly efficient since not only the discretizations are localized, all the parameters of a subproblem (for example, the solver, PDE, type of mesh, etc.) can be tailored to the specific needs of each individual subproblem.

Finally I have to point out that my prototyping of survey decomposition was done on the desktop computer or small cluster. The speed-up can be almost linearly increased if more processors can be used. This makes it an ideal application for the emerging massive parallel computing devices, including GPUs and cloud computing. Infrastructure as a service (IaaS) in cloud computing offers virtually unlimited computing resources in a scalable manner, and thus could be a good platform to implement the framework. Once the massive parallelization of the survey decomposition is established, the paradigm shift of using 3D modeling as a routine would be just a fingertip away.

Bibliography

Abraitis, P., Pattrick, R., & Vaughan, D., 2004. Variations in the compositional, textural and electrical properties of natural pyrite: A review, *International Journal of Mineral Processing*, **74**(1), 41–59. → pages 2

Alumbaugh, D. L., Newman, G. A., Prevost, L., & Shadid, J. N., 1996. Three-dimensional wideband electromagnetic modeling on massively parallel computers, *Radio Science*, **31**(1), 1–23. → pages 17

Amestoy, P. R., Guermouche, A., L'Excellent, J. Y., & Pralet, S., 2006. Hybrid scheduling for the parallel solution of linear systems, *Parallel Computing*, **32**(2), 136–156. → pages 18, 50

Annan, A. P., Smith, R. S., Lemieux, J., O'Connell, M. D., & Pedersen, R. N., 1996. Resistive-limit, time-domain AEM apparent conductivity, *Geophysics*, **61**(1), 93–99. → pages 9

Aurenhammer, F., 1991. Voronoi diagrams - a survey of a fundamental geometric data structure, *ACM Computing Surveys (CSUR)*, **23**(3), 345–405. → pages 103

Badea, E. A., Everett, M. E., Newman, G. A., & Biro, O., 2001. Finite-element analysis of controlled-source electromagnetic induction using Coulomb-gauged potentials, *Geophysics*, **66**(3), 786–799. → pages 12

Beamish, D., 2003. Airborne EM footprints, *Geophysical Prospecting*, **51**(1), 49–60. → pages 42

Benkabbour, B., Toto, E., & Fakir, Y., 2004. Using DC resistivity method to characterize the geometry and the salinity of the Plioquaternary consolidated coastal aquifer of the Mamora plain, Morocco, *Environmental Geology*, **45**(4), 518–526. → pages 6

Bernstone, C., Dahlin, T., Ohlsson, T., & Hogland, H., 2000. DC-resistivity mapping of internal landfill structures: Two pre-excavation surveys, *Environmental Geology*, **39**(3-4), 360–371. → pages 6

Börner, R.-U., Ernst, O. G., & Spitzer, K., 2008. Fast 3-D simulation of transient electromagnetic fields by model reduction in the frequency domain using Krylov subspace projection, *Geophysical Journal International*, **173**(3), 766–780. → pages 52

Brodie, R. & Sambridge, M., 2006. A holistic approach to inversion of time-domain airborne EM, *ASEG Extended Abstracts*, **2006**, 1–4. → pages 11

Chambers, J. E., Kuras, O., Meldrum, P. I., Ogilvy, R. D., & Hollands, J., 2006. Electrical resistivity tomography applied to geologic, hydrogeologic, and engineering investigations at a former waste-disposal site, *Geophysics*, **71**(6), B231–B239. → pages 2

Chwala, A., Schultze, V., Stolz, R., Ramos, J., IJsselsteijn, R., Meyer, H.-G., & Kretzschmar, D., 2001. An HTS dc SQUID system in competition with induction coils for TEM applications, *Physica C: Superconductivity*, **354**(1), 45–48. → pages 98

Commer, M. & Newman, G., 2004. A parallel finite-difference approach for 3D transient electromagnetic modeling with galvanic sources, *Geophysics*, **69**(5), 1192–1202. → pages 12, 17, 18

Commer, M. & Newman, G. a., 2006. An accelerated time domain finite difference simulation scheme for three-dimensional transient electromagnetic modeling using geometric multigrid concepts, *Radio Science*, **41**(3), RS3007. → pages 18

Commer, M. & Newman, G. a., 2008. New advances in three-dimensional controlled-source electromagnetic inversion, *Geophysical Journal International*, **172**(2), 513–535. → pages 18, 66, 151

Commer, M., Newman, G. A., Carazzone, J. J., Dickens, T. A., Green, K. E., Wahrmund, L. A., Willen, D. E., & Shiu, J., 2008. Massively parallel electrical-conductivity imaging of hydrocarbons using the IBM Blue Gene/L supercomputer, *IBM Journal of Research and Development*, **52**(1.2), 93–103. → pages 17

Constable, S., 2010. Ten years of marine CSEM for hydrocarbon exploration, *Geophysics*, **75**(5), 75A67–75A81. → pages 2, 6

Constable, S. C., Parker, R. L., & Constable, C. G., 1987. Occam's inversion: A practical algorithm for generating smooth models from electromagnetic sounding data, *Geophysics*, **52**(3), 289–300. → pages 16

Cox, L. H., Wilson, G. a., & Zhdanov, M. S., 2010. 3D inversion of airborne electromagnetic data using a moving footprint, *Exploration Geophysics*, **41**(4), 250–259. → pages 19

Cox, L. H., Wilson, G. a., & Zhdanov, M. S., 2012. 3D inversion of airborne electromagnetic data, *Geophysics*, **77**(4), WB59–WB69. → pages 19

De Jong, E., Ballantyne, A., Cameron, D., & Read, D., 1979. Measurement of apparent electrical conductivity of soils by an electromagnetic induction probe to aid salinity surveys, *Soil Science Society of America Journal*, **43**(4), 810–812. → pages 6

Eaton, P. A., 1998. Application of an improved technique for interpreting transient electromagnetic data, *Exploration Geophysics*, **29**(2), 175–183. → pages 10

Edwards, R., 1974. The magnetometric resistivity method and its application to the mapping of a fault, *Canadian Journal of Earth Sciences*, **11**(8), 1136–1156. → pages 62

Edwards, R. N., 1997. On the resource evaluation of marine gas hydrate deposits using sea-floor transient electric dipole-dipole methods, *Geophysics*, **62**(1), 63–74. → pages 6

Farquharson, C. G., 2007. Constructing piecewise-constant models in multidimensional minimum-structure inversions, *Geophysics*, **73**(1), K1–K9. → pages 14

Farquharson, C. G. & Oldenburg, D. W., 1993. Inversion of time-domain electromagnetic data for a horizontally layered earth, *Geophysical Journal International*, **114**(3), 433–442. → pages 11, 27

Fincham, A., Christensen, J., Barker, J., Samier, P., et al., 2004. Up-gridding from geological model to simulation model: Review applications and limitations, in *SPE Annual Technical Conference and Exhibition*, Society of Petroleum Engineers. → pages 66

Fitterman, D. V. & Deszcz-Pan, M., 1998. Helicopter EM mapping of saltwater intrusion in Everglades National Park, Florida, *Exploration Geophysics*, **29**(1/2), 240–243. → pages 6

Fitterman, D. V. & Stewart, M. T., 1986. Transient electromagnetic sounding for groundwater, *Geophysics*, **51**(4), 995–1005. → pages 2

Foks, N. L., Krahenbuhl, R., & Li, Y., 2014. Adaptive sampling of potential-field data: A direct approach to compressive inversion, *Geophysics*, **79**(1), IM1–IM9. → pages 76

Fraser, D. C., 1978. Resistivity mapping with an airborne multi-coil electromagnetic system, *Geophysics*, **43**(1), 144–172. → pages 5, 9

Fullagar, P. K., Vrbancich, J., & Pears, G., 2010. Geologically-constrained 1D TEM inversion, *ASEG Extended Abstracts*, **2010**(1), 1–4. → pages 11

Grayver, A. V., Streich, R., & Ritter, O., 2013. Three-dimensional parallel distributed inversion of CSEM data using a direct forward solver, *Geophysical Journal International*, **193**(3), 1432–1446. → pages 18, 50

Guitton, A. & Symes, W. W., 2003. Robust inversion of seismic data using the Huber norm, *Geophysics*, **68**(4), 1310–1319. → pages 14

Haber, E., 1997. *Numerical strategies for the solution of inverse problems*, Ph.D. thesis, University of British Columbia. → pages 16

Haber, E., 2004. A multilevel, level-set method for optimizing eigenvalues in shape design problems, *Journal of Computational Physics*, **198**(2), 518–534. → pages 14

Haber, E. & Ascher, U. M., 2001. Fast finite volume simulation of 3D electromagnetic problems with highly discontinuous coefficients, *SIAM Journal on Scientific Computing*, **22**(6), 1943–1961. → pages 12

Haber, E. & Oldenburg, D., 2000. A GCV based method for nonlinear ill-posed problems, *Computational Geosciences*, **4**(1), 41–63. → pages 16

Haber, E., Ascher, U., Aruliah, D., & Oldenburg, D., 2000. Fast simulation of 3D electromagnetic problems using potentials, *Journal of Computational Physics*, **163**(1), 150–171. → pages 12, 43

Haber, E., Oldenburg, D. W., & Shekhtman, R., 2007. Inversion of time domain three-dimensional electromagnetic data, *Geophysical Journal International*, **171**(2), 550–564. → pages 18, 51

Hansen, P. C. & O'Leary, D. P., 1993. The use of the L-curve in the regularization of discrete ill-posed problems, *SIAM Journal on Scientific Computing*, **14**(6), 1487–1503. → pages 16

Herrmann, F. J., 2010. Randomized sampling and sparsity: Getting more information from fewer samples, *Geophysics*, **75**(6), WB173–WB187. → pages 77

Hestenes, M. R. & Stiefel, E., 1952. Methods of conjugate gradients for solving linear systems, *Journal of research of the National Bureau of Standards*, **49**(6), 409–436. → pages 16

Hohmann, G. W., 1975. Three-dimensional induced polarization and electromagnetic modeling, *Geophysics*, **40**(2), 309–324. → pages 12, 17

Holtham, E. & Oldenburg, D. W., 2012. Large-scale inversion of ZTEM data, *Geophysics*, **77**(4), WB37–WB45. → pages 17

Huang, H. & Won, I., 2003. Characterization of UXO-like targets using broadband electromagnetic induction sensors, *Geoscience and Remote Sensing, IEEE Transactions on*, **41**(3), 652–663. → pages 2

Hyndman, R. D. & Hyndman, D. W., 1968. Water saturation and high electrical conductivity in the lower continental crust, *Earth and Planetary Science Letters*, **4**(6), 427 – 432. → pages 2

Jago, C. P., 2008. *Metal- and alteration-zoning, and hydrothermal flow paths at the moderately-tilted, silica-saturated Mt. Milligan copper-gold alkalic porphyry deposit*, Master's thesis, University of British Columbia. → pages 25

Jardani, A., Revil, A., Santos, F., Fauchard, C., & Dupont, J., 2007. Detection of preferential infiltration pathways in sinkholes using joint inversion of self-potential and EM-34 conductivity data, *Geophysical Prospecting*, **55**(5), 749–760. → pages 6

Johnson, W. J., 2003. Case histories of dc resistivity measurements to map shallow coal mine workings, *The Leading Edge*, **22**(6), 571–573. → pages 6

Keating, P. B. & Crossley, D. J., 1990. The inversion of time-domain airborne electromagnetic data using the plate model, *Geophysics*, **55**(6), 705–711. → pages 10

Kovacs, A., Holladay, J. S., & Bergeron Jr, C. J., 1995. The footprint/altitude ratio for helicopter electromagnetic sounding of sea-ice thickness: comparison of theoretical and field estimates, *Geophysics*, **60**(2), 374–380. → pages 42

Lane, R., Green, A., Golding, C., Owers, M., Pik, P., Plunkett, C., Sattel, D., & Thorn, B., 2000. An example of 3D conductivity mapping using the TEMPEST airborne electromagnetic system, *Exploration Geophysics*, **31**(2), 162–172. → pages 11

Legault, J. M., Carriere, D., & Petrie, L., 2008. Synthetic model testing and distributed acquisition dc resistivity results over an unconformity uranium target from the Athabasca Basin, northern Saskatchewan, *The Leading Edge*, **27**(1), 46–51. → pages 6

Leslie, K., Binks, R., Lam, S., Sullivan, P., Tilbrook, D., Thorn, R., & Foley, C., 2008. Application of high-temperature superconductor SQUIDs for ground-based TEM, *The Leading Edge*, **27**(1), 70–74. → pages 98

Li, D., Beckner, B., Kumar, A., et al., 2001. A new efficient averaging technique for scaleup of multimillion-cell geologic models, *SPE Reservoir Evaluation & Engineering*, **4**(4), 297–307. → pages 66

Li, Y. & Oldenburg, D. W., 1996. 3-D inversion of magnetic data, *Geophysics*, **61**(2), 394–408. → pages 14

Liu, G. & Becker, A., 1990. Two-dimensional mapping of sea-ice keels with airborne electromagnetics, *Geophysics*, **55**(2), 239–248. → pages 42

MacGregor, L. & Sinha, M., 2000. Use of marine controlled-source electromagnetic sounding for sub-basalt exploration, *Geophysical Prospecting*, **48**(6), 1091–1106. → pages 6

Macnae, J., 1998. Fast AEM data processing and inversion, *Exploration Geophysics*, **29**(2), 163–169. → pages 9, 10

Macnae, J., Mortimer, R., & Gilgallon, K., 2010. Deep conductor delineation through improved EMFlow data processing, *ASEG Extended Abstracts*, **2010**, 1–4. → pages 10

Marchant, D., Haber, E., & Oldenburg, D. W., 2013. Inductive source induced polarization, *Geophysical Journal International*, **192**(2), 602–612. → pages 42

Mitchinson, D. & Enkin, R., 2011. Continued investigations of physical property-geology relationships in porphyry-deposit settings in the Quest and Quest-West project area, Central British Columbia (NTS 093E, K, L, M, N), Tech. rep., Geoscience BC. → pages 23, 25

Mogi, T., Tanaka, Y., Kusunoki, K., Morikawa, T., & Jomori, N., 1998. Development of grounded electrical source airborne transient EM (GREATEM), *Exploration Geophysics*, **29**(1/2), 61–64. → pages 62

Mwenifumbo, C., Elliott, B., Jefferson, C., Bernius, G., & Pflug, K., 2004. Physical rock properties from the Athabasca Group: Designing geophysical exploration models for unconformity uranium deposits, *Journal of Applied Geophysics*, **55**(1), 117–135. → pages 2

Nabighian, M. N. & Macnae, J. C., 1991. Time domain electromagnetic prospecting methods, in *Electromagnetic Methods in Applied Geophysics, Volume 2. Applications*, pp. 427–520, ed. Nabighian, M. N., Society of Exploration Geophysicists. → pages 9, 27

Newman, G. A. & Alumbaugh, D. L., 1997. Three-dimensional massively parallel electromagnetic inversion - I. Theory, *Geophysical Journal International*, **128**(2), 345–354. → pages 17, 50

Newman, G. A., Hohmann, G. W., & Anderson, W. L., 1986. Transient electromagnetic response of a three-dimensional body in a layered earth, *Geophysics*, **51**(8), 1608–1627. → pages 12, 17

Oldenburg, D. W. & Li, Y., 1994a. Inversion of induced polarization data, *Geophysics*, **59**(9), 1327–1341. → pages 50

Oldenburg, D. W. & Li, Y., 1994b. Subspace linear inverse method, *Inverse Problems*, **10**(4), 915. → pages 50

Oldenburg, D. W., McGillivray, P. R., & Ellis, R. G., 1993. Generalized subspace methods for large-scale inverse problems, *Geophysical Journal International*, **114**(1), 12–20. → pages 50

Oldenburg, D. W., Li, Y., & Ellis, R. G., 1997. Inversion of geophysical data over a copper gold porphyry deposit: A case history for Mt. Milligan, *Geophysics*, **62**(5), 1419–1431. → pages 40, 42

Oldenburg, D. W., Haber, E., & Shekhtman, R., 2008. Forward modelling and inversion of multisource TEM data, in *SEG Technical Program Expanded Abstracts 2008*, pp. 559–563, Society of Exploration Geophysicists. → pages 18

Oldenburg, D. W., Haber, E., & Shekhtman, R., 2013. Three-dimensional inversion of multisource time domain electromagnetic data, *Geophysics*, **78**(1), E47–E57. → pages 18, 29, 32, 35, 43, 50, 52, 67, 92, 148

Osmond, R. T., Watts, A. H., Ravenhurst, W. R., Foley, C. P., & Leslie, K. E., 2002. Finding nickel from the B-field at Raglan - 'To B or not DB', in *SEG*

Technical Program Expanded Abstracts 2002, pp. 404–407, Society of Exploration Geophysicists. → pages 98

Palacky, G., 1981. The airborne electromagnetic method as a tool of geological mapping, *Geophysical Prospecting*, **29**(1), 60–88. → pages 5

Palacky, G. J., 1988. Resistivity characteristics of geologic targets, in *Electromagnetic Methods in Applied Geophysics, Volume 1. Theory*, pp. 52–129, ed. Nabighian, M. N., Society of Exploration Geophysicists. → pages 1

Palacky, G. J., 1993. Use of airborne electromagnetic methods for resource mapping, *Observations of Earth from Space*, **13**, 5–14. → pages 5, 9

Palacky, G. J. & West, G. F., 1973. Quantitative interpretation of INPUT AEM measurements, *Geophysics*, **38**(6), 1145–1158. → pages 9

Palacky, G. J. & West, G. F., 1991. Airborne electromagnetic methods, in *Electromagnetic Methods in Applied Geophysics, Volume 2. Applications*, pp. 811–877, ed. Nabighian, M. N., Society of Exploration Geophysicists. → pages 9

Pfaffhuber, A., Grimstad, E., Domaas, U., Auken, E., Foged, N., & Halkjær, M., 2010. Airborne EM mapping of rockslides and tunneling hazards, *The Leading Edge*, **29**(8), 956–959. → pages 6

Portniaguine, O. & Zhdanov, M. S., 1999. Focusing geophysical inversion images, *Geophysics*, **64**(3), 874–887. → pages 14

Pridmore, D., Hohmann, G., Ward, S., & Sill, W., 1981. An investigation of finite-element modeling for electrical and electromagnetic data in three dimensions, *Geophysics*, **46**(7), 1009–1024. → pages 12

Raiche, A., 2004. Practical 3D airborne EM inversion in complex terranes, *ASEG Extended Abstracts*, **2004**, 1–4. → pages 10

Raiche, A., Jupp, D., Rutter, H., & Vozoff, K., 1985. The joint use of coincident loop transient electromagnetic and Schlumberger sounding to resolve layered structures, *Geophysics*, **50**(10), 1618–1627. → pages 11

Reid, J. & Fullagar, P., 1998. Conductivity-depth transformation of slingram transient electromagnetic data, *Exploration Geophysics*, **29**(4), 570–576. → pages 10

Reid, J. & Vrbancich, J., 2004. A comparison of the inductive-limit footprints of airborne electromagnetic configurations, *Geophysics*, **69**(5), 1229–1239. → pages 42

Reid, J., Pfaffling, A., & Vrbancich, J., 2006. Airborne electromagnetic footprints in 1D earths, *Geophysics*, **71**(2), G63–G72. → pages 42

Ritter, O., Hoffmann-Rothe, A., Bedrosian, P., Weckmann, U., & Haak, V., 2005. Electrical conductivity images of active and fossil fault zones, *Geological Society, London, Special Publications*, **245**(1), 165–186. → pages 2

SanFilipo, W. A. & Hohmann, G. W., 1985. Integral equation solution for the transient electromagnetic response of a three-dimensional body in a conductive half-space, *Geophysics*, **50**(5), 798–809. → pages 17

Sattel, D., 2005. Inverting airborne electromagnetic (AEM) data with Zohdy's method, *Geophysics*, **70**(4), G77–G85. → pages 11, 26

Schwarz, G., Haak, V., & Rath, V., 1985. Electrical conductivity studies in the Travale geothermal field, Italy, *Geothermics*, **14**(5), 653–661. → pages 2

Schwarzbach, C. & Haber, E., 2013. Finite element based inversion for time-harmonic electromagnetic problems, *Geophysical Journal International*, **193**(2), 615–634. → pages 12, 50

Seigel, H. O., 1974. The magnetic induced polarization (MIP) method, *Geophysics*, **39**(3), 321–339. → pages 62

Siripunvaraporn, W. & Egbert, G., 2000. An efficient data subspace inversion method for 2-D magnetotelluric data, *Geophysics*, **65**(3), 791–803. → pages 76

Smith, R. & Annan, P., 1998. The use of B-field measurements in an airborne time-domain system: Part I. benefits of B-field versus dB/dt data, *Exploration Geophysics*, **29**(1/2), 24–29. → pages 98

Smith, R. S. & Lee, T. J., 2002. The moments of the impulse response: a new paradigm for the interpretation of transient electromagnetic data, *Geophysics*, **67**(4), 1095–1103. → pages 10

Smith, R. S. & Wasylechko, R., 2012. Sensitivity cross-sections in airborne electromagnetic methods using discrete conductors, *Exploration Geophysics*, **43**(2), 95–103. → pages 42

Spitzer, K. & Chouteau, M., 2003. A dc resistivity and IP borehole survey at the Casa Berardi gold mine in northwestern Quebec, *Geophysics*, **68**(2), 453–463. → pages 6

Stolz, E. M., 2000. Electromagnetic methods applied to exploration for deep nickel sulphides in the Leinster area, Western Australia, *Exploration Geophysics*, **31**(1/2), 222–228. → pages 2

Strack, K.-M., Hanstein, T., LeBrocq, K., Moss, D., Vozoff, K., & Wolfgram, P., 1989. Case histories of LOTEM surveys in hydrocarbon prospective area, *First Break*, **7**(12). → pages 6

Strack, K.-M., Lüschen, E., & Kötz, A., 1990. Long-offset transient electromagnetic (LOTEM) depth soundings applied to crustal studies in the Black Forest and Swabian Alb, Federal Republic of Germany, *Geophysics*, **55**(7), 834–842. → pages 6, 62

Sun, J. & Li, Y., 2014. Adaptive l_p inversion for simultaneous recovery of both blocky and smooth features in a geophysical model, *Geophysical Journal International*, **197**(2), 882–899. → pages 14

Tarantola, A., 2005. *Inverse problem theory and methods for model parameter estimation*, SIAM. → pages 13

Tikhonov, A. & Arsenin, V. Y., 1977. *Methods for solving ill-posed problems*, John Wiley and Sons, Inc. → pages 14

Um, E. S., Harris, J. M., & Alumbaugh, D. L., 2010. 3D time-domain simulation of electromagnetic diffusion phenomena: A finite-element electric-field approach, *Geophysics*, **75**(4), F115–F126. → pages 52

Vallee, M. A. & Smith, R. S., 2009. Inversion of airborne time-domain electromagnetic data to a 1D structure using lateral constraints, *Near Surface Geophysics*, **7**(1), 63–71. → pages 11

van Schoor, M., 2002. Detection of sinkholes using 2D electrical resistivity imaging, *Journal of Applied Geophysics*, **50**(4), 393–399. → pages 2

Waff, H. S., 1974. Theoretical considerations of electrical conductivity in a partially molten mantle and implications for geothermometry, *Journal of Geophysical Research*, **79**(26), 4003–4010. → pages 2

Wang, G. L., Torres-Verdín, C., Salazar, J. M., & Voss, B., 2009. Fast 2D inversion of large borehole EM induction data sets with an efficient Fréchet-derivative approximation, *Geophysics*, **74**(1), E75–E91. → pages 18

Wang, T. & Hohmann, G. W., 1993. A finite-difference, time-domain solution for three-dimensional electromagnetic modeling, *Geophysics*, **58**(6), 797–809. → pages 12

Wang, T., Oristaglio, M., Tripp, A., & Hohmann, G., 1994. Inversion of diffusive transient electromagnetic data by a conjugate-gradient method, *Radio science*, **29**(4), 1143–1156. → pages 17

Waxman, M. H., Thomas, E., et al., 1974. Electrical conductivities in Shaly Sands - I. The relation between hydrocarbon saturation and resistivity index; II. The temperature coefficient of electrical conductivity, *Journal of Petroleum Technology*, **26**(02), 213–225. → pages 2

Welhener, H. E., Labrenz, D. M., & Huang, J., 2007. Mt. Milligan project - resource report, Omenica Mining District, British Columbia, Tech. rep., Independent Mining Consultants, Inc. → pages 22

Wilson, G., Raiche, A., & Sugeng, F., 2006. 2.5D inversion of airborne electromagnetic data, *Exploration Geophysics*, **37**(4), 363–371. → pages 11

Wolfgram, P. & Karlik, G., 1995. Conductivity-depth transform of GEOTEM data, *Exploration Geophysics*, **26**(3), 179–185. → pages 10

Wolfgram, P., Sattel, D., & Christensen, N. B., 2003. Approximate 2D inversion of AEM data, *Exploration Geophysics*, **34**(2), 29–33. → pages 11

Xie, G., Li, J., Majer, E. L., Zuo, D., & Oristaglio, M. L., 2000. 3-D electromagnetic modeling and nonlinear inversion, *Geophysics*, **65**(3), 804–822. → pages 17

Yee, K. S., 1966. Numerical solution of initial boundary value problems involving Maxwells equations, *IEEE Transactions on Antennas and Propagation*, **14**(3), 302–307. → pages 147

Yin, C., Huang, X., Liu, Y., & Cai, J., 2014. Footprint for frequency-domain airborne electromagnetic systems, *Geophysics*, **79**(6), E243–E254. → pages 42

Yu, W. W., 2012. *Inversion of airborne electromagnetic data in 2.5D*, Master's thesis, University of British Columbia. → pages 11

Zang, M., 1993. UTEM case history of a base metal prospect Goianesia, Brazil, *Exploration Geophysics*, **24**(3/4), 859–862. → pages 6

Zaslavsky, M., Druskin, V., & Knizhnerman, L., 2011. Solution of 3D time-domain electromagnetic problems using optimal subspace projection, *Geophysics*, **76**(6), F339–F351. → pages 52

Zaslavsky, M., Druskin, V., & Abubakar, A., 2013. Large-scale Gauss-Newton inversion of transient controlled-source electromagnetic measurement data using the model reduction framework, *Geophysics*, **78**(4), E161–E171. → pages 52

Appendix A

3D Forward Modeling Using Finite Volume Method

Using equation 1.1c, 1.1d, 1.2b and 1.2c and neglecting the displacement current, we can derive Maxwell's equations with the quasi-static approximation in time domain as

$$\nabla \times \mathbf{E} + \mu \frac{\partial \mathbf{H}}{\partial t} = 0, \tag{A.1a}$$

$$\nabla \times \mathbf{H} - \sigma \mathbf{E} = \mathbf{J}_s(t). \tag{A.1b}$$

Approximating the time derivative with finite differences using the backward Euler (implicit) method, equation A.1 becomes

$$\nabla \times \mathbf{E}^{i+1} + \mu_0 \frac{\mathbf{H}^{i+1} - \mathbf{H}^{i}}{\delta t} = 0, \tag{A.2a}$$

$$\nabla \times \mathbf{H}^{i+1} - \sigma \mathbf{E}^{i+1} = \mathbf{J}_s^{i+1}, \tag{A.2b}$$

where super-script i denotes the field at time i, and δt is a small time step. By eliminating the electric field we obtain a second-order system for the magnetic field

$$\nabla \times \sigma^{-1} \nabla \times \mathbf{H}^{i+1} + \gamma \mathbf{H}^{i+1} = \nabla \times \sigma^{-1} \mathbf{J}_s^{i+1} + \gamma \mathbf{H}^{i}, \tag{A.3}$$

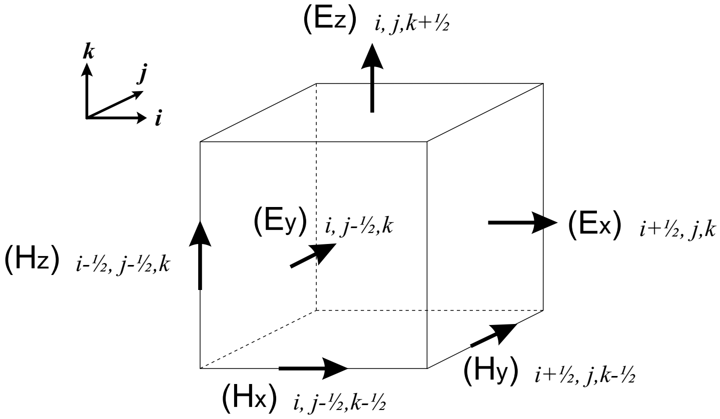

Figure A.1: Staggered discretization in 3D.

where $\gamma = \mu_0/\delta t$.

The finite volume method integrates the differential Maxwell's equations over elemental control volumes aligned with the mesh grid. The modeling domain is discretized by a rectilinear mesh consisting of many rectangular cells. I use a staggered grid (Yee, 1966) with H fields on the edges, E fields on the faces and the conductivity at the cell centres (Figure A.1).

Upon discretization, equation A.3 can be written in a discrete form

$$\left[\mathbf{C}^\top \mathbf{diag}\left(\mathbf{A}_v \exp(\mathbf{m})\right)^{-1}\mathbf{C} + \gamma^{i+1}\mathbf{I}\right]\mathbf{h}^{i+1} = \mathbf{C}^\top \mathbf{diag}\left(\mathbf{A}_v \exp(\mathbf{m})\right)^{-1}\mathbf{j}_s^{i+1} + \gamma^{i+1}\mathbf{h}^i, \tag{A.4}$$

where $\mathbf{C}$ is a curl operator mapping a field from edges to faces, $\mathbf{C}^\top$ is also a curl operator but from faces to edges, $\mathbf{A}_v$ is a harmonic averaging matrix mapping the conductivity values from cell centers to faces, $\exp(\mathbf{m})$ is a vector of conductivities ($\mathbf{m} = \ln\sigma$) for all the cell centers, $\gamma^{i+1} = \mu_0/\delta t$ at time step $i+1$, $\mathbf{h}$ is a vector of H fields in x, y, z directions, and $\mathbf{j}_s$ is a vector of the current density field $\mathbf{J}_s$ due to the source. If a closed transmitter loop is modeled, I first analytically compute the magnetic vector potential $\mathbb{A}$ of the loop source to ensure the source is divergencefree, then obtain the primary current density field

$$\mathbf{J}_s = \nabla \times \mu_0^{-1} \nabla \times \mathbb{A}. \tag{A.5}$$

Equation A.4 is for the H field at a particular time step. Expressing the Maxwell matrix as $\mathbf{A}(\mathbf{m}, \delta t)$ and the first term in the right-hand side of equation A.4 as $\mathbf{q}$, we have the system of equations for the entire modeling in time

$$\begin{pmatrix} \mathbf{A}(\mathbf{m}, \delta t_1) & & & \\ -\gamma(\delta t_2)\mathbf{I} & \mathbf{A}(\mathbf{m}, \delta t_2) & & \\ & \ddots & \ddots & \\ & & \ddots & \ddots \end{pmatrix} \begin{pmatrix} \mathbf{h}^1 \\ \mathbf{h}^2 \\ \vdots \\ \vdots \end{pmatrix} = \begin{pmatrix} \mathbf{q}^1(\mathbf{m}) + \gamma(\delta t_1)\mathbf{h}^0 \\ \mathbf{q}^2(\mathbf{m}) \\ \vdots \\ \vdots \end{pmatrix} \tag{A.6}$$

or in a symbolic form

$$\mathbf{Bh} = \mathbf{rhs}. \tag{A.7}$$

The initial field $\mathbf{h}^0$ can be obtained by solving a static problem.

If the symmetric positive definite matrices $\mathbf{A}(\mathbf{m}, \delta t_i)$ in equation A.6 are factorized into $\mathbf{L}$ and $\mathbf{L}^\top$ and stored, the operator $\mathbf{B}^{-1}$ is available by forward time stepping and $\mathbf{B}^\top$ by backward time stepping. The forward modeled data is

$$F(\mathbf{m}) = \mathbf{QB}^{-1}\mathbf{rhs}, \tag{A.8}$$

where $\mathbf{Q}$ is a sparse space-time interpolation matrix mapping the fields from cell edges to the receiver location and from time steps to the measurement times. I also note $\mathbf{A}$ only changes if δt changes, so in practice, if many delay times from early to late are modeled, the rule of thumb is to adjust the length of δt for each decade to match the scale of time dynamics. In this way, only a few $\mathbf{A}$ matrices need to be factorized.

A detailed description of the algorithm can be found in Oldenburg et al. (2013).

Appendix B

Calculation of Sensitivity

Differentiating both sides of equation A.7 with respect to **m** yields

$$\frac{\partial \mathbf{h}}{\partial \mathbf{m}} = \mathbf{B}^{-1}\left(\frac{\partial \mathbf{rhs}}{\partial \mathbf{m}} - \frac{\partial \mathbf{B}}{\partial \mathbf{m}}\mathbf{h}\right) = \mathbf{B}^{-1}\mathbf{G}. \tag{B.1}$$

Then using the same interpolation matrix **Q** in equation A.8, the sensitivity matrix is

$$\mathbf{J} = \mathbf{Q}\frac{\partial \mathbf{h}}{\partial \mathbf{m}} = \mathbf{Q}\mathbf{B}^{-1}\mathbf{G}, \tag{B.2}$$

and its transpose is

$$\mathbf{J}^{\top} = \mathbf{G}^{\top}\mathbf{B}^{-T}\mathbf{Q}^{\top}, \tag{B.3}$$

where **Q** is again a space-time interpolation matrix. As the operators $\mathbf{B}^{-1}$ and $\mathbf{B}^{-\top}$ are already available in forward modeling (Appendix A), I derive other components in $\mathbf{G}$ as

$$\frac{\partial \mathbf{rhs}}{\partial \mathbf{m}} = \begin{pmatrix} \frac{\partial \mathbf{q}^1}{\partial \mathbf{m}} \\ \frac{\partial \mathbf{q}^2}{\partial \mathbf{m}} \\ \vdots \\ \vdots \end{pmatrix} = \begin{pmatrix} -\mathbf{C}^\top \mathbf{diag}(\mathbf{j}_s^1)\mathbf{diag}[(\mathbf{A}_v \exp(\mathbf{m}))^{-2}]\mathbf{A}_v \mathbf{diag}(\exp(\mathbf{m})) \\ -\mathbf{C}^\top \mathbf{diag}(\mathbf{j}_s^2)\mathbf{diag}[(\mathbf{A}_v \exp(\mathbf{m}))^{-2}]\mathbf{A}_v \mathbf{diag}(\exp(\mathbf{m})) \\ \vdots \\ \vdots \end{pmatrix}, \tag{B.4}$$

and

$$\frac{\partial \mathbf{B}}{\partial \mathbf{m}}\mathbf{h} = \begin{pmatrix} \frac{\partial}{\partial \mathbf{m}}\left[\mathbf{A}(\mathbf{m})\mathbf{h}^1\right] \\ \frac{\partial}{\partial \mathbf{m}}\left[\mathbf{A}(\mathbf{m})\mathbf{h}^2\right] \\ \vdots \\ \vdots \end{pmatrix} = \begin{pmatrix} -\mathbf{C}^\top \mathbf{diag}(\mathbf{C}\mathbf{h}^1)\mathbf{diag}[(\mathbf{A}_v \exp(\mathbf{m}))^{-2}]\mathbf{A}_v \mathbf{diag}(\exp(\mathbf{m})) \\ -\mathbf{C}^\top \mathbf{diag}(\mathbf{C}\mathbf{h}^2)\mathbf{diag}[(\mathbf{A}_v \exp(\mathbf{m}))^{-2}]\mathbf{A}_v \mathbf{diag}(\exp(\mathbf{m})) \\ \vdots \\ \vdots \end{pmatrix}. \tag{B.5}$$

Now we have options of implicit or explicit sensitivity. The implicit method stores $\mathbf{Q}$, $\mathbf{B}^{-1}$ (essentially several factorized $\mathbf{A}$) and $\mathbf{G}$ for the operation of $\mathbf{J}$ times a vector during inversion, eliminating the necessity of forming $\mathbf{J}$. However, $\mathbf{G}$ requires the field $\mathbf{h}$ everywhere in the mesh at every time, which actually consumes more memory than storing $\mathbf{J}$ in EM modeling.

The explicit method computes and stores the sensitivity matrix $\mathbf{J}$. Since the number of data is always much less than the number of model parameters, I compute $\mathbf{J}^\top$ in equation B.3 column by column. An explicit $\mathbf{J}$ can be formed using equation B.3 and solving for $\mathbf{J}^\top \mathbf{e}_i$ for the ith datum where $\mathbf{e}_i$ is the zero vector with a unit value in the ith entry. Computation of each column of $\mathbf{J}^\top$ is equivalent to one complete forward modeling involving backward time stepping ($\mathbf{B}^{-\top}$ operation). After $\mathbf{J}$ is stored, the matrices and vectors associated with $\mathbf{B}$, $\mathbf{G}$, and $\mathbf{Q}$ can all be deleted from the memory.

Appendix C

Fast Discrete Model Conversion Between Arbitrary Meshes

The key of survey decomposition is to be able to solve many subproblems that effectively represent the original problem for specific source-receiver pairs. The representation of the global model on a local mesh is crucial in this process. One popular and proven method is intersecting-volume weighted material averaging similar to Commer & Newman (2008). Here I document an approach that implements fast model conversion between two meshes embedded in different coordinate systems using the proxy of point cloud. I assume the following information are known: the original model $\mathbf{m}_1$, a vector, on mesh 1 in coordinate system 1, the destination mesh 2 in coordinate system 2, and the coordinate transform parameters, which, in case of two Cartesian systems, are a 3-by-3 rotation matrix $\mathbf{M}$ and a displacement vector $\mathbf{t}$ so that for a point $\mathbf{p}_1 = (x_1, y_1, z_1)$ in system 1 and its coordinate $\mathbf{p}_2 = (x_2, y_2, z_2)$ in system 2, we have

$$\mathbf{p}_2 = \mathbf{M}\,\mathbf{p}_1 + \mathbf{t}. \tag{C.1}$$

In addition, there is no restrictions on the type of the meshes (regular or irregular), but it requires that a point can be located in a mesh in terms of cell index. The goal is to find an averaging matrix $\mathbf{A}$, with which we can get the model on the second mesh m_2 by applying the averaging matrix

$$\mathbf{m}_2 = \mathbf{A}\,\mathbf{m}_1. \tag{C.2}$$

An entry a_{ij} in $\mathbf{A}$ represents the weight, essentially the normalized intersecting volume, when calculating the relative contribution of cell j in the original model to cell i in the converted model. Summing elements of $\mathbf{A}$ along rows yields a vector of all ones

$$\sum_{j=1}^{N_{\mathbf{m}_1}} a_{ij} = \mathbf{1}. \tag{C.3}$$

Matrix $\mathbf{A}$ is usually sparse because a cell in $\mathbf{m}_2$ may intersect only a few cells in m_1. If inactive cells are involved, the corresponding rows and columns of $\mathbf{A}$ must be removed.

The point cloud approach first initiates N_p random points in every cell in $\mathbf{m}_2$ knowing the cell index of every point in $\mathbf{m}_2$; then the points are transformed to system 1 and their cell indices in $\mathbf{m}_1$ are calculated; finally $\mathbf{A}$ is assembled by counting the number of points: if there is a point from cell i in $\mathbf{m}_2$ is found in cell j in $\mathbf{m}_1$, then a_{ij} adds one. Each row of $\mathbf{A}$ is normalized by its sum before $\mathbf{A}$ is applied to $\mathbf{m}_1$. When N_p approaches infinity, the estimated intersecting volumes are identical with the results computed using the rigorous analytic geometry; however, I have noticed even a small N_p, for example 50, is adequate to quickly produce a good converted model for the purpose of EM modeling. The point cloud approach is especially useful when using a rigorous method is expensive for complex meshes.