Holistic Inversion
of
Airborne Electromagnetic Data

Ross Colin Brodie

A thesis submitted for the degree of

Doctor of Philosophy

of

The Australian National University

May 2010

Declaration

This thesis is the result of research undertaken while I was a student in the Research School of Earth Sciences at the Australian National University. Except as otherwise stated in the text, the work described is original and my own. The thesis has never been submitted to another university or similar institution.

Ross Colin Brodie

4 May 2010

Acknowledgements

Malcolm Sambridge has been a truly excellent supervisor to me throughout my studies at the Research School of Earth Sciences. Malcolm's mastery of all-things-inversion was inspiring. Most important to me was his ability to teach— to explain with clarity. Malcolm was always a pleasure to work with, he always provided me with the right balance between freedom and focus, and provided a great amount of support when I had to juggle competing workloads. Let us make sure that our collaborations do continue.

I thank my employer, Geoscience Australia, for its support in providing the sponsorship to undertake this research. It has been a marvellous opportunity and I have truly appreciated it. Special thanks go to Barry Drummond for instigating the whole thing. I am glad you nudged me out of my comfort zone and gave me the encouragement to undertake a PhD.

I would like to show my gratitude toward some colleagues who were important leading up to this research. When I first began working in the field of airborne electromagnetics, but really did not yet know a great deal about the subject, Andy Green was of great practical assistance an of immense educational value to me. It was from a collaborative project with Andy that the work in this thesis all began really. Thanks to James Reid who set me off on a path toward quantitative modelling by suppling me with Guptasarma and Singh's Hankel transform coefficients. Although these are now a distant memory, it was a formative early step. I have always enjoyed and learned from the many collaborations and discussions with Richard Lane. Since he arrived at Geoscience Australia, Richard has been a great mentor and has taught me so much about airborne electromagnetics that one simply does not learn from the literature. Other important scientific mentors at Geoscience Australia, from whom I have learnt so much about airborne geophysics over the years, have been Brian Minty, Peter Milligan,

Murray Richardson and formerly Ian Hone. Extra thanks go to Brian, who diligently proofread the bulk of this thesis.

Also deserving acknowledgement are, my mother Helen, my late father John, and all of my eight siblings, who have always been of great support. They provided the educational opportunities and right family environment for achievement— something that should not be underestimated. Gabrielle Sheen has been a wonderfully close personal friend to me throughout the entirety of these studies. Gabrielle's support and encouragement is truly appreciated.

I thank the South Australian Salinity Mapping and Management Support Project, which was jointly funded by the Australian and South Australian Governments under the National Action Plan for Salinity and Water Quality, for permission to use and publish the Riverland dataset in this research. I also thank the Bureau of Rural Sciences, an agency of the Department of Agriculture, Fisheries and Forestry, for permission to use and publish parts of the Lower Macquarie dataset.

Abstract

A holistic method for simultaneously calibrating, processing, and inverting frequency-domain airborne electromagnetic data has been developed. A spline-based, 3D, layered conductivity model covering a complete survey area is recovered through inversion of an entire raw airborne data set and available independent geoelectric and interface-depth data. The holistic inversion formulation includes a mathematical model to account for systematic calibration errors such as incorrect gain, phase and zero-level. By taking these elements into account in the inversion, the need to pre-process the airborne data prior to inversion is eliminated.

Conventional processing schemes involve the sequential application of a number of calibration corrections, with data from each frequency being treated separately. This is followed by inversion of each multi-frequency airborne sample in isolation from other samples. By simultaneously considering all of the available information in a holistic inversion, the inter-frequency and spatial coherency characteristics of the data are able to be exploited. The formulation ensures that the conductivity and calibration models are optimal with respect to the airborne data and prior information. Introduction of inter-frequency inconsistency and multistage error propagation stemming from the sequential nature of conventional processing schemes is also avoided.

It is confirmed that accurate conductivity and calibration parameter values are recovered from holistic inversion of synthetic data sets. It is also demonstrated that the results from holistic inversion of raw survey data are superior to the output of conventional 1D inversion of final processed contractor delivered data. In addition to the technical benefits, it is expected that holistic inversion will reduce costs by avoiding the expensive calibration→processing→recalibration paradigm. Furthermore, savings may

also be made because specific high altitude zero-level observations, needed for conventional processing, may not be required.

The same philosophy is also applied to the inversion of time-domain data acquired by fixed-wing towed-bird systems. A spline-based, 2D, layered conductivity model covering a complete survey line is recovered along with a calibrations model. In this instance, the calibration model is a spline based representation of three unmeasured elements of the system geometry. By inverting the less processed total field data, the procedure is able to prevent incorrect assumptions made in conventional primary field removal from being propagated into the inversion stage. Furthermore, by inverting a complete line of data at once the along-line spatial coherency of the geology and the geometry variations is exploited.

Using real survey data, it was demonstrated that all components of the data could be simultaneously and satisfactorily fitted and that the resulting conductivity model was consistent with independent prior information. This was an improvement over the conventional approach, in which the data could not be satisfactorily fitted, nor was the conductivity model consistent with prior information. It was further established that by using the holistic inversion spline parameterization, the resulting conductivity model was more continuous and interpretable than if the conventional style discrete parameterization was used.

If adopted, the holistic approach, could reduce survey costs, reduce data processing turnaround times, and improve the quantitative information that can be extracted from data, and hence, increase the value of airborne electromagnetics for mineral exploration and environmental mapping applications.

Table of Contents

Table of Appendices

Chapter 1

Introduction

1.1 Outline

In this introductory chapter to the thesis I will provide an overview of the airborne electromagnetic method. The overview will introduce the basic principles of electromagnetic geophysics, its applications, the governing theory, some useful rules of thumb, and an explanation of acquisition systems and the data presented to the quantitative interpreter. The overview material will be familiar to the experienced electromagnetic geophysicist, however it provides necessary background for those readers who are unfamiliar with the technique.

Following the overview, I discuss the factors that limit the accuracy of subsurface electrical conductivity estimates derived from airborne electromagnetic data through geophysical inversion. These include the non-uniqueness of the inverse problem, systematic calibration errors that contaminate the data, and the inconsistencies that may be created in the data processing steps taken to mitigate the systematic errors. This sets out the motivations for the research that I have pursued in an effort to address these important issues.

I then propose a new framework for the calibration, processing and inversion of airborne electromagnetic data. The overarching proposition is that entire airborne datasets be simultaneously inverted to solve for a continuous 3D conductivity model and a calibration model that mathematically parameterizes systematic error in the data. Following that, a broad outline of the remainder of the thesis is set out.

1.2 The airborne electromagnetic method

1.2.1 Basic principles

Electromagnetic methods are one group of geophysical techniques that allow estimation of the distribution and electromagnetic properties of subsurface materials via non-invasive remotely-measured observations. Electromagnetic methods are founded upon the phenomenon of electromagnetic induction, which is governed by the famous Maxwell Equations (Maxwell, 1892). The method involves the transmission of electromagnetic energy via a wire loop carrying a time-varying current which has associated magnetic and electric fields (Figure 1.1). These primary fields induce eddy currents to flow in electrically conductive subsurface material. The subsurface currents in turn have their own associated secondary magnetic and electric fields that may be detected on or above the surface. Since the induced subsurface currents and hence their associated secondary fields are influenced by the electromagnetic properties of the subsurface, analysis of the detected fields enables inferences to be made about those subsurface materials. Knowledge of the electromagnetic properties of subsurface materials, the electrical conductivity (σ), the magnetic permeability (μ) and the dielectric permittivity (ε), has significance in geoscience, environmental science and geotechnical sciences.

There are several different modes of operation for electromagnetic methods. They are classified according to attributes like: dimensions, orientation and spectral content of the transmitting source; the relative position and orientation of the source and receiver; how the receiver detects the fields; and the manner in which a survey is carried out. This work deals only with the airborne electromagnetic (AEM) method, which is one particular class of electromagnetics in which an aircraft carries and/or tows the transmitting and receiving instrumentation through the air. The two main subcategories of AEM systems are frequency-domain and time-domain systems. These will be

described in detail later in the chapter. However, for now it is only necessary to note that frequency-domain systems transmit a continuous sinusoidal current waveform through multiple loops at different discrete frequencies. In contrast, time-domain systems transmit a pulsed (i.e. on then off) current waveform through one loop.

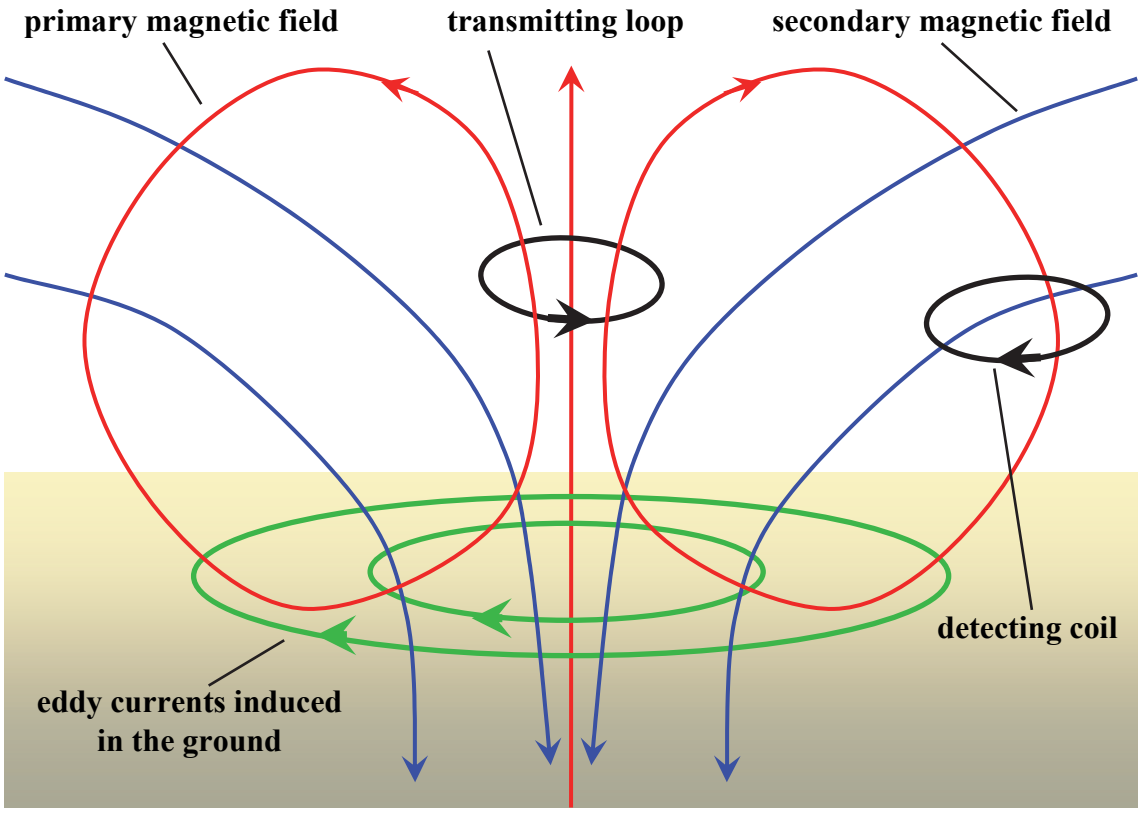

Figure 1.1 *The principal concepts of AEM: Above the ground, a time varying current is driven through a wire loop, causing eddy currents to flow in the ground. The eddy currents have an associated varying magnetic field which induces a voltage in a receiving coil that is recorded and analysed.*

AEM was first trialled in Canada in 1946 (Palacky and West, 1991) and was further developed from 1948 by a small exploration company when they adapted a ground based tractor-towed electromagnetic system so that it could be carried by a wooden Anson aircraft (Fountain, 1998). In 1954 an AEM survey led to the discovery of the Heath Steel zinc-copper-lead-silver deposit, which triggered a wider take up of the new technology. No doubt the motivation for the development of AEM was its potential to rapidly survey large areas in a regular spatial fashion, without the impediments of topographic and cultural features (e.g. hills, rivers, roads and fences) that obstructed

ground based surveys. Because of the limitations of operating from a moving airborne platform, AEM data cannot match the resolution and accuracy that ground based surveys are capable of. However, dense along line sampling and quasi-continuous spatial coverage, which is typical of modern AEM surveys, enables detailed and meaningful geological signatures to be distinguished using the full spatial context of the AEM data. Ground based methods are typified by relatively few transects and/or point soundings, from which important signatures may not be distinguished due to the lack of complete spatial context. AEM is therefore often used as a tool to assess a broad area in a cost and time efficient manner, thereby providing a means of identifying and prioritizing specific features that can be followed up with more detailed and discriminating ground based techniques.

Initially, AEM was primarily used in the exploration for base metal ore bodies. Some ore bodies are extremely conductive, and often reside in a resistive host rock, which means the large contrast can give rise to a localized anomaly when an AEM system is flown over them. Since the anomaly appears as bump on a plotted profile of AEM data against an otherwise bland background, the exploration for these discrete conductive bodies is often referred to as 'bump-finding'. The presence of thick conductive regolith (the soil and weathered material between the surface and fresh bedrock) in Australia, and many other parts of the world, has been an impediment to the application of AEM because it masks and complicates the anomaly caused by discrete conductors. Nevertheless, AEM has steadily progressed over the 56 years following the first mineral discovery and is now a widely used tool, not only for bump-finding, but ironically, also for environmental and geotechnical applications that are solely focused on mapping the regolith that masked the bumps.

1.2.2 Applications

The purpose of AEM is to enhance our understanding of the subsurface by gaining knowledge of its electromagnetic properties. It is applicable where there is sufficient contrast between the electromagnetic properties of the various subsurface units of interest. In these scenarios the electromagnetic properties can act as a surrogate parameter for subsurface mapping. The method cannot be used to definitively determine the composition of mineral ores or lithological units because the bulk electromagnetic properties of rocks are a complex function of multiple variables (e.g. mineral content, porosity, pore fluid conductivity and saturation), and hence they are not prescriptive.

Electrical conductivity is the most feasible of the electromagnetic properties that can be resolved by AEM. Mapping of magnetic permeability and dielectric permittivity may be possible under some limited circumstances. The depth of investigation and the resolution of AEM are strongly dependent on the conductivity of the subsurface and the AEM system being used. However, broadly speaking, AEM methods are suitable for investigation in the top 600 m and when the vertical and horizontal resolution is not required to be better than 4 m and 40 m respectively.

AEM was originally developed as a mineral exploration tool under the understanding that mineral ores, particularly massive sulphides, have a large electrical conductivity contrast with their host rock (Palacky and West, 1991). Numerous case studies have been presented in the literature showing how AEM can be used to delineate base metal ore bodies (e.g. Wolfgram and Golden, 2001, Smith et al., 2003, Witherly and Irvine, 2006). Other mineral commodities, for example diamondiferous kimberlite and lamproite bodies (Macnae, 1995; Cunion, 2009) and uranium (Palacky, 1990; Reid and Viezzoli, 2007; Sorensen et al., 2009), have been explored for using AEM. Unlike the base metals, these commodities themselves are not especially conductive, but they may for example be detectable by virtue of distinctive weathering signatures of their hosts.

Further application has been found in geological mapping where subtle contrasts between various geological units are exploited (Schaefer et al., 1998; Worrall et al., 1998; Lawrie et al., 2000b). AEM is used in hydrogeological investigations for detection of water resources and water quality characterisation (Sengpiel, 1983; Fitterman and Deszcz-Pan, 1998; Sattel and Kgotlhang, 2004; Auken et al., 2007). Mapping of saline soils is a particularly widespread application in Australia (Anderson et al., 1993; Street et al., 1998; Lawrie et al., 2000a; Brodie et al., 2004b).

Other applications include: mapping of sea ice thickness (Kovacs and Valleau, 1990; Liu and Becker, 1990; Reid et al., 2003); bathymetry mapping (Zollinger et al., 1987; Vrbancich et al., 2000; Wolfgram and Vrbancich, 2006); unexploded ordinance detection (Gamey et al., 2000); and characterisation of landslides (Konishi, 1998). Further applications have been reported by Hodges (1999), Pellerin (2002) and Ackman (2003), some of which include detection of buried objects, mine discharge contamination and soil characterisation for pipeline construction.

1.2.3 Electromagnetic properties of rocks

There are three physical properties of rocks that are important in electromagnetics. They are the electrical conductivity σ (S/m), magnetic permeability μ (H/m) and dielectric permittivity ε (F/m). Natural materials are usually non-symmetric in crystalline structure and their electromagnetic properties depend on the direction in which they are measured. Thus the quantities σ, μ, and ε are anisotropic and can only correctly be represented as tensor quantities (Keller, 1988). The properties may also be dispersive or non-linear, meaning that their values may be a function of the frequency and amplitudes, respectively, of the applied magnetic or electric field. However in AEM there is rarely enough information content in the data to resolve isotropic values let alone the complete tensor, or dispersive or non-linear effects. So

practically dictates that we assume the electromagnetic properties to be scalar and isotropic in almost all cases.

As we will see in the next section, it turns out that over the frequencies used in AEM ($\leq 10^5$ Hz) dielectric permittivity usually has negligible influence on observed electromagnetic fields compared to conductivity and magnetic permeability and errors in AEM measurements. It has non-negligible effect only in resistive terrains and only at frequencies in the upper end of the AEM operating range. Thus dielectric permittivity is unlikely to be important or resolvable in the conductive Australian setting. However under the more resistive settings discussed by Huang and Fraser (2001) and Hodges (2004) dielectric permittivity is important and may even be resolvable by AEM.

Magnetic permeability is related to the more fundamental property, magnetic susceptibility (k_μ) through the expression,

$$\mu = \mu_0(1 + k_\mu), \tag{1-1}$$

where $\mu_0 = 4\pi \times 10^{-7} H/m$ is the magnetic permeability of free space. The magnetic susceptibility of rock forming minerals can vary over several orders of magnitude. Magnetic susceptibility and hence permeability will only become increasing important as it moves away from zero. The only minerals abundant enough in nature to significantly influence the bulk magnetic permeability of rock units (i.e. at the macroscopic scale relevant to AEM) is magnetite and to a lesser extend a few other ferromagnetic minerals such as hematite, ilmenite, pyrrhotite and maghaemite. In many situations the magnetic permeability is taken to be the free space value. This is an especially good assumption in sedimentary environments, but important exceptions do exist and need to be considered. In fact Huang and Fraser (2001) and Hodges (2004) also show that magnetic permeability can be interpreted from AEM data in some cases.

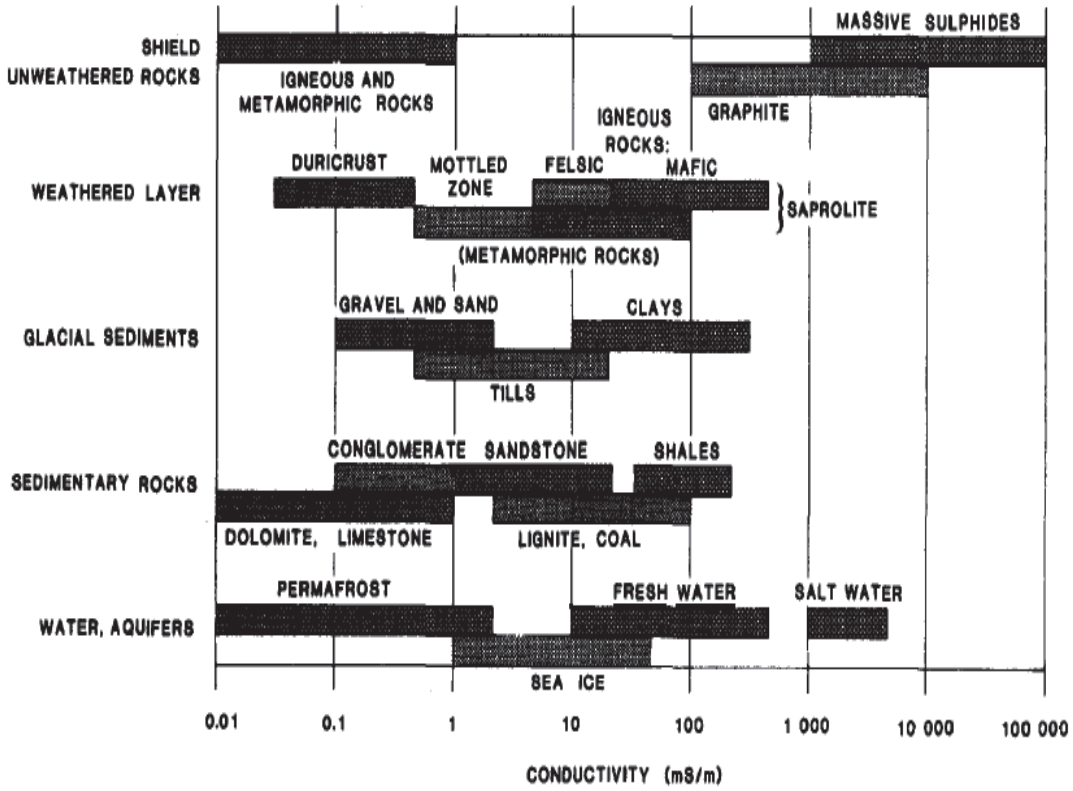

Figure 1.2 A diagram after Palacky (1993) showing the typical range of conductivities of common earth materials.

Electrical conductivity is the most important and most feasible of the electromagnetic properties to be estimated from AEM data. Like magnetic susceptibility, conductivity also ranges over several orders of magnitude in naturally occurring materials. Figure 1.1 is a diagram after Palacky (1993) that shows the typical range of conductivities of common earth materials. Base metal sulphides are typically very conductive and are of course important targets in mineral exploration applications. Native metals and graphite are also conductive rock constituents. Almost all other rock forming minerals act as insulators. Important exceptions are some clay minerals that serve to increase conductivity via ion exchange processes if they are wet (Keller, 1988). This is an important factor in AEM mapping of the regolith zone. Bulk rock conductivity is influenced by the electrical conductivity of the constituent mineral grains; their shape, connectivity and concentration; the volumetric amount of pore fluid; the concentration and mobility of dissolved ions in the pore fluid; and connectivity and shape of the pores. These factors have been described in detail by Archie (1942) and

McNeill (1980). Fresh bedrock without large concentrations of the highly conductive minerals noted above (i.e. metals, graphite or sulphides) and with few interconnecting pores tends to have a very low conductivity. The volumetric content and conductivity of the pore fluids are the dominating factors influencing bulk rock conductivity in the regolith. This is an especially important factor in Australia where saline groundwater is abundant.

1.2.4 The physics of electromagnetics

In electromagnetic induction surveys a transmitter generates a time varying current in a single or multi-turn loop of wire. The current has an associated time varying magnetic field, which is called the primary field, and it is related to the current in the loop through Ampere's law. The primary magnetic field propagates from the transmitter and induces eddy currents to flow in subsurface conductive material. These are also called secondary or induced currents, and they have an associated secondary magnetic field. The time rate of change of the total magnetic field, the sum of the primary and secondary fields, can be detected at a receiver consisting of one or more wire coils. It is detected by measuring the voltage that the time varying total magnetic field induces in the receiver coils according to Faraday's law. More than one receiver coil may be used in order to measure separate directional components of the magnetic field. The electromagnetic properties of materials that the electromagnetic energy penetrates influence the amplitude, direction and relative phases of the measured fields. Thus, by analysis of the voltages measured in the receiver coil(s), inferences can be made about the electromagnetic properties of the subsurface materials.

Most readers will be familiar with the physical principles of induction, eddy currents, Ampere's law and Faraday's law that are used in the summary above. However making accurate inferences about subsurface electromagnetic properties is not trivial because it requires consideration of electromagnetic theory. Thorough examinations of

electromagnetic theory relevant to geophysical applications are provided, for example, by Ward (1967), Wait (1982), Ward and Hohmann (1988), and West and Macnae (1991). These authors explain how the relationships between the current in the transmitter loop, the electromagnetic fields, the subsurface electromagnetic properties and the voltages measured in the receiver coils are governed by Maxwell's equations (Maxwell, 1892).

Maxwell elegantly related Faraday's Law, Ampere's Law and Gauss's Laws for electricity and magnetism into a single unifying theory of electromagnetism through the expression of four fundamental uncoupled differential equations. In their differential time-domain form they are,

$$\nabla \times \mathbf{e} + \frac{\partial \mathbf{b}}{\partial t} = 0, \tag{1-2}$$

$$\nabla \times \mathbf{h} - \frac{\partial \mathbf{d}}{\partial t} = \mathbf{j}, \tag{1-3}$$

$$\nabla \cdot \mathbf{b} = 0, \tag{1-4}$$

$$\nabla \cdot \mathbf{d} = \rho. \tag{1-5}$$

In these equations ρ (C/m^3) represents electric charge density, and the vector quantities $\mathbf{b}$ (Wb/m^2) is the magnetic flux density or magnetic induction, $\mathbf{h}$ (A/m) is the magnetic field intensity, $\mathbf{d}$ (C/m^2) is the electric flux density or dielectric displacement, $\mathbf{e}$ (V/m) is the electric field intensity, $\mathbf{j}$ (A/m^2) represents electric current density, and $\partial/\partial t$ signifies a time rate of change.

To be of use in geophysics the Maxwell's equations must of course be related somehow to the electromagnetic properties of the subsurface, electrical conductivity (σ) magnetic permeability (μ) and dielectric permittivity (ε) that were introduced in the previous section. They couple Maxwell's equations through the so called constitutive relations,

$$\mathbf{b} = \mu\,\mathbf{h}\,, \tag{1-6}$$

$$\mathbf{d} = \varepsilon\,\mathbf{e}\,, \tag{1-7}$$

$$\mathbf{j} = \sigma\,\mathbf{e}\,. \tag{1-8}$$

As noted in the previous section, σ, μ and ε may in general be anisotropic, nonlinear and dispersive. However in all but a few geophysical applications they are considered to be scalar, isotropic, linear and non-dispersive.

One step toward making inferences about the subsurface, in a quantitative sense, is to solve Maxwell's equations so that the electromagnetic field components at the receiver can be expressed in terms of the subsurface electromagnetic properties. In Chapter 2 it is shown how for fields in homogenous source free regions, that vary harmonically in time t with angular frequency ω in the form $e^{i\omega t}$ ($i = \sqrt{-1}$), the electric and magnetic fields be expressed as the frequency-domain wave equations,

$$\nabla^2\mathbf{E} + (\mu\varepsilon\omega^2 - i\mu\sigma\omega)\mathbf{E} = 0\,, \tag{1-9}$$

and,

$$\nabla^2\mathbf{H} + (\mu\varepsilon\omega^2 - i\mu\sigma\omega)\mathbf{H} = 0\,, \tag{1-10}$$

where $\mathbf{E}$ and $\mathbf{H}$ are the frequency-domain equivalents of $\mathbf{e}$ and $\mathbf{h}$.

The wave equations, also known as the Helmholtz wave equations, describe the propagation of electromagnetic energy with propagation constant $k = \sqrt{\mu\varepsilon\omega^2 - i\mu\sigma\omega}$. For low frequencies ($\leq 10^5$ Hz) and real earth materials it is usually true that $\mu\varepsilon\omega^2 << \mu\sigma\omega$. Known as the quasi-static case, this assumption is usually valid for the frequencies employed in AEM. Physically the quasi-static assumption means that conduction currents dominate over displacement currents. It also means that the Helmholtz equations reduce to diffusion equations and thus diffusion is the dominant process, which results in the resultant lack of resolution in electromagnetic methods (Ward and Hohmann, 1988).

It is only for certain rather simple subsurface distributions of the electromagnetic properties that closed form analytic solutions of Maxwell's Equations can be derived. Otherwise one must resort to numerical solutions of the differential equations involving finite difference or finite element methods. The theory and codes for forward modelling of general 2D and 3D distributions of the electromagnetic properties do exist (e.g. Ellis, 1995; Sugeng and Raiche, 2004), however they are not routinely used for interpretation of complete AEM datasets at this point in time. This is primarily due to the large computational cost associated with forward modelling of general 2D and 3D distributions.

Due to the vast amounts of data to be interpreted in AEM, practicality dictates that simplifying assumptions about the spatial distribution of the properties and/or approximations in the forward modelling must be made. One such assumption is the so called layered-earth or 1D approximation in which the electromagnetic properties are assumed to be constant within a series of vertically stacked layers. The 1D layered-earth assumption was used for the inversion work carried out in this research. The details of the 1D forward modelling routine used for this work, which is based on the solution to the Helmholtz equations described in Wait (1982), are presented in Chapter 2.

For the simplest case of a homogeneous earth the solution of the Helmholtz wave equations provides insight into some general principles of electromagnetics. Ward and Hohmann (1988) show how the solution of the wave equation under the quasi-static assumption for a harmonic vertical dipole (horizontal loop) source over a homogeneous earth, leads to a much used parameter in electromagnetics. The skin depth, which is given by,

$$\delta = \sqrt{\frac{2}{\omega\mu\sigma}}, \tag{1-11}$$

is the depth at which the primary fields decay to $1/e$ of their value at the surface. The skin depth is often used as an indication of the depth of investigation of electromagnetic methods. Skin depth cannot be used to definitively determine or compare the depth of investigation of electromagnetic systems as it takes no account of the source power or geometry and the sensitivity and noise levels of the receiver. However, skin depth does mean that, all other things being equal, as the frequency or conductivity increase the primary field decay is more rapid. This leads to the general rule of thumb that for harmonic sources (e.g. frequency-domain AEM systems) the measured electromagnetic fields are more sensitive to shallow conductive features at high source frequencies, and are more sensitive to deep features at low source frequencies.

A related principle is the diffusion length (Nabighian and Macnae, 1991) which applies, not to harmonic, but to step-function sources in a conducting homogeneous medium. The diffusion length is,

$$d = \sqrt{\frac{2t}{\mu\sigma}}\,, \tag{1-12}$$

and it is the distance at which the electric field and the currents are at their maximum for a given delay time t after the current step. Another useful approximation for a step-function excitation was presented by Nabighian (1979). He showed that the measured response can be approximately represented by a downward and outward moving equivalent current filament, of diminishing amplitude and having the same shape as the transmitter loop. The general rule of thumb to be gleaned is that for step-function sources (e.g. time-domain AEM systems), measurements at early time after switch off will be most sensitive to shallow conductive features, and that late time measurements will be most sensitive to deeper conductive features.

The variation in sensitivity of measurements at different frequencies and delay times, which is described by the skin depth and diffusion depth concepts, is what gives

electromagnetics the ability to perform conductivity-depth soundings. By transmitting and measuring at several frequencies (in a frequency-domain system) or measuring at several delay times (in a time-domain system), an AEM system can potentially discriminate a conductivity-depth curve.

1.2.5 Frequency domain AEM systems

Frequency-domain AEM systems consist of multiple transmitter-receiver-bucking coil triplets or 'coilsets'. Contemporary systems combine four to six coilsets into the one system. The transmitter and receiver coils are sometimes housed inside a cigar shaped tube, known as a 'bird', which is towed beneath a helicopter as shown in Figure 1.3. Otherwise they may be housed in pods mounted on the wing tips of a fixed-wing aircraft as shown in Figure 1.4. The transmitter and receiver coils are separated by approximately 8 m and 21 m in the helicopter and fixed-wing systems shown respectively. Systems usually have coilsets whose coil axes are nominally orientated vertically (horizontal-coplanar or HCP loops) or in the horizontal flight line direction (vertical-coaxial or VCX loops). The choice of orientation depends on the orientation of the targets to be detected, because HCP coilsets are more sensitive to horizontal conductors and VCX coilsets are more sensitive to dipping conductors.

A continuous sinusoidal current is passed through the transmitter coil of each coilset at a single fixed discrete frequency. Typically the frequencies used are in the range 300 Hz to 100 kHz. In each coilset there is one receiver coil tuned to receive at the specific frequency of its paired transmitter. The quantity that is measured is the voltage induced in the receiver coil. The voltage is proportional to the time rate of change of the magnetic flux threading the coil (i.e. the area of the coil × times the magnetic flux density $\mathbf{B}$ in the direction of the coil's axis) and the number of coil turns. The magnetic field at the receiver is the vector sum of the primary magnetic field emanating

directly from the transmitter coil, and the secondary magnetic field due to the eddy currents flowing in the subsurface.

Figure 1.3 Photograph of the DIGHEM frequency-domain AEM system that shows a magnetometer bird (top) and electromagnetic bird (bottom) that is towed approximately 30 m below the helicopter.

It is the only the secondary field that is instructive about the subsurface, however it is typically only a small fraction (e.g. 10^{-4}) of the primary magnetic field strength. Therefore the influence of the primary field must somehow be suppressed before accurate measurements of the secondary field can be made. To deal with this primary field contamination problem a 'bucking' coil is used. The bucking coil is the third member of each coilset which feeds the negative of the voltage that is induced in it into the receiver coil. Its position, area and number of turns are chosen such that it theoretically cancels out the voltage induced in the receiver coil by the primary field. Bucking reduces the dynamic range over which the receiver must measure and thus improves the resolution achievable by the analogue to digital converters.

Figure 1.4 The GTK-AEM-05 frequency-domain AEM system which has pods mounted on the wing tips that house the transmitter and receiver coils.

The effectiveness of the bucking system is highly dependent on the stability of the electronic circuitry. It depends also on the stability of the distance between the transmitter, receiver and bucking coils. Accordingly the relative orientations and positions of the transmitter, receiver and bucking coils must be kept as rigid as possible. Due to the inverse-cubed fall off in the primary field strength, changes in the transmitter-receiver separation of the order of 0.1 mm, caused for example by thermal expansion of flexure of the 8 m long structure, are significant (Huang and Fraser, 1999).

The measured receiver coil voltage is transformed to components that are in phase and at quadrature (90° out of phase) with the current in the transmitter coil. Data are presented as the ratio, in units of parts per million (ppm), of the measured secondary to theoretical primary field induced voltages. Typically a measurement or sample is output every 0.1 s (or ~3 to 7 m) along a flight line.

To give the reader an appreciation of the data that is acquired, two synthetically computed forward models responses are shown in Figure 1.5 for a six horizontal-coplanar coilset, RESOLVE frequency-domain system. Figure 1.5a shows the response for a 0.01 S/m halfspace, and Figure 1.5b shows the response for the same model but with a more conductive 2 m thick layer at surface. It can be seen how, as the skin depth calculation would suggest, the presence of the conductive surface layer has greatest effect on the data in the highest frequency coilsets. In these examples the response for a 20 frequency system was actually computed (and plotted as the solid curves) to

demonstrate the continuity of the response across the frequency spectrum. The circles represent the 12 data (6 inphase and 6 quadrature) that would actually be recorded every 0.1 s by a RESOLVE system flying over this model.

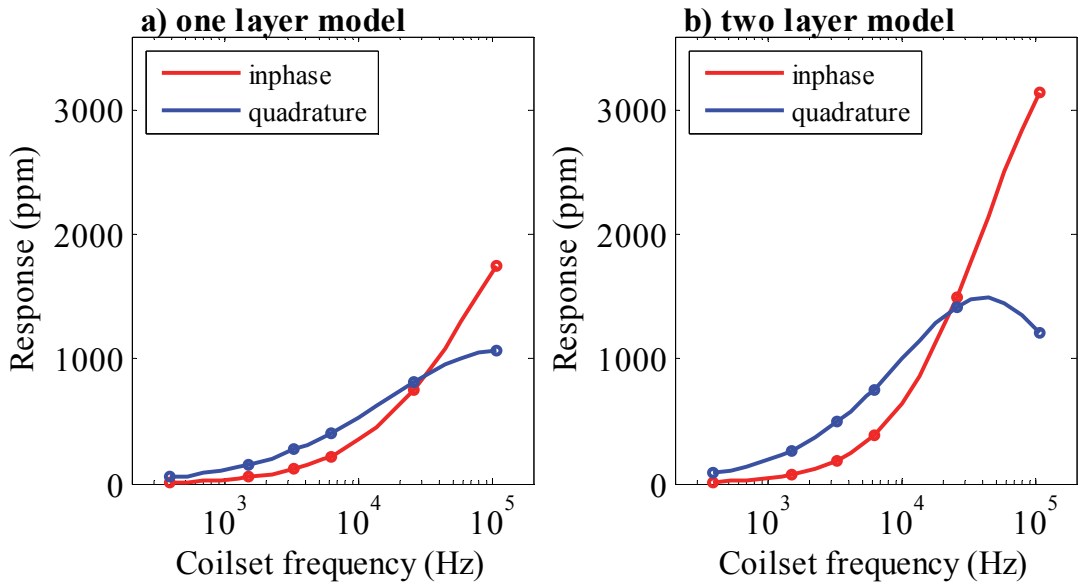

Figure 1.5 *The computed forward model response of a six frequency RESOLVE system flying at 30 m height over (a) a 0.01 S/m halfspace, and (b) a more conductive 0.1 S/m layer that is 2 m thick and overlying the same 0.01 S/m halfspace.*

1.2.6 Time domain AEM systems

In time-domain electromagnetics the general procedure is to transmit pulses of electromagnetic energy interleaved with intervals of zero transmission. An idealized representation of a possible current waveform is shown in the top panel of Figure 1.6. Current flows in the transmitter loop for a period of time (the on-time) and is then switched off and is followed by a period of nil current flow in the transmitter loop (the off-time), and then another current pulse of opposite polarity is transmitted, and so on. The rationale for this mode of transmission is that the small voltages that are induced in the receiver coil(s), due to the eddy currents dissipating after the current is switched off, can be measured in the off-time when the large primary field is not present. This is the time-domain method of addressing the primary field contamination problem that was discussed in the previous section in relation to frequency-domain systems.

Usually only one, nominally horizontal, transmitter loop is used. In fixed-wing installations the transmitter loop is slung around the extremities of the aircraft as shown in the left hand panel of Figure 1.7. For helicopter systems the transmitter loop and receiver coils are attached to some form of supporting structure and is towed below the aircraft as shown in the right hand panel of Figure 1.7.

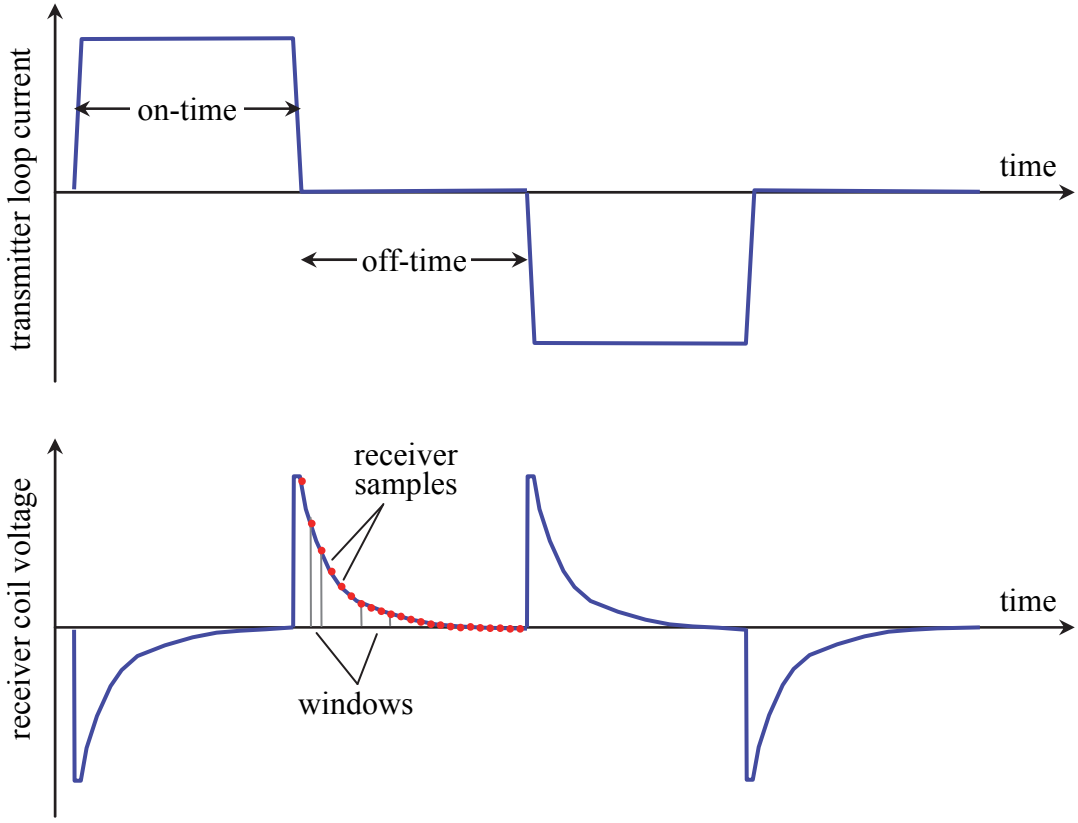

Figure 1.6 *Schematic representation of a time-domain current waveform (top) and the corresponding receiver coil voltage (bottom).*

A variety of transmitter current waveforms shapes are used, the choice of which depends on the application. The shape of the current waveform determines the amount of power that energises the subsurface conductors at various spectral frequencies. A current pulse that is slowly switched off does not generate as much power at high frequencies as a pulse that is abruptly switched off. Also a faster switch off time is achievable with smaller peak currents. In combination, these two factors mean that approximately square waveforms are typically favoured for shallow geological mapping applications, and are employed, for example, by the SKYTEM (Sorensen and Auken, 2004) and TEMPEST (Lane et al., 2000) systems. In contrast, systems that are

primarily focused on detection of deep conductors, for example GEOTEM (Annan and Lockwood, 1991) and VTEM (Witherly et al., 2004), high power is required but high frequencies are not necessary. Therefore, in these systems slower switch off waveforms (e.g. half-sine) are usually employed.

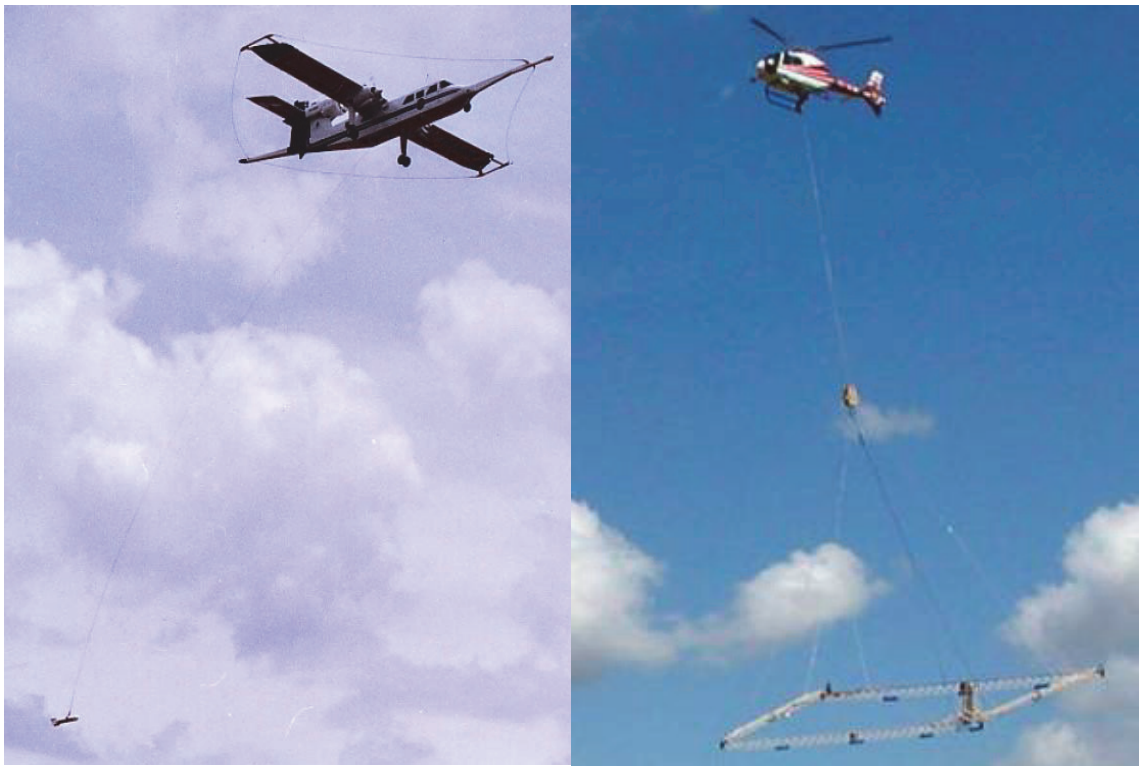

Figure 1.7 *Examples of time-domain AEM systems. The left hand panel shows an early TEMPEST system aircraft carrying a wire transmitter loop (slung between the wingtips, tail and nose), and towing the receiver bird (bottom left of photograph). The right hand panel shows the SKYTEM system in which a frame, that carries both the transmitter loop and receiver coils, is towed below a helicopter.*

The waveform is transmitted in alternating polarity pulses whose periods (including the interlacing periods of off time) are of the order of 4 to 40 ms. The so called base frequency, the reciprocal of the period, is usually in the range 25 to 250 Hz. The base-frequency is the lowest frequency that can be transmitted and detected by the system. The highest frequency that can be can be transmitted and detected by the system is dictated by both, the highest frequencies that are transmitted (the waveform shape) and the response characteristics of the receiver coils. The range from lowest to highest frequency is called the system bandwidth. Large bandwidth systems are the most

desirable as they have a greater resolving power if all other system characteristics are equal.

In fixed-wing installations the receiver coils are housed inside a shell known as the bird, which is towed behind (~120 m) and below (~40 m) the aircraft. The receiver bird can be seen in the lower left corner of Figure 1.7, attached to the end of a tow cable which is faintly visible as well. In helicopter installations the receiver coils are attached to some form of support structure that also carries the transmitter loop as shown in the right hand panel of Figure 1.7. The receivers often consist of two or three sensors, arranged in orthogonal fashion, so that directional components of the fields can be measured.

The sensors are multiple turn coils that measure the voltage induced in the receiver coil. The voltage induced in each coil is proportional to the time rate of change of the magnetic flux threading the coil (i.e. the area of the coil × times the magnetic flux density $\mathbf{B}$ in the direction of the coil's axis) and the number of coil turns. Unlike frequency-domain system receiver coils, the time-domain receiver coils are not tuned to receive at a specific frequency. Instead, the measurement involves sampling, or integrating in some cases, the voltage in the receiver coils at multiple delay times during each half-cycle (i.e. the time period relating to one pulse). The receiver sampling is schematically illustrated by the red circle in the lower panel of Figure 1.6.

The corresponding delay-time receiver samples from several time-adjacent half-cycles are stacked together into one airborne sample (after reversing the sign of data associated with negative polarity pulses). Stacking is a form of synchronous signal detection used to significantly increase the signal to noise ratio (Macnae et al., 1984). The number of half-cycles that are stacked into one output sample (the stacking length) varies from system to system. However, to demonstrate by way of example, for the TEMPEST system the stacking length is 3.04 s, because 152 of the 0.02 s half-cycles are stacked into one output sample. The output sample is calculated (i.e. drawn from the stacking

filter) for intervals spaced 0.2 s (~12.5 m) along the flight line. In other words, the data in each 12.5 m sample are effectively a filtered version of the data acquired while the aircraft travelled ~190 m (3.04 s). It is clear then that while stacking improves the signal to noise ratio, it also has the effect of reducing lateral resolution.

The delay-time samples from one stacked output sample are then binned or windowed. The binning process combines several adjacent delay-time samples from the stacked half-cycle into wider time 'windows' for further signal to noise enhancement (cf. lower panel of Figure 1.6). This may, for example, be achieved via a simple box-car shaped filter (Lane et al., 2000) or a linear-tapered shaped filter (Macnae et al., 1984). For example, the receiver in the VTEM system samples the receiver coil voltage at 96 kHz to produce 1920 delay-time samples per half-cycle spaced at 10.4166 μs intervals. However, these are binned into 26 or 30 logarithmically spaced delay-time windows using a linear-tapered averaging scheme. We might concisely summarize stacking as filtering across half-cycles and binning as filtering within a half-cycle.

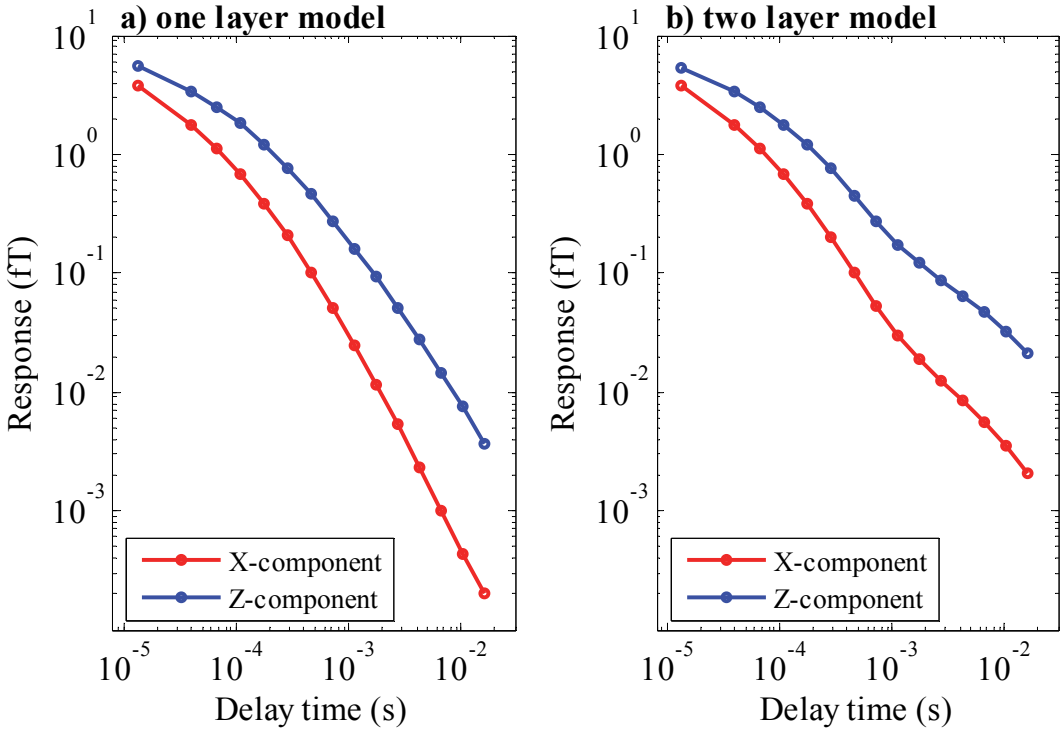

Figure 1.8 *The computed forward model response for the TEMPEST system flying over (a) a 0.01 S/m halfspace, and (b) a 400 m thick 0.01 S/m layer that is underlain by more conductive a 0.1 S/m halfspace.*

To give the reader an appreciation for the data that are presented to the interpreter, two synthetically computed forward models responses are shown in Figure 1.8 for the TEMPEST system. The system has and X- and Z-component receiver coil and the processed data are binned into 15 delay-time windows (hence the 15 circles on each response curve). Therefore each stacked and binned airborne sample is comprised of 30 data. Figure 1.8a shows the response for a 0.01 S/m halfspace. Figure 1.8b shows the response for a two layer model that has the same 0.01 S/m conductivity in the top 400 m, but with a more conductive 0.1 S/m halfspace below that. By comparing the two responses, it can be seen how the presence of the conductive layer at 400 m depth does not have significant influence on the response until a delay time of approximately 10^{-3} s. This result would be predicted by the diffusion depth approximation (Equation 1-12) presented earlier.

A short outline of time-domain systems, like that provided above, cannot exhaustively describe the diverse array of time-domain AEM systems that exist.

1.2.7 AEM system footprint

AEM systems have differing lateral resolutions, which was described by Liu and Becker (1990) as the system footprint. They defined the footprint of an AEM system as '*the side of a square surface, centred directly below the transmitter coil, that contains the induced currents which account for 90% of the observed secondary magnetic field*'. The calculations of Liu and Becker (1990) were based on an inductive limit approximation (i.e. for an infinitely conductive this sheet at the surface and infinite frequency) and it provides a minimum estimate of the true footprint size. The true footprint is partially influenced by the height, separation and orientations of the transmitter loop and receiver coils. It also depends on the conductivity of the ground and the source frequency and delay-time. More recent analysis (Reid et al., 2006) concluded that for more general 1D earths at finite frequencies, the footprint may be

several times the inductive limit approximation. Based on tables presented in Reid and Vrbancich (2004) the minimum footprints for various AEM systems are presented below in Table 1.1. Their calculations were derived in terms of the ratio between a square footprint's side length and the transmitter loop height, which is shown in the third column of the table.

Table 1.1 Minimum footprint sizes for various AEM systems.

System	Transmitter loop height (m)	Footprint side length / transmitter loop height	Minimum footprint size (m)
RESOLVE	30	3.72* / 1.34**	112* / 40.2**
GTK-AEM-05	30	1.49	45
HOISTEM	30	3.68	110
SKYTEM	30	3.68	110
VTEM	30	3.68	110
SPECTREM	95	3.93† / 2.57‡	373† / 244‡
TEMPEST	120	3.93† / 2.57‡	471† / 308‡

†Z component data: ‡X component data; *Horizontal coplanar configuration; **Vertical coaxial configuration

The effective footprint of an AEM system is further influenced by the amount of spatial filtering or stacking that is applied to the data during its processing. These operations must increase the minimum footprint size by approximately one half of the filter width.

1.3 Motivation for the research

1.3.1 The call for accurate conductivity estimates

In the search for discrete conductors, interpretation of the data often consisted of little more than classifying responses as anomalous or background (Lane, 2002). Qualitative, assessment of the amplitude and shape of anomalies plotted in profile form, and their visual comparison to pre-computed anomaly responses of elementary bodies (e.g. thin plates or spheres), has been a typical means of interpretation of isolated conductors. Palacky and West (1991) describe schemes for the interpretation of a number of elementary bodies using such techniques.

The uptake of image processing of airborne geophysical data in the late 1980's offered new possibilities for the application of AEM (Anderson et al., 1993). Image processing allowed one of the most valuable attributes of airborne methods, the spatial context of the data, to be more readily exploited. It allowed geologically distinguishing patterns to be interpreted. Until this time the signal from the regolith was considered by explorers as a nuisance that masked the anomalies of discrete conductors, and was largely ignored (Green, 1998a).

Another impetus for the increased uptake of AEM for geological mapping applications was the emergence of methods that made the transformation of AEM data into estimates of subsurface conductivity feasible for entire surveys. These were based on a method developed by Macnae and Lamontagne (1987) that allowed approximate transformation of ground electromagnetic data into conductivity estimates for a quasi-layered-earth. The method was further developed into a fast approximate inversion algorithm that could rapidly process data acquired by any AEM system (Macnae et al., 1998).

The new applications were more concerned with shallower depths of investigation and the conductivity estimates could be readily followed up with ground truthing. It was soon recognised that AEM data were not sufficiently accurate to allow the desired level of quantitative interpretation for these applications (Deszcz-Pan et al., 1998; George et al., 1998). The numerous practical difficulties in calibrating AEM systems played a large part in this problem. It was stated in a summary, discussion and future trends forum at the International Conference on Airborne Electromagnetics, Sydney, 1998, that '*Calibration issues are becoming increasingly important as we move to quantitative interpretation. Optimising current waveforms, calibrating sensors for amplitude and phase drift, for system geometry and temperature are all crucial to best practice. The potential to improve hardware is still significant, but designers must concentrate their efforts on better positioning, geometry and calibration',* (Spies et al., 1998).

The industry has to some degree responded to the challenge with the development of AEM systems that feature superior monitoring, calibration and near-surface discrimination (e.g. Lane et al., 2000; Sorensen and Auken, 2004; Fugro Airborne Surveys, 2009a). However it has been reported by Brodie et al. (2004a), Ley-Cooper and Macnae (2004), and Lane et al. (2004a) that data from the most sophisticated AEM systems are still subject to substantial calibration error.

Another reason for the difficulty in attaining accurate conductivity models from AEM data is the ambiguity or non-uniqueness of geophysical inversion. Non-uniqueness means that there may be an infinite number of models that fit the observed data equally well. Therefore, even if the data are error free, a particular model that is estimated by an inversion algorithm is unlikely to be the true model.

It is clear that in order to improve conductivity estimates from AEM data, the influence of both systematic calibration errors and non-uniqueness must somehow be reduced. In the following three subsections I will demonstrate the effect that non-uniqueness and systematic calibration errors have on conductivity estimates, and how those errors are typically addressed in conventional data processing.

1.3.2 Non-uniqueness

Non-uniqueness is the fundamental property of geophysical inverse problems meaning that, if any model can be found to fit the observed data, then an infinite number of them exist (Backus and Gilbert, 1967; Sambridge and Mosegaard, 2002). Or as described by Ellis (1998) in the context of AEM inversion, the result of non-uniqueness is that subsurface conductivity cannot be uniquely determined from a finite number of noisy data.

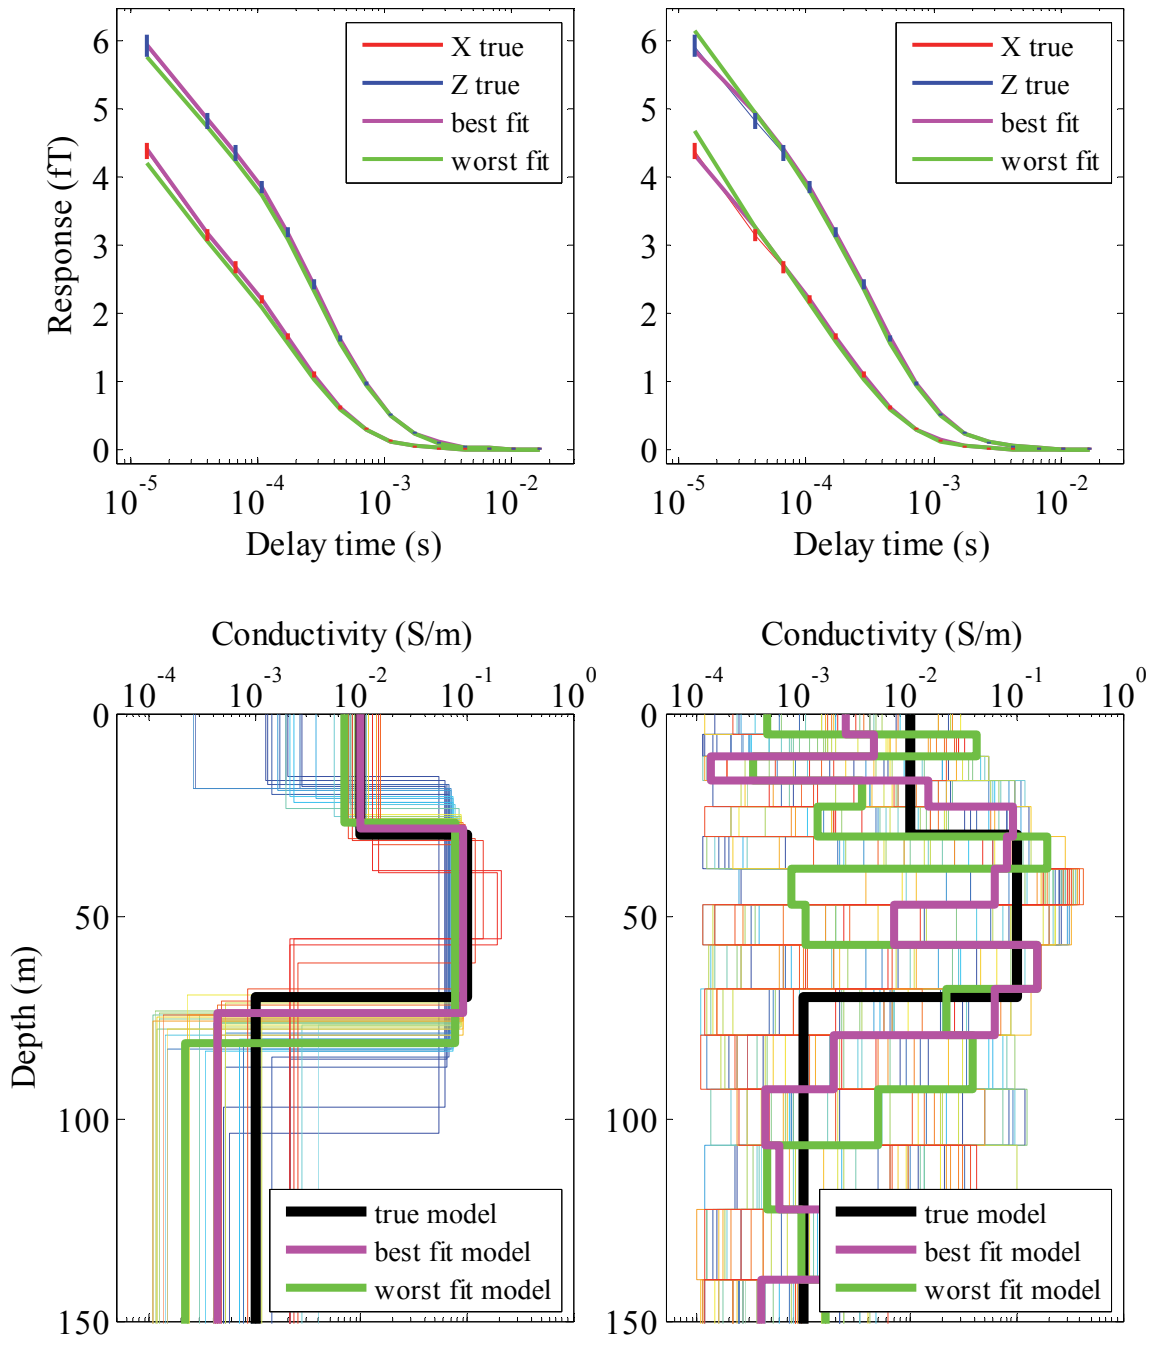

Figure 1.9 *An illustration of non-uniqueness in the inversion of AEM data: The bottom left panel shows a synthetic three layer model (thick black line) and an ensemble of three layer models (coloured lines), found using the Neighbourhood Algorithm inversion method, and whose forward model response fits the true model response satisfactorily. Plotted in the top left panel is the true model's forward response with accompanying error bars and the ensemble's best and worst fitting forward responses. The right hand panels show the same information, except that in this case an ensemble models with 15 fixed-thickness layers was generated.*

Figure 1.9 demonstrates non-uniqueness in the inversion of synthetic TEMPEST time-domain AEM system data that has been computed from a three layer conductivity model. The same three layer conductivity models is shown in the bottom two panels of

the figure as a thick black line. The forward model response of the three layer model, the synthetic X- and Z-component data, are shown in the top two panels in the red and blue curves respectively. Noise estimates that would realistically be expected in typical survey data have also been generated and these are shown as error bars on the data. The synthetic data shown in the top left and top right panels are identical because they are generated from the same three layer model.

The synthetic data have been inverted using the Monte Carlo style Neighbourhood Algorithm inversion method (Sambridge, 1999) to find an ensemble of models that all satisfy the data within the assigned noise levels. The inversion has been run twice, firstly using a three layer parameterization where both the layer conductivities and thicknesses were solved for (few-layer inversion), and again using a 15-layer parameterization where the thicknesses are kept fixed and the layer conductivities are solved for (multi-layer inversion).

The ensemble of models that satisfactorily fitted the synthetic data (i.e. whose forward response matches the synthetic data within the noise levels) from the few- and multi-layer inversions are shown as the coloured models in the bottom left and bottom right panels respectively. The best (i.e. with the lowest data misfit) and worst (i.e. with the highest acceptable data misfit) fitting models from each ensemble have been plotted as thicker magenta and green models respectively.

The forward response of the best and worst fitting models in each ensemble are shown in the top left and top right panels. It can be seen how each ensemble's worst fitting forward responses (green curves) plot on average within the error bars of the true synthetic response. This shows that all of the models in both ensembles must fit the data satisfactorily. The great variety of models that satisfactorily fit the data in these examples clearly demonstrates the non-uniqueness in the inversion of AEM data.

The problem of non-uniqueness of geophysical inversion cannot be entirely circumvented. However, the range of possible models may be narrowed down or constrained by introducing assumptions and prior information about the likely form of the plausible earth models or by adding additional independent data. Indeed, the choice to parameterize the subsurface by a certain number of discrete layers (e.g. 3 or 15 as in the example above) is itself an assumption about the likely form of plausible earth models. Another way is to restrict the set of fitting models to a more geologically reasonable set via regularization (Tikhonov and Arsenin, 1977). Regularization may be used, for example, to limit solutions to only smooth models (Constable et al., 1987), or those models that are acceptably close to a reference model constructed from prior information (Farquharson and Oldenburg, 1998). Additional data may be explicitly added to the inverse problem to further constrain the solution. For example Auken and Christiansen (2004) used estimates of the depth to a specific layer gleaned from seismic refraction data to constrain the 2D inversion of electrical resistivity data.

1.3.3 Systematic calibration errors

Errors in AEM data can usually be subdivided into 'random noise' and 'systematic error' components. Random noise is caused by largely unpredictable events such as, vibration of the transmitter-receiver assembly; atmospheric sferic events; powerline interference; and other electronic interference. It is well recognised in the geophysical industry and, to some degree, its reduction was the main focus of attempts to improve AEM systems over many years. It is extremely important in the bump finding applications because it often presents in data at the same high spatial frequencies as the anomalies of discrete conductors, and thus makes their differentiation difficult. Random error is typically zero-centred and Gaussian in nature. This type of noise can often be successfully reduced with filtering or can simply be excised from the dataset. Most geophysical inversion algorithms are designed to cater for random noise.

Systematic errors are caused by some form of miscalibration of, or imprecise knowledge of the characteristics of, the transmission and measurement system. Comprehensive analyses of the possible causes of errors in calibration of frequency-domain systems were published in Fitterman (1997) and Fitterman (1998). Some examples of systematic error in frequency-domain systems are: imperfect cancellation of the primary field; incorrect knowledge of the system gain; incorrect synchronization between the transmitter and receiver time bases; and incorrect altimeter measurements of the height of the transmitter-receiver assembly. They may be constant or vary slowly spatially and temporally, typically at scale lengths greater than the geological signal.

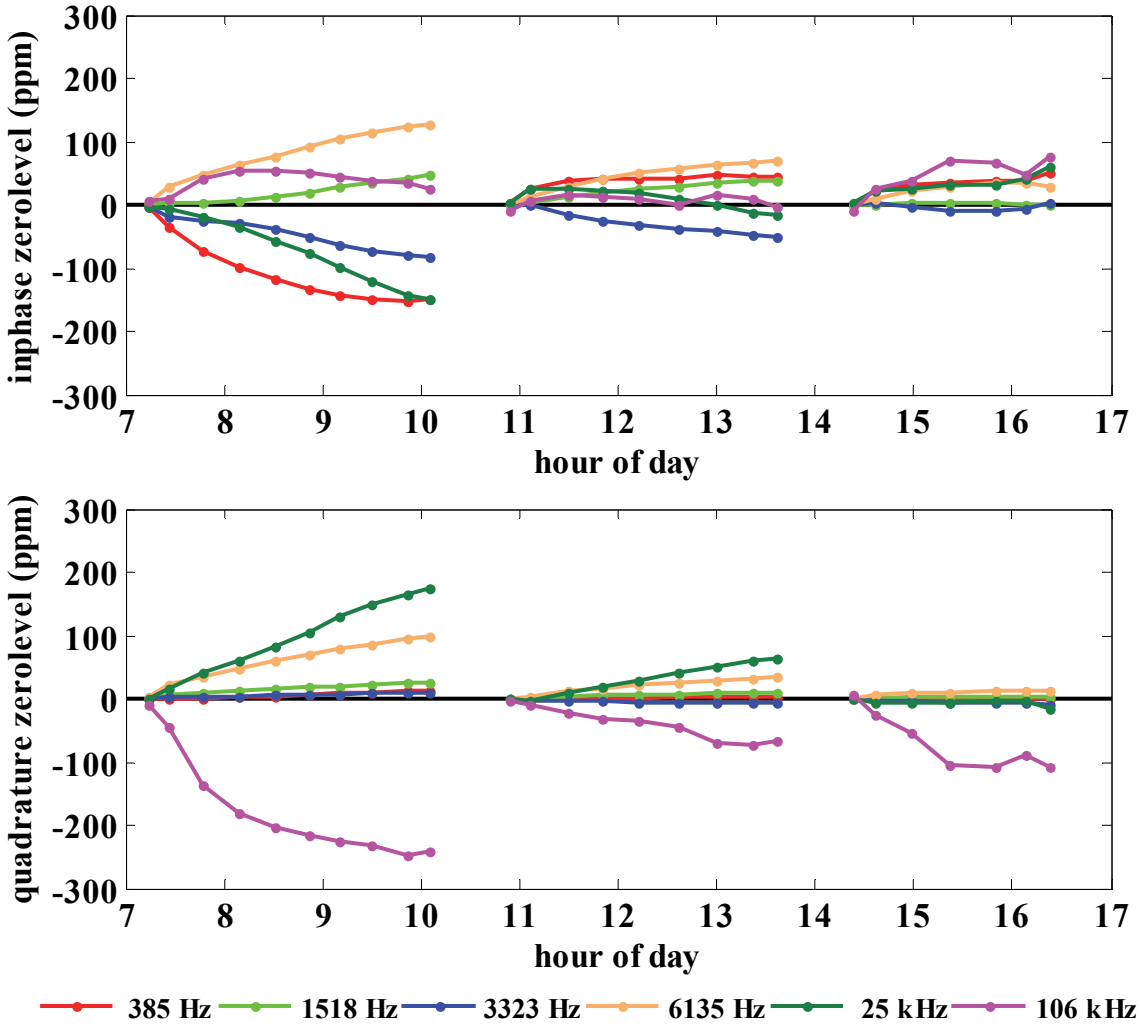

Figure 1.10 *Inphase (top) and quadrature (bottom) zero-level measurements made by the RESOLVE system during three separate flights on one day of a survey. Each individual dot on the curves represents a measurement made at ~20 minute intervals during excursions to high altitude.*

An example of one form of calibration problem that is prominent in frequency-domain systems, zero-level drift (or bias), is shown in Figure 1.10. It shows the measurements of the inphase and quadrature system response made at high altitude (>500 m) by the RESOLVE system at ~20 minutes intervals during the course of three separate flights on one day of a survey. At high altitude the response due to the conductivity of the ground is negligible (<1 ppm for a 2 S/m halfspace). However the measured inphase response is significant, implying that the primary field due to direct coupling with the transmitter coils have not been perfectly cancelled by the bucking system. Furthermore, since any primary field contamination would be expected to be purely inphase, it appears that there is some phase error due to incorrect synchronization between the transmitter and receiver time bases. Although, in this case the apparent phase errors would have to be unusually large to explain the non-zero quadrature response and they may therefore have a different cause (e.g. the zero-levels were not measured at sufficiently high enough altitude). These zero-level drift errors are usually attributed to thermal expansion and contraction of the bird that houses the transmitter and receiver coils (Huang and Fraser, 1999; Valleau, 2000). This is because any small changes in the transmitter-receiver-bucking coils separations or orientations change the effectiveness of the primary field cancellation by the bucking coils. Temperature changes also affect the stability of the electronics.

Similarly to frequency-domain AEM systems, there are several calibration problems that limit the quantitative analysis of time-domain AEM data. Some of these include; the unmeasured position and orientation of the receiver bird in fixed-wing systems (Smith, 2001a); unmeasured transmitter-receiver assembly height and tilt in helicopter systems (Davis et al., 2006); improper amplitude calibration of the measured signal (Vrbancich and Fullagar, 2007); unknown or inaccurate knowledge of the waveform (Auken et al., 2008); imperfect bucking of the primary field in the closely coupled

helicopter systems (Walker and Rudd, 2008). The first of these calibration problems is addressed in the time-domain holistic inversion method presented in Chapter 5.

1.3.4 Limitations of conventional data processing

The method of dealing with zero-level drift in conventional frequency-domain data processing is to subtract out the zero-level drift measured at high altitude from the survey altitude data (Valleau, 2000). This is a justifiable approach if it is assumed that the zero-level measured at high altitude is representative of that at survey altitude and that it drifts linearly in time between excursions to altitude. However, these assumptions may not be entirely valid due to temperature difference between survey and high altitudes or if there are rapid fluctuations in temperature (Huang and Fraser, 1999). Typically it is found that after zero-level subtraction, residual errors remain in the dataset and further processing is required.

The well-established techniques of tie line cross-over levelling (Luyendyk, 1997) and micro-levelling (Minty, 1991) that are used for dealing with systematic calibration error in airborne magnetic and gamma-ray spectrometric surveys are generally not useful for AEM survey data. This is because the underlying assumption that the response should be equal on a flight-line and a tie-line at their cross-over (intersection) is less valid for AEM because it is much more strongly dependent the height of the system above the source. Furthermore, if the data are to be quantitatively interpreted the inter-channel (i.e. between coilset frequencies/delay-times and components) integrity of the data must remain consistent with the physics of electromagnetic induction, which is not necessarily the case with these techniques. Green (2003) described a method of reducing residual drift errors by a least squares minimization of long wavelength line to line data differences. This method is prone to the same problems as the crossover and micro-levelling methods as it works on each channel individually and does not account for height differences.

Commenting on frequency-domain data processing techniques, Huang and Fraser (1999) stated, *"There is a paucity of references on airborne resistivity levelling as the subject falls within the 'tricks of the trade' of geophysical contractors"*. Later, Valleau (2000) and more recently Siemon (2009) have described iterative manual levelling techniques that are commonly used to address residual errors left after zero-level subtraction. These are based on the pseudolayer-halfspace method (Fraser, 1978). The pseudolayer-halfspace method is a lookup-table transformation of the observed data from a single coilset into apparent conductivity (σ_{app}) and apparent distance (d_{app}) parameters. The parameter σ_{app} represents the conductivity of a homogeneous halfspace situated at distance h_{app} below the transmitter-receiver bird that satisfies the data for that coilset frequency. Using the measured altimeter data (h_{alt}), the halfspace's apparent depth (d_{app}) below the true ground surface is then calculated (i.e. $d_{app} = h_{app} - h_{alt}$ is the thickness of an infinitely resistive pseudolayer between the ground surface and the top of the conducting halfspace).

The method described by Valleau (2000) is primarily guided by the aim of generating coherent apparent conductivity maps that are free of artefacts that the processor interprets to be indicative of the presence of residual error. Such artefacts are typically those features correlated with a particular acquisition entity (e.g. a flight line or complete flight). After inspecting the apparent conductivity map, the processor must decide on and then apply corrections to the inphase and quadrature, hoping that the apparent conductivity map recalculated from the corrected data will be artefact free. The procedure is iteratively applied until the processor is satisfied with the quality of the result.

Figure 1.11 shows images of the apparent conductivity and depth parameters calculated from final processed RESOLVE survey data. The survey was flown along north–south orientated flight lines. It can be seen in Figure 1.11a that the image of apparent conductivity is coherent and largely free of artefact aligned parallel with the flight lines.

However, the apparent depth image (Figure 1.11b) is not geologically plausible because of the significant striping in the flight line direction. This suggests that in levelling the conductivity the systematic calibration errors were not eliminated, but were simply 'disguised' by the apparent depth parameter. An alternate explanation might be that there was systematic altimeter error that is affecting the apparent depth parameter. However, if that was the case similar artefacts would be expected to be visible on apparent depth parameter images for the other high frequency coilsets along the same flight lines. That is not the case for this dataset.

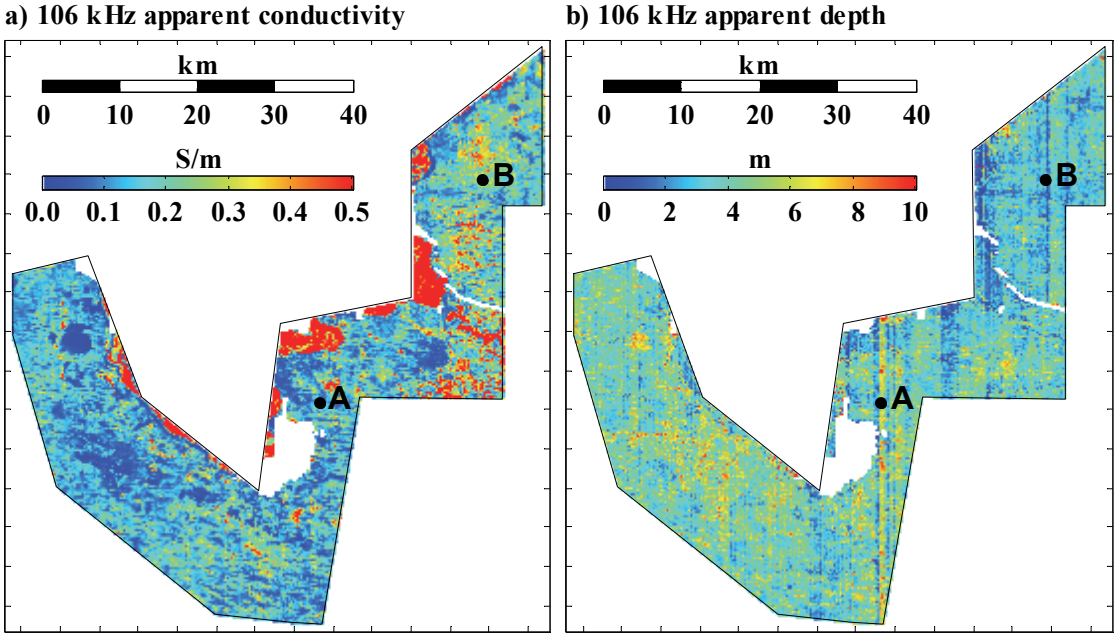

Figure 1.11 Pseudolayer-halfspace (a) apparent conductivity, and (b) apparent depth parameters calculated from final processed RESOLVE 106 kHz data.

Identification of the features that are considered to be artefacts is an inherently subjective process. Having selected a feature, the processor must choose which of several possible classes of corrections (e.g. gain, phase, bias or height) to apply. Due to the complex and nonlinear relationship between the calibration parameters and the inphase and quadrature data and thence the apparent conductivity, it is unlikely that even an experienced processor can choose the correct combination of classes and sizes of corrections to apply.

Although a degree of optimization may be achieved through several time-consuming, sequentially applied correction iterations (Valleau, 2000), the final result is far from the best that might result from a more objective and systematic approach. Also, since levelling is carried out independently for each frequency, it is impossible to identify and correct any inter-frequency inconsistencies that may exist or to ensure that additional inconsistencies are not introduced. Furthermore, sequentially applied corrections allow the propagation of errors from one processing step to the next. Huang and Fraser (1999) have noted that poor levelling decisions may generate false features and eliminate real features. In the following section I will demonstrate how these frequency-domain data processing limitations impact upon quantitative inversions, with specific attention on paid to the points on the images labelled A and B.

In the processing of time-domain AEM data, zero-level drift and manually derived levelling corrections are also applied (e.g. Carter et al., 2009). These procedures have the same limitations as the processing of frequency-domain data. A separate issue for fixed wing systems is that some elements of the system geometry are not known and must be estimated from the data during the processing. This requires assumptions to be made about the conductivity of the ground and the orientation of the receiver bird (Smith, 2001a; Lane et al., 2004a). This particular limitation is discussed in detail in Chapter 5.

1.3.5 Effect of systematic error on inversions

In contrast to random errors, systematic errors are not well handled by algorithms that estimate conductivities from AEM data. Systematic errors are not necessarily zero-centred and they tend to be correlated in time and correlated across two or more channels of data. Furthermore, unlike random noise, due to their slowly varying nature systematic errors may not be readily identifiable as errors when considered as individual channels of data. However when multiple channels of data are considered together in a

quantitative inversion, the influence that systematic errors have on the absolute numeric data values can lead to mutual inconsistency between the channels of data, the physics of electromagnetic induction, and a plausible conductivity model.

The most obvious impact of systematic calibration errors is that they will translate into systematic errors or biases in conductivity estimates derived from the data through either inversions or conductivity imaging routines. A perhaps less well appreciated impact is that mutual inconsistency between the channels of data may actually prevent the data from being fitted to any plausible conductivity model at all.

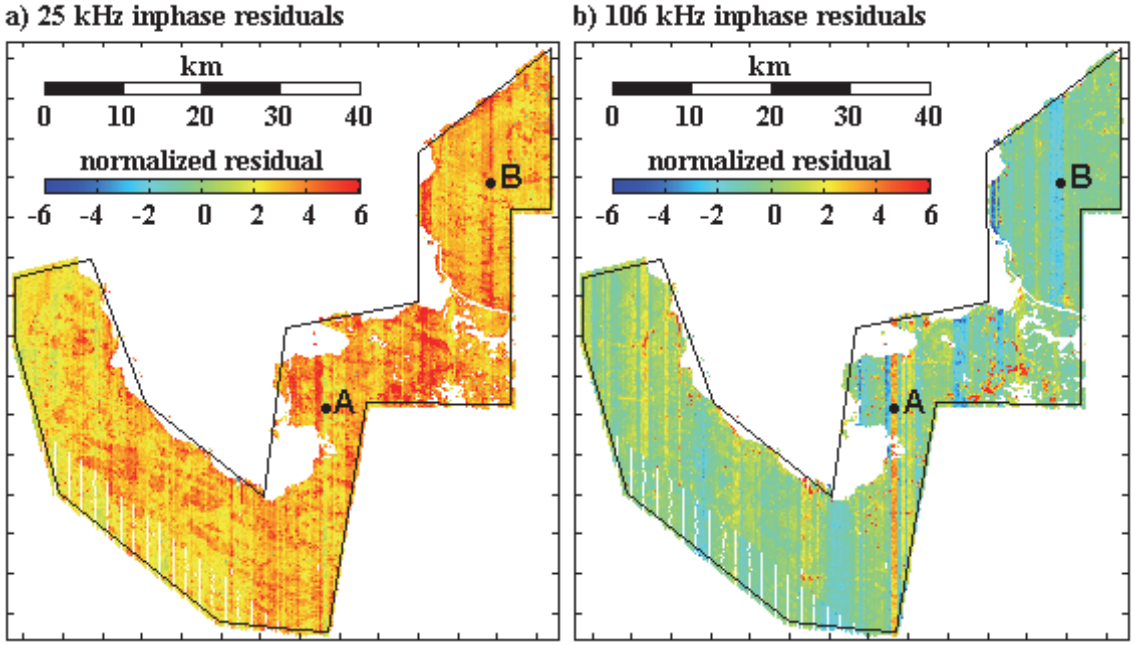

Figure 1.12 Noise-normalized residuals from the inversion of conventionally processed (a) 25 kHz inphase, and (b) 106 kHz inphase RESOLVE data.

Figure 1.12 shows an example of this for the same RESOLVE survey data as was presented in the previous section. The conventionally processed final data were inverted using the 1D sample-by-sample inversion algorithm described by Brodie et al. (2004b). They found that the inversion could not find any plausible model that satisfied the data adequately and that the AEM data were inconsistent with downhole conductivity logs. The noise-normalized residuals of the inversion (i.e. the difference between the observed data and the forward response of the best fit model normalized by the estimated noise) are shown in Figure 1.12a and b for the 25 kHz and 106 kHz coilset

inphase data respectively. It can be seen that the 25 kHz residuals are systematically large and positive.

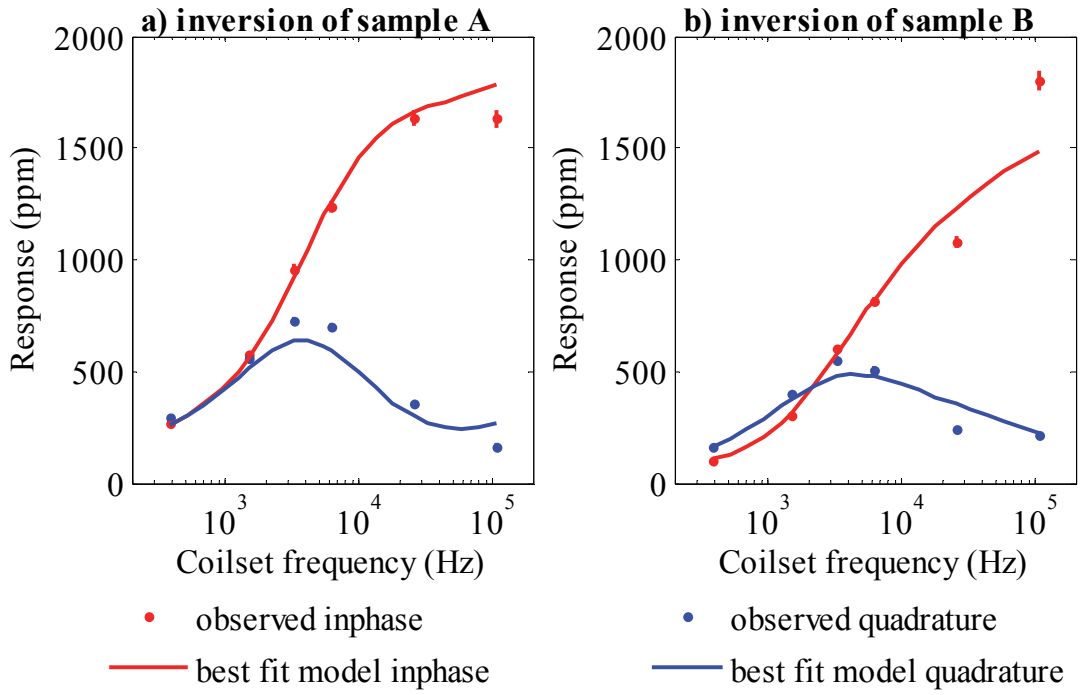

Figure 1.13 *The observed data and the best fitting inversion models for the two airborne samples labelled (a) A, and (b) B on Figure 1.11 and Figure 1.12. Note that for the purposes of this plot only, the vertical coaxial orientated coilset (3323 Hz) data have been scaled so that their amplitudes are approximately consistent with the data from other coilsets, all of which were horizontal coplanar orientated.*

The observed data and the best fitting models from the inversion of two specific airborne samples, labelled A and B on Figure 1.11 and Figure 1.12, are detailed in Figure 1.13a and Figure 1.13b respectively. Note for instance that the 106 kHz best fit model response is considerably larger than the observed data for Sample A (Figure 1.13a) but the opposite is the case for Sample B (Figure 1.13b). Inspection of the 106 kHz inphase inversion residuals in Figure 1.12b shows that this was more or less the case for the entire duration of both flights. It is clear that both sets of inversion residuals spatially correlated with particular flights or individual north–south flight lines.

There is also a strong correlation between the inversion residuals and the artefacts shown in the 106 kHz apparent depth parameter images shown in Figure 1.11b.

Together, all these characteristics suggest some form of residual or introduced systematic calibration error is present in the final processed dataset. If there were no systematic calibration errors present, we would expect the inversion residuals to be much more randomly distributed and zero-centred. Brodie et al. (2004b) found that to fit these data to plausible layered conductivity models, it was necessary to rescale the data by factors derived by Brodie et al. (2004a) from a regression of the data against forward models of downhole log data.

Deszcz-Pan et al. (1998) had previously recognized that systematic calibration error was the cause of unacceptably large misfits in the inversion of DIGHEM data. Using an independent method, Ley-Cooper and Macnae (2004) and Ley-Cooper et al. (2006) also found compelling evidence of systematic scaling errors in RESOLVE and DIGHEM datasets. Vrbancich and Fullagar (2007) also found that data scaling errors in time-domain HOISTEM data prevented the data from being fitted using plausible conductivity models. Lane et al. (2004a) demonstrated how the problem of primary field removal and system geometry estimation in fixed-wing TEMPEST data led to systematic biases in conductivity estimates derived from conductivity depth imaging algorithms.

1.4 The holistic inversion framework

In the previous section I discussed the motivation for this research which stems from the increasing demand by AEM survey stakeholders for more accurate conductivity estimates. I demonstrated how non-uniqueness, systematic calibration errors, and the limitations of how they are dealt with in conventional data processing, all impact negatively on the outcome of interpretation of AEM data.

In this section I will outline the proposed strategy for addressing these problems. The overarching proposition is that entire airborne dataset be inverted, along with prior information, to solve for a continuous 3D conductivity model and a calibration model

that mathematically parameterizes systematic error in the data. The strategy is dubbed a 'holistic approach' because, in effect, it simultaneously calibrates, processes and inverts entire datasets.

The motivation behind combining calibration, processing and inversion into one step is three-fold. Firstly, it circumvents the all-too-common problem when data are finally presented to an inversion algorithm it is found that they cannot be fitted because of systematic calibration errors and/or assumptions made during the data processing. Secondly, it is a mechanism of ensuring that the assumptions made in the calibration, processing, and inversion plus any prior information are all, at a minimum, mutually consistent. Finally, it aims to avoid the time consuming and expensive circularly iterative paradigm of calibration→processing→inversion→validation→recalibration.

The inversion model is comprised of a continuous 3D conductivity model and a calibration model. The role of the calibration model is to mathematically describe systematic calibration errors. The idea is to simulate calibration errors in the forward model response calculations of the inversion, rather than to remove them per se. In doing this calibration errors are accounted for, and thus there is no need to separately calibrate and process the data. The formulation of the calibration model is based on a sound physical understanding the sources of calibration error (e.g. Fitterman, 1998). It is an objective means of accounting for calibration error that was motivated by the work of (Deszcz-Pan et al., 1998).

The rationale for simultaneously inverting an entire dataset, rather than the conventional approach of independently inverting each airborne sample, is two-fold. Firstly, it permits the spatial coherency of the geology be capitalized upon by allowing spatial constraints to be placed on the conductivity model in all horizontal directions. This concept was first published as part of this research (Brodie and Sambridge, 2004; Brodie and Sambridge, 2006), and has since been successfully used by Viezzoli et

al. (2008). It is a natural extension of previous methods that have gained advantage by placing vertical (e.g. Constable et al., 1987) and along-line lateral constraints (e.g. Gyulai and Ormos, 1999; Auken et al., 2005) on the resultant conductivity model. Secondly, simultaneous inversion of an entire dataset has added advantage in the holistic approach because the temporal coherency of systematic calibration error (e.g. smooth zero-level drift throughout a flight shown in Figure 1.10) can also be capitalized upon. This would not be possible if only one sample or one flight line were being inverted. Overall then, the inversion of an entire dataset provides the opportunity to apply additional constraint, thereby helping to reduce difficulties associated with non-uniqueness.

A single continuously defined 3D model is used to parameterize the subsurface conductivity. Local 1D models are extracted from the 3D model for the purpose of the electromagnetic forward model calculations. This is a fundamental rearrangement of the conventional 1D inversion notion of parameterizing the subsurface with many discrete 1D models that are subsequently stitched together to form a quasi-3D model. The 3D parameterization is used because it conceptualises the conductivity structure in precisely the way one envisages a layered-earth to be.

It also provides a natural mechanism for incorporating independent prior information into the inversion which helps further reduce the influence of non-uniqueness. This is because the conductivity model is continuously defined (i.e. even between flight lines), thereby allowing prior information (e.g. downhole logs, ground electromagnetic soundings) that in general does not lie on the airborne survey flight lines to be used without any need for extrapolation. To help clarify this point, consider for example a borehole that lies partway between two flight lines. If a conventional discrete parameterization is used one must choose to either; use the borehole information to constrain the discrete 1D model associated with the nearest airborne sample; or to have it constrain multiple discrete 1D models using some ad hoc distance threshold or

weighting. However, using a continuous model the borehole information can simply and naturally constrain the conductivity model where the borehole data was actually observed.

The holistic approach therefore provides a framework for considering all of the available information together in a single self-consistent treatment of the data.

1.5 Outline of the thesis

The remainder of the thesis is set out as follows. Chapter 2 describes the details of the electromagnetic 1D forward modelling routine that has been used in the research. Theoretical material on the quasi-static solution of Maxwell's equations is provided, leading to the magnetic field expressions for a dipole source above a layered-earth. It then shows how these expressions are used to build the complete AEM system response. It covers the analytic partial derivatives of the magnetic field expressions required in the holistic inversion, which were derived and are provided in appendices.

Chapter 3 is a major chapter of the thesis. It sets out the formal theoretical development of the holistic inversion technique for frequency-domain AEM data. Material covered includes: the conductivity and calibration model parameterizations; the three different input data types; the calibration equations; and the forward modelling expressions for each data type and their partial derivatives. A description of the inversion minimization scheme concludes the chapter.

Chapter 4 presents three different applications of the frequency-domain holistic inversion. The first is a small synthetic example to test that the theory has been implemented properly. The second example is a few-layer model inversion of real data including substantial prior information. The final example is a multi-layer inversion of an entire dataset which requires parallel computations. Here I use only minimal prior

information to test if the method can be applied in 'greenfields' areas where prior constraints may not be available.

Chapter 5 presents a holistic inversion framework for data acquired by fixed-wing time-domain AEM systems. It discusses the calibration problems associated with these systems relating to primary field and system geometry estimation. It sets out the formal theoretical development for the holistic inversion based on a 2D (rather than 3D) conductivity model. Material covered includes: the conductivity and calibration model parameterizations; the input data; the forward modelling expressions and their partial derivatives; and the inversion minimization scheme. A real data example is then presented. The chapter concludes with a discussion on how the method might be applied to other time-domain platforms.

Chapter 6 contains a summary of the material covered, along with accompanying conclusions, and a discussion of likely future directions for the holistic approach.

Chapter 2

Layered-earth forward modelling

2.1 Outline

This chapter covers the forward modelling of the response of an AEM system positioned above a layered-earth. A layered-earth, otherwise known as a 1D earth, is an idealised model of the real earth in which we consider it to consist of several electrically homogeneous isotropic horizontal layers that are stacked on top of each other.

This simplification of the real earth is used because it allows analytic closed form solutions to Maxwell's equations to be derived. Layered-earths allow relatively rapid forward modelling and inversion compared to the more general 2D or 3D earth models, which typically require numerical solutions. Layered-earth models are good approximations of the real earth for modelling in stratified geological environments where the scale length of lateral variation is greater than the AEM system's footprint (Liu and Becker, 1990; Reid et al., 2006). A further argument for the use of 1D models is that the electromagnetic data are inherently ambiguous even when 1D models are used. Since 2D and 3D conductivity distributions are by definition more complex, then inverting for such distributions must therefore be even more ambiguous.

Amongst others, Wait (1982), Ward and Hohmann (1988), and Frischknecht (1967), have derived equivalent analytic expressions for the resultant magnetic field above a layered-earth in the presence of infinitesimal vertical and horizontal magnetic dipole sources. Solutions of this type form the basis of layered-earth modelling in airborne electromagnetics.

The expressions provided in this chapter are based on the conventions and derivation developed by Wait (1982). The derivation is rather long and complicated so it will not be reproduced here, however a brief outline of the derivation is provided in Section 2.5.1. The relevant final expressions, summarised on page 113 of Wait's book, are reproduced in Section 2.5.3. Note that the terms have been rearranged into the more convenient matrix notation provided by Fitterman and Yin (2004).

A complete AEM system response is modelled by: (1) representing the current in the actual airborne system's transmitter loop(s) with an equivalent assembly of infinitesimal magnetic dipoles whose strength oscillate at discrete frequencies; (2) computing each of their resultant magnetic fields at equivalent infinitesimal receiver loops; (3) then combining the results of all dipoles appropriately to replicate the actual measured system response.

Treating the finite sized transmitter loop(s) as a magnetic dipole source(s) and finite sized receiver coil(s) as infinitesimal is quite adequate if the height of the transmitter loop above the ground and the horizontal distance between transmitter loop and receiver coils is approximately three times greater than the dimensions of the transmitter loop and receiver coils. This is the case in all the airborne systems considered herein.

The necessary number of infinitesimal dipoles that are required depends on the style of the airborne system. A vertical and a horizontal transmitter dipole are required to represent arbitrarily orientated transmitter loops for each necessary frequency. Then, in frequency-domain systems, transmitter dipoles are required for each of the coilset transmitting frequencies. For time-domain systems, the magnetic field computations are carried out in the frequency-domain at multiple discrete frequencies (e.g., 25, or 5 per decade of frequency) and are then transformed to the time-domain after multiplications with the frequency spectrum of the transmitter loop's current waveform.

In all commercially available airborne electromagnetic systems the receiver sensors do not actually measure the magnetic field itself. Rather they sense voltages induced in receiver induction coils due to the negative time rate of change of the magnetic field flux threading the coil. Notwithstanding this, all the fundamental calculations can be carried out in the 'magnetic field domain' and then analytically transformed where necessary to required voltages or ratios of secondary to primary field induced voltages.

In this chapter the coordinate system is defined, a convention for defining the orientation of the transmitter loop and receiver coils is provided, and the geoelectric properties of the layers are defined. The expressions for the magnetic fields due to an infinitesimal dipole source are then provided, followed by the procedure for combining those results into a complete system response. Partial derivatives of the forward model that have been analytically derived for use in the inversion algorithms are discussed, however these derivatives are shown in full in Appendix B.

2.2 Coordinate system

Figure 2.1 schematically depicts a layered-earth model as well as some aspects of the coordinate system that is used in this chapter. The coordinate system in which the forward modelling expressions are defined is a right handed Cartesian coordinate system. The positive x-axis is horizontal and in the direction of flight of the aircraft, the positive y-axis is directed toward the left hand side of the aircraft's flight and the positive z-axis is directed vertically upwards. Following the convention of Fitterman and Yin (2004), this is called the inertial coordinate system. The origin is directly below the transmitter on the surface of the earth. The centre of the transmitter loop is at height h above the earth and is thus located at the coordinate $(0,0,h)$. The receiver coils' assembly is centred at the coordinate (x,y,z). To simplify the magnetic field expressions presented later, it is convenient here to define two purely geometric scalar quantities,

$$r = \sqrt{x^2 + y^2}\,, \text{ and} \tag{2-1}$$

$$R = \sqrt{x^2 + y^2 + (z-h)^2}\,, \tag{2-2}$$

which respectively represent the horizontal and total distances between transmitter loop and the receiver coils' assembly.

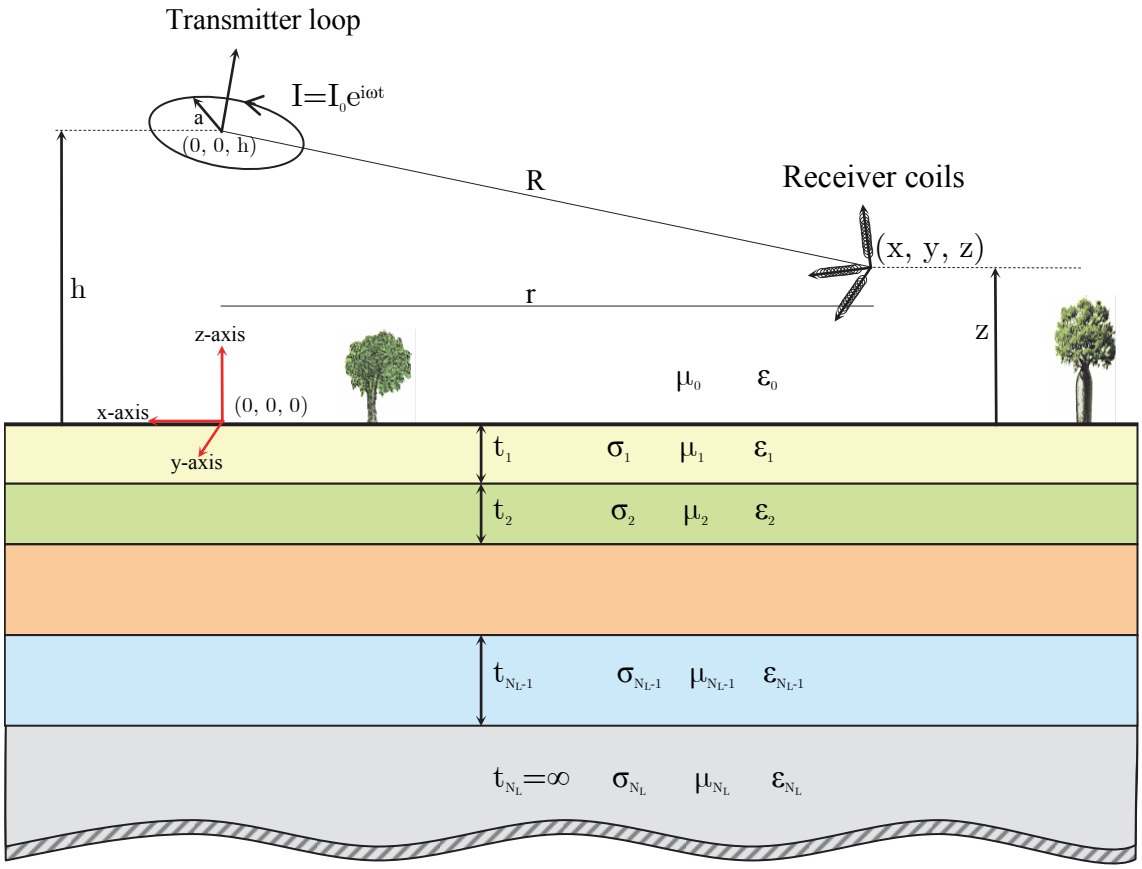

Figure 2.1 Schematic representation of a magnetic dipole transmitter and a three axis receiver above an N_L-layered-earth.

2.3 Orientation of transmitter loop and receiver coils

In general, the transmitter and receiver loops in an AEM system will be in some arbitrary orientation due to the flight manoeuvres that the aircraft, towed bird or towed transmitter/receiver assembly undergo during flight. The orientation may be defined by the unit normal vector to the plane of the loop, or the unit vector parallel to the coil's axis. Let us say its reference orientation $\mathbf{v}_0$ is its orientation when in straight and level flight. For example, we would say that the transmitter loop on a fixed wing towed bird

system has a reference orientation equal to the unit vertical vector in the inertial coordinate system $\mathbf{v}_0 = [0,0,1]^T$. Similarly, we would define the reference orientation of a vertical coaxial loop in a frequency-domain system as $\mathbf{v}_0 = [1,0,0]^T$.

Typically, the orientation of a transmitter loop or receiver coil is expressed as roll, pitch and yaw rotation angles with respect to its reference orientation. Fitterman and Yin (2004) provide a compact and convenient treatment of the subject of roll, pitch and yaw, on which the convention in this thesis is based. Details on the subject are provided in Appendix A, however the basic equations follow. Roll, (ψ_R), pitch (ψ_P) and yaw (ψ_Y) are defined as counter-clockwise rotations about the inertial coordinate system's z-axis, y-axis and x-axis respectively, for an observer looking toward the origin. After a loop or coil undergoes successive yaw then pitch then roll manoeuvres it will have a general orientation $\mathbf{v}$. The two vectors $\mathbf{v}_0$ and $\mathbf{v}$ are related through the orthogonal rotation matrix equations,

$$\mathbf{v}_0 = \mathbf{R}\mathbf{v}, \text{ and} \tag{2-3}$$

$$\mathbf{v} = \mathbf{R}^T\mathbf{v}_0, \tag{2-4}$$

where $\mathbf{R}$ is the rotation matrix that defines the orientation of the loop or coil given by,

$$\mathbf{R} = \begin{bmatrix} \mathrm{c}\psi_P\,\mathrm{c}\psi_Y & \mathrm{c}\psi_P\,\mathrm{s}\psi_Y & -\mathrm{s}\psi_P \\ \mathrm{s}\psi_R\,\mathrm{s}\psi_P\,\mathrm{c}\psi_Y - \mathrm{c}\psi_R\,\mathrm{s}\psi_Y & \mathrm{s}\psi_R\,\mathrm{s}\psi_P\,\mathrm{s}\psi_Y + \mathrm{c}\psi_R\,\mathrm{c}\psi_Y & \mathrm{s}\psi_R\,\mathrm{c}\psi_P \\ \mathrm{c}\psi_R\,\mathrm{s}\psi_P\,\mathrm{c}\psi_Y + \mathrm{s}\psi_R\,\mathrm{s}\psi_Y & \mathrm{c}\psi_R\,\mathrm{s}\psi_P\,\mathrm{s}\psi_Y - \mathrm{s}\psi_R\,\mathrm{c}\psi_Y & \mathrm{c}\psi_R\,\mathrm{c}\psi_P \end{bmatrix}, \tag{2-5}$$

using the shorthand $\cos\psi = \mathrm{c}\psi$ and $\sin\psi = s\psi$.

The orientations of the loops are not measured in all AEM systems. For such systems, the loops must be assumed to be in some reference orientation for modelling purposes. In some fixed-wing systems the orientation of the transmitter loop is measured via an inertial gyroscope mounted on the aircraft that tracks the loop's roll, pitch and yaw rotations with respect to some reference orientation (Lane et al., 2000). In some

helicopter systems the orientation of the towed transmitter/receiver assembly is measured by inclinometers (Sorensen and Auken, 2004) or with three GPS units mounted on the towed assembly (Fitzpatrick and Munday, 2007).

2.4 Geoelectric properties of the layered-earth

Layered-earth models have N_L discrete horizontal layers stacked vertically on top of each other. Within any one layer the electrical properties are uniform and isotropic. The k^{th} layer from the top is assigned a thickness t_k, electrical conductivity σ_k, a magnetic permeability μ_k and dielectric permittivity ε_k. The bottom layer, also called the basement or halfspace layer, is infinitely thick so that $t_{N_L} = \infty$. For later convenience we define a vector containing all of the layer thicknesses $\mathbf{t} = \left[t_1, t_2, \cdots, t_{N_L-1}\right]^{\mathbf{T}}$. In an analogous manner we define the vectors $\boldsymbol{\sigma} = \left[\sigma_1, \sigma_2, \cdots, \sigma_{N_L}\right]^{\mathbf{T}}$, $\boldsymbol{\mu} = \left[\mu_1, \mu_2, \cdots, \mu_{N_L-1}\right]^{\mathbf{T}}$ and $\boldsymbol{\varepsilon} = \left[\varepsilon_1, \varepsilon_2, \cdots, \varepsilon_{N_L-1}\right]^{\mathbf{T}}$.

2.5 Magnetic field expressions for an infinitesimal dipole source

2.5.1 Outline of the solution of the Maxwell's equations

I begin here by showing the derivation of the frequency-domain Helmholtz wave equations in $\mathbf{E}$ and $\mathbf{H}$ (Equations 1-9 and 1-10) that were given in Chapter 1. They are derived from the time-domain Maxwell's Equations (Equations 1-2 to 1-5) and the constitutive relations (Equations 1-6 to 1-8). We consider sources whose moment varies harmonically in time t with angular frequency ω in the form $m = m_0 e^{i\omega t}$ so that $\partial/\partial t = i\omega$ ($i = \sqrt{-1}$). Following Ward and Hohmann (1988), using uppercase letters for the frequency-domain vector quantities, and making use of the constitutive relations, we may rewrite Equation 1-2 in the frequency-domain as,

$$\nabla \times \mathbf{E} + \frac{\partial \mathbf{B}}{\partial t} = \nabla \times \mathbf{E} + i\mu\omega\mathbf{H} = 0\,, \tag{2-6}$$

and Equation 1-3 as,

$$\nabla \times \mathbf{H} - \frac{\partial \mathbf{D}}{\partial t} - \mathbf{J} = \nabla \times \mathbf{H} - i\omega\varepsilon\mathbf{E} - \sigma\mathbf{E} = \nabla \times \mathbf{H} - (\sigma + i\varepsilon\omega)\mathbf{E} = 0\,. \qquad (2\text{-}7)$$

Then taking the curl ($\nabla\times$) of Equations 2-6 and making use of the vector identity $\nabla \times \nabla \times \mathbf{A} = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2\mathbf{A}$, we attain,

$$\nabla \times [\nabla \times \mathbf{E} + i\mu\omega\mathbf{H}] = \nabla(\nabla \cdot \mathbf{E}) - \nabla^2\mathbf{E} + i\mu\omega\nabla \times \mathbf{H} = 0\,, \qquad (2\text{-}8)$$

and similarly the curl of Equation 2-7 is given by,

$$\nabla \times [\nabla \times \mathbf{H} - (\sigma + i\varepsilon\omega\mathbf{E})] = \nabla(\nabla \cdot \mathbf{H}) - \nabla^2\mathbf{H} - \nabla \times (\sigma + i\varepsilon\omega)\mathbf{E} = 0\,. \qquad (2\text{-}9)$$

For homogenous regions $\nabla \cdot \mathbf{H} = 0$ because $\nabla \cdot \mathbf{B} = \nabla \cdot \mu\mathbf{H} = \mu\nabla \cdot \mathbf{H} = 0$ in accordance with Equations 1-4 and 1-6.

A further simplification can be made for conductive earth materials ($\geq 10^{-4}$ S/m) because free electric charge in homogenous regions dissipates in $< 10^{-6}$ s (Ward and Hohmann, 1988). So for the frequencies used in AEM ($\leq 10^5$ Hz) we can say $\nabla \cdot \mathbf{E} = 0$ because $\nabla \cdot \mathbf{D} = \nabla \cdot \varepsilon\mathbf{E} = \varepsilon\nabla \cdot \mathbf{E} = \rho \sim 0$ in accordance with Equations 1-5 and 1-7. Therefore substituting $\nabla \times \mathbf{H} = (\sigma + i\varepsilon\omega)\mathbf{E}$ from Equation 2-7 into Equation 2-8, and $\nabla \times \mathbf{E} = -i\mu\varpi\mathbf{H}$ from Equation 2-6 into Equation 2-9, yields the frequency-domain wave equations,

$$\nabla^2\mathbf{E} + (\mu\varepsilon\omega^2 - i\mu\sigma\omega)\mathbf{E} = 0\,, \qquad (2\text{-}10)$$

and,

$$\nabla^2\mathbf{H} + (\mu\varepsilon\omega^2 - i\mu\sigma\omega)\mathbf{H} = 0\,. \qquad (2\text{-}11)$$

These are the Helmholtz wave equations in $\mathbf{E}$ and $\mathbf{H}$ (Equation 1-9 and 1-10) which were given in Section 1.2.4. These equations describe the propagation of electromagnetic energy with propagation constant,

$$k = \sqrt{\mu\varepsilon\omega^2 - i\mu\sigma\omega}\,. \qquad (2\text{-}12)$$

For low frequencies ($\leq 10^5$ Hz) and real earth materials it is usually true that $\mu\varepsilon\omega^2 << \mu\sigma\omega$. For the frequencies employed in airborne geophysics we usually make this assumption, called the quasi-static assumption, and rewrite Equation 2-12 as,

$$k^2 = -i\mu\sigma\omega\,. \tag{2-13}$$

Physically, the quasi-static assumption means that conduction currents dominate over displacement currents. It also means that the Helmholtz equations reduce to diffusion equations and thus diffusion is the dominant process, which results in the lack of resolution in electromagnetic methods (Ward and Hohmann, 1988). The magnetic field expressions shown herein were derived by Wait under this quasi-static assumption.

For a vertical magnetic dipole, working in circular cylindrical coordinates, Wait (1982) defined a scalar magnetic Hertz potential from which the electric and magnetic field are derivable. It turns out that the scalar Hertz potential also satisfies the scalar Helmholtz equation and is more readily solved than solving for the magnetic fields directly. In the air (the region containing the source transmitter) the general solution for the Hertz potential is found to be,

$$F_0 = \frac{Ia}{2}\int_0^\infty \frac{J_1(\lambda a)}{\lambda}\left[e^{-\lambda|z-h|} + \mathsf{R}_0 e^{-\lambda(z+h)}\right]\, J_0(\lambda r)d\lambda\,, \tag{2-14}$$

and in the kth homogenous layer it is found to be,

$$F_k = \frac{Ia}{2}\int_0^\infty \frac{J_1(\lambda a)}{\lambda}\left[A_k e^{u_k z} + B_k e^{-u_k z}\right]\, J_0(\lambda r)d\lambda\,, \tag{2-15}$$

where a is the radius of the infinitesimal current loop carrying current I, $u_k = \sqrt{\lambda^2 - k_k^2} = \sqrt{\lambda^2 + i\mu_k\sigma_k\omega}$, J_0 and J_1 are Bessel functions of order 0 and 1 respectively and λ is the integration variable. The terms $\mathsf{R}_0(\boldsymbol{\sigma},\boldsymbol{\mu},\mathbf{t},\omega,\lambda)$, $A_k(\boldsymbol{\sigma},\boldsymbol{\mu},\mathbf{t},\omega,\lambda)$ and $B_k(\boldsymbol{\sigma},\boldsymbol{\mu},\mathbf{t},\omega,\lambda)$ are coefficients of up-going and down-going wave fields in each region that must be found through the application of boundary

conditions. The suitable boundary conditions insist that the normal magnetic flux density and tangential magnetic fields are continuous at the surface and layer interfaces. It is also conditional that the coefficient of the down-going wave field in the bottom layer must be zero ($B_{N_L} = 0$) since the fields must vanish as $z \to -\infty$. Application of these conditions yields a set of equations from which the coefficients can be successively eliminated until the required coefficient R_0 is determined. A recursive procedure for determining R_0 and an alternative propagation matrix method are shown in Section 2.5.2 following.

According to (Ward and Hohmann (1988)) the solution is in the form of a superposition of plane waves with continuously varying complex angles of incidence. They also refer to the term R_0 as a reflection coefficient. To find the magnetic fields in the air, as is required for airborne modelling, it is then a matter of substituting R_0 into Equation 2-14 and using the definition of the Hertz potential to derive expressions for the magnetic field components at the receiver.

The procedure for a horizontal dipole is more complicated due to lack of cylindrical symmetry but follows a similar path. The final resulting expressions for both vertical and horizontal dipoles are provided in Section 2.5.3.

2.5.2 Complex reflection coefficient

Recursive method

Wait (1982) derived a recursive procedure for determining the complex reflection coefficient R_0 that is required for evaluation of the magnetic fields. He found that,

$$\mathsf{R}_0 = \frac{N_0 - Y_1}{N_0 + Y_1}, \tag{2-16}$$

where $N_0 = \lambda / i\mu_0\omega$ and Y_1 is the admittance at the surface and is yet to be determined. The admittance at the interface between the k^{th} and $(k+1)^{th}$ layers is related through the expression,

$$Y_k = N_k \frac{Y_{k+1} + N_k \tanh(u_k t_k)}{N_k + Y_{k+1} \tanh(u_k t_k)}, \tag{2-17}$$

where,

$$N_k = \frac{u_k}{i\mu_k\omega}, \text{ and} \tag{2-18}$$

$$u_k = \sqrt{(\lambda^2 + i\mu_k\sigma_k\omega)}. \tag{2-19}$$

In the basement halfspace $t_{N_L} = \infty$ and $\tanh(u_{N_L} t_{N_L}) \to 1$, Equation 2-17 reduces to $Y_{N_L} = N_{N_L}$ which is a known quantity as it can be calculated from Equation 2-18. Thus we have the basis for a recursive algorithm in which the unknown Y_1 is determined recursively by computing the Y_k's beginning at the bottom most interface (Y_{N_L}), and then using Equation 2-17 successively compute $Y_{N_L-1}, Y_{N_L-2}, \cdots, Y_2, Y_1$. The R_0 term, required in the Hankel transforms, can then be calculated from Equation 2-16.

Propagation matrix method

Farquharson et al. (2003) and Farquharson (2000) have described a convenient method of computing the term P_{21}/P_{11}, that they use in their magnetic field expressions, and is identical to the complex reflections R_0 term used herein. They call this the propagation matrix method. The procedure is to define individual matrices for each layer, which for the *k*th layer is,

$$\mathbf{M}_k = \frac{1}{2}\begin{bmatrix} \left(1 + \frac{\mu_{k-1}u_k}{\mu_k u_{k-1}}\right) & \left(1 - \frac{\mu_{k-1}u_k}{\mu_k u_{k-1}}\right) \\ \left(1 - \frac{\mu_{k-1}u_k}{\mu_k u_{k-1}}\right) e^{-2u_{k-1}t_{k-1}} & \left(1 + \frac{\mu_{k-1}u_k}{\mu_k u_{k-1}}\right) e^{-2u_{k-1}t_{k-1}} \end{bmatrix}, \; k = 2, 3, \cdots, N_L. \tag{2-20}$$

For the special case $k = 1$ the layer matrix is simplified to,

$$\mathbf{M}_1 = \frac{1}{2}\begin{bmatrix} \left(1 + \frac{\mu_0 u_1}{\mu_1 u_0}\right) & \left(1 - \frac{\mu_0 u_1}{\mu_1 u_0}\right) \\ \left(1 - \frac{\mu_0 u_1}{\mu_1 u_0}\right) & \left(1 + \frac{\mu_0 u_1}{\mu_1 u_0}\right) \end{bmatrix}. \tag{2-21}$$

The composite propagation matrix, the product of all the individual layer matrices, is then given by,

$$\mathbf{P} = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} = \mathbf{M}_1 \prod_{k=2}^{N_L} \mathbf{M}_k \,, \tag{2-22}$$

and the complex reflection coefficient R_0 is,

$$\mathsf{R}_0 = \frac{P_{21}}{P_{11}}. \tag{2-23}$$

The propagation matrix method is more computationally efficient, and convenient in a programming sense, than the recursive method when partial derivatives of the reflection coefficient with respect to properties of individual layers are required in an inversion algorithm. Analytic partial derivatives of R_0 have been derived for the layer conductivities and thicknesses. The derivation and resulting expressions are given in Appendix B. Farquharson (2000) described how the layer matrices and partial propagation matrices can be stored for each layer during the forward modelling stage and then later be reused to efficiently calculate partial derivatives. For the research reported herein, each layer matrix $\mathbf{M}_k$ is calculated and stored and the corresponding partial pre- and post-propagation matrices are accumulated. These partial propagation matrices are defined as,

$$\mathbf{P}_k^{pre} = \prod_{j=1}^{k-1} \mathbf{M}_j \text{ and } \mathbf{P}_k^{post} = \prod_{j=k+1}^{N_L} \mathbf{M}_j \,, \tag{2-24}$$

and are subsequently reused in Equation B-57 in Appendix B for efficient calculation of partial derivatives.

2.5.3 Magnetic field expressions

Consider an infinitesimal dipole source located at position $(0,0,h)$, arbitrarily orientated in the direction of the unit vector $\hat{\mathbf{m}}$, and whose moment varies harmonically in time. Its moment is given by,

$$\mathbf{m} = m_0 e^{i\omega t}\hat{\mathbf{m}}, \tag{2-25}$$

where m_0 is a scalar constant, ω represents angular frequency and t represents time. The total magnetic field $(\mathbf{H})$ at position (x,y,z) is,

$$\mathbf{H} = \mathbf{H}^P + \mathbf{H}^S, \tag{2-26}$$

where $\mathbf{H}^P$ is the primary field and $\mathbf{H}^S$ is the secondary field. By definition, the primary field is caused only by currents flowing in the transmitter loop, that is, as if the conducting medium of the subsurface was not present. The primary field is in phase with the transmitter current and is a function only of the system geometry and transmitter dipole moment. The primary field is given by,

$$\mathbf{H}^P = \mathbf{G}^P\mathbf{m} = \frac{1}{4\pi}\begin{bmatrix} \dfrac{3x^2-R^2}{R^5} & \dfrac{3xy}{R^5} & \dfrac{3x(z-h)}{R^5} \\ \dfrac{3xy}{R^5} & \dfrac{3y^2-R^2}{R^5} & \dfrac{3y(z-h)}{R^5} \\ \dfrac{3x(z-h)}{R^5} & \dfrac{3y(z-h)}{R^5} & \dfrac{3(z-h)^2-R^2}{R^5} \end{bmatrix}\mathbf{m}, \tag{2-27}$$

where the matrix $\mathbf{G}^P$ represents the primary field Green's tensor (Fitterman and Yin, 2004).

In contrast, the secondary field results from the currents induced in the subsurface by the time varying primary excitation. The secondary field is out of phase with the transmitter current and are thus represented by complex quantities. They are a function

of the transmitter moment, system geometry, frequency and subsurface electrical properties. The secondary field is given by,

$$\mathbf{H}^S = \mathbf{G}^S\mathbf{m} = \frac{1}{4\pi}\begin{bmatrix} \frac{\left(x^2-y^2\right)T_2}{r^3} - \frac{T_0 x^2}{r^2} & 2\frac{xyT_2}{r^3} - \frac{xyT_0}{r^2} & -\frac{xT_1}{r} \\ 2\frac{xyT_2}{r^3} - \frac{xyT_0}{r^2} & \frac{\left(y^2-x^2\right)T_2}{r^3} - \frac{T_0 y^2}{r^2} & -\frac{yT_1}{r} \\ \frac{xT_1}{r} & \frac{yT_1}{r} & -T_0 \end{bmatrix}\mathbf{m}, \quad (2\text{-}28)$$

where the matrix $\mathbf{G}^S$ represents the secondary field Green's tensor. The electrical property and frequency dependence of the expressions enter through the terms T_0, T_1 and T_2 which are Hankel Transform integrals. The definition and evaluation of the Hankel Transform integrals is discussed in Section 2.5.4 below.

The columns of the tensors $\mathbf{G}^P$ and $\mathbf{G}^S$ can be physically interpreted as being the contribution of a unit moment transmitter dipole in the x-axis, y-axis and z-axis directions. Similarly, each row relates to the resultant field at the receiver in each of the axes directions.

Note that in Fitterman and Yin (2004) negative signs should **not** be shown in front of the term in the first column of the second row of both the primary and secondary field Green's tensors. Accordingly, the error is rectified in equations 2-27 and 2-28.

2.5.4 Hankel transforms

The terms T_0, T_1 and T_2 that appear in the secondary field Green's tensor are Hankel Transform integrals of the form,

$$T_0 = -\int_0^\infty \mathsf{R}_0(\lambda)\, \lambda^2 e^{-\lambda(z+h)}\, J_0(\lambda r) d\lambda\,, \tag{2-29}$$

$$T_1 = -\int_0^\infty \mathsf{R}_0(\lambda)\, \lambda^2 e^{-\lambda(z+h)}\, J_1(\lambda r) d\lambda\,, \tag{2-30}$$

and,

$$T_2 = -\int_0^\infty \mathsf{R}_0(\lambda)\, \lambda e^{-\lambda(z+h)}\, J_1(\lambda r) d\lambda\,. \tag{2-31}$$

Inside the integrals the complex reflection coefficient term R_0 is a complex function of the layer conductivities ($\boldsymbol{\sigma}$), magnetic permeabilities ($\boldsymbol{\mu}$), thicknesses ($\mathbf{t}$), transmitter frequency (ω), and integration variable (λ). The terms $J_0(\lambda r)$ and $J_1(\lambda r)$ are Bessel functions of order 0 and 1 whose argument is the product of the transmitter-receiver horizontal separation (r) and integration variable. These integrals do not have analytic solutions for the general case and must be evaluated numerically.

The repeated calculation of R_0 for several values of the integration variable in the numerical evaluation of the integrals, and its repetition for each transmitter frequency, is the most computationally expensive part of layered-earth forward modelling. The difficulty in the numerical integration of the Hankel transform exists for two reasons—the infinite range of integration, and the oscillatory nature of the Bessel function for large values of λr. There is no shortage of publications on Hankel transform evaluation in the literature. Notable examples relevant to electrical geophysical methods would include, Ghosh (1971), Koefoed (1972), Johansen and Soerensen (1979), Anderson (1982), Chave (1983), Christensen (1990) and Guptasarma and Singh (1997). Except for Chave (1983), all of these authors favour the use of linear digital filters.

In the linear filtering techniques the integral is transformed into a convolution integral via a judicious substitution of variables. It is then approximated by a discrete

convolution summation, in which the summands are the product of kernel function evaluations with predetermined filter weights. The summation can be calculated using various efficient adaptive and lagged convolution schemes. According to Ward and Hohmann (1988), the digital filtering techniques are useful when there is a requirement to evaluate the integral for many different values of r, in which case the results of expensive evaluations of R_0 may be stored and later reused for other values of r. Chave (1983) opted for a direct quadrature approach with continued fraction expansion of the partial integrals evaluated between successive Bessel function zero crossings. These publications tend to be generalist, meaning that the methods are designed to cater for a wide range of kernel functions and a wide range of the horizontal transmitter to receiver distances r.

In layered-earth forward modelling of airborne electromagnetic systems there is no need to cater for a wide variety of integrands because there are only three and they are closely related. Also, unlike ground based systems there is no need to evaluate the integrands for many receiver positions because, for any one layered-earth model, there is only one receiver position.

In this work I have opted to take the specialist approach. A direct quadrature method specific to the airborne case and to layered-earths was developed. The emphasis is on minimising the range of the integration, and hence the number of expensive evaluations of R_0, by analytically estimating the peak position and decay of the integrands. This is made possible because $\mathsf{R}_0(\lambda)$ and $\lambda^n e^{-\lambda(z+h)}$ $(n = 1, 2)$ are smooth functions of λ, and when dealing specifically with the airborne case both the transmitter height (h) and receiver height (z) are positive values. Because of this, the exponential terms $\lambda^n e^{-\lambda(z+h)}$ rapidly decay, with increasing values of λ before the Bessel functions become oscillatory.

The peak position and rate of decay away from the peak of the term $\lambda^n e^{-\lambda(z+h)}$ can be analytically determined, and the peak position of $\mathsf{R}_0(\lambda)$ can also be estimated relatively accurately with just one evaluation of R_0 for a halfspace that approximates the true earth model. This means that the expensive evaluations of R_0 can be immediately restricted to an area immediately about the integrand peaks. Furthermore, the abscissa can be chosen to be the same for all three integrands.

Extensive experimentation, on multilayered models, has found that by using this quadrature integration scheme the transforms can generally be evaluated to a relative accuracy of less than 0.1% with only 17 evaluations of R_0. This level of accuracy is considered sufficient since it is much better that the two to five percent noise that is typical in airborne electromagnetic data. Prior to developing this method, I had used the digital filtering method and coefficients described by Guptasarma and Singh (1997). In their scheme there are 120 and 140 coefficients for the $J_0()$ and $J_1()$ Bessel function kernels respectively. Since the abscissas were not coincident, a total of 260 evaluations of R_0 were required to compute all three integrals.

Consequently, the quadrature Hankel transform integration method developed as part of this research is approximately 15 times faster than the digital filtering method of Guptasarma and Singh (1997) because only 17 abscissa are used. I acknowledge however, that this is for the specific airborne case, and for the 0.1% required level of accuracy. The filters of Guptasarma and Singh (1997) are more widely applicable and offer a high degree of accuracy. No further trials were carried out to compare the performance to other published digital filtering techniques.

2.6 Building of the complete system response

To build the complete system response we must first compute magnetic fields of equivalent infinitesimal dipoles from Equations 2-27 or 2-28, for each transmitter loop, receiver coil position and frequency separately.

Consider a n_{TX} turn transmitter loop of area A_{TX} centred at $(0,0,h)$. It carries a time varying current $I = I_0 e^{i\omega t}$. The loop is approximated by an equivalent infinitesimal magnetic dipole directed perpendicular to the plane of the loop with dipole moment,

$$\mathbf{m} = n_{TX}\, I A_{TX} \mathbf{R}_{TX}^{T}\, \hat{\mathbf{m}}_{TX}, \tag{2-32}$$

where the unit vector $\hat{\mathbf{m}}_{TX}$ is the reference orientation vector of the loop when it is in straight and level flight. For example, $\hat{\mathbf{m}}_{TX} = [0,0,1]^{\mathbf{T}}$ for the transmitter loop on a time-domain fixed-wing system, and $\hat{\mathbf{m}}_{TX} = [1,0,0]^{\mathbf{T}}$ for the transmitter coil in a vertical-coaxial coilset on a frequency-domain system. The loop's roll, pitch and yaw are more generally defined by the rotation matrix $\mathbf{R}_{TX}$.

We can substitute Equation 2-32 directly into the magnetic field expressions (Equations 2-27 and 2-28) to calculate the fields due to an arbitrary oriented transmitter dipole. The expressions given in Equations 2-27 and 2-28 are the directional magnetic field components along each of the coordinate system axes. However the magnetic field components that we wish to model are the components actually sensed along the axis of each of the receiver coils, which may, in general, be rotated.

Consider a three component receiver assembly that has three orthogonal sensor coils with n_{RX} turns of area A_{RX} located at position (x,y,z). These are typically called the X-component, Y-component and Z-components receiver coils whose axes, when the receiver assembly is in straight and level flight, are parallel to the x-axis, y-axis and z-axis of the inertial coordinate system respectively. However, in general, the receiver assembly is rotated such that it has roll, pitch and yaw defined by the rotation matrix $\mathbf{R}_{RX}$. A three component vector $\mathbf{H}_{RX}$ whose elements are the magnetic fields sensed by the X-component, Y-component and Z-component receiver coils can be defined and expressed as,

$$\mathbf{H}_{RX} = \mathbf{R}_{RX}\mathbf{H}. \tag{2-33}$$

Therefore, after collecting previous terms we may write,

$$\mathbf{H}_{RX} = \mathbf{R}_{RX}\left[(\mathbf{G}^{P} + \mathbf{G}^{S})I\ n_{TX}\ A_{TX}\ \mathbf{R}_{TX}^{T}\hat{\mathbf{m}}_{TX}\right]\ . \tag{2-34}$$

Alternatively we may wish to calculate the primary field component alone as,

$$\mathbf{H}_{RX}^{P} = \mathbf{R}_{RX}\left[\mathbf{G}^{P}I\ n_{TX}\ A_{TX}\ \mathbf{R}_{TX}^{T}\hat{\mathbf{m}}_{TX}\right], \tag{2-35}$$

or the secondary field alone as,

$$\mathbf{H}_{RX}^{S} = \mathbf{R}_{RX}\left[\mathbf{G}^{S}I\ n_{TX}\ A_{TX}\ \mathbf{R}_{TX}^{T}\hat{\mathbf{m}}_{TX}\right]. \tag{2-36}$$

Equations 2-35 and 2-36 are fundamental equations that express the primary and secondary magnetic fields along the axes of an arbitrarily oriented receiver assembly in the presence of an arbitrarily oriented transmitter loop that transmits at a single frequency.

The receiver actually measures the voltage induced in the coil rather than the magnetic field itself. From Faraday's law the induced voltage is the negative of the time rate of change of the total magnetic flux threading the loop. For the harmonic time variation of the dipole source considered here $\partial/\partial t = i\varpi$. Therefore the total induced voltage in the receiver coil is,

$$\mathbf{V}_{RX} = -n_{RX}\frac{\partial(\mathbf{B}_{RX}A_{RX})}{\partial t} = -n_{RX}\frac{\partial(\mu_{0}\mathbf{H}_{RX}A_{RX})}{\partial t} = -i\ \mu_{0}\,\omega\ n_{RX}A_{RX}\,\mathbf{H}_{RX}, \tag{2-37}$$

and the primary and secondary components are respectively,

$$\mathbf{V}_{RX}^{P} = -i\ \omega\,\mu_{0}\,n_{RX}A_{RX}\ \mathbf{R}_{RX}\left[\mathbf{G}^{P}I\ n_{TX}\ A_{TX}\ \mathbf{R}_{TX}^{T}\hat{\mathbf{m}}_{TX}\right] \tag{2-38}$$

and,

$$\mathbf{V}_{RX}^{S} = -i\ \omega\,\mu_{0}\,n_{RX}A_{RX}\ \mathbf{R}_{RX}\left[\mathbf{G}^{S}I\ n_{TX}\ A_{TX}\ \mathbf{R}_{TX}^{T}\hat{\mathbf{m}}_{TX}\right]. \tag{2-39}$$

It is possible then to compute the response of a general airborne electromagnetic system using appropriate combinations of Equation 2-37. From here the modelling becomes dependent on the actual system being modelled. The possibilities are too numerous to cover here. Suffice to say that one must be careful to compute quantities equivalent to the data that have been recorded, normalized and transformed. For example, for some systems the results are presented in terms of the magnetic field at the receiver rather than the induced voltage.

To model most frequency-domain systems that use several coilsets, each having one transmitter loop, one receiver loop and operating at one specific frequency, it is only necessary to have one equivalent dipole source for each coilset.

To model time-domain systems one can make use of the analytic frequency-domain expressions. The procedure used for computing a time-domain system response is as follows:

a) digitise the transmitter's current waveform (or the processed data equivalent waveform in the case of TEMPEST and SPECTREM) to produce a transmitter current time series w_t and compute its discrete frequency spectrum $W_f = \mathsf{F}(w_t)$ via Fast Fourier Transform (FFT);

b) compute the required magnetic field $B_f^{\log}$ components at the receiver for several logarithmically spaced discrete frequencies across the full bandwidth of the system;

c) spline interpolate the logarithmically spaced $B_{f^{\log}}$ to generate a linearly spaced frequency spectrum B_f at the same discrete frequencies that exist in W_f;

d) depending on whether a magnetic B field or voltage response is required, compute $b_t = \mathsf{F}^{-1}(B_f W_f)$ or $v_t = \mathsf{F}^{-1}(-iwB_f W_f)$ by inverse FFT to yield the response time series;

e) window the time series b_t or v_t as necessary to properly simulate the given system's windowing method (e.g. boxcar, linear taper, area under curve) to yield the required time-domain window response.

f) apply normalizations and unit conversions to the window responses to appropriately simulate any such adjustments made to the acquired data being simulated.

The number of logarithmically spaced frequencies at which the expensive magnetic field calculations need to be made depends on the system bandwidth and the desired accuracy. Essentially, one must use a sufficient density of frequencies to enable accurate spline representation of the continuous frequency spectrum from the discrete samples. Six frequencies per decade are typically required to be computed over four to five decades (Raiche, 1998).

2.7 Analytic partial derivatives

The holistic inversion algorithms for both frequency- and time-domain airborne electromagnetic data, which are discussed in Chapter 3 and Chapter 5, use linearized gradient-based minimization techniques. This requires that partial derivatives of the forward modelled data, with respect to each the unknown inversion model parameters, to be calculated. It would be simple to calculate the required derivatives by finite difference methods, however that would be computationally prohibitive due to the large number of parameters to be solved for (Sambridge et al., 2007). Since it is much more efficient and accurate to use analytically calculated partial derivatives, expressions for these have been derived.

Expressions for the partial derivatives of the forward model with respect to: individual layer parameters of the 1D model (σ_k, t_k) and the system geometry location parameters (h, x, y, z and r) are presented in Appendix B. Derivatives with respect to the

system geometry orientation parameters, (ψ_R, ψ_P and ψ_Y) for both the transmitter and receiver loops are presented in Appendix A. These are in turn used in the holistic inversion algorithms (see Sections 3.5.4 and Section 5.5.4) when compiling the derivatives with respect to the underlying inversion model parameters.

Chapter 3

Frequency-domain holistic inversion

3.1 Outline

The holistic approach jointly inverts raw *airborne* electromagnetic frequency-domain data, *geoelectric* data and *interface-depth* data to simultaneously estimate a *calibration* model and a *conductivity* model. It is the conductivity model that is ultimately of interest to the end-user. The purpose of the calibration model is to account for systematic calibration errors in the airborne electromagnetic data. Figure 3.1 summarises the principal elements of the holistic inversion scheme.

The conceptual 3D model of the Earth is one in which the subsurface is comprised of a series of layers vertically stacked on top of each other. The conductivity and thickness of each layer varies laterally (in the plane of the x and y axes) in a smooth and continuous fashion. The logarithm of the layer conductivities and thicknesses are parameterized by separate independent uniform bicubic B-spline meshes. It is the node coefficients of the splines that are solved for in the inversion to ultimately yield the conductivity model of interest. Once solved for, the spline coefficients can be used to calculate the conductivity anywhere within the model bounds.

The choice exists to solve for both layer conductivities and thicknesses, which will be referred to as a '*variable thickness inversion*'. Otherwise, the layer thicknesses may be set and kept fixed while the layer conductivities are solved for. This will be referred to as a '*fixed thickness inversion*'.

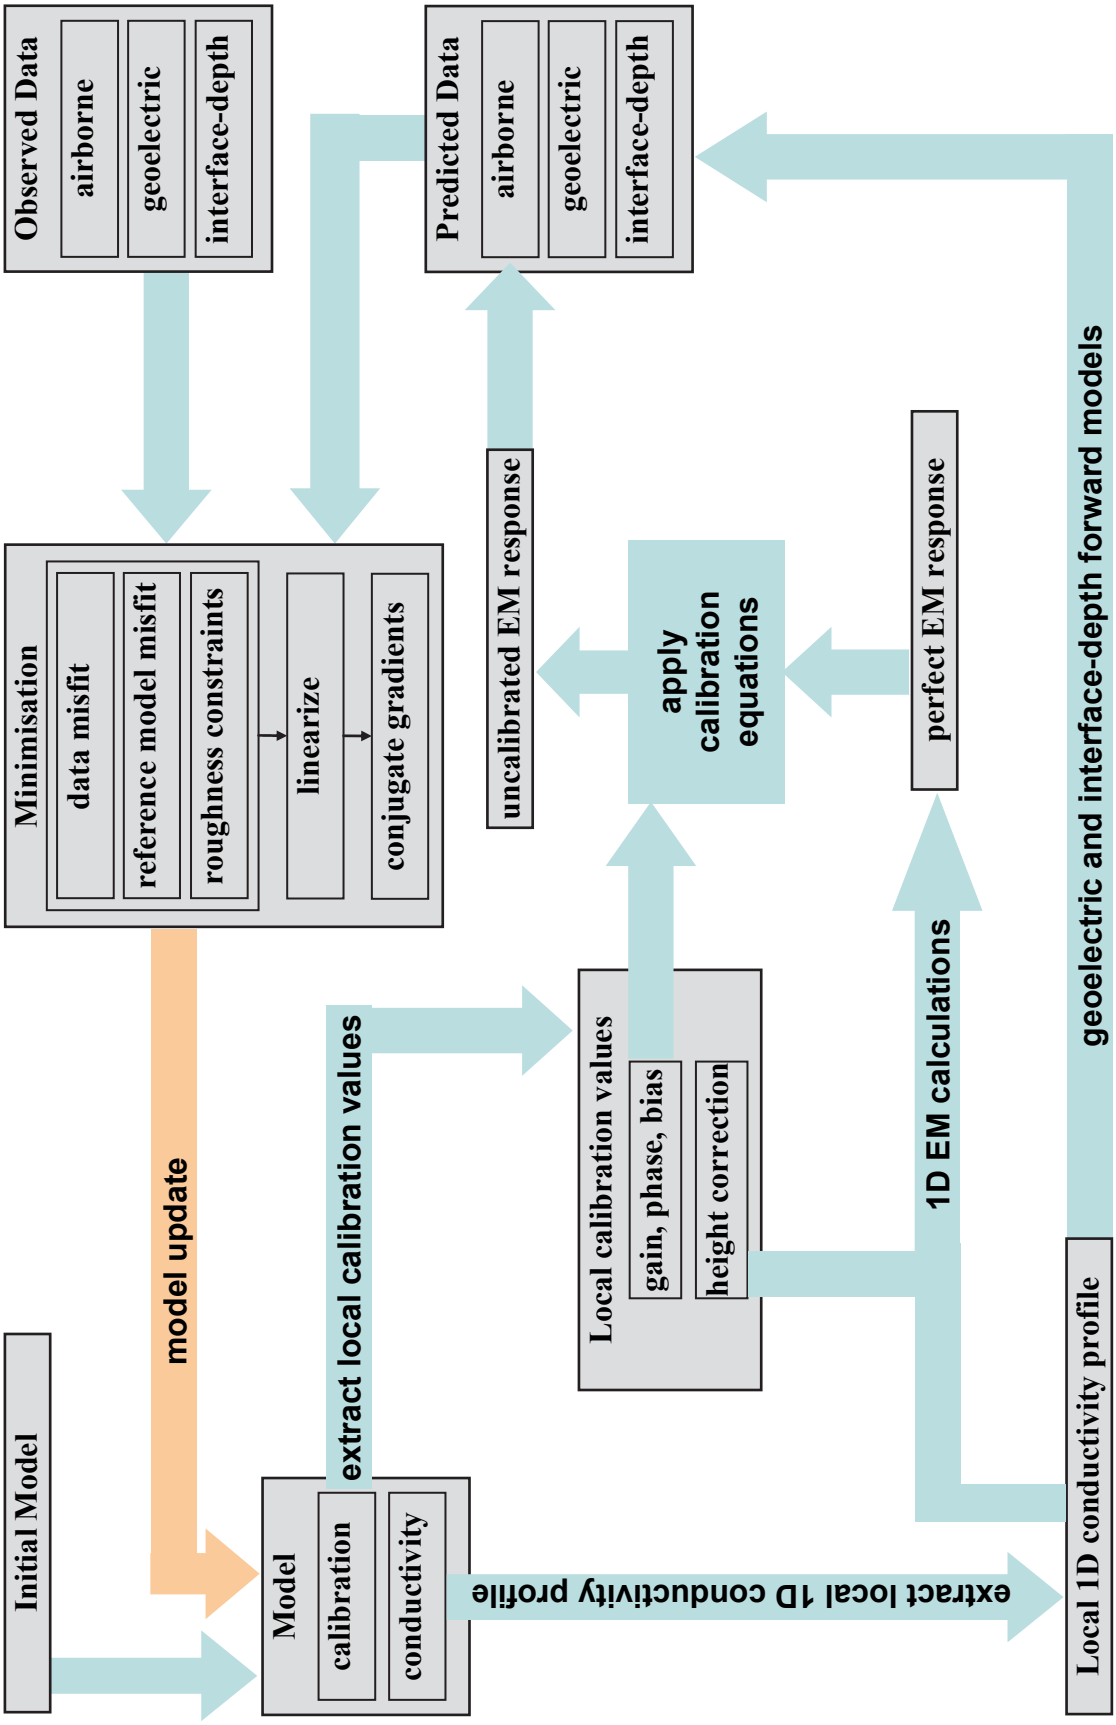

Figure 3.1 *Schematic flowchart of the principal elements of the holistic inversion scheme.*

The systematic calibration errors, which include gain, phase, bias and height errors, are represented by either discrete values (for individual days of flying, individual flights or individual lines) or piecewise linear functions which, in combination, make up the calibration model.

Geoelectric data are independent direct or inferred measurements of subsurface conductivity at specific locations. Interface-depth data are independent measurements of the depth to particular interfaces in the conceptual conductivity model at specific locations. The airborne data are the inphase and quadrature measurements for all of the coilsets and airborne samples in the dataset.

Although the inversion employs a 3D representation of subsurface conductivity, 1D layered-earth electromagnetic forward modelling theory is used to approximate the airborne response. In principle, full 3D modelling could be used for the holistic inversion, and would indeed be more consistent with the model used. However, simultaneous 3D inversion of datasets in excess of a few thousand line kilometres would not be computationally feasible. Comparison of the difference in computation load between 1D and 3D forward modelling is not straightforward because it depends on the complexity of the models and the discretization of the meshes involved. However, as a guide, Ellis (1995) found that for a 12×12×4 mesh, an approximate 3D finite element forward model took 6,900 times longer to compute than a four layer 1D forward model. He found that a more accurate hybrid 3D finite element forward model took 32,900 times longer than the 1D model. In a more recent case study (Annetts et al., 2003) computed a 3D finite element forward model for a 2.4×1.0×0.75 km domain (60×29×26 mesh) in 32 s per frequency per sample. They did not compare times with their own 1D modelling, however a 26 layer 1D forward model can be computed in 0.001 s per frequency per sample using code developed for this thesis. These results suggest that full 3D takes 10^4 longer and hence is impractical for the holistic inversion.

Therefore the procedure that I have developed for forward modelling (prediction) of any particular datum is a two-stage process. Firstly, it requires extraction of the *local 1D conductivity profile* from the 3D model at the location of the datum, plus, for an airborne datum, it also requires extraction of the *local calibration values* for that datum. Secondly, for an airborne datum, the local 1D conductivity profile and local height calibration value are used to compute a theoretical airborne response, to which the local gain, phase and bias calibration values are applied through a set of calibration equations to predict the airborne datum. In the case of a geoelectric or interface-depth datum the local 1D conductivity profile can be used on its own to predict the datum.

It is important to realise that uncalibrated airborne data are input into the inversion. This requires that uncalibrated data be simulated in the forward modelling. In light of this, the calibration equations mentioned above may well be better described as *'uncalibration'* equations because, in effect, they transform the theoretical response for a perfect AEM system into the theoretical response for an uncalibrated system.

The inversion is based on the minimisation of an objective function comprised of data misfit, reference model misfit, spatial conductivity roughness and temporal roughness of some calibration errors. Since the objective function is nonlinear with respect to the model parameters, it is minimised via an iterative scheme where an initial set of model parameters are chosen and then within each iteration the current set of model parameters are updated to a new set. The new model parameters are determined by using the conjugate gradient method to solve a system of linear equations that are locally-linearised about the current set of model parameters.

3.1.1 Nomenclature

Nomenclature that appears throughout this chapter is listed below. Terms are described as they are encountered, but are provided here for the reader's convenience.

N_s number of samples in the airborne dataset to be inverted

N_c number of coilsets (frequencies) in the airborne system

$N_a = N_s \times N_c \times 2$ number of airborne data

N_g number of geoelectric data

N_i number of interface-depth data

$N_d = N_a + N_g + N_i$ total number of data

N_L number of layers in the earth model

N_m number of unknown model parameters to be solved for

N_{lx} number of reference east-west roughness constraints

N_{ly} number of reference north-south roughness constraints

N_{lz} number of vertical roughness constraints

N_{lb} number of bias roughness constraints

c^{σ}_{ijk} coefficient of the ijth node of kth conductivity spline mesh

c^{t}_{ijk} coefficient of the ijth node of kth thickness spline mesh

$\sigma_k(x,y)$ local 1D model kth layer conductivity evaluated from splines at x, y

$t_k(x,y)$ local 1D model kth layer thickness evaluated from splines at x, y

$w^{\sigma}_{ijk}(x,y)$ weight of ijth node coefficient used in evaluating $\sigma_k(x,y)$

$w^{t}_{ijk}(x,y)$ weight of ijth node coefficient used in evaluating $t_k(x,y)$

$\mathbf{s}_{sc}$ local gain calibration value for cth coilset at the sth sample

θ_{sc} local phase calibration value for cth coilset at the sth sample

b^{ip}_{sc} local inphase bias calibration value for cth coilset at the sth sample

b^{q}_{sc} local quadrature bias calibration value for cth coilset at the sth sample

$\triangle h_s$ local height calibration value for sth sample

$\boldsymbol{f}^{pft}(\ )$ 1D forwrad model function for perfectly calibrated AEM system

$\boldsymbol{f}^{obs}(\ )$ 1D forwrad model function for miscalibrated AEM system

$\boldsymbol{g}(\ )$ nonlinear vector valued forward model function that maps $\mathbf{m}$ to $\boldsymbol{g}$

$\mathbf{d} = [\mathbf{d}^{air} \mid \mathbf{d}^{geo} \mid \mathbf{d}^{\mathrm{int}}]^{\mathbf{T}}$ observed airborne, geoelectric and interface-depth data

$\boldsymbol{g} = [\boldsymbol{g}^{air} \mid \boldsymbol{g}^{geo} \mid \boldsymbol{g}^{\mathrm{int}}]^{\mathbf{T}}$ predicted airborne, geoelectric and interface-depth data

$\mathbf{m} = [\mathbf{m}^{cal} \mid \mathbf{m}^{ear}]^{\mathbf{T}}$ calibration and earth model parameters to be solved for

$\mathbf{G}_{\mathbf{m}}$ Jacobian matrix where $\mathbf{G}_{\mathbf{m}}\Big|_{ij} = \partial \boldsymbol{g}_i(\mathbf{m}) \Big/ \partial \mathbf{m}_j$

3.2 Conductivity model

3.2.1 Choice of conductivity model style

There is no one holistic inversion conductivity model parameterization that will be optimal for all electromagnetic surveys. A parameterization choice must be made based on: (a) its ability to adequately represent the conductivity structure of the survey area; (b) its compatibility with 1D layered-earth modelling; (c) its ability to satisfy the requirements of the holistic approach; and (d) practical implementation considerations.

In conventional 1D AEM inversions algorithms, every multi-frequency measurement, that is, every airborne sample acquired at different spatial locations, is inverted independently to estimate a 1D conductivity models (Ellis, 1998; Sattel, 1998; Sengpiel and Siemon, 1998; Huang and Fraser, 2003; Farquharson et al., 2003; Brodie et al., 2004b; Lane et al., 2004a). Since each sample is inverted separately, I call this sample-by-sample or SBS inversion. The conductivity model for the complete survey area is subsequently compiled post-inversion by 'stitching' together all of the independently inverted 1D models in some fashion, for example by gridding or kriging.

AEM data are highly coherent because of the spatial continuity of the geology and the temporal continuity of the systematic calibration errors. This coherency is additional available information which could potentially be used to constrain conductivity predictions and identify calibration errors. Clearly, it is not possible for SBS style algorithms, which invert each airborne sample independently of all the other samples, to benefit from this extra coherency information.

In the laterally constrained inversion (LCI) of Auken et al. (2005) as well as a similar method described by Sasaki (2004), the SBS style of parameterization is used. However unlike SBS inversion, the LCI inverts multiple adjacent samples along a profile at once. This approach is able to exploit the spatial coherency in the direction of

the profile because it explicitly ties together the separate 1D models in the inversion by applying conductivity model covariance or roughness constraints that operate between adjacent 1D models along the profile. Viezzoli et al. (2008) have recently extended this concept to spatially constrained inversion (SCI). The SCI exploits spatial coherency in all lateral directions by inverting data from several profiles at once and applies covariance constraints to adjacent models in all directions. In the LCI and SCI methods, the individual 1D models are compiled post-inversion into a model for the complete survey area, in the same way as for the SBS method. Figure 3.2 illustrates how different conceptual styles of model, which have been used in 1D AEM inversion, are able to make use of spatial coherency information.

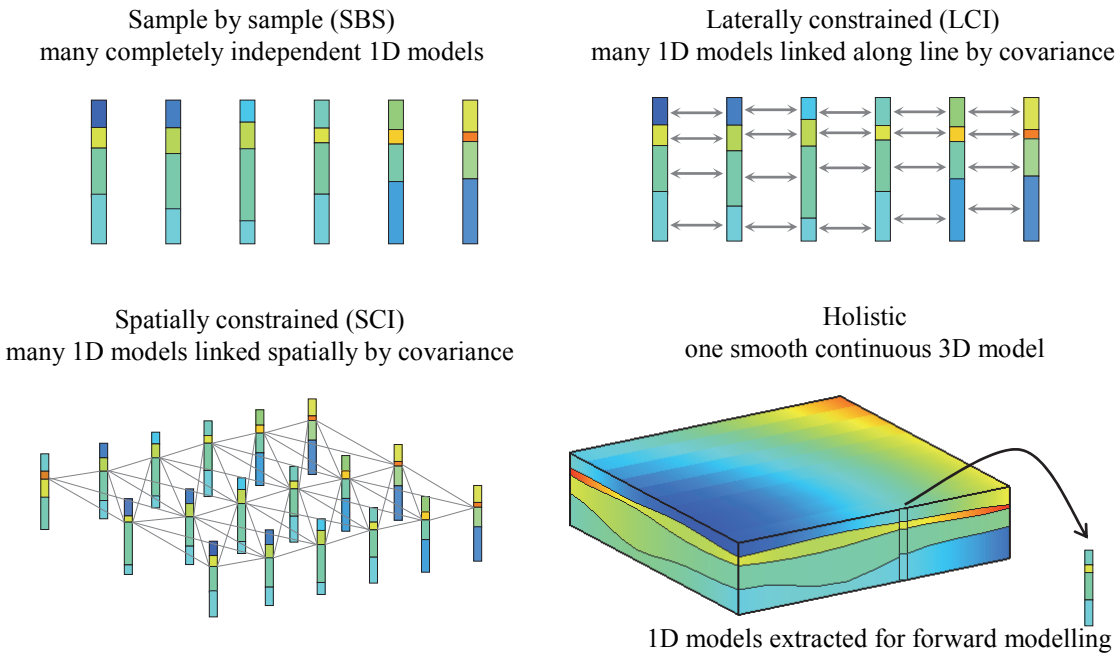

Figure 3.2 *Schematic illustration of the different conceptual styles of conductivity model used in the sample-by-sample, laterally constrained, spatially constrained and holistic inversion methods.*

The holistic inversion can also take advantage of spatial coherency, but achieves this in a subtly different manner. For the holistic inversion, I choose to use a vertically layered 3D conductivity model in which the conductivity and thickness of each layer is directly parameterized by smooth and continuous mathematical basis functions. Since the model is laterally smooth and continuous, exploitation of the aforementioned valuable

spatial coherency information is automatically ensured. There is no need for explicit covariance coupling of independent 1D models.

This single 3D model style is physically appealing because it conceptualises the conductivity structure in precisely the way one envisages a layered-earth to be. It seems somewhat more intuitive or straightforward than conceptualising the subsurface as some combination of 1D models. Another attractive feature of the chosen model style is that the conductivity value at any point in the model is uniquely defined by the mathematical basis functions, including between flight lines where AEM data samples do not exist. This allows the independent geoelectric and interface depth data, which in general do not lie exactly on flight lines, to be naturally incorporated in the inversion without any requirement to extrapolate them to nearby flight lines. To help clarify this point, consider for example a borehole that lies partway between two flight lines. If a conventional discrete parameterization is used one must choose to either; use the borehole information to constrain the discrete 1D model associated with the nearest airborne sample; or to have it constrain multiple discrete 1D models using some ad hoc distance threshold or weighting. However, using a continuous model the borehole information can simply and naturally constrain the conductivity model where the borehole data was actually observed.

A further benefit of this particular property is that examination and visualisation of the model is unique. This is because there is no need for gridding, kriging or projection when generating images and cross sections since conductivity values can be exactly evaluated where the information is required.

The layer properties must obviously be laterally variable to accommodate geological variations across the survey area. The layer properties are chosen to be smooth and continuous because it is expected that the lateral variations of the Earth's electrical properties are typically gradational rather than abrupt. This assertion does not always

hold true, but it is the case in the many layered geological environments, particularly at the scale length of the AEM system footprint. Clearly there will be geological environments where abrupt variations are present, for example faulted and folded terrains, in which case the laterally smooth and continuous model will not be able to adequately represent the geology. However, like all other 1D inversion methods, the holistic inversion is generally not suitable to those geological environments. Smoothness and continuity will also minimise the shortcoming of the 1D modelling approximations that must be made. A vertically layered model that has constant conductivity from the top to the bottom of each layer, was chosen for its compatibility with 1D forward modelling theory, which requires that the model have such distinct layering.

To correctly compute the exact forward response of a model containing lateral variations it would be necessary to use 3D electromagnetic calculations. However that would not be feasible, from a practical implementation viewpoint, due to the computational burden. The solution to this is to use 1D modelling approximations which are of the order of 10^4 times faster as noted earlier. Forward modelling is carried out by extracting a local 1D model, by drilling down vertically into the 3D model, wherever a forward response is required.

A contrast can now be seen between the various methods. The SBS, LCI and SCI methods parameterize conductivity by **many 1D** models that are stitched together to form a **single 3D** model which is ultimately desired for interpretation. On the other hand, in the holistic inversion a **single 3D** conductivity model is employed, from which **many 1D** models must be extracted for forward modelling purposes.

3.2.2 Choice of layer property basis function

There are several types of smooth 2D functions that could possibly be used to represent the required conductivity or thickness of each layer in construction of the chosen style

of model. In the simpler 1D case, Gyulai and Ormos (1999) used polynomial and periodic basis functions to represent smooth changes in layer conductivity and thickness along a profile in inversion of DC-resistivity data. This concept could be extended to analogous 2D basis functions. Alternatively, piecewise planar or bilinear basis functions could also be used. A further candidate, and the one chosen for the holistic inversion because of its suitable properties is the uniform bicubic B-spline surface (Bartels et al., 1987).

Uniform bicubic B-spline surfaces are widely used in computer graphics applications, particularly for the approximation of surfaces in 3D. They have also been used in seismic tomography for the parameterization of geological interfaces (e.g. Rawlinson, 2000). These spline surfaces are defined by a topologically rectangular mesh which has a spline node and associated coefficient at each mesh intersection. The coefficients are the unknowns that are solved for in the inversion. An in-depth description is provided in Appendix C, but here I discuss the properties that make them particularly suitable for the holistic inversion.

The mathematical formulation ensures that the surfaces are continuous and that their first and second derivatives, with respect to distance in any direction, are also continuous (i.e. they are C^2 functions). Our requirement for representation of a smooth and laterally continuous property distribution is implicit in the parameterization. Therefore it does not have to be explicitly enforced by smoothness constraints or post-inversion filtering.

These surfaces are locally-supported, because the value of the surface at any location is only dependent upon the 16 adjacent surrounding spline node coefficients, rather than all of the mesh's coefficients. Local-support is an important property because it means that when the Jacobian matrices are calculated for the inversion, other than for the entries corresponding to the 16 adjacent coefficients, the vast majority of entries are

zero and hence the matrices are sparse. The sparsity is vital for computational efficiency and is the major reason that the goal of inverting complete datasets to take advantage of spatial coherency in the geological signal and temporal coherency in the calibration error is actually feasible. The local-support has a further intuitive appeal in that the computed airborne response becomes a function of a weighted average of local Earth properties rather than distant properties. This is akin to the true airborne response being a weighted function of the Earth properties primarily within the system's footprint.

Bicubic B-spline surfaces are popular in computer graphics because of their ability to represent complicated surfaces with as few parameters as possible. This is also attractive for the holistic inversion because the data can be satisfied without resort to unnecessarily fine discretization or the need for a separate set of model parameters for every airborne sample. This also reduces the total number of unknowns, thereby lowering the computational load and improving the condition of the inversion problem.

3.2.3 Conductivity model parameterization

The subsurface is deemed to be comprised of discrete layers stacked in layer-cake fashion. The bottommost layer, or the basement halfspace, is infinitely thick. Figure 3.3 illustrates a possible four layer model, showing how both the conductivity (represented by colour) and thickness of each discrete layer, vary smoothly and continuously in the lateral sense. However at any one horizontal position the conductivity is constant from the top to the bottom of the layer.

The logarithm of conductivity and logarithm of thickness of each layer is parameterized by a uniform bicubic B-spline surface. The B-spline parameterization allows a node spacing to be chosen such that the subsurface conductivity can be adequately represented with as few model parameters as possible. The chosen node spacing is dependent on the minimum of the expected scale length of lateral conductivity

variations and the airborne system's footprint. Each of the spline surfaces are completely independent. Different node spacings may be chosen for each layer, and different node spacings may be chosen in each orthogonal coordinate system axes.

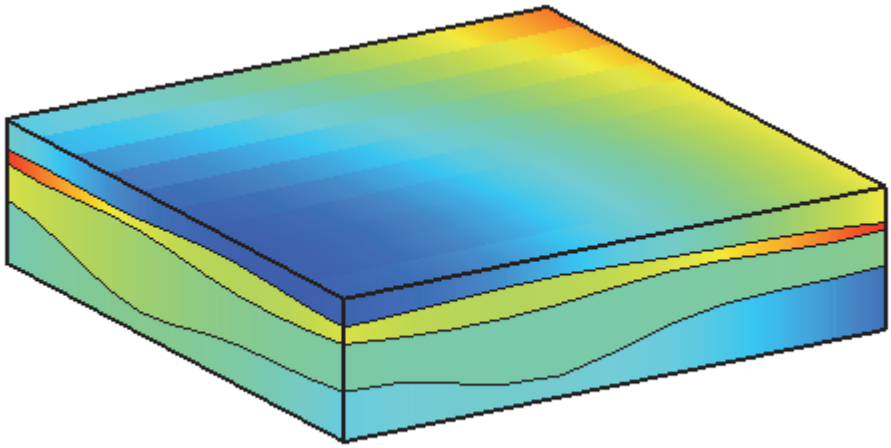

Figure 3.3 The layered and laterally continuous conductivity model.

In general, to parameterize an N_L layer earth $2N_L - 1$ spline surfaces are required—N_L for the conductivities and $N_L - 1$ for the layer thicknesses. Since the bottom layer is infinitely thick its thickness does not need to be parameterized. In the variable thickness inversion, where we choose to invert for both layer thickness and conductivity, all $2N_L - 1$ splines are required. However, in the fixed thickness inversion we may use splines to represent the layer thicknesses, or alternatively the layer thicknesses can be spatially constant. In the latter case only N_L splines are required in total. However, for the ensuing development we will consider the more general case where both conductivities and thicknesses are parameterized.

The spline surface representing the kth layer conductivity and thickness is denoted S_k^{σ} and S_k^{t} respectively. To avoid repetition, for the moment consider the spline S_k^{v} where v may represent either σ or t. As detailed in Appendix C, S_k^{v} is comprised of a topologically rectangular array, of $N_{kx}^{v} - 1$ columns by $N_{ky}^{v} - 1$ rows of surface patches (the area bounded by four spline nodes, *cf.* Appendix C), which, when mosaiced, form a spline control mesh having N_{kx}^{v} columns by N_{ky}^{v} rows by mesh intersections or spline

nodes. The i,jth spline node has a corresponding coefficient value c^v_{ijk}. It is these coefficients that are solved for in the inversion. The total number of coefficients is,

$$N_c = \sum_{k=1}^{N_L} N^{\sigma}_{ky} N^{\sigma}_{kx} + \sum_{k=1}^{N_L-1} N^{t}_{ky} N^{t}_{kx} \,. \tag{3-1}$$

At any given lateral position x, y that lies on the i,jth surface patch of the spline mesh S^v_k, the conductivity or thickness parameter $v_k(x,y)$ is parameterized by,

$$l\,\mathrm{og}(v_k(x,y)) = S^v_{ijk}(x,y) = \sum_{a=i-1}^{i+2} \sum_{b=j-1}^{j+2} w^v_{abk}(x,y)\; c^v_{abk} \,, \tag{3-2}$$

where $w^v_{abk}(x,y)$ is the known weight that determines the contribution of the a,bth spline node coefficient on the kth mesh, c^k_{abk} to the spline surface value at point x, y. The weights are a function only of the position x, y on the surface patch. For details of determining the weights see Equation C-4 in Appendix C. The positions of the spline nodes are set at the beginning of the inversion and remain fixed throughout. The weights therefore remain fixed for each x, y position as well.

It can be seen then from Equation 3-2 that the spline evaluation is a simple weighted sum of 16 (4×4) coefficients. An important point is that the conductivity or thickness of any one layer at any one location is influenced by only the coefficients of the sixteen immediately surrounding nodes. This is what is meant by 'locally supported'. Accordingly, the vast majority of the derivatives in the inversion's Jacobian matrix are zero, which allows sparse matrix computational methods in the inversion.

Note that there is no general requirement that all spline nodes lie on a regularly spaced rectangular mesh, but they must maintain a rectangular topology. Likewise, since each spline mesh is independent, there is no requirement that the nodes have the same spatial location on each of the meshes. Furthermore, only those nodes that are required for computation of a local 1D conductivity structure need to have 'defined' coefficients. This means that irregularly shaped survey areas can be accommodated.

3.3 Calibration model

3.3.1 Background

Errors in a frequency-domain airborne electromagnetic system, and many other devices, can usually be subdivided into 'random noise' and 'systematic error' components. The meaning of random noise here, are those errors caused by relatively unpredictable events, although they may not be truly random in the strict sense of the word. Some examples are: vibration of the transmitter-receiver assembly; atmospheric sferic events; powerline interference; and other electronic interference. Random noise is well recognised in the geophysical industry and has to some degree been the main focus of attempts to improve AEM systems over many years. This is because it often presents in data at the same high spatial frequencies as the anomalies caused by discrete geological conductors, which mineral explorers are interested in, thereby making it difficult to discriminate between the random noise and the conductors. One aim of the holistic inversion is to deterministically model errors and thereby remove any detrimental effect they may otherwise have on the conductivity estimates. However, since random noise cannot be deterministically modelled it cannot be targeted by the holistic inversion. Random noise is thus treated in the same way as it is in conventional inversion algorithms.

Systematic errors, on the other hand, are not random because they are caused by some form of miscalibration of the measurement system. Comprehensive analyses of the possible causes of miscalibration are published in Fitterman (1997) and Fitterman (1998). Some examples of systematic error in a frequency-domain system are: incorrect cancellation of the primary field; incorrect knowledge of the system gain; incorrect synchronization between the transmitter and receiver time bases; and incorrect altimeter measurements of the height of the transmitter-receiver assembly. They may be constant or vary slowly spatially and temporally, typically at scale lengths greater than the

geological signal. Systematic errors in a frequency-domain dataset are not well accommodated by the typical assumptions made about errors in inversion algorithms because they are not necessarily zero-centred, and they tend to be correlated in time and across two or more channels of data. Furthermore, unlike random noise, due to their slowly varying nature systematic errors may not be readily identifiable as errors when viewed as individual channels of data. However, when multiple channels of data are considered together in a quantitative inversion, the influence that systematic errors have on the absolute numeric values of the dataset can lead to mutual inconsistency between the channels of data and a plausible earth model.

Since systematic errors have postulated non-random physical causes they can be expressed deterministically rather than just stochastically. In the holistic inversion, this allows us to classify systematic errors as being a signal component, additional to the true ground response signal, of the measured AEM data. That is, they are not considered as error in the same sense as random noise. The calibration model is used to simulate the systematic error signal. In so doing, rather than transforming systematic error to artefacts in the conductivity model, it is transformed to the parameters of the calibration model.

The purpose of the calibration model is thus to parameterise the systematic errors. The calibration model used for the research reported here has been adopted from the work described by Deszcz-Pan et al. (1998), which included *gain*, *phase* and *bias* errors categories. In the present work a further category, that of *height* error, has been included.

3.3.2 Systematic error classes

To explain what is meant by these different classes of errors, consider a perfectly calibrated single-frequency AEM system that, at height h above a hypothetical layered-earth, that has conductivities $\boldsymbol{\sigma}$ and thicknesses $\mathbf{t}$. The perfect system would observe a

complex response $\boldsymbol{f}^{pft}(\boldsymbol{\sigma}, \mathbf{t}, h)$. However, above the same earth, a miscalibrated system would observe a response given by,

$$\boldsymbol{f}^{obs} = se^{i\theta}\left[\boldsymbol{f}^{pft}(\boldsymbol{\sigma}, \mathbf{t}, h^{obs} + \triangle h) + b\right]. \tag{3-3}$$

In Equation 3-3 s is a multiplicative real amplitude scaling factor, known as the system gain, which for a perfectly calibrated system would have the value $\mathrm{s} = 1.0$. The real valued phase θ is the phase difference between the transmitter and receiver time references, which for a perfectly calibrated system would have the value $\theta = 0.0$. The complex valued bias b, otherwise known as the zero-level, is the response that is observed when no secondary field is present (e.g., at high altitude), which for the ideal system should be $b = 0.0 + 0.0i$. In the miscalibrated system, the altimeter measurement of the true transmitter-receiver assembly height h, is incorrectly measured as h^{obs}, thus an additive calibration factor $\triangle h$ is required to correct the height such that $h = h^{obs} + \triangle h$.

The gain, phase, bias and height errors will typically vary as a function of time throughout the survey. However, in some cases they may not vary at all. Ideally, the parameterization of these errors would be based on a sound understanding of how each class of error varies and what the causes are. However, there are unfortunately no publicly documented reports of studies where tests flights of airborne frequency-domain systems have been made to specifically investigate this matter.

Some insight into the variation of calibration errors can be gained by inspection of the spatial patterns of artefacts caused by systematic calibration errors in images of survey data. An example can be seen in Figure 1.12 in Chapter 1, where it was demonstrated how artefacts in a dataset corresponded to individual flights, and in some cases with individual flight lines. Other examples of artefacts that are more or less constant over blocks of flight lines have been published by Huang and Fraser (1999), Huang (2008) and Siemon (2009). Gain and phase calibration procedures are usually carried out at the

beginning of each survey, the beginning of each day of flying, or at the beginning of each flight (Deszcz-Pan et al., 1998). It is reasonable then to expect that the systematic calibration errors would change from survey to survey, from day to day, or from flight to flight because of the different errors that could possibly be introduced on each occasion (Fitterman, 1998). Therefore, the parameterization should allow for calibration errors that have a constant value for all samples within each survey, day or flight, but that have a different value for each of survey, day and flight. Furthermore because the calibrations are carried out separately for each coilset, separate parameters are required for each coilset.

Further insight can be gained by inspection of the data that are recorded at high altitude at regular intervals during a survey flight. At high altitude the effect of the ground is negligible and thus the measured secondary field would ideally be zero. However, a non-zero value, the bias or zero-level, is generally observed. As shown in Figure 1.10 of Chapter 1, the bias tends to drift slowly throughout a survey flight. This is consistent with a gradual change in temperature during the flight, which is the postulated cause of bias errors (Huang and Fraser, 1999; Valleau, 2000). This suggests that an appropriate bias parameterization could represent a slow variation with time over a flight.

In the examples discussed later, I illustrate some possibilities where gain calibrations errors are parameterized to be different for each frequency but are constant for an entire survey. Phase is different for each frequency, as well as different for each day of flying, to reflect the possibility of different phase errors resulting from different daily calibrations. The inphase and quadrature bias for each frequency is deemed to vary piecewise linearly over the duration of each flight. The height error is parameterised to be constant over all coilsets and over the complete survey.

3.3.3 Parameterization styles

The general parameterization that has been formulated incorporates the following entities: (a) survey; (b) day of flying; (c) flight; (d) line; (e) coilset; and (f) inphase or quadrature channel. Calibration errors may be considered constant over one such entity, over a group of entities, or alternatively, they can be considered to vary piecewise linearly over any one entity. The specific choice will be system and survey specific, guided by a trade-off between adequate representation of the systematic calibration errors, inversion stability and size. The number of possible ways to parameterize the calibration model is large, but the best results will be achieved when the parameterization accurately describes the characteristics of the systematic calibration error. For example, it is highly unlikely that one would choose to parameterize gain to vary piecewise-linearly throughout the survey because, based on understanding of the causes for gain error, there is no physical justification to do so. Thus at least two styles of parameterization are required.

Piecewise constant parameterization

The first is a simple piecewise-constant parameterization. For each class of error a set of parameters is created. A list of all survey indices, day indices, flight indices, line indices, coilset indices and channel indices that are associated with each parameter is created and maintained. An airborne datum can only belong to one parameter from each calibration error category. This ensures that a unique set of parameters are defined for each datum. The piecewise-constant calibration values for each airborne datum are then simply 'looked-up' through a hierarchal index table. The partial derivative of a piecewise-constant calibration value for a given airborne datum with respect to a given parameter is unity if that datum is associated with that parameter, and is zero otherwise.

Piecewise linear parameterization

The second style of parameterization is the piecewise-linear which is illustrated in Figure 3.4. The figure shows how the inphase or quadrature bias in any channel is represented by a piecewise-linear curve throughout the flight. A set of N_b fiducials, denoted $f_1, f_2, \ldots, f_{N_b}$, ranging from the beginning to the end of the flight are chosen. These fiducials remain fixed throughout the inversion. Their spacing may be variable, but would normally be chosen at equal intervals so that the likely bias variation can be adequately represented (e.g. 20 minutes). At the top of Figure 3.4 the time-spans of seven hypothetical individual flight lines are shown. Note that there is no specific correspondence between the start and end positions of flight lines and the chosen fixed fiducials.

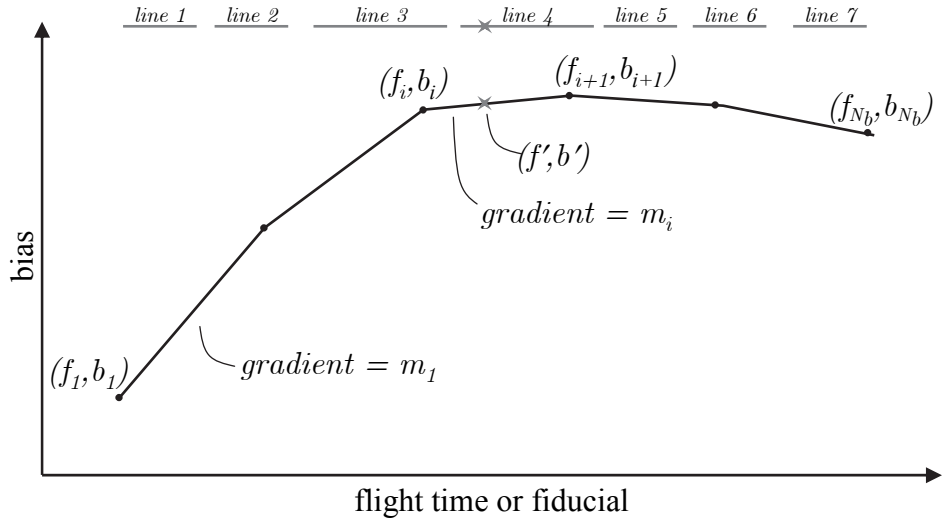

Figure 3.4 Illustration of a piecewise-linear parameterization of bias that varies slowly throughout a survey flight. The time-spans of a series of seven survey flight lines flown during the flight are shown at the top of the plot.

For the ensuing discussion, the bias at the ith fixed fiducial is denoted b_i, and the gradient of the bias between the ith and $(i+1)$th fixed fiducial is denoted m_i. Also, the bias at fiducial f', lying on the ith linear segment ($f_i \le f' < f_{i+1}$), has bias value denoted b'. Two variations on the piecewise-linear parameterization have been

formulated. Structurally, both formulations are identical, however the actual parameters solved for are different.

Bias parameterization by start value and gradients

In the first formulation the inversion parameters are the bias at the first fixed fiducial (b_1) and the gradients of each linear segment ($m_1, m_2, \ldots, m_{N_b-1}$). Here the bias can then be evaluated at any fiducial f' on the ith linear segment by the expression,

$$b' = b_1 + \left[\sum_{p=1}^{i-1} m_p (f_{p+1} - f_p)\right] + m_i (f' - f_i). \tag{3-4}$$

In Section 3.5.4, we will need to evaluate expressions for the partial derivatives of the bias values with respect to the parameters to be solved for. These are needed to construct the Jacobian matrix used in the minimization. It is convenient to write these here as follows,

$$\frac{\partial b'}{\partial b_1} = 1, \tag{3-5}$$

and,

$$\frac{\partial b'}{\partial m_k} = \begin{cases} f_{k+1} - f_k; & k < i \\ f' - f_k; & k = i \\ 0; & k > i \end{cases}. \tag{3-6}$$

Bias parameterization by values

In the second formulation the inversion parameters are simply the biases at each of the fixed fiducials ($f_1, f_2, \ldots, f_{N_b}$). The bias can be evaluated at any fiducial f' on the ith linear segment by the expression,

$$b' = b_i + (b_{i+1} - b_i)\left[\frac{f' - f_i}{f_{i+1} - f_i}\right] = b_i \left[\frac{f_{i+1} - f'}{f_{i+1} - f_i}\right] + b_{i+1}\left[\frac{f' - f_i}{f_{i+1} - f_i}\right]. \tag{3-7}$$

Again it is convenient to write the partial derivatives with respect to the inversion parameters here for later use in Section 3.5.4. They are given by,

$$\frac{\partial b'}{\partial b_k} = \begin{cases} \dfrac{f_{k+1} - f'}{f_{k+1} - f_k} & ;k = i \\ \dfrac{f' - f_k}{f_{k+1} - f_k} & ;k = i+1 \\ 0 & ;k < i \ and \ k >= i+1 \end{cases} . \tag{3-8}$$

Although both these forms can represent identical bias variations, each has its own advantages. The first mixes parameter types which leads to an extra level of algorithmic complexity. However it may be an ideal choice if the bias does not vary significantly once the transmitter-receiver assembly and the electronics reach a thermally stable state. In this case the prior on the gradients would be set to zero. Temporal roughness constraints can be applied to either of these alternatives by minimizing the first finite differences of adjacent gradient parameters, in the first case, or minimizing second finite differences of bias parameters in the second case.

In would also be possible to parameterize the bias by a linear variation over each individual survey line, and for example, solve for the bias at the beginning and end of each line. This line-wise method would be inherently less stable than the flight-wise methods described above. The flight-wise method can take advantage of the fact that the bias at the end of one survey line will be similar to the bias at the beginning of the following survey line. In essence, the flight-wise method is able to take advantage of the temporal continuity of the bias signal. However an important exception occurs when the system bias is reset to zero between survey lines during high altitude calibrations. This results in the bias becoming discontinuous, and the flight-wise parameterization will not be suitable. However, if recorded, these resets can be identified and reversed as noted in Section 4.2.4 (*cf.* Figure 4.7*c*).

3.4 Observed data

The observed data is comprised of up to three separate data types, or classes, that are called *airborne* data ($\mathbf{d}^{air}$), *geoelectric* data ($\mathbf{d}^{geo}$) and *interface-depth* data ($\mathbf{d}^{\mathrm{int}}$). The complete data vector can be written as,

$$\mathbf{d} = \left[\mathbf{d}^{air} \mid \mathbf{d}^{geo} \mid \mathbf{d}^{\mathrm{int}}\right]^{\mathbf{T}}. \tag{3-9}$$

The meaning and construction of each class of data are detailed in Sections 3.4.1, 3.4.2, and 3.4.3 that follow.

3.4.1 Observed airborne data

Since the aim of the holistic inversion is to processes and calibrate the data within the inversion, the airborne data are a substantially raw version of the frequency-domain electromagnetic data recorded at survey altitude. Substantially raw means the acquired data prior to any post-flight calibration (e.g. scaling, phase adjustments and drift) or levelling corrections having been applied, but possibly with high frequency noise rejection filtering applied. Tie-line and daily-repeat-line data may also be included in the inversion without any special treatment so long as they have been acquired in the same manner as the survey data.

The airborne data include the inphase (real) and quadrature (imaginary) parts of the ratio of the secondary to primary magnetic fields, in units of parts per million, for all coilsets. The inphase and quadrature channel data corresponding to the cth coilset of the sth airborne sample are denoted $d^{ip}_{s,c}$ and $d^{q}_{s,c}$ respectively. We then write the complex quantity as $d^{air}_{s,c} = d^{ip}_{s,c} + j d^{q}_{s,c}$ where, $j = \sqrt{-1}$. Complex vector and matrix arithmetic are not used in the inversion, so the real and imaginary parts are treated as two data. The complete vector of airborne data can thus be expressed as,

$$\mathbf{d^{air}} = \left[d^{ip}_{1,1}, d^{q}_{1,1}, \cdots, d^{ip}_{1,Nc}, d^{q}_{1,Nc}, \cdots, d^{ip}_{Ns,1}, d^{q}_{Ns,1}, \cdots, d^{ip}_{Ns,Nc}, d^{q}_{Ns,Nc}\right]^{\mathbf{T}}, \tag{3-10}$$

where N_s is the number of airborne samples in the inversion, and the airborne system has N_c coilsets. The total number of airborne data is $N_s \times N_c \times 2$.

The inversion uses airborne data that are 'substantially raw', i.e. the observed data without bias, gain, phase or other systematic error corrections having been applied. The reason why these data are used is that I wish to avoid any errors or inconsistencies that may have been introduced into the dataset during the standard data processing. Since the calibration does not simulate random errors, the data may however have been filtered to remove short wavelength effects of random noise sources such as powerline noise and sferic events. In many systems each channel is reset, or zeroed, during high altitude calibrations. As discussed in Section 3.3, this practice introduces instantaneous level shifts in the recorded data, which should be reversed if the piecewise-linear parameterization of bias over an entire flight is to be used.

3.4.2 Observed geoelectric data

The geoelectric data are the natural logarithm of conductivity 'measurements' at specific locations and over specific depth intervals. In general there are N_G subsurface conductivity measurements. The sth conductivity measurement is denoted σ_s^{geo}, and the geoelectric data vector can be written as,

$$\mathbf{d}^{\mathbf{geo}} = \left[\log(\sigma_1^{geo}), \log(\sigma_2^{geo}), \cdots, \log(\sigma_{N_G}^{geo})\right]^T. \tag{3-11}$$

As shown in Figure 3.5, each conductivity measurement has an associated known horizontal position (x_s^{geo}, y_s^{geo}), depth below the surface to the top of the measurement interval ($z_{s,top}^{geo}$), and depth to the bottom of the interval ($z_{s,bot}^{geo}$). Several measurements may share the same horizontal position.

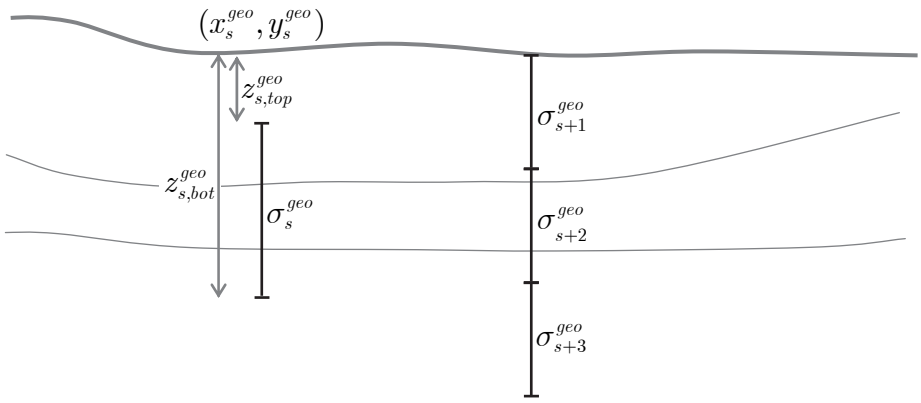

Figure 3.5 An illustration of four geoelectric data measured at different horizontal positions and/or depth intervals. Note that there is no particular relationship between layer interfaces in the conductivity model and the depth intervals

Possible sources of geoelectric data include downhole conductivity logs (e.g., EM39), surface electromagnetic surveys (e.g., EM-31), ground electromagnetic surveys (e.g., NanoTEM) and laboratory analysis (e.g., pore fluid conductivity and EC1:5 analyses). Although the term 'conductivity measurements' is used above, a more rigorously description might be 'interpretations of conductivity' because the conductivity values are predicted from geophysical data or laboratory analysis. However, the circumstances surrounding their measurement are typically more favourable than for airborne data. So in the holistic inversion scheme geoelectric data are considered to provide a form of 'ground truth' for constraining conductivity predictions. Nevertheless, the geoelectric data are assigned an associated uncertainty to reflect the level of confidence in them or lack thereof, because ground based data can also be noisy or miscalibrated.

Downhole conductivity logs typically measure the conductivity at 0.02 to 0.10 m depth intervals. This is certainly beyond the resolving capability of any airborne system. It is thus pragmatic to average the downhole data over depth intervals roughly equivalent to the vertical resolving ability of the airborne system, for example over 1 to 5 m intervals, before they are included in the holistic inversion.

3.4.3 Observed interface-depth data

Interface-depth data may include any information that is available regarding the depth, at a specific location, to a specific interface in the conceptual layered geological model. In general there are N_I interface-depth measurements. If the sth interface depth measurement is denoted d_s^{int}, the geoelectric data vector can be written as,

$$\mathbf{d}^{\text{int}} = \left[d_1^{\text{int}}, d_2^{\text{int}}, \cdots, d_{N_I}^{\text{int}}\right]^T . \tag{3-12}$$

As shown in Figure 3.6, each interface-depth measurement has an associated known horizontal position (x_s^{int}, y_s^{int}), and integer index (i_s^{int}) of the interface to which it refers. Different datum may refer to different interfaces, and several data may share the same horizontal position.

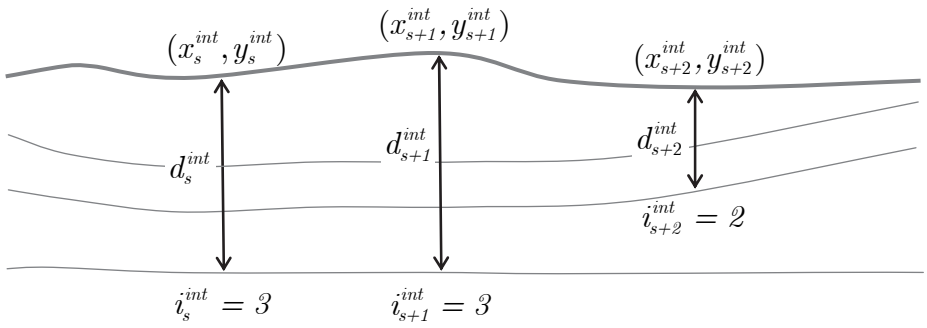

Figure 3.6 An illustration of three interface-depth data measured at different horizontal positions. Note that each datum may refer to different interfaces.

Interface-depth data are only befitting when a few-layer model is being used, in which case each layer in the conductivity model corresponds to a particular layer in the conceptual geological model. Interface-depth data might include the depth to the watertable measured in boreholes as used by Brodie et al. (2004b). It also might include observations of depths to geological boundaries observed in lithological logs.

3.5 Forward modelling and partial derivative calculations

3.5.1 Outline

In this section the mathematical simulation or prediction of a set of data that would be expected to be observed for a given set of model parameters is described. As explained in Section 3.1, the holistic inversion employs a 3D representation of subsurface conductivity, but for computational reasons 1D layered-earth electromagnetic forward modelling theory is used to approximate the airborne response. Because of this, the procedure for forward modelling of any particular datum is a two-stage process.

The first stage is the extraction of the ***local 1D conductivity profile*** from the 3D model at the location of the datum. If the datum is an airborne datum, this step also involves extraction of its ***local calibration values***. The procedure for this step and the governing equations are detailed in Sections 3.5.2 and 3.5.3.

The second stage takes the local 1D conductivity profile and the local calibration values and uses them to simulate the observed airborne, geoelectric and interface-depth data. Sections 3.5.4, 3.5.5 and 3.5.6, separately detail the forward modelling for each of the three data types. The predicted data for a given set of model parameters $\mathbf{m}$ is the combined result of three vector-valued nonlinear forward model functions, denoted as,

$$\boldsymbol{g}(\mathbf{m}) = \left[\boldsymbol{g}^{air}(\mathbf{m}) \mid \boldsymbol{g}^{geo}(\mathbf{m}) \mid \boldsymbol{g}^{\mathrm{int}}(\mathbf{m})\right]^{\mathbf{T}}, \tag{3-13}$$

that separately simulate the airborne, geoelectric and interface-depth data.

Since a gradient based scheme is used to iteratively minimize an objective function in the inversion, a Jacobian matrix must be constructed. It is therefore necessary to determine the partial derivatives of each datum with respect to each inversion model parameter. Computation of the derivatives firstly requires the derivatives of the data with respect to the local 1D profile layer properties and local calibration values. Then

the partial derivatives of the local layer properties and local calibration values with respect to the underlying conductivity and calibration model parameters are computed. These must then be combined, using the chain rule, to attain the derivatives of the data with respect to the actual inversion model parameters. Analytic expressions have been derived for all required derivative combinations. These are described in the relevant section for each class of data.

Shown at the top of Figure 3.7 is a hypothetical set of survey lines and a spline mesh. The location of two airborne samples that have 12 associated data from 6 coilsets and 2 channels (A and B) and a geoelectric datum (C) are shown on the mesh. The grid/matrix at the bottom of the figure schematically illustrates the influence of inversion model parameters on the predicted data for A, B and C. Where a grid square is coloured, that parameter has nonzero influence on the corresponding datum. Non-coloured squares indicate that a parameter has no influence on the datum, and hence the corresponding entry in the Jacobian matrix is zero.

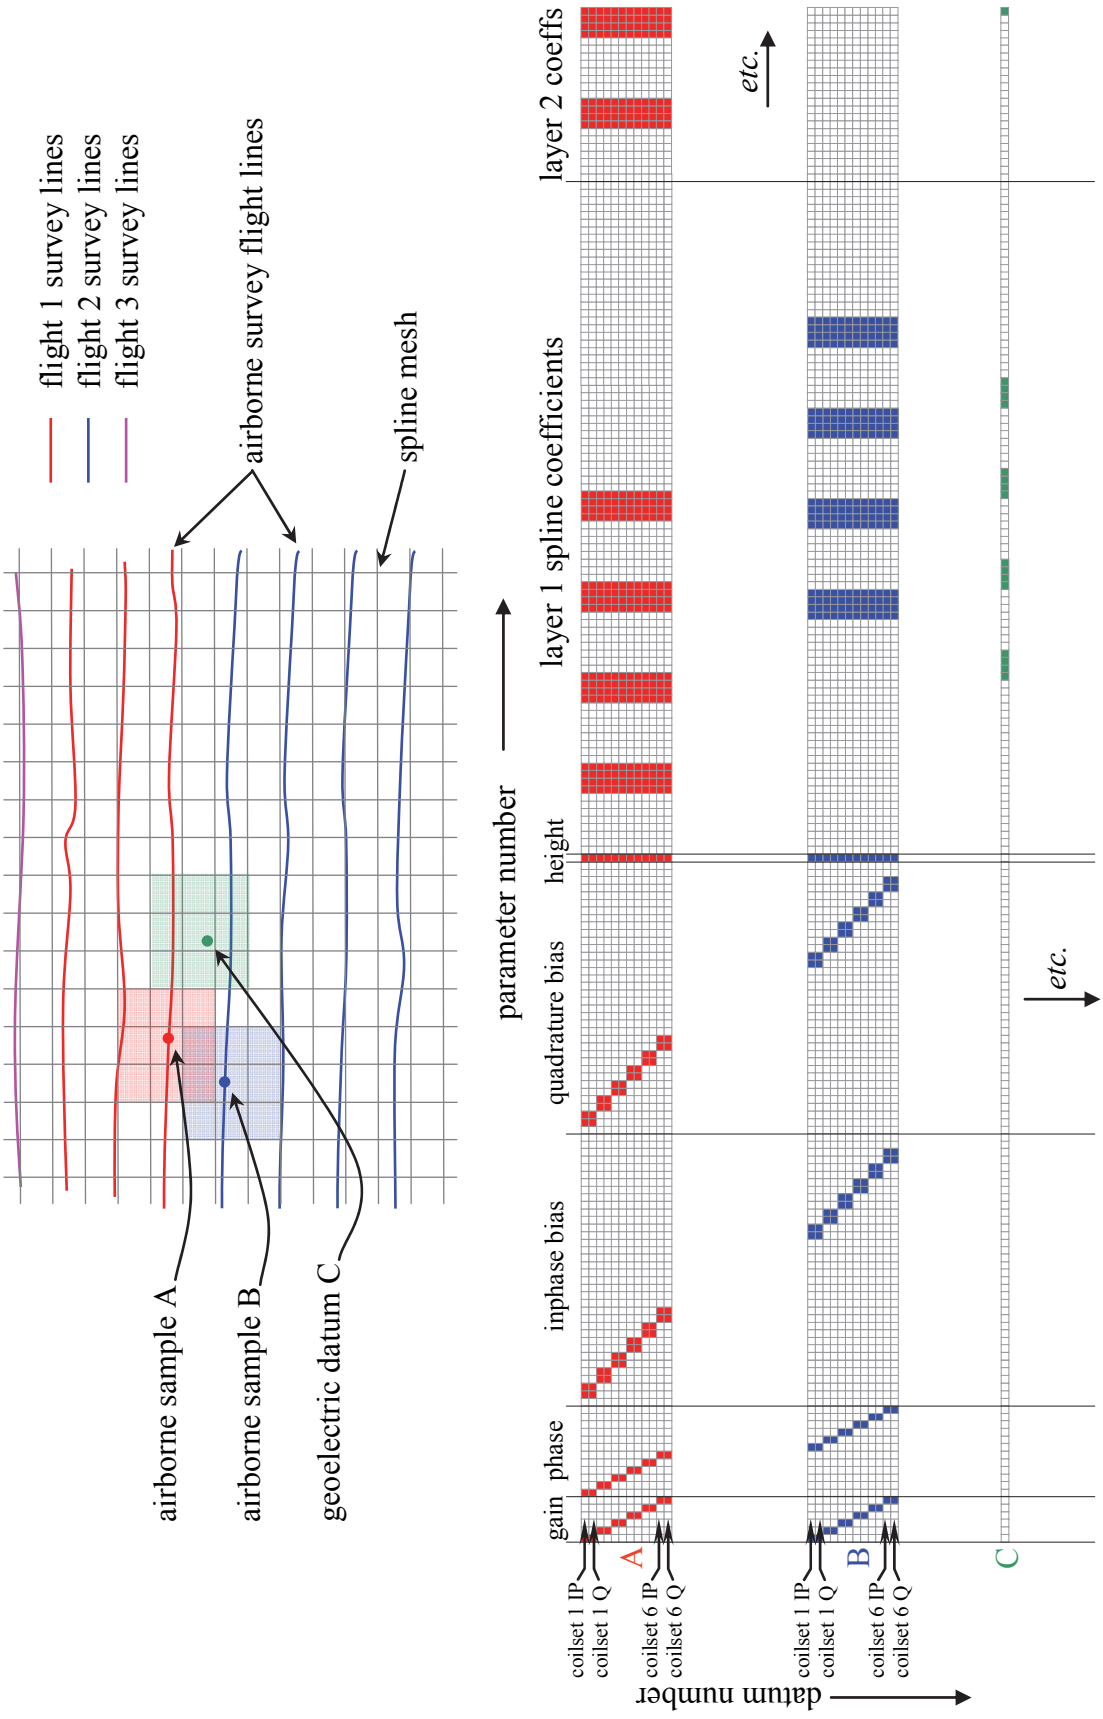

Figure 3.7 *Schematic illustration of how the parameterization cross-links data from different airborne samples, airborne flights and geoelectric measurement stations.*

3.5.2 Extraction of local 1D conductivity profile

The local 1D conductivity profile, corresponding to a given datum located at lateral position x,y and is on the ith column and jth row of surface patches, is the conductivities and thicknesses of the 3D model at the lateral position of that datum. It can be thought of as the 1D conductivity profile that would be encountered if you drill vertically down through the 3D model at the datum's position. The layer conductivities and thicknesses are denoted $\boldsymbol{\sigma}(x,y) = [\sigma_1(x,y), \sigma_2(x,y), \cdots, \sigma_{N_L}(x,y)]^{\mathrm{T}}$ and $\mathbf{t}(x,y) = [t_1(x,y), t_2(x,y), \cdots, t_{N_L}(x,y)]^{\mathrm{T}}$. The kth layer conductivity is calculated by evaluating the bicubic spline formula Equation 3-2, as,

$$\sigma_k(x,y) = \exp\left\{ \sum_{a=i-1}^{i+2} \sum_{b=j-1}^{j+2} w^{\sigma}_{abk}(x,y)\; c^{\sigma}_{abk} \right\}, \tag{3-14}$$

where c^{σ}_{abk} are the spline node coefficients and $w^{\sigma}_{abk}(x,y)$ are the known corresponding weights for the datum position. Similarly we can write $t_k(x,y)$ as,

$$t_k(x,y) = \exp\left\{ \sum_{a=i-1}^{i+2} \sum_{b=j-1}^{j+2} w^{t}_{abk}(x,y)\; c^{t}_{abk} \right\}. \tag{3-15}$$

The spline node coefficients c^{σ}_{abk} and c^{t}_{abk} are elements of the model parameter vector which are solved for in the inversion. The exponentials are required because the conductivities and thicknesses are parameterized by their logarithms. All datum that share a common lateral position, for example all the channels of data for an airborne sample, will share the same local 1D conductivity profile (i.e. it only has to be computed once per airborne sample).

In the following sections partial derivatives of the local layer conductivities and thicknesses, with respect to the underlying spline node coefficients, will be required. These are given by,

$$\frac{\partial \sigma_k(x,y)}{\partial c^{\sigma}_{uvk}} = \exp\left\{ \sum_{a=i-1}^{i+2} \sum_{b=j-1}^{j+2} w^{\sigma}_{abk}(x,y)\; c^{\sigma}_{abk} \right\} \times w^{\sigma}_{uvk}(x,y) = \sigma_k(x,y) w^{\sigma}_{uvk}(x,y) \tag{3-16}$$

and,

$$\frac{\partial t_k(x,y)}{\partial c^t_{uvk}} = \exp\left\{\sum_{a=i-1}^{i+2}\sum_{b=j-1}^{j+2} w^t_{abk}(x,y)\; c^t_{abk}\right\} \times w^t_{uvk}(x,y) = t_k(x,y) w^t_{uvk}(x,y)\,, \quad \text{(3-17)}$$

respectively. An important note regarding the spline node coefficient derivative expressions is that the weights $w^{\sigma}_{uvk}(x,y)$ and $w^{t}_{uvk}(x,y)$, and hence the derivatives, are nearly always zero because of the local-support of the bi-cubic spline basis functions. The weights are non-zero only when $i-1 \le u \le i+2$ and $j-1 \le v \le j+2$. Additionally, derivatives with respect to coefficients from any layer other than the *k*th layer are always zero. It is these properties that lead to a Jacobian matrix that is extremely sparse.

3.5.3 Extraction of local calibration values

Calibration values only apply to airborne data. They are used in the mathematical expression (Equation 3-19) that transforms the theoretical (calibrated) airborne response to the simulated observed (miscalibrated) airborne response. The local calibration values, corresponding to a particular airborne sample being forward modelled, can be thought of as the values that describe the state of the AEM system's calibration at the instant in time when the sample was measured, that is for the specific survey, day, flight, line, and fiducial. Depending on the chosen parameterization, there may or may not be different values for each sample, coilset and inphase/quadrature data channel.

In the ensuing development we will denote the local calibrations values for the *s*th airborne sample and the *c*th coilset with the subscripts s and c respectively. Thus, for the *s*th sample and the *c*th coilset we define the local gain calibration value as $s_{s,c}$, the local phase calibration value as $\theta_{s,c}$, the local complex bias calibration values as $b_{s,c} = b_{s,c,ip} + jb_{s,c,ip}$, and the local height calibration value Δh_s. Here j refers to the complex constant $\sqrt{-1}$.

The local calibration values must be computed from the parameters of the underlying calibration model. To do this, the model parameters that have non-zero effect on the particular airborne sample being forward modelled are selected from a hierarchal lookup-table. The lookup-table indexes the elements of the model parameter vector $\mathbf{m}$ against calibration values for each survey, day, flight, line, fiducial and coilset. For the calibration error classes that are parameterized in piecewise-constant fashion the local calibration values are just the selected model parameters. However, for the calibration errors that are parameterized in piecewise-linear fashion, two or more model parameters must be selected from the lookup-table. The local calibration values are then calculated from Equation 3-4 or Equation 3-7.

One calibration model parameter may apply to many airborne samples. For example, the phase calibration parameter $\theta_{s,c}$ for one coilset might apply to all the airborne samples acquired on a particular day of flying. It is through this mechanism, where many airborne samples are linked back to the same calibration model parameter, that the temporal coherency of the systematic calibration errors is exploited.

3.5.4 Airborne data forward model and partial derivatives

Here we consider the task of simulating the inphase g_{sc}^{ip} and quadrature g_{sc}^{q} datum corresponding to the cth coilset and the sth airborne sample. The sample was measured at lateral position x, y which lies on the ith column and jth row of surface patches. It was acquired when the transmitter receiver assembly was actually at height h_s^{pft} but was incorrectly measured to be at height h_s^{obs}.

The local 1D conductivity profile and the local calibration values are first computed as described in the preceding sections. Then, using the local conductivities $\boldsymbol{\sigma} = \boldsymbol{\sigma}(x, y)$, thicknesses $\mathbf{t} = \mathbf{t}(x, y)$ and height calibration value $\triangle h_s$, the layered-earth forward modelling routine (Chapter 2) is used to calculate the theoretical forward model. This

simulates the response for the system as if it were at the height $h_s^{obs} + \triangle h_s$ above ground, and for perfectly calibrated gain, phase and bias. This can be written as,

$$f_{sc}^{pft}(\boldsymbol{\sigma}, \mathbf{t}, h_s^{obs} + \triangle h_s) = f_{sc}^{ip}() + j f_{sc}^{q}(), \tag{3-18}$$

where $f_{sc}^{pft}()$ denotes the complex valued forward modelling function.

The calibration model equation (Equation 3-3) is then applied to simulate the transformation of the perfect system response into the simulated uncalibrated system response. The combined complex response is given by,

$$g_{sc}^{ip} + j g_{sc}^{q} = s_{sc} e^{j\theta_{sc}} \left[f_{sc}^{pft}(\boldsymbol{\sigma}, \mathbf{t}, h_s^{obs} + \triangle h_s) + b_{sc}^{ip} + j b_{sc}^{q} \right], \tag{3-19}$$

which can be separated into its constituent inphase and quadrature parts,

$$g_{sc}^{ip} = s_{sc} \left[(f_{sc}^{ip} + b_{sc}^{ip}) \cos\theta_{sc} - (f_{sc}^{q} + b_{sc}^{q}) \sin\theta_{sc} \right], \tag{3-20}$$

and,

$$g_{sc}^{q} = s_{sc} \left[(f_{sc}^{ip} + b_{sc}^{ip}) \sin\theta_{sc} + (f_{sc}^{q} + b_{sc}^{q}) \cos\theta_{sc} \right]. \tag{3-21}$$

In the following sections we derive the partial derivatives of the two airborne data g_{sc}^{ip} and d_{sc}^{q} with respect to the inversion model parameters of each class.

Conductivity and thickness parameter derivatives

To attain the derivative with respect to the conductivity spline coefficient located at the u,vth intersection of the kth layer mesh, we use Equations 3-20 and 3-21 with 3-16. After using the chain rule and dropping the (x,y) notation, this yields,

$$\frac{\partial g_{sc}^{ip}}{\partial c_{uvk}^{\sigma}} = \frac{\partial g_{sc}^{ip}}{\partial \sigma_k} \times \frac{\partial \sigma_k}{\partial c_{uvk}^{\sigma}} = s_{sc} \left[\frac{\partial f_{sc}^{ip}}{\partial \sigma_k} \cos\theta_{sc} - \frac{\partial f_{sc}^{q}}{\partial \sigma_k} \sin\theta_{sc} \right] \times \sigma_k \times w_{uvk}^{\sigma} \tag{3-22}$$

and,

$$\frac{\partial g_{sc}^{q}}{\partial c_{uvk}^{\sigma}} = \frac{\partial g_{sc}^{q}}{\partial \sigma_k} \times \frac{\partial \sigma_k}{\partial c_{uvk}^{\sigma}} = s_{sc} \left[\frac{\partial f_{sc}^{ip}}{\partial \sigma_k} \sin\theta_{sc} + \frac{\partial f_{sc}^{q}}{\partial \sigma_k} \cos\theta_{sc} \right] \times \sigma_k \times w_{uvk}^{\sigma} . \tag{3-23}$$

Expressions for derivatives with respect to the thickness spline coefficients can be derived in analogous fashion, this time using Equation 3-17, to give,

$$\frac{\partial g_{sc}^{ip}}{\partial c_{uvk}^{t}} = \frac{\partial g_{sc}^{ip}}{\partial t_k} \times \frac{\partial t_k}{\partial c_{uvk}^{t}} = s_{sc} \left[\frac{\partial f_{sc}^{ip}}{\partial t_k} \cos\theta_{sc} - \frac{\partial f_{sc}^{q}}{\partial t_k} \sin\theta_{sc} \right] \times t_k \times w_{uvk}^{t} , \tag{3-24}$$

and,

$$\frac{\partial g_{sc}^{q}}{\partial c_{uvk}^{t}} = \frac{\partial g_{sc}^{q}}{\partial t_k} \times \frac{\partial t_k}{\partial c_{uvk}^{t}} = s_{sc} \left[\frac{\partial f_{sc}^{ip}}{\partial t_k} \sin\theta_{sc} + \frac{\partial f_{sc}^{q}}{\partial t_k} \cos\theta_{sc} \right] \times t_k \times w_{uvk}^{t} . \tag{3-25}$$

The analytic expressions for the partial derivatives of the 1D electromagnetic forward model functions ($\partial f_{sc}^{ip} / \partial \sigma_k$, $\partial f_{sc}^{q} / \partial \sigma_k$, $\partial f_{sc}^{ip} / \partial t_k$ and $\partial f_{sc}^{q} / \partial t_k$) that are required above are derived in Appendix B.

Gain parameter derivatives

If p is a gain parameter, we can write the derivatives of the modelled inphase and quadrature data with respect to p as,

$$\frac{\partial g_{sc}^{ip}}{\partial p} = \frac{\partial g_{sc}^{ip}}{\partial s_{sc}} \frac{\partial s_{sc}}{\partial p} = \left[(f_{sc}^{ip} + b_{sc}^{ip}) \cos\theta_{sc} - (f_{sc}^{q} + b_{sc}^{q}) \sin\theta_{sc} \right] \frac{\partial s_{sc}}{\partial p} , \tag{3-26}$$

and,

$$\frac{\partial g_{sc}^{q}}{\partial p} = \frac{\partial g_{sc}^{q}}{\partial s_{sc}} \frac{\partial s_{sc}}{\partial p} = \left[(f_{sc}^{ip} + b_{sc}^{ip}) \sin\theta_{sc} + (f_{sc}^{q} + b_{sc}^{q}) \cos\theta_{sc} \right] \frac{\partial s_{sc}}{\partial p} . \tag{3-27}$$

If p is a parameter that influences the gain for the cth coilset and the sth airborne sample, the term $\partial s_{sc} / \partial p = 1$, otherwise $\partial s_{sc} / \partial p = 0$.

Phase parameter derivatives

Similarly, if p is a phase parameter, the derivatives with respect to p are,

$$\frac{\partial g_{sc}^{ip}}{\partial p} = \frac{\partial g_{sc}^{ip}}{\partial \theta_{sc}} \frac{\partial \theta_{sc}}{\partial p} = -s_{sc} \left[(f_{sc}^{ip} + b_{sc}^{ip}) \sin \theta_{sc} + (f_{sc}^{q} + b_{sc}^{q}) \cos \theta_{sc} \right] \frac{\partial \theta_{sc}}{\partial p}, \tag{3-28}$$

and,

$$\frac{\partial g_{sc}^{q}}{\partial p} = \frac{\partial g_{sc}^{q}}{\partial \theta_{sc}} \frac{\partial \theta_{sc}}{\partial p} = s_{sc} \left[(f_{sc}^{ip} + b_{sc}^{ip}) \cos \theta_{sc} - (f_{sc}^{q} + b_{sc}^{q}) \sin \theta_{sc} \right] \frac{\partial \theta_{sc}}{\partial p}. \tag{3-29}$$

If p is a parameter that influences the phase for the cth coilset and the sth airborne sample, the term $\partial \theta_{sc} / \partial p = 1$, otherwise $\partial \theta_{sc} / \partial p = 0$.

Inphase bias parameter derivatives

If p is an inphase bias parameter the derivatives with respect to p are,

$$\frac{\partial g_{sc}^{ip}}{\partial p} = \frac{\partial g_{sc}^{ip}}{\partial b_{sc}^{ip}} \frac{\partial b_{sc}^{ip}}{\partial p} = s_{sc} \cos \theta_{sc} \frac{\partial b_{sc}^{ip}}{\partial p}, \tag{3-30}$$

and,

$$\frac{\partial g_{sc}^{ip}}{\partial p} = \frac{\partial g_{sc}^{ip}}{\partial b_{sc}^{ip}} \frac{\partial b_{sc}^{ip}}{\partial p} = s_{sc} \sin \theta_{sc} \frac{\partial b_{sc}^{ip}}{\partial p}. \tag{3-31}$$

Since the local inphase bias calibration values may depend on two or more bias model parameters the term $\partial b_{sc}^{ip} / \partial p$ must be evaluated in alternate ways depending on the style of parameterization. If the start-bias/gradient style of parameterization (Equation 3-4) is used, we use Equations 3-5 and 3-6. Otherwise, if the alternate style of bias parameterization (Equation 3-7) is used, then $\partial b_{sc}^{ip} / \partial p$ is evaluated from Equation 3-8.

Quadrature bias parameter derivatives

If p is a quadrature bias parameter, the derivatives with respect to p are,

$$\frac{\partial g_{sc}^{q}}{\partial p} = \frac{\partial g_{sc}^{q}}{\partial b_{sc}^{q}} \frac{\partial b_{sc}^{q}}{\partial p} = -s_{sc} \sin \theta_{sc} \frac{\partial b_{sc}^{q}}{\partial p}, \tag{3-32}$$

and,

$$\frac{\partial g_{sc}^{q}}{\partial p} = \frac{\partial g_{sc}^{q}}{\partial b_{sc}^{q}} \frac{\partial b_{sc}^{q}}{\partial p} = s_{sc} \cos\theta_{sc} \frac{\partial b_{sc}^{q}}{\partial p}. \tag{3-33}$$

In analogous fashion to the inphase bias parameters, depending on the parameterization, the term $\partial b_{sc}^{q} \big/ \partial p$ must be evaluated from either Equations 3-5 and 3-6 or otherwise from Equation 3-8.

Height calibration parameter derivatives

If m_n is a height correction parameter, we can write the derivatives of the modelled inphase and quadrature data with respect to m_n as,

$$\frac{\partial g_{sc}^{ip}}{\partial p} = \frac{\partial g_{sc}^{ip}}{\partial \triangle h_s} \times \frac{\partial \triangle h_s}{\partial p} = s_{sc} \left[\frac{\partial f_{sc}^{ip}}{\partial \triangle h_s} \cos\theta_{sc} - \frac{f_{sc}^{q}}{\partial \triangle h_s} \sin\theta_{sc} \right] \times \frac{\partial \triangle h_s}{\partial p}, \tag{3-34}$$

and,

$$\frac{\partial g_{sc}^{q}}{\partial p} = \frac{\partial g_{sc}^{q}}{\partial \triangle h_s} \times \frac{\partial \triangle h_s}{\partial p} = s_{sc} \left[\frac{\partial f_{sc}^{ip}}{\partial \triangle h_s} \sin\theta_{sc} + \frac{\partial f_{sc}^{q}}{\partial \triangle h_s} \cos\theta_{sc} \right] \times \frac{\partial \triangle h_s}{\partial p}. \tag{3-35}$$

The analytic expressions for the partial derivatives of the 1D electromagnetic forward model functions ($\partial f_{sc}^{ip} \big/ \partial \triangle h_s$ and $\partial f_{sc}^{q} \big/ \partial \triangle h_s$), that are required above, are derived Appendix B. If m_n is a parameter that influences the height correction value for the sth airborne sample, the term $\partial \triangle h_s \big/ \partial p = 1$, otherwise $\partial \triangle h_s \big/ \partial p = 0$.

3.5.5 Geoelectric data forward model and partial derivatives

Here we consider the task of simulating a geoelectric datum g^{geo} which corresponds to a conductivity measurement made at position x,y that lies on the ith column and jth row of surface patches. The measurement is made over the depth interval from z_{top} to z_{bot}. The local 1D conductivity profile ($\boldsymbol{\sigma} = \boldsymbol{\sigma}(x,y)$ and $\mathbf{t} = \mathbf{t}(x,y)$) is first evaluated from the spline meshes as described in Section 3.5.2. The local thickness values are interrogated to determine the uppermost and lowermost layers, of the local 1D model, into which the measurement interval extends. The uppermost is denoted the pth layer

and the lowermost is denoted the qth layer. If the top of the interval coincides with the bottom interface of the pth layer, then it is said to begin in the pth layer.

The depth interval has a total length $L = z_{bot} - z_{top}$. The amount of overlap between the geoelectric depth interval and the kth layer in the local 1D model is determined, and are denoted l_k. The average conductivity over the depth interval is,

$$\overline{\sigma} = \frac{1}{L}\sum_{a=p}^{q}\sigma_a l_a \,. \tag{3-36}$$

Since the forward model of the geoelectric datum is the natural logarithm of the average conductivity, g^{geo} can be expressed as,

$$g^{geo} = \log\overline{\sigma} \,. \tag{3-37}$$

To derive the partial derivative of the geoelectric datum with respect to the conductivity spline coefficient at the u,vth intersection of the kth spline mesh, we first take the derivative with respect to the local 1D profile's layer conductivity,

$$\frac{\partial g^{geo}}{\partial \sigma_k} = \frac{\partial \log\overline{\sigma}}{\partial \overline{\sigma}} \times \frac{\partial \overline{\sigma}}{\partial \sigma_k} = \frac{1}{\overline{\sigma}}\frac{l_k}{L} \,. \tag{3-38}$$

Then using the chain rule and Equations 3-38 and 3-16, the derivative with respect to the spline node coefficient model parameter is,

$$\frac{\partial g^{geo}}{\partial c^{\sigma}_{uvk}} = \frac{\partial g^{geo}}{\partial \sigma_k} \times \frac{\partial \sigma_k}{\partial c^{\sigma}_{uvk}} = \frac{1}{\overline{\sigma}}\frac{l_k}{L} \times \sigma_k \times w^{\sigma}_{uvk} \,. \tag{3-39}$$

To derive the partial derivative of the geoelectric datum with respect to the thickness spline coefficient, we first take the derivative with respect to the local 1D profile's layer thickness. By direct inspection it is easy to see that,

$$\frac{\partial g^{geo}}{\partial t_k} = \begin{cases} \frac{1}{\overline{\sigma}L}(\sigma_p - \sigma_q); & k < p \\ \frac{1}{\overline{\sigma}L}(\sigma_k - \sigma_q); & p \le k < q \\ 0; & k \ge q \end{cases} , \tag{3-40}$$

which is to be used to evaluate the derivative with respect to the thickness spline coefficient given by the expression,

$$\frac{\partial g^{geo}}{\partial c^t_{uvk}} = \frac{\partial g^{geo}}{\partial t_k} \times \frac{\partial t_k}{\partial c^t_{uvk}} = \frac{\partial g^{geo}}{\partial t_k} \times t_k \times w^t_{uvk} \,. \tag{3-41}$$

As with the airborne data, the vast majority of the geoelectric data derivatives with respect to conductivity and thickness spline coefficients are zero because the weights w^{σ}_{uvk} and w^{t}_{uvk} are only non-zero when $i-1 \le u \le i+2$ and $j-1 \le v \le j+2$. All derivatives with respect to calibration model parameters are zero as they do not influence geoelectric data.

3.5.6 Interface-depth data forward model and partial derivatives

Here we consider the task of simulating an interface-depth datum g^{int} which corresponds to a depth measurement made at position x,y which lies on the ith column and jth row of surface patches. The measurement refers to the depth to the pth interface (i.e. bottom of the pth layer). The local 1D thicknesses ($\mathbf{t} = \mathbf{t}(x,y)$) are evaluated from the spline meshes as described in Section 3.5.2. The forward model is the cumulative thickness of all layers above the pth interface,

$$g^{\text{int}} = \sum_{a=1}^{p} t_a \,. \tag{3-42}$$

The derivative of g^{int} with respect to the local 1D model kth layer thickness is,

$$\frac{\partial g^{\text{int}}}{\partial t_k} = \begin{cases} 1; & k \le p \\ 0; & k > p \end{cases}, \tag{3-43}$$

which is to be used to evaluate the derivative with respect to the kth layer thickness spline coefficients given by the expression,

$$\frac{\partial g^{\text{int}}}{\partial c^t_{uvk}} = \frac{\partial g^{\text{int}}}{\partial t_k} \times \frac{\partial t_k}{\partial c^t_{uvk}} = \frac{\partial g^{\text{int}}}{\partial t_k} \times t_k \times w^t_{uvk} \,. \tag{3-44}$$

Again the vast majority of the interface-depth data derivatives with respect to thickness spline coefficients are zero because the weights w^t_{uvk} are only non-zero when $i-1 \leq u \leq i+2$ and $j-1 \leq v \leq j+2$. Also, the conductivity spline coefficient parameters and all of the calibration model parameters have no influence on interface-depth, and hence their corresponding derivatives are zero.

3.6 Inversion scheme

3.6.1 Outline

During the inversion a set of model parameters is iteratively updated until there is an acceptable level of agreement between the observed data and the data predicted from the model. The procedure is guided by the minimization of an objective function that is comprised of data misfit and model regularization terms. The data misfit term guides the inversion toward the principal goal of matching the observed and predicted data to within the expected noise levels. However because there is no unique solution, additional constraints are added through the model regularization terms. The model regularization penalizes the difference between the model and a priori reference values as well as spatial roughness of the conductivity and temporal roughness of the bias. Physically, this means that in cases where many alternate models satisfy the data equally well, the model with smaller penalty terms will be preferred, i.e. the model closer to the reference model and spatially smoother. Lagrange multipliers control the relative weight of the individual terms.

The iterative procedure begins with an initial set of inversion model parameters (spline coefficients and calibration parameters) that map to a set of reference values (layer properties and calibration values). The procedure consists of two logical loops (Figure 3.8). The outer loop constructs a linearized system of equations which are designed so that, when solved, they yield a new set of model parameters whose corresponding objective function value is smaller than that corresponding to the original

parameters at the start of the loop. It is in the inner loop that the linear system is actually solved via an iterative preconditioned conjugate gradient algorithm.

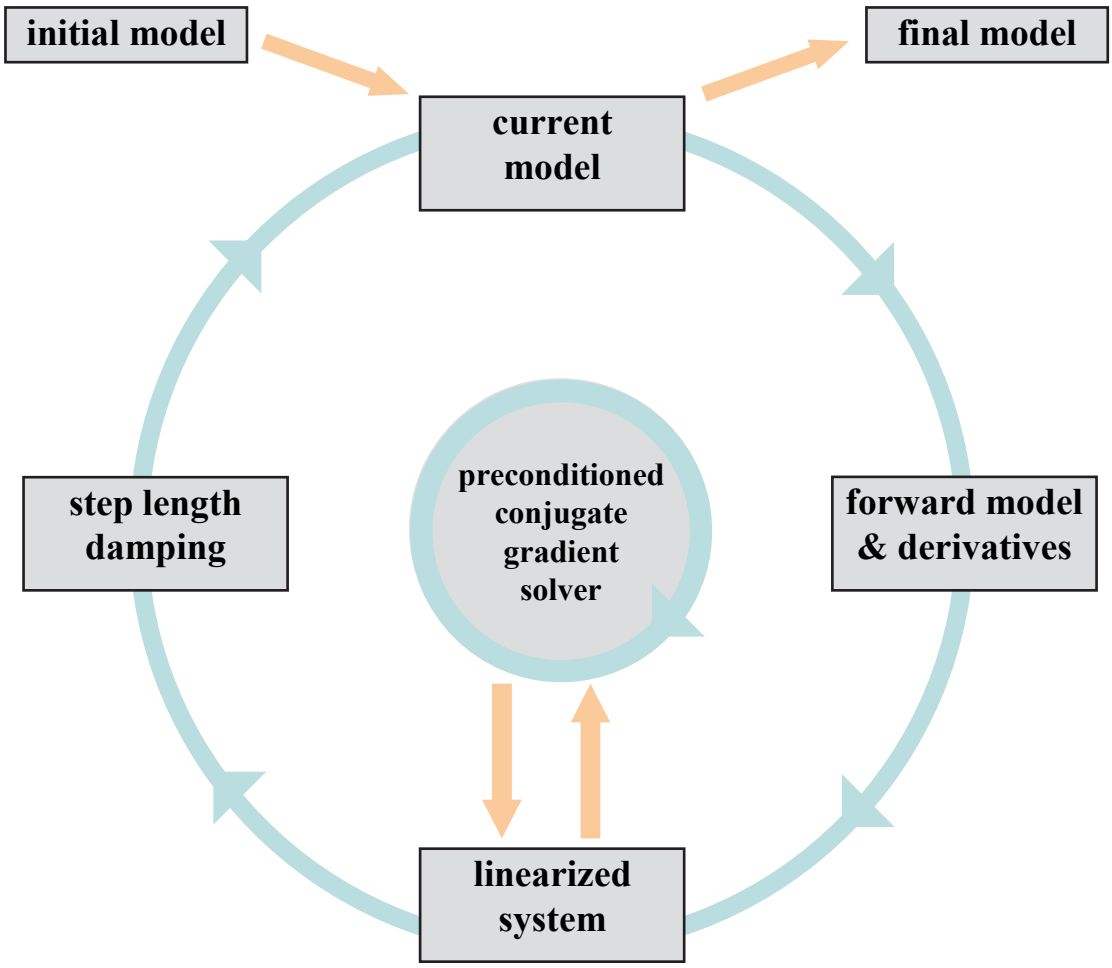

Figure 3.8 *Schematic illustration of the outer and inner loops of the iterative minimization scheme.*

To design the linear system within each outer loop iteration, knowledge that the objective function will be minimized when its derivative with respect to the new model parameters is zero, is capitalized upon. Thereupon, by differentiating the nonlinear objective function with respect to the new model parameters and equating the result to zero, equations that minimize the objective function are defined. These nonlinear equations are linearized via a two term Taylor series approximation, which leaves a set of linear equations to be solved within the inner loop.

The iterative procedure terminates when an acceptable level of data misfit is reached or the objective function can no longer be reduced. The following subsections describe the objective function and the iterative procedure used to minimize it.

3.6.2 Objective function definition

The inversion scheme seeks to minimize an objective function of the form,

$$\Phi(\mathbf{m}) = \Phi_d(\mathbf{m}) + \lambda_r\Phi_r(\mathbf{m}) + \lambda_x\Phi_x(\mathbf{m}) + \lambda_y\Phi_y(\mathbf{m}) + \lambda_z\Phi_z(\mathbf{m}) + \lambda_b\Phi_b(\mathbf{m}), \quad (3\text{-}45)$$

where $\Phi_d(\mathbf{m})$ is a data misfit term, $\Phi_r(\mathbf{m})$ is a reference-model misfit term, $\Phi_x(\mathbf{m})$, $\Phi_y(\mathbf{m})$ and $\Phi_z(\mathbf{m})$ are east–west, north–south and vertical model roughness terms respectively, and $\Phi_b(\mathbf{m})$ is a bias roughness term. The $\lambda's$ are multipliers that weight the relative importance of each term.

The model regularization objective function terms used here are a subset of the broader class of L_2 model-structure measures, which are widely used in geophysical inversion to constrain solutions (e.g. Constable et al., 1987; Menke, 1989; Farquharson and Oldenburg, 1998). A nuance in the way they are used here, is that the constraints are not imposed directly on the inversion model parameters. The justification for this is that it is not particularly physically intuitive to directly place constraints on the somewhat abstract spline coefficient parameters. Instead, I impose the constraints on the more physically meaningful layer properties that the coefficients parameterize.

Data misfit

The data misfit Φ_d is defined in the usual noise weighted least-squares sense,

$$\Phi_d(\mathbf{m}) = \left[\mathbf{d} - \boldsymbol{g}(\mathbf{m})\right]^{\mathbf{T}} \mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m})\right]. \quad (3\text{-}46)$$

Here $\mathbf{d} = (\mathbf{d}^{air} \mid \mathbf{d}^{geo} \mid \mathbf{d}^{\mathrm{int}})^{\mathbf{T}}$ is the N_d length vector of data made up of the observed airborne, geoelectric and interface-depth data. The vector $\mathbf{m} = (\mathbf{m}^{cal} \mid \mathbf{m}^{ear})^{\mathbf{T}}$ is the N_m length vector of unknown calibration and conductivity model parameters to be

estimated. The function $\boldsymbol{g}(\mathbf{m})$ is the nonlinear vector-valued forward model function that maps the model parameters to predictions of data (i.e. Equations 3-19, 3-37, and 3-42).

The $N_d \times N_d$ matrix $\mathbf{C}_d$ expresses covariance of the errors expected in the observed data. If the errors are assumed to be uncorrelated and normally distributed then $\mathbf{C}_d\big|_{ii}$ is the variance of the error associated with the ith datum, and all off-diagonal entries of $\mathbf{C}_d$ are zero. It is instructive to reiterate here that, as explained in Section 3.3.1, the errors are the random errors only and not the systematic errors, which in the holistic inversion are parameterized and solved for in the inversion.

Reference model misfit

The reference model misfit part of the objective function Φ_r allows constraint to be imposed on the solution via prior expectation of the conductivity structure that is likely to be encountered in the survey area, and the likely state of calibration of the AEM system. These are probabilistic, or soft, constraints because they do not set hard bounds on model parameters. Nor are they intended to mimic hard evidence (i.e. measurements). In the holistic inversion, information from specific located measurements is utilised through the geoelectric and interface-depth data.

The regularisation encourages the solution toward the replication of a reference set of layer conductivity, layer thickness, and calibration values which might be expected if the airborne, geoelectric and interface-depth data were not available. The prior expected values might be derived, for example, from statistical analysis of downhole conductivity data from the survey area, or from a geological environment known to be similar. If the prior information available is limited (e.g. just a few nearby bores), then the reference values will necessarily be spatially constant for conductivity and thickness values and temporarily constant for calibration values. On the other hand, the prior information may be abundant enough to generate a reference model with spatially

varying layer properties (e.g. by gridding of prior conductivity information) and temporarily varying calibration values (e.g. from high altitude zero-level measurements).

To implement the constraints, a set of N_m reference values denoted $\mathbf{r}^0$ are selected from the prior information. For reasons mentioned earlier, the reference values are not the somewhat abstract spline coefficients or bias gradients. Instead they are the more physically meaningful logarithms of layer conductivities, logarithms of layer thicknesses, and calibration values. Each reference value has an associated position or acquisition time, depending on whether it is a conductivity, thickness or calibration value. For simplicity, the conductivity and thickness reference values are always chosen to coincide with the spline mesh intersections and with individual calibration parameters. A prior level of uncertainty can be placed on the reference values via a $N_r \times N_r$ covariance matrix $\mathbf{C}_r$. This would typically be a diagonal matrix in which $\mathbf{C}_r\big|_{ii}$ is the variance of the uncertainty on the ith reference value.

A $N_m \times N_m$ linear operator matrix $\mathbf{S}$ that maps the underlying model parameters $\mathbf{m}^0$ to the elements of $\mathbf{r}^0$ can be defined such that,

$$\mathbf{S}\mathbf{m}^0 = \mathbf{r}^0. \tag{3-47}$$

The operator $\mathbf{S}$ is sparse and known through definition of the model. That is, the columns of $\mathbf{S}$ corresponding to the conductivity model spline coefficients will be filled partly with the weights required for the evaluation of the spline functions at the locations corresponding to the elements of $\mathbf{r}^0$. The columns of $\mathbf{S}$ corresponding to calibration model inversion parameters will hold the value unity on the diagonal and zero elsewhere, except when the parameterization is piecewise linear, in which case two nonzero values will be needed.

Since the aim is to encourage the reference values to be replicated by the model, an appropriate function to be minimised is a least squares difference measure weighted by the prior uncertainty on the reference values, which is written as,

$$\Phi_r = \left[\mathbf{r}_0 - \mathbf{Sm}\right]^{\mathrm{T}} \mathbf{C}_r^{-1} \left[\mathbf{r}_0 - \mathbf{Sm}\right]. \tag{3-48}$$

Conductivity model roughness

Although the 3D conductivity model is locally smooth in lateral directions by virtue of the spline parameterization, a greater degree of smoothness can be enforced by introducing the roughness terms Φ_x, Φ_y and Φ_z into the objective function. By minimising these measures of conductivity model roughness in each direction, a longer wavelength smoothness in the model is encouraged. Each term quantifies the cumulative roughness over the whole model by the norm of the second derivatives of the layer conductivities logarithms and/or thicknesses calculated at spline mesh intersections. Again, because it is the layer properties and not the coefficients that we wish to be smooth, the roughness measure is based on the logarithms of the layer properties that the model coefficients represent, rather than the coefficients themselves.

The second derivatives are approximated by second finite difference calculations between property value triplets at adjacent intersections. Considering now the second derivative calculation in the east–west direction (for Φ_x) at the i,jth intersection of the kth mesh. The logarithm of the property value at that intersection is just one of the elements of the vector $\mathbf{r} = \mathbf{Sm}$, which we will denote $r_{i,j,k}$. The second derivative approximation is then,

$$l_{ijk}^{x} = \frac{r_{i+1,j,k} - r_{i,j,k}}{\triangle x} - \frac{r_{i,j,k} - r_{i-1,j,k}}{\triangle x} = \frac{1}{\triangle x}\left[r_{i-1,j,k} - 2r_{i,j,k} + r_{i+1,j,k}\right], \tag{3-49}$$

where $\triangle x$ is the node spacing. Therefore a vector of derivatives $\mathbf{l}_x$ can be built up by the product,

$$\mathbf{l}_x = \mathbf{L}_x\mathbf{r} = \mathbf{L}_x\mathbf{S}\mathbf{m}\,, \tag{3-50}$$

where $\mathbf{L}_x$ is a $N_{lx}\times N_m$ roughening matrix of the form,

$$\mathbf{L}_x = \frac{1}{\triangle x}\begin{bmatrix} \cdots & 1 & \cdots & -2 & \cdots & 1 & & & \\ & & & 1 & \cdots & -2 & \cdots & 1 & \\ & & & & \cdots & 1 & \cdots & -2 & \cdots & 1 \\ & & & & & & & \vdots & & \vdots \end{bmatrix}. \tag{3-51}$$

Each row of $\mathbf{L}_x$ corresponds to a triplet of adjacent spline nodes along the east–west direction. The number of rows is equal to the total number of mesh intersections less those on the eastern and western borders, because differences cannot be applied on edge nodes. If there is no motivation to penalise roughness of the layer thicknesses, then the corresponding rows are simply omitted from the roughening matrix. The precise column positioning of the nonzero $\begin{bmatrix}1 & \cdots & -2 & \cdots & 1\end{bmatrix}$ entries in each row depends on the ordering of the coefficients in the model parameter vector $\mathbf{m}$. The overall model roughness in the east–west direction to be minimised is the norm of $\mathbf{l}_x$,

$$\Phi_x = \mathbf{l}_x^{\mathrm{T}}\mathbf{l}_x = \mathbf{m}^{\mathrm{T}}\mathbf{S}^{\mathrm{T}}\mathbf{L}_x^{\mathrm{T}}\mathbf{L}_x\mathbf{S}\mathbf{m}\,. \tag{3-52}$$

Through the same line of reasoning the roughness in the north–south and vertical directions are given by,

$$\Phi_y = \mathbf{m}^{\mathrm{T}}\mathbf{S}^{\mathrm{T}}\mathbf{L}_y^{\mathrm{T}}\mathbf{L}_y\mathbf{S}\mathbf{m}\,, \tag{3-53}$$

and,

$$\Phi_z = \mathbf{m}^{\mathrm{T}}\mathbf{S}^{\mathrm{T}}\mathbf{L}_z^{\mathrm{T}}\mathbf{L}_z\mathbf{S}\mathbf{m}\,. \tag{3-54}$$

Since it does not make sense to apply vertical roughness constraints to layer thicknesses the rows of $\mathbf{L}_z$ that correspond to thickness spline meshes parameters are omitted.

There are a number of alternative measures of model-structure that could be used if it was deemed appropriate for a particular survey area. For example, it may be appropriate to promote flatness rather than smoothness, in which case the first derivatives would be minimised and the $\begin{bmatrix}1 & \cdots & -2 & \cdots & 1\end{bmatrix}$ entries of the $\mathbf{L}$ matrices

would become $\begin{bmatrix} -1 & \cdots & 1 & \cdots \end{bmatrix}$ (*cf.* Menke, 1989, p. 53). In some circumstances it might be considered appropriate to encourage smoothness or flatness of the model's deviation from the reference model. In this case, expressions of the form $\Phi_x = (\mathbf{m} - \mathbf{m}^0)^{\mathrm{T}} \mathbf{S}^{\mathrm{T}} \mathbf{L}_x^{\mathrm{T}} \mathbf{L}_x \mathbf{S} (\mathbf{m} - \mathbf{m}^0)$ would be used in Equation 3-52. Additional terms could be added, for example, to encourage smoothness and flatness simultaneously.

Conductivity model roughness with respect to elevation

In some geological settings, for example where conductivity is controlled by weathering, it is quite reasonable to expect conductivity variations to follow the topography as is illustrated by the colours on the left hand side of Figure 3.9. However, in many other geological environments, for example where conductivity is hydrogeologically controlled, it is more reasonable to expect the conductivity to vary according to elevation above sea level as shown by colours on the right hand side of Figure 3.9. In the latter case it makes sense to apply lateral roughness constraints that promote smoothness with respect to elevation.

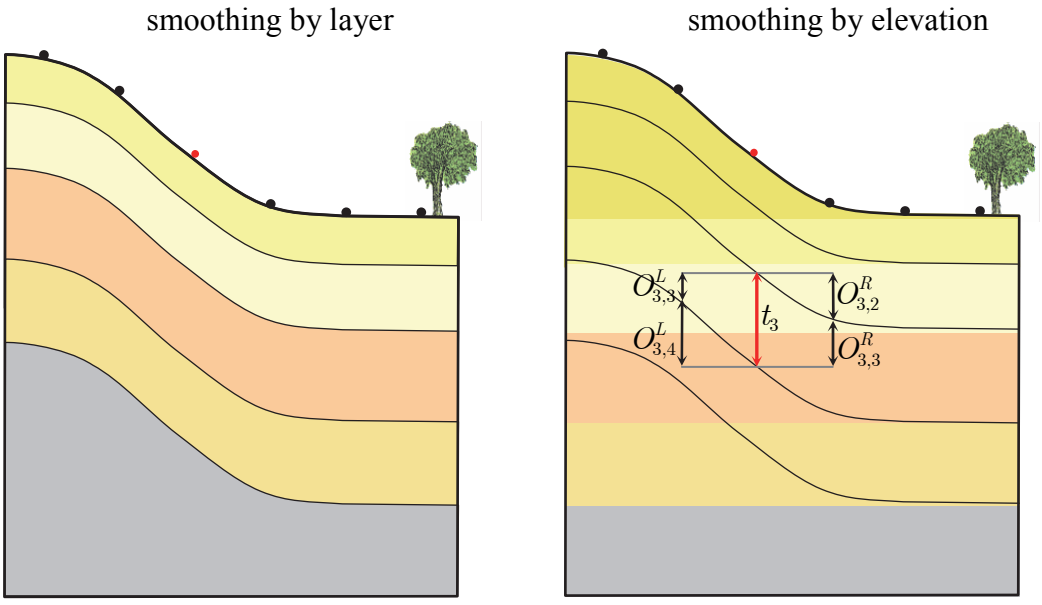

Figure 3.9 *Lateral roughness constraints can be imposed to encourage the conductivity model to be smooth with respect to either depth (left) or elevation (right).*

However, the objective function terms Φ_x and Φ_y discussed so far only promote smoothness of layer conductivities, irrespective of changes in the layer's elevation. Therefore, to promote smoothness with respect to elevation an alternative form of the lateral roughening matrices needs to be used.

To understand how this is achieved, refer to the left hand side of Figure 3.9 where spline nodes are marked on an uneven topographic surface. Consider the determination of the entries in the roughening matrix corresponding to the third layer spline mesh at the intersection shown by the red node marker. The amount of overlap between the elevation ranges of the *i*th layer at the red node, and the *j*th layer at the nodes to its immediate left and right are determined and denoted as $o_{i,j}^{L}$ and $o_{i,j}^{R}$ respectively. The nonzero overlaps are shown in Figure 3.9 for the third layer at the red node. The corresponding roughening matrix will have the form,

$$\mathbf{L}_x = \frac{1}{\triangle x}\begin{bmatrix} \ddots & & \ddots & & \ddots & & \ddots & & \ddots & & \ddots \\ \cdots & \frac{o_{3,3}^{L}}{t_3} & \cdots & \frac{o_{3,4}^{L}}{t_3} & \cdots & -2 & \cdots & \frac{o_{3,2}^{L}}{t_3} & \cdots & \frac{o_{3,3}^{L}}{t_3} & \cdots \\ & \ddots & & \ddots & & & & \ddots & & \ddots & \end{bmatrix}. \tag{3-55}$$

The entries on each sides of the *-2* entry sum to unity, and they have the effect of averaging the conductivity at the adjacent left and right nodes over the same elevation range as the central node. Typically, there will be five non-zero entries in each row of the roughness matrices. However, in general, there may be more or less nonzero entries depending on the topographic variations and layer thicknesses. This alternative formulation is not recommended for variable layer thickness inversions because the roughening matrices would need to be recalculated at each iteration due to the changing thicknesses.

Bias roughness

Since the postulated cause of the variation in bias is thought to be gradual change in temperature during the flight (Huang and Fraser, 1999; Valleau, 2000), it is expected

that the bias varies gradually during a survey flight. Observations of zero-level made at high altitude during survey flights discussed in Chapter 1 (Figure 1.10) confirm this behaviour. To take advantage of this temporal coherency an additional Φ_b regularization term is included.

Similarly to the conductivity roughness terms discussed above, the bias roughness term quantifies the total roughness, over all flights and channels, via the norm of the second derivatives approximated by finite difference operators. This is expressed as,

$$\Phi_b = \mathbf{m}^{\mathrm{T}}\mathbf{S}^{\mathrm{T}}\mathbf{L}_b^{\mathrm{T}}\mathbf{L}_b\mathbf{S}\mathbf{m}\,, \tag{3-56}$$

where $\mathbf{L}_b$ is a finite difference operator. If the bias is parameterized by gradients, the rows of $\mathbf{L}_b$ contain entries in the form $\begin{bmatrix}-1 & \cdots & 1 & \cdots\end{bmatrix}$, and if parameterized by bias values directly, the entries are in the form $\begin{bmatrix}1 & \cdots & -2 & \cdots & 1\end{bmatrix}$.

It is expected that gain and phase calibration errors are largely related to the on-ground location at which they were calibrated, mispositioning of calibration instruments and human error (Fitterman, 1998). Since such errors are likely to be independent from day to day rather than smoothly varying, there is no motivation for penalizing their roughness in the parameterization scheme used here. The parameterization of height calibration values yields only one parameter, so regularization is not relevant in this case. If, for example, height was parameterized by a separate value at each airborne sample, an along line roughness penalty term would be an appropriate way to take advantage of the knowledge that the transmitter-receiver assembly's height always varies smoothly.

3.6.3 Objective function minimization

Iterative procedure

If the objective function Φ was quadratic, with respect to the model parameters, it would be possible to directly solve for the desired model by classical linear inverse

theory techniques (e.g. Menke, 1989; Lawson and Hanson, 1974). However, since the relationship between the model parameters and the data is nonlinear, the forward model function $\boldsymbol{g}(\mathbf{m})$ and $\Phi(\mathbf{m})$ are also nonlinear, and hence $\Phi(\mathbf{m})$ is not quadratic in $\mathbf{m}$. This means we must resort to iterative nonlinear inverse techniques (e.g. Marquardt, 1963; Lines and Treitel, 1984).

An initial estimate of the model parameters $\mathbf{m} = \mathbf{m}^0$ is made by using the conjugate gradient algorithm (Section 3.6.4) to solve the linear system,

$$\mathbf{r}^0 = \mathbf{S}\mathbf{m}^0, \tag{3-57}$$

which finds the initial model parameters that would replicate the reference values. Then an iterative loop is set up, which in the nth iteration aims to update the current model parameters $\mathbf{m}^n$ to a new estimate $\mathbf{m}^{n+1}$, such that the objective function is reduced (i.e. $\Phi(\mathbf{m}^{n+1}) < \Phi(\mathbf{m}^n)$).

The problem of deciding on a new $\mathbf{m}^{n+1}$ can be posed in terms of solving for a model perturbation $\triangle\mathbf{m}^n = \mathbf{m}^{n+1} - \mathbf{m}^n$ or of solving directly for $\mathbf{m}^{n+1}$. These are known as '*creeping'* and '*jumping'* methods respectively (Parker, 1994). An in depth discussion of the differences is provided in Scales et al. (1990). In this thesis, the jumping method is used because it is more intuitive, and it is simpler to impose the regularisation constraints in terms of the model rather than the model perturbations.

Determining the new model

To make an appropriate choice of the so far unknown $\mathbf{m}^{n+1}$ we make use of the knowledge that the derivative of the objective function with respect to the new model parameters will be zero at the minimum. That condition is,

$$\frac{\partial\Phi}{\partial\mathbf{m}^{n+1}} = \frac{\partial\Phi_d}{\partial\mathbf{m}^{n+1}} + \lambda_r\frac{\partial\Phi_r}{\partial\mathbf{m}^{n+1}} + \lambda_x\frac{\partial\Phi_x}{\partial\mathbf{m}^{n+1}} + \lambda_y\frac{\partial\Phi_y}{\partial\mathbf{m}^{n+1}} + \lambda_z\frac{\partial\Phi_z}{\partial\mathbf{m}^{n+1}} + \lambda_b\frac{\partial\Phi_b}{\partial\mathbf{m}^{n+1}} = 0\ . \tag{3-58}$$

To proceed, the individual parts of Equation 3-58 must be determined. Using Equation 3-46 the derivative of the first term is,

$$\frac{\partial \Phi_d}{\partial \mathbf{m}^{n+1}} = \frac{\partial}{\partial \mathbf{m}^{n+1}}\left[\left[\mathbf{d}^{obs} - \boldsymbol{g}(\mathbf{m}^{n+1})\right]^{\mathbf{T}} \mathbf{C}_d^{-1}\left[\mathbf{d}^{obs} - \boldsymbol{g}(\mathbf{m}^{n+1})\right]\right]. \tag{3-59}$$

So far $\boldsymbol{g}(\mathbf{m}^{n+1})$ is unknown, however it can be approximated by assuming that $\boldsymbol{g}()$ is linear in the vicinity of $\mathbf{m}^n$ and expanding it as a two term Taylor series,

$$\boldsymbol{g}(\mathbf{m}^{n+1}) \cong \boldsymbol{g}(\mathbf{m}^n) + \mathbf{G}_n(\mathbf{m}^{n+1} - \mathbf{m}^n). \tag{3-60}$$

Here $\mathbf{G}_n$ is the Jacobian matrix,

$$\mathbf{G}_n\Big|_{ij} = \frac{\partial g_i}{\partial m_j^n} = \frac{\partial \boldsymbol{g}_i(\mathbf{m}^n)}{\partial m_j^n}, \tag{3-61}$$

or, in other words, the entry at the ith row and jth column of $\mathbf{G}_n$ is the partial derivative of the ith predicted datum with respect to the jth model parameter evaluated at point $\mathbf{m}^n$ in model space. All the expression required for computing the entries of the Jacobian matrix have been provided in Section 3.5.

Making use of Equation 3-60, we can now rewrite Equation 3-59 as,

$$\frac{\partial \Phi_d}{\partial \mathbf{m}^{n+1}} = \frac{\partial\left[\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^n) - \mathbf{G}_n(\mathbf{m}^{n+1} - \mathbf{m}^n)\right]^{\mathbf{T}} \mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^n) - \mathbf{G}_n(\mathbf{m}^{n+1} - \mathbf{m}^n)\right]\right]}{\partial \mathbf{m}^{n+1}} \tag{3-62}$$

Since neither $\mathbf{d}$, $\boldsymbol{g}(\mathbf{m}^n)$ nor $\mathbf{G}_n$ depend on $\mathbf{m}^{n+1}$, and noting the identity

$$\frac{\partial \mathbf{a}(x)^T \mathbf{B}\mathbf{a}(x)}{\partial x} = 2 \times \frac{\partial \mathbf{a}(x)^{\mathbf{T}}}{\partial x} \times \mathbf{B} \times \mathbf{a}(x),$$

Equation 3-62 reduces to,

$$\frac{\partial \Phi_d}{\partial \mathbf{m}^{n+1}} = -2\mathbf{G}_n^{\mathbf{T}}\mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^n) - \mathbf{G}_n(\mathbf{m}^{n+1} - \mathbf{m}^n)\right]. \tag{3-63}$$

Similarly, noting that $\mathbf{r}^0$, $\mathbf{S}$, $\mathbf{L}_x$, $\mathbf{L}_y$, $\mathbf{L}_z$ and $\mathbf{L}_b$ are all independent of $\mathbf{m}^{n+1}$, differentiating the remainder of the terms in Equation 3-58 yields,

$$\frac{\partial \Phi_r}{\partial \mathbf{m}^{n+1}} = \frac{\partial \left[\mathbf{r}^0 - \mathbf{S}\mathbf{m}^{n+1}\right]^{\mathrm{T}} \mathbf{C}_r^{-1} \left[\mathbf{r}^0 - \mathbf{S}\mathbf{m}^{n+1}\right]}{\partial \mathbf{m}^{n+1}} = -2\mathbf{S}^{\mathrm{T}}\mathbf{C}_r^{-1}\left[\mathbf{r}^0 - \mathbf{S}\mathbf{m}^{n+1}\right], \tag{3-64}$$

$$\frac{\partial \Phi_x}{\partial \mathbf{m}^{n+1}} = \frac{\partial (\mathbf{m}^{n+1})^{\mathrm{T}} \mathbf{S}^{\mathrm{T}} \mathbf{L}_x^{\mathrm{T}} \mathbf{L}_x \mathbf{S}\mathbf{m}^{n+1}}{\partial \mathbf{m}^{n+1}} = 2\mathbf{S}^{\mathrm{T}} \mathbf{L}_x^{\mathrm{T}} \mathbf{L}_x \mathbf{S}\mathbf{m}^{n+1}, \tag{3-65}$$

$$\frac{\partial \Phi_y}{\partial \mathbf{m}^{n+1}} = \frac{\partial (\mathbf{m}^{n+1})^{\mathrm{T}} \mathbf{S}^{\mathrm{T}} \mathbf{L}_y^{\mathrm{T}} \mathbf{L}_y \mathbf{S}\mathbf{m}^{n+1}}{\partial \mathbf{m}^{n+1}} = 2\mathbf{S}^{\mathrm{T}} \mathbf{L}_y^{\mathrm{T}} \mathbf{L}_y \mathbf{S}\mathbf{m}^{n+1}, \tag{3-66}$$

$$\frac{\partial \Phi_z}{\partial \mathbf{m}^{n+1}} = \frac{\partial (\mathbf{m}^{n+1})^{\mathrm{T}} \mathbf{S}^{\mathrm{T}} \mathbf{L}_z^{\mathrm{T}} \mathbf{L}_z \mathbf{S}\mathbf{m}^{n+1}}{\partial \mathbf{m}^{n+1}} = 2\mathbf{S}^{\mathrm{T}} \mathbf{L}_z^{\mathrm{T}} \mathbf{L}_z \mathbf{S}\mathbf{m}^{n+1}. \tag{3-67}$$

$$\frac{\partial \Phi_b}{\partial \mathbf{m}^{n+1}} = \frac{\partial (\mathbf{m}^{n+1})^{\mathrm{T}} \mathbf{S}^{\mathrm{T}} \mathbf{L}_b^{\mathrm{T}} \mathbf{L}_b \mathbf{S}\mathbf{m}^{n+1}}{\partial \mathbf{m}^{n+1}} = 2\mathbf{S}^{\mathrm{T}} \mathbf{L}_b^{\mathrm{T}} \mathbf{L}_b \mathbf{S}\mathbf{m}^{n+1}. \tag{3-68}$$

Collecting the results from Equations 3-63 to 3-68 and substituting into Equation 3-58 results in,

$$\begin{aligned} &-2\mathbf{G}_n^{\mathrm{T}}\mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^n) - \mathbf{G}_n(\mathbf{m}^{n+1} - \mathbf{m}^n)\right] - 2\lambda_r \mathbf{S}^{\mathrm{T}}\mathbf{C}_r^{-1}\left[\mathbf{r}_0 - \mathbf{S}\mathbf{m}^{n+1}\right] + 2\lambda_x \mathbf{S}^{\mathrm{T}}\mathbf{L}_x^{\mathrm{T}}\mathbf{L}_x\mathbf{S}\mathbf{m}^{n+1} \\ &\quad + 2\lambda_y \mathbf{S}^{\mathrm{T}}\mathbf{L}_y^{\mathrm{T}}\mathbf{L}_y\mathbf{S}\mathbf{m}^{n+1} + 2\lambda_z \mathbf{S}^{\mathrm{T}}\mathbf{L}_z^{\mathrm{T}}\mathbf{L}_z\mathbf{S}\mathbf{m}^{n+1} + 2\lambda_b \mathbf{S}^{\mathrm{T}}\mathbf{L}_b^{\mathrm{T}}\mathbf{L}_b\mathbf{S}\mathbf{m}^{n+1} = 0 \end{aligned} \tag{3-69}$$

Then collecting the terms in the unknown vector $\mathbf{m}^{n+1}$ onto the left hand side yields,

$$\begin{aligned} &\left[\mathbf{G}_n^{\mathrm{T}}\mathbf{C}_d^{-1}\mathbf{G}_n + \lambda_r \mathbf{S}^{\mathrm{T}}\mathbf{C}_r^{-1}\mathbf{S} + \lambda_x \mathbf{S}^{\mathrm{T}}\mathbf{L}_x^{\mathrm{T}}\mathbf{L}_x\mathbf{S} + \lambda_y \mathbf{S}^{\mathrm{T}}\mathbf{L}_y^{\mathrm{T}}\mathbf{L}_y\mathbf{S} + \lambda_z \mathbf{S}^{\mathrm{T}}\mathbf{L}_z^{\mathrm{T}}\mathbf{L}_z\mathbf{S} + \lambda_b \mathbf{S}^{\mathrm{T}}\mathbf{L}_b^{\mathrm{T}}\mathbf{L}_b\mathbf{S}\right]\mathbf{m}^{n+1} \\ &\qquad = \mathbf{G}_n^{\mathrm{T}}\mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^n) + \mathbf{G}_n\mathbf{m}^n\right] + \lambda_r \mathbf{S}^{\mathrm{T}}\mathbf{C}_r^{-1}\mathbf{r}^0 \end{aligned}. \tag{3-70}$$

This may be recast into the form,

$$\mathbf{A}\mathbf{m}^{n+1} = \mathbf{b}, \tag{3-71}$$

where,

$$\mathbf{A} = \left[\mathbf{G}_n^{\mathrm{T}}\mathbf{C}_d^{-1}\mathbf{G}_n + \lambda_r \mathbf{S}^{\mathrm{T}}\mathbf{C}_r^{-1}\mathbf{S} + \lambda_x \mathbf{S}^{\mathrm{T}}\mathbf{L}_x^{\mathrm{T}}\mathbf{L}_x\mathbf{S} + \lambda_y \mathbf{S}^{\mathrm{T}}\mathbf{L}_y^{\mathrm{T}}\mathbf{L}_y\mathbf{S} + \lambda_z \mathbf{S}^{\mathrm{T}}\mathbf{L}_z^{\mathrm{T}}\mathbf{L}_z\mathbf{S} + \lambda_b \mathbf{S}^{\mathrm{T}}\mathbf{L}_b^{\mathrm{T}}\mathbf{L}_b\mathbf{S}\right], \tag{3-72}$$

and,

$$\mathbf{b} = \mathbf{G}_n^{\mathrm{T}}\mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^n) + \mathbf{G}_n\mathbf{m}^n\right] + \lambda_r \mathbf{S}^{\mathrm{T}}\mathbf{C}_r^{-1}\mathbf{r}^0. \tag{3-73}$$

Equation 3-71 thus represents a linear system of equations which can be solved to yield the required vector $\mathbf{m}_{n+1}$. This is achieved by using a preconditioned conjugate gradient method as described in Section 3.6.4.

Step length damping

Due to the linear approximation of the nonlinear forward modelling function $\boldsymbol{g}()$ made above, the $\mathbf{m}^{n+1}$ will almost certainly not be located at the minimum. It is only guaranteed that the path from $\mathbf{m}^{n}$ to $\mathbf{m}^{n+1}$ across the objective function surface departs $\mathbf{m}^{n}$ going downhill. The path may cross a valley and then go uphill so that in fact it is possible that $\Phi(\mathbf{m}^{n+1}) > \Phi(\mathbf{m}^{n})$. In this case it is necessary to apply step-length damping (Parker, 1994) to ensure the objective function decreases at every iteration. Two forms of step-length damping have been used in the current work.

In the first method, once a new $\mathbf{m}_{n+1}$ is generated, $\Phi_d(\mathbf{m}^{n+1})$ is calculated to check for improvement. If there is no improvement, then the parameter change is iteratively reduced via a construct of the form,

$$
\begin{aligned}
&k = 1 \\
&while \quad (k < ntrys \quad and \quad \Phi_d(\mathbf{m}^{n+1}) > \Phi_d(\mathbf{m}^{n})) \\
&\qquad \mathbf{m}^{n+1} = \mathbf{m}^{n} + \beta \times (\mathbf{m}^{n+1} - \mathbf{m}^{n}) \\
&\qquad k = k + 1 \\
&end
\end{aligned}
\qquad ,
$$

where, typically $\beta = 0.75$ and $ntrys = 10$.

The second method is conservative in that at each iteration the aim is to reduce the data misfit to a value that is a large fraction of its current value. This method is designed to prevent the above situation from occurring in the first place and is therefore applied within every iteration. To implement this a line search is carried out on β over the interval $[0,1]$ to find a revised solution $\mathbf{m}^{n+1}_{revised}$ such that,

$$\mathbf{m}^{n+1}_{revised} = \mathbf{m}^{n} + \beta \times (\mathbf{m}^{n+1} - \mathbf{m}^{n}), \tag{3-74}$$

and,

$$0.5 < \Phi_d(\mathbf{m}^{n+1}_{revised}) \Big/ \Phi_d(\mathbf{m}^n) < 0.7\,. \tag{3-75}$$

In both methods forward modelling is required for each set of trial parameters, however it is not necessary to compute derivatives. The second method tends to be more robust because of the conservative rate of convergence. However it is more difficult to implement and is generally more computationally expensive because it is applied at every iteration. Nevertheless, the second method is preferred because the extra expense is not prohibitive.

Termination conditions

The iterative procedure continues until an acceptable level of misfit between the observed and predicted data is achieved. An acceptable level of misfit occurs when the chi-squared misfit,

$$\chi^2 = \Phi_d(\mathbf{m}^{n+1}) = \Big[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^{n+1})\Big]^{\mathbf{T}} \mathbf{C}_d^{-1} \Big[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^{n+1})\Big], \tag{3-76}$$

is less than the number of data. The inversion will thus terminate if the normalized data misfit,

$$\Phi_d(\mathbf{m}^{n+1}) \Big/ N_d \leq 1\,. \tag{3-77}$$

A further termination condition is triggered when the improvement in the data misfit between successive iterations is less than a specified percentage minimum $p^{\min}$, i.e.,

$$100 \times \left[\frac{\Phi_d(\mathbf{m}^n) - \Phi_d(\mathbf{m}^{n+1})}{\Phi_d(\mathbf{m}^n)}\right] < p^{\min}\,.$$

Typically a value of $p^{\min} = 1\%$ is used.

Selection of weighting parameters

The choice of values of the λ weighting factors in Equation 3-45 is made at the beginning of the inversion and they remain fixed throughout the inversion. Their values are determined by trial and error after an assessment of the resulting models. The selection procedure uses the highest value of the weighting factors that still allows the data to be fitted. This requires several runs, initially setting the weighting values high then progressively reducing them. This is the most subjective element of the holistic inversion. Unfortunately, to some degree this has been unavoidable because more automated and objective methods, such as generalized cross validation, are not practical due to the size of the problem.

3.6.4 Solution of linear system

It is now possible to get an insight into the size and sparsity structure of the matrices that together make up the linear system (Equations 3-71) to be solved. To do this we will examine the matrices involved in the holistic inversion of airborne survey data shown in Section 4.4 of Chapter 4. In this example 2590 line kilometres of data, which were acquired over a 406 km^2 irregularly shaped area, were inverted. This represented the largest portion of a survey that could be inverted on a standalone computer with 2 gigabytes of random-access memory available. A five layer conductivity model, parameterized by just 8 spline meshes because the fourth layer thickness was kept fixed, was used. There were 689,499 data in total, of these, 575,208 were airborne data, 330 were geoelectric data and 113,961 were interface-depth data. There were 148,674 unknown parameters to be estimated in total. Of these, 146,574 were spline node coefficients and 2,100 were calibration model parameters.

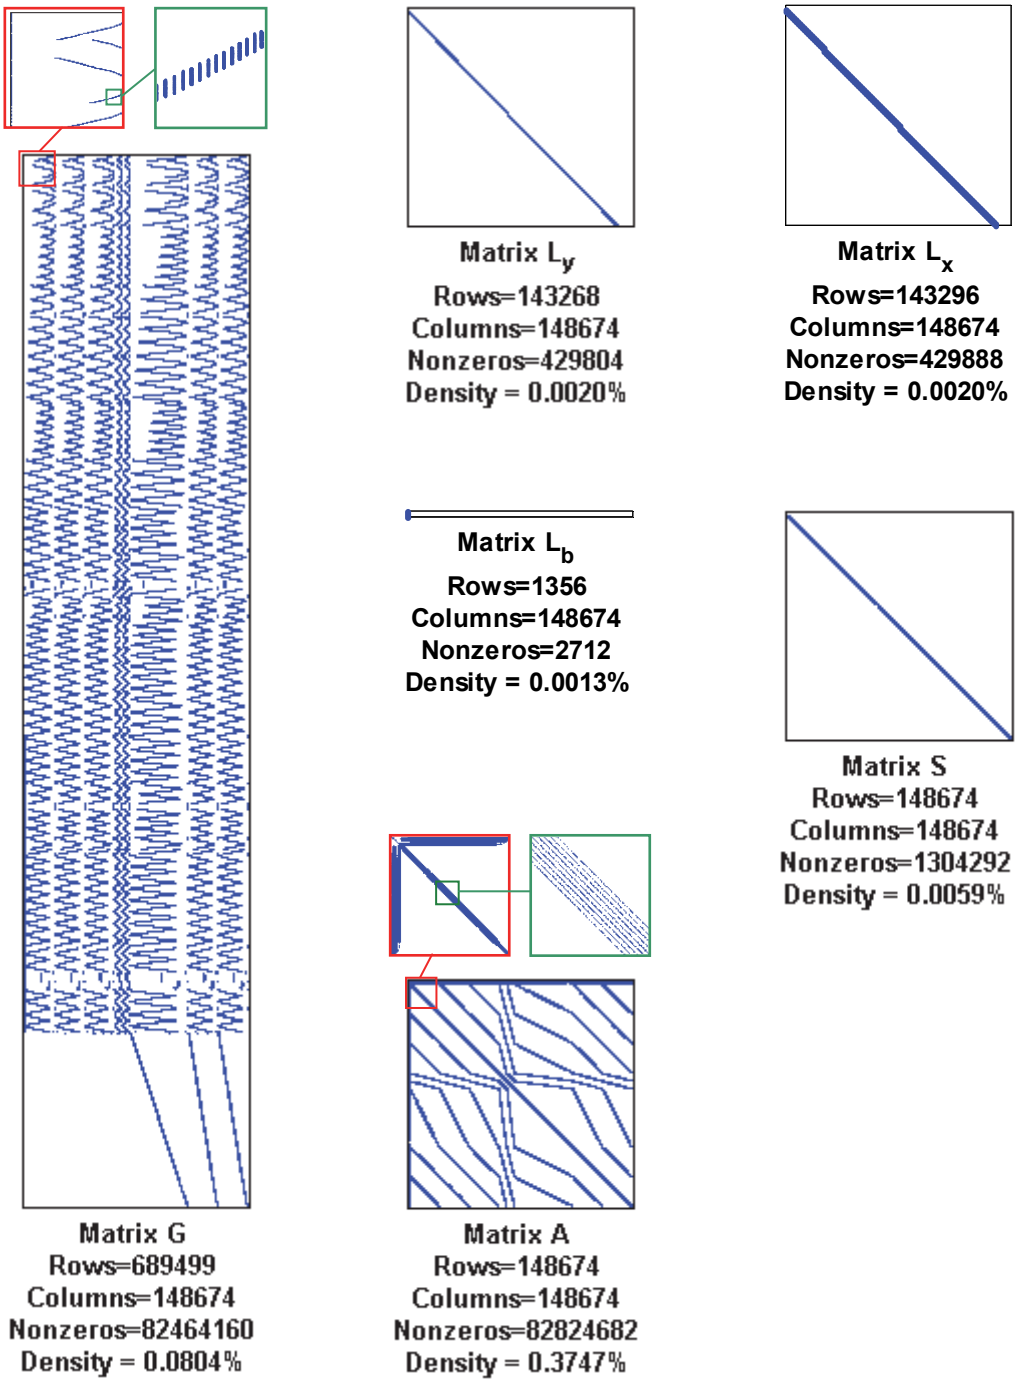

Figure 3.10 *An example of the sparsity structure of the matrices that go toward creating the linear system to be solved. Due to the drawing resolution the matrices appear more dense than they truly are.*

Figure 3.10 shows the structure of the matrices involved in this example. Each airborne datum is influenced by only the 16 adjacent nodes on 8 spline meshes plus 6 calibration parameters. Therefore just 134 $(16\times(5+3)+6)$ of the 148,674 column entries on each row of the Jacobian matrix corresponding to airborne data are nonzero. There are fewer for the interface-depth data as they are not affected by conductivity meshes.

Overall, the **G** matrix is approximately 0.08% dense and had 82.5×10^6 nonzero entries and takes 1.0 gigabytes to store in memory in an efficient compressed format where only the nonzero entries are stored. There are typically only 16 nonzero entries per row in $\mathbf{S}$, and 3 to 5 nonzero entries per row in the **L** matrices, so they are even sparser. In general the matrix structure depends on the relative proportions of calibration and conductivity model parameters and the chosen parameterizations.

The linear system matrix $\mathbf{A}$ in Equation 3-71 is thus also sparse. To quantify how sparse, consider that the i,jth entry of $\mathbf{G}^{\mathbf{T}}C_d^{-1}\mathbf{G}$, and hence of $\mathbf{A}$, is nonzero if any one datum is influenced by both the ith and the jth parameter. For the case where the meshes on each layer have the same spacing, an approximate lower bound on the ratio of the density of the $\mathbf{A}$ to the **G** matrix will be 3.06. The factor of 3.06 ($^{49}/_{16}$) results from the fact that one parameter can influence data lying on a 4×4 region of surface patches (bounded by 5×5 nodes) which can in turn be effected by a wider group of 7×7 nodes. Consequently the A matrix is substantially denser than the **G** matrix or any of the other underlying matrices. The situation is complicated by varying mesh sizes and the proportion of calibration model parameters. The **A** matrix in the example above had an overall density of 0.37%, (4.6 times that of **G**) and had 82.8×10^6 nonzero entries, which would also take 1.0 gigabytes to store in memory in the efficient compressed format. In the portion of the A matrix that was effected by calibration parameters (i.e. the first 2,100 rows and columns) the average density was 4.02%. Despite their different densities, the similar number of nonzero entries only occurs coincidentally in this example because there are ~4.6 times more rows in **G** than there are in **A**.

Clearly, the size of the system of equations to be solved can be substantial. It is thus necessary to use an iterative method, rather than a direct method, to solve the system. A suitable choice for the symmetric positive-definite matrix **A** is the conjugate gradient method (Barrett et al., 1994). The conjugate gradient algorithm solves the linear system by making a series of approximations to the solution based on orthogonal search

directions, which successively reduce the residual vector. Each iteration requires a matrix-vector product and two inner products. However, only a small number of vectors need be kept in memory.

The only access the algorithm needs to $\mathbf{A}$ is in computing the product $\mathbf{Ax}$ where $\mathbf{x}$ is an approximation to the solution. Thus, the real advantage for the current application is that the matrix $\mathbf{A}$ does not need to be explicitly formed and stored in memory. Instead, the individual components that make up the matrix $\mathbf{A}$ (i.e. $\mathbf{G}$, $\mathbf{S}$ $\mathbf{L}_x$ etc.) are stored in compressed form and are used in calculating the product $\mathbf{Ax}$. This requires substantially less memory as discussed above. For the holistic inversion, the individual matrices are stored in a virtual C++ matrix object which has a member function that performs the matrix product $\mathbf{Ax}$ without explicitly forming $\mathbf{A}$.

The conjugate gradient algorithm has been implemented through two different publicly available code libraries. In the first generation of the holistic inversion code, which was for a standalone computer, it was implemented via the routine *"CG"* described by (Dongarra et al., 1996). This is the method used in the example shown in Section 4.4 of Chapter 4. In the second generation the code is parallelized and the routine *"KSPSolve"* from the PETSc library is used (Balay et al., 2005). This parallel version is used for the much larger problem discussed in Section 4.5 of Chapter 4.

In both methods the diagonal of the $\mathbf{A}$ matrix is used as the preconditioning matrix. It can also be computed efficiently without explicitly forming the whole of $\mathbf{A}$. The algorithm is supplied with the $\mathbf{b}$ vector and an initial guess $\mathbf{x}^0 = \mathbf{m}^n$ of the solution. The routines computes the norm of the initial vector of residuals ($|| \mathbf{Ax}^0 - \mathbf{b} ||$) and then iterates until the revised norm meets specified criteria. In the first method, iterations continue until the revised norm reached a specified fraction of the norm of $\mathbf{b}$,

$$\frac{|| \mathbf{Ax} - \mathbf{b} ||}{|| \mathbf{b} ||} < tol, \tag{3-78}$$

for which a values of $tol = 10^{-14}$ have been used.

In the second method (the parallel code) an additional criteria is used where iterations continue until the revised norm is less than a specified fraction of the norm of the initial vector of residuals,

$$\frac{|| \mathbf{Ax} - \mathbf{b} ||}{|| \mathbf{Ax}^0 - \mathbf{b} ||} < rtol, \tag{3-79}$$

for which values of $rtol = 10^{-5}$ have been used.

If convergence is not reached within $\mathbf{N}_m$ (i.e. the number of parameters) iterations, the procedure terminates and returns the current solution. It has been found in practice that, for the holistic inversion, the conjugate gradient algorithm converges well before the maximum number of iterations is reached. Details are provided in the relevant sections of Chapter 4.

Chapter 4

Frequency-domain applications

4.1 Outline

In this chapter three applications of the holistic inversion to frequency-domain AEM data are presented. The first example (Section 4.3) is an inversion of synthetically generated data which confirms that the formal theory has been correctly implemented. It also establishes that the method recovers the known model parameters correctly.

In the second example (Section 4.4) real AEM survey data from the Riverland Airborne Electromagnetic Survey are inverted. In Chapter 1 the Riverland survey data was shown to be negatively affected by systematic calibration errors. The first generation of code, developed for a single-processor computer, was used and hence only a one-quarter subset of the total survey area could be inverted. The inversion employed a five layer variable-thickness conductivity model in which each layer corresponded to a specific unit in the well understood conceptual (hydro)geological model. This model was based on sound prior knowledge of the geology of the survey area, and geoelectric and interface depth data were explicitly included in the inversion. The example demonstrates that the holistic inversion can be successfully applied to real data since the estimated calibration model parameters are consistent with independently derived parameters. Furthermore, unlike in the conventional sample-by-sample inversion, the holistic inversion is able to fit the observed data and, the estimated conductivity model is free of the artefacts caused by calibration errors.

One question arising from the second example was whether or not its success could be attributed to the use of good prior information and well-developed knowledge of the geology of the survey area. This prompted an investigation into how the method would perform if 'survey-area-specific' prior information was unavailable, for example in so called 'greenfields' mapping surveys. When reliable prior information is available, tight constraints can be imposed on the layer properties, which make it possible to solve for the layer thicknesses in a stable manner. In such circumstances, a few-layer parameterization is a good choice because the model can be tailored to adequately represent the true conductivity variations with a minimal number of unknowns. Furthermore, if the inversion parameters are direct proxies for (hydro)geological features there is no need to carry out additional post-inversion manipulations to extract features which may be required downstream.

A shortage of prior information limits the understanding of the geological layering and how it should translate to a conceptual conductivity model. There is no way to know how to choose a generic few-layered model that can adequately represent the large (but unknown) range of vertical conductivity profile shapes that are possible. Without prior knowledge of the geoelectric properties of the (hydro)geological units of the survey area, it is difficult to build an accurate reference model with associated uncertainties. In such cases, a multi-layer fixed-thickness parameterization is a suitable choice because it can represent a wide range of profile shapes. To improve stability it is necessary to fix the layer thicknesses and to regularise the profile shape, with, for example, vertical smoothness constraints (Constable et al., 1987). In this type of parameterization, particular layers in the model do not correspond to particular geological features. Instead the features must be extracted by some post-inversion manipulation or interpretation.

In the third example (Section 4.5) the case of minimal prior information is simulated, and thus geoelectric and interface depth data are not included in the inversion. For the

reasons stated above, a multi-layer fixed-thickness model was used. Since this parameterization requires many more model parameters, a second generation of the code was developed to allow the method to be implemented on a parallel computer. This allowed the complete Riverland survey to be inverted in one inversion. The example demonstrates that the method can be used when only minimal prior information is available. Without using explicit prior information, the recovered conductivity model is geologically plausible and consistent with downhole logs and groundwater depth data. Zero-levels predicted from the calibration model correlate well with zero-level measurements recorded at high altitude. Gain values are also similar to gain values predicted from downhole log data.

Since the application examples are based on the inversion of data from the Riverland Survey, it is practical to first present some background information on the survey area and objectives, the airborne electromagnetic system used, survey procedure, data processing, and prior information. This is presented in the following subsection before moving on to the synthetic and real data examples.

4.2 The Riverland survey

4.2.1 Background

The Riverland Survey covers a 10 to 15 km wide corridor following the southern bank of the River Murray in South Australia. The zone stretches between Kingston-On-Murray in the west to Wilperna Island northeast of Renmark in the east (Figure 4.1). The survey was flown as part of a project conducted under the auspices of the South Australia Salinity Mapping and Management Support Project, and funded under the National Action Plan for Salinity and Water Quality.

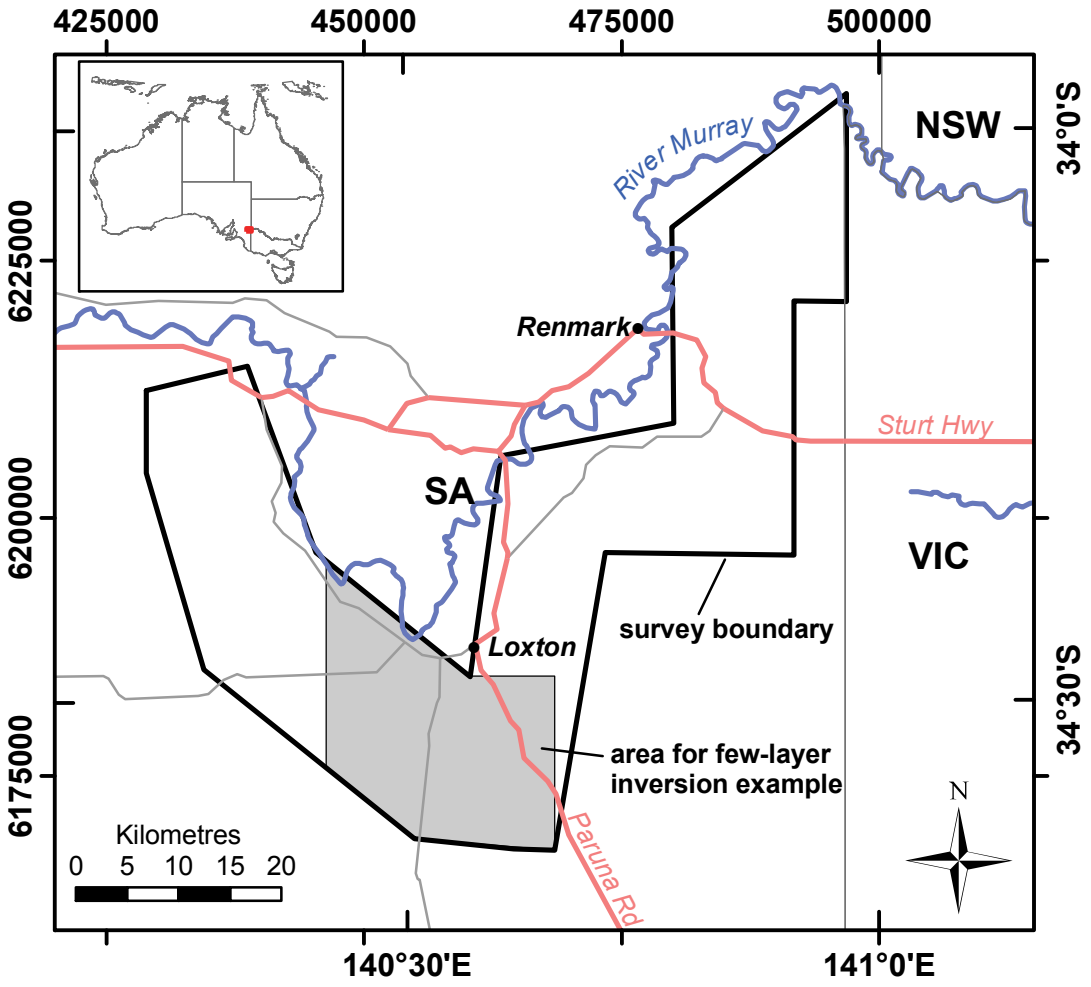

Figure 4.1 *Locality map of the Riverland Airborne Electromagnetic Survey showing the boundary of the survey area and the subset area used for the few-layer inversion example.*

The aim of the project was to map the spatial distribution and thickness of near surface clay rich units, which are important because they are perhaps the only sediments in the area that can impede groundwater recharge. Such recharge causes the groundwater, which in many places is very saline, to rise and increases its flow towards the River Murray. The survey results could thus be used in conjunction with soil hydrological models to estimate rates of groundwater recharge, which were required to help develop natural resource management plans for the area (Green et al., 2004; Munday et al., 2004). To give the reader a feeling for the land use and topography of the survey area, Figure 4.2 shows satellite imagery and surface elevation data.

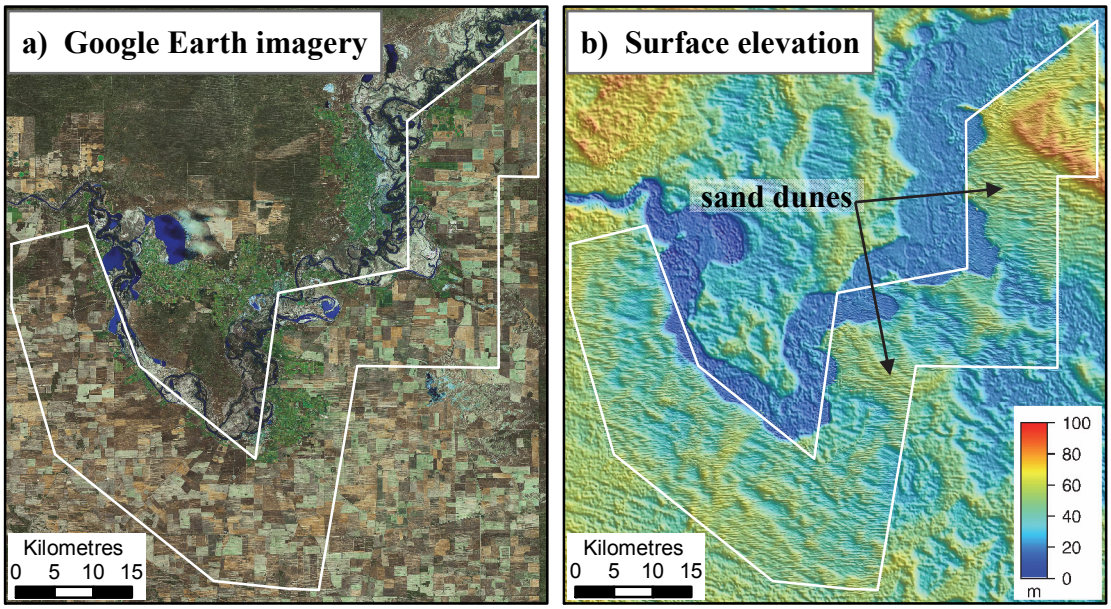

Figure 4.2 Google Earth imagery and surface elevation data in the Riverland area.

The survey data were acquired and processed under a contract between the Government of Australia and Fugro Airborne Surveys (FAS) between June and September 2002. Full details of the data acquisition and processing is provided in the survey operations report (Cowey et al., 2003).

This dataset has been the subject of several previous studies. Brodie et al. (2004a) and Ley-Cooper and Macnae (2004) identified calibration problems in the Riverland dataset. Brodie et al. (2004b) inverted the dataset with conventional 1D sample-by-sample inversion. The results of that inversion work was interpreted by Green et al. (2004) and has also been assessed by several other authors (Munday et al., 2004; Tan et al., 2004).

4.2.2 RESOLVE airborne electromagnetic system

The data were acquired by FAS using its frequency-domain airborne electromagnetic system known as RESOLVE. The system is comprised of six coilsets which are housed in an instrument pod, known as a bird, towed below a helicopter (Figure 4.3). The RESOLVE system has not been formally described in the literature, however an article describing the system is available on FAS's internet site (Fugro Airborne Surveys,

2009a). The RESOLVE system is a more recent generation of the well know DIGHEM systems that were described in Fraser (1978) and Fraser (1990).

The new developments in the RESOLVE system included replacement of analogue components with digital components, digitising of the signal in the bird and adding a digital receiver in the aircraft. A further important difference was the inclusion of tuned calibration coils inside the bird so that the system could be calibrated in the air (Fugro Airborne Surveys, 2009b). This was to avoid the errors introduced in ground based procedures by conductive terrain at the calibration site (Fitterman, 1998).

Figure 4.3 The RESOLVE system bird and helicopter during the Riverland survey (photograph courtesy of Tim Munday).

For the Riverland survey the system was configured as shown in Table 4.2. The bird was slung via a 30 m tow cable hitched below the helicopter. It was fitted with a dual-frequency Global Positioning System (GPS) receiver. The GPS data were post processed against a ground base station to allow accurate positioning of the bird. Note that the more recently developed three-sensor GPS attitude monitoring equipment (Hodges et al., 2007) was not available at the time of the survey.

The bird was fitted with a laser altimeter to measure its height above ground level. A radar altimeter was also fitted on the helicopter. The GPS sensor was located in the centre of the towed bird, 3 metres behind the laser altimeter sensor that was positioned at the forward flying portion of the bird (Cowey et al., 2003).

Table 4.2 RESOLVE coilset configuration for the Riverland survey.

Coilset	**1**	**2**	**3**	**4**	**5**	**6**
Frequency (Hz)	385	1,518	3,323	6,135	25,380	106,140
Separation (m)	7.86	7.86	8.99	7.86	7.86	7.86
Orientation	HCP	HCP	VCX	HCP	HCP	HCP

HCP means "horizontal coplanar" and VCP means "vertical coaxial"

4.2.3 Data acquisition

The helicopter was flown so that the bird would nominally maintain a height of 30 m above ground level. A total of 11,476 line kilometres of production data were acquired. Flight lines were oriented north–south and, for most of the survey, were spaced 150 m apart. However lines were spaced 300 m apart in part of the area, as is shown in Figure 4.4. East–west orientated tie lines were flown at approximately 6 km spacing. Repeat lines were flown on a daily basis during the survey over a designated 5 km long transects as a quality control check on repeatability and for noise estimation. Due to operational reasons, the two different locations shown on Figure 4.4 were used for this purpose. A substantial part of the survey area could not be flown due to safety reasons. Referring back to Figure 4.1 it can be seen that these were mainly around the built-up areas of the township of Loxton, irrigation precincts, and along the Sturt Highway.

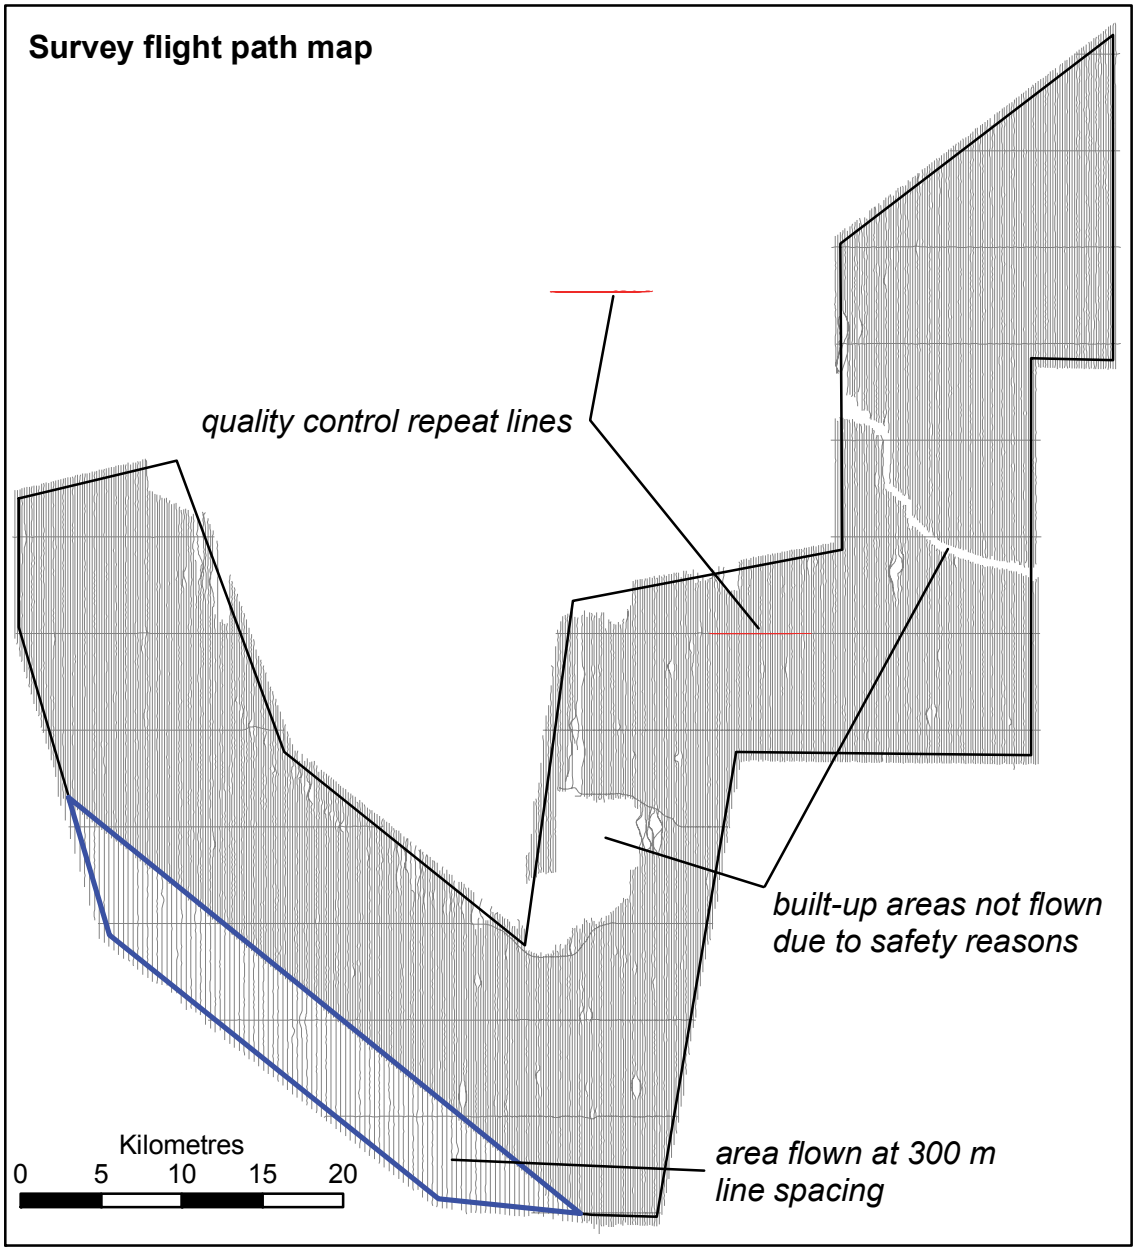

Figure 4.4 *Flight path map of the Riverland survey showing two sets of repeat lines flown for quality control checking, the portion of the area flown at 300 m line spacing, and areas that were not flown due to safety reasons.*

Figure 4.5 shows the overall temporal pattern of data acquisition. There was a total of 37 days of production flying, and generally two to three flights were conducted on each day. There were 89 flights in total. There were 650 lines in total, of which 17 were east–west tie lines and 46 were repeat lines. The RESOLVE system measures an airborne sample every 0.1 s as it traverses along the flight lines. At an average speed of approximately 33 metres per second, this resulted in an airborne sample every 3.3 m along the flight path. Over the entire survey, 3.7 million samples were collected, which

equates to some 44.4 million data since there are 6 inphase and 6 quadrature data per sample.

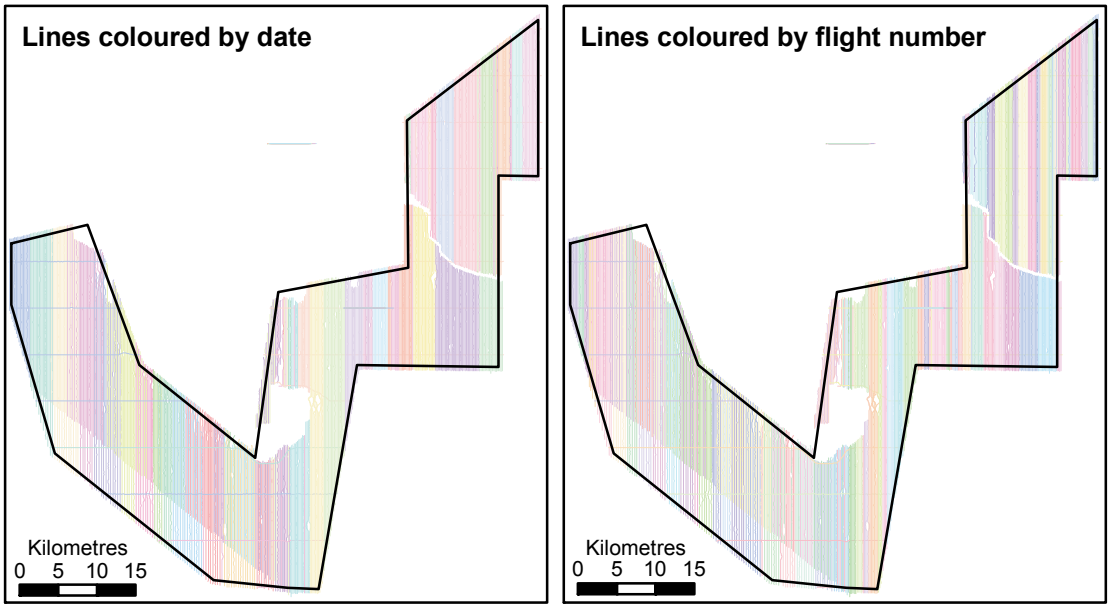

Figure 4.5 *Flight path map illustrating the temporal pattern of data acquisition. The lines are coloured according to the day of flying (left) and the flight number (right).*

4.2.4 Calibrations

Calibrations were carried out in-flight using FAS's *AutoCal* automatic, internal calibration process (Cowey et al., 2003). The internal calibration coils are factory calibrated at FAS's calibration site at Mountsberg, Ontario, chosen because of its resistive geology. The coils are used to provide a known and stable reference signal. At the beginning and end of each flight, and at intervals of approximately 20 minutes (Figure 4.6a) during each survey flight, the system was flown to a height (>500m above ground level) such that the secondary field from the ground was negligible. Any signal measured in the receiver coils, the so called zero-level, was recorded so that it could later be corrected in the data processing. On some occasions, usually when it was large, the zero-level was reset to zero by adjustments within the system (Figure 4.7b). Following the zero-level measurement and reset, the internal calibration coils were consecutively triggered for each frequency for sufficient time to determine an accurate

response through any ambient noise (Figure 4.6b and Figure 4.7b). The receiver response to each calibration coil event was compared to the expected response, known from the factory calibration, for both amplitude and phase. A gain and phase correction was then applied to match the measured signal to the expected reference signal.

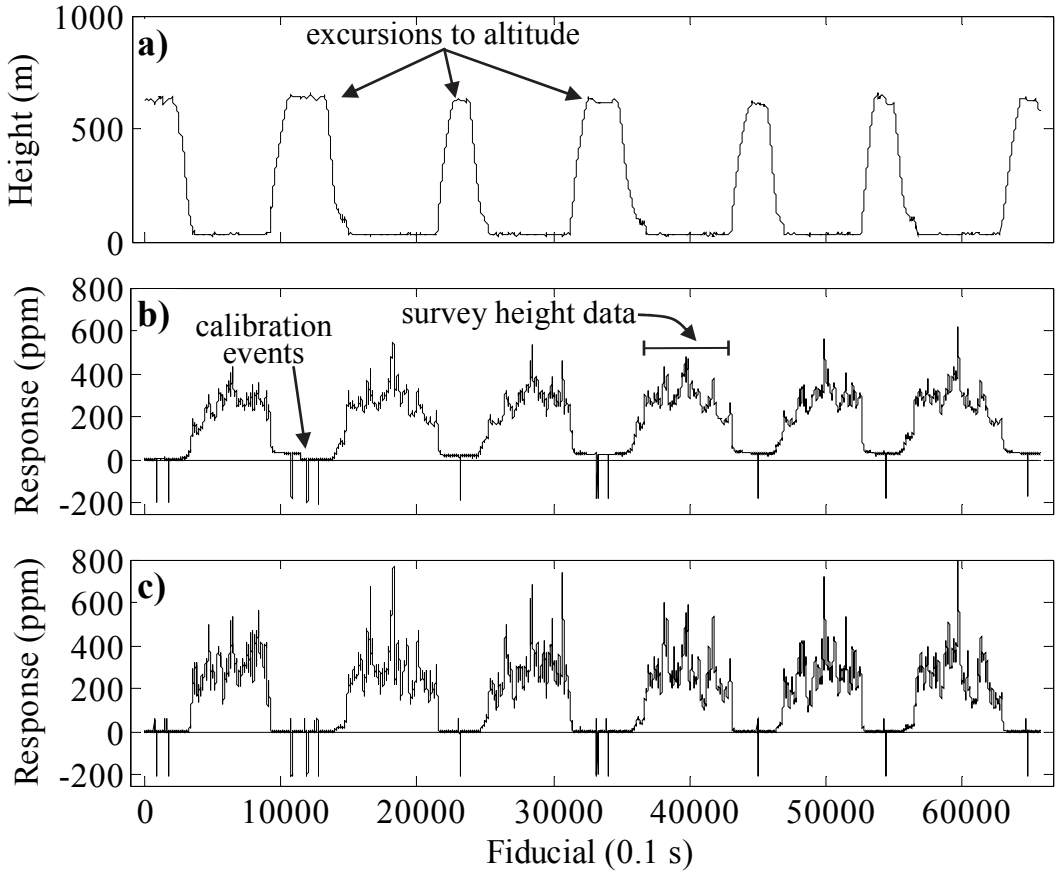

Figure 4.6 *Profiles showing (a) bird height above ground, (b) 385 Hz inphase and (c) 385 Hz quadrature data, for an entire flight including the regular excursions to high altitude to reset the zero-level and perform calibrations with the internal calibration coils.*

Figure 4.7 shows a detailed view of (a) bird height and (b) the inphase channel of the 385 Hz coilset data during one excursion to altitude in which the zero-level was measured and reset during the *AutoCal* procedure. In the holistic inversion I wish to take advantage of the fact that the bias changes slowly and smoothly over time (i.e. of its temporal coherency) and have therefore parameterized it as a continuous piecewise linear function of time as shown in Figure 3.4. However the zero-level resets shown in Figure 4.7b create discontinuities in the bias. Therefore it is necessary to reverse the zero-level resets as shown in Figure 4.7c to produce a continuous bias. This was done

for the Riverland data before it was inverted in the examples shown in Sections 4.4 and 4.5. This step would not be necessary if the bias was parameterized more simply, for example as a constant for each flight line or as a linear variation over each flight line, because the discontinuity does not occur while survey altitude flight lines data are being acquired. Notwithstanding this, the former parameterization is still preferred because it takes greater advantage of coherency and there are fewer parameters to be estimated. It should therefore be better conditioned.

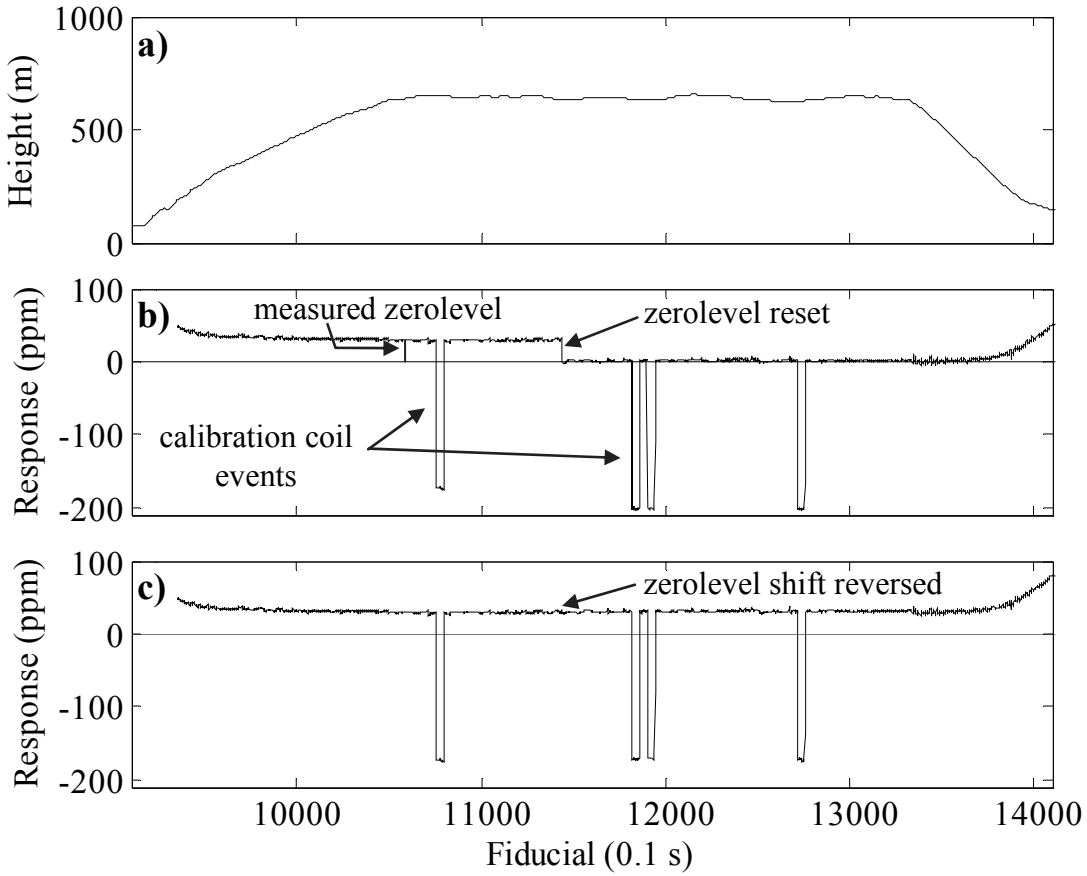

Figure 4.7 Detailed view of one high altitude excursion showing (a) bird height above ground, and 385 Hz inphase data (b) before and (c) after reversal of the level shift created during the calibration.

4.2.5 Data processing

This section explains the post-flight data processing carried out on the electromagnetic data by Fugro Airborne Surveys, and then explains which aspects were used for the examples in this thesis. Gain corrections were made to the data for the two highest frequency coilsets. This involved scaling the inphase and quadrature channels from the

25,380 Hz and 106,150 Hz coilsets by the factors 1.133 and 1.189 respectively. According to the report of Cowey et al. (2003) the reason was to improve the match between the measured response from *AutoCal* high altitude events and the expected factory-calibrated reference response. No explanation was given as to why the *AutoCal* procedure had not made this adjustment automatically.

The signal measured at high altitude excursion, as shown in Figure 4.7b and marked as 'measured zero-level', were used to calculate a zero-level or drift correction (Valleau, 2000). For each high altitude excursion a small sample of data is selected and averaged to make a so called zero-level 'pick' (i.e. a fid, value pair). A pick is made on either side of a level shift if the zero-level is reset. Between each set of picks the value is linearly interpolated at the fiducial of each survey altitude flight line sample and is then subtracted from the survey data. This procedure is designed to remove bias, assuming that the bias measured at altitude is representative of the bias at survey altitude, and that it varies linearly in between.

The zero-level corrected data were then filtered to reduce the effect of high frequency random noise caused by factors such as atmospheric sferic activity (e.g. distant lightning strikes), powerline noise, electronic noise and high frequency mechanical vibration of the bird. The filters were a 9-point median filter followed by a 9-point Hanning (raised cosine) shaped convolution filter. This procedure only effects spatial wavelengths less than ~30 m, which is much less the system's minimum inductive limit footprint size of around 112 m (Reid and Vrbancich, 2004; Reid et al., 2006). The filters are effective at reducing the noise but do not eliminate it entirely, especially in close proximity to powerlines.

Further data processing by the contractor FAS included '*levelling*' of the dataset and production of apparent resistivity data (Cowey et al., 2003). The levelling procedure identified flights, individual lines or line segments which required further adjustments

based on visual inspection of apparent resistivity maps generated for each coilset from the inphase and quadrature data channels. The reason was attributed to subtle changes between in-flight calibrations of the system causing flight-to-flight and line-to-line differences in apparent resistivity. The report does not specifically say what the criteria were for deciding when an adjustment was necessary or what the adjustment should be. Presumably, the decision was based on identifying and eliminating 'levelling busts' (Valleau, 2000). In other words, eliminating spatially incoherent parts of the maps or artefacts that are in some way correlated with certain flights or lines and can be recognised as not being geologically plausible. Flight-wise constant phase changes were applied to certain flights, in the range -3.5º to +2.0º for the 106,104 Hz coilset, and +3.5º to all flights for the 385 Hz coilset. Base level (constant) or tilt (linear) adjustments were made to a total of 1,187 individual line/channels combinations.

It is important to recognise that the holistic inversion seeks to avoid making the somewhat subjective, and most probably time consuming, data processing steps designed to remove systematic calibration error that have been described above. Instead, the holistic approach seeks to account for the systematic errors by inverting for them, along with the conductivity model, in a more objective and efficient procedure. Therefore raw filtered data (without the application of the other data processing corrections) are inverted in the holistic inversion examples in this thesis. The same 9-point median and Hanning filters are applied, since these are necessary to remove the occasional very high amplitude spikes associated with anthropogenic anomalies.

4.2.6 Noise estimates

The noise in the RESOLVE system can be considered to have a component that is independent of the secondary field ground response and a component that is related to ground response, and hence ground conductivity. These have been named 'additive' and 'multiplicative' noise components, respectively, in the method of estimating noise

described by Green and Lane (2003). Additive noise is typically the same high frequency noise that the 9-point filters described above are designed to reduce (i.e. sferic activity, powerline, electronic and vibrational). Multiplicative noise is related to ground conductivity and is caused primarily by variations in system geometry that are not measured and not accounted for. In the Green and Lane (2003) noise model the standard deviation of the noise for a given coilset and channel is,

$$d^{err} = \sqrt{a^2 + (0.01 \times m \times d^{obs})^2}\,, \tag{4-1}$$

where a is the standard deviation of the additive noise, m is the relative or multiplicative noise expressed as a percentage, and d^{obs} is the observed response. For this thesis the same noise model is employed. I use the additive and multiplicative noise estimates that were made for inversion of the Riverland data by Brodie et al. (2004a). They calculated additive noise from the standard deviation of the filtered high altitude data acquired before and after each flight. The percentage multiplicative noise was calculated from the filtered daily repeat calibration line data mentioned earlier in this section. The parameters of the noise model (i.e. a and m), are given in Table 4.3 below.

Table 4.3 Estimated noise model parameters.

Coilset	**1**	**2**	**3**	**4**	**5**	**6**
Frequency (Hz)	385	1,518	3,323	6,135	25,380	106,140
Inphase additive (ppm)	2.55	4.15	2.90	5.15	8.5	13.8
Quadrature additive(ppm)	1.50	1.90	1.50	3.20	6.65	10.4
Inphase multiplicative (%)	1.20	1.60	1.90	1.85	2.10	2.15
Quadrature multiplicative (%)	1.85	2.35	2.70	2.60	2.70	2.45

4.2.7 Prior information

The location of 49 downhole conductivity logs inside the survey area that were available to be used as prior information in the inversion are shown in Figure 4.8a. The data were acquired using an AUSLOG A034 inductive downhole conductivity probe (Jones and Spring, 2003). The data were used to calibrate the Riverland data by

comparing forward models of the logs to the observed data at the nearest airborne samples (Brodie et al., 2004a). They were also statistically analysed in preparing the reference model constraints for the sample-by-sample inversion of the dataset (Brodie et al., 2004b).

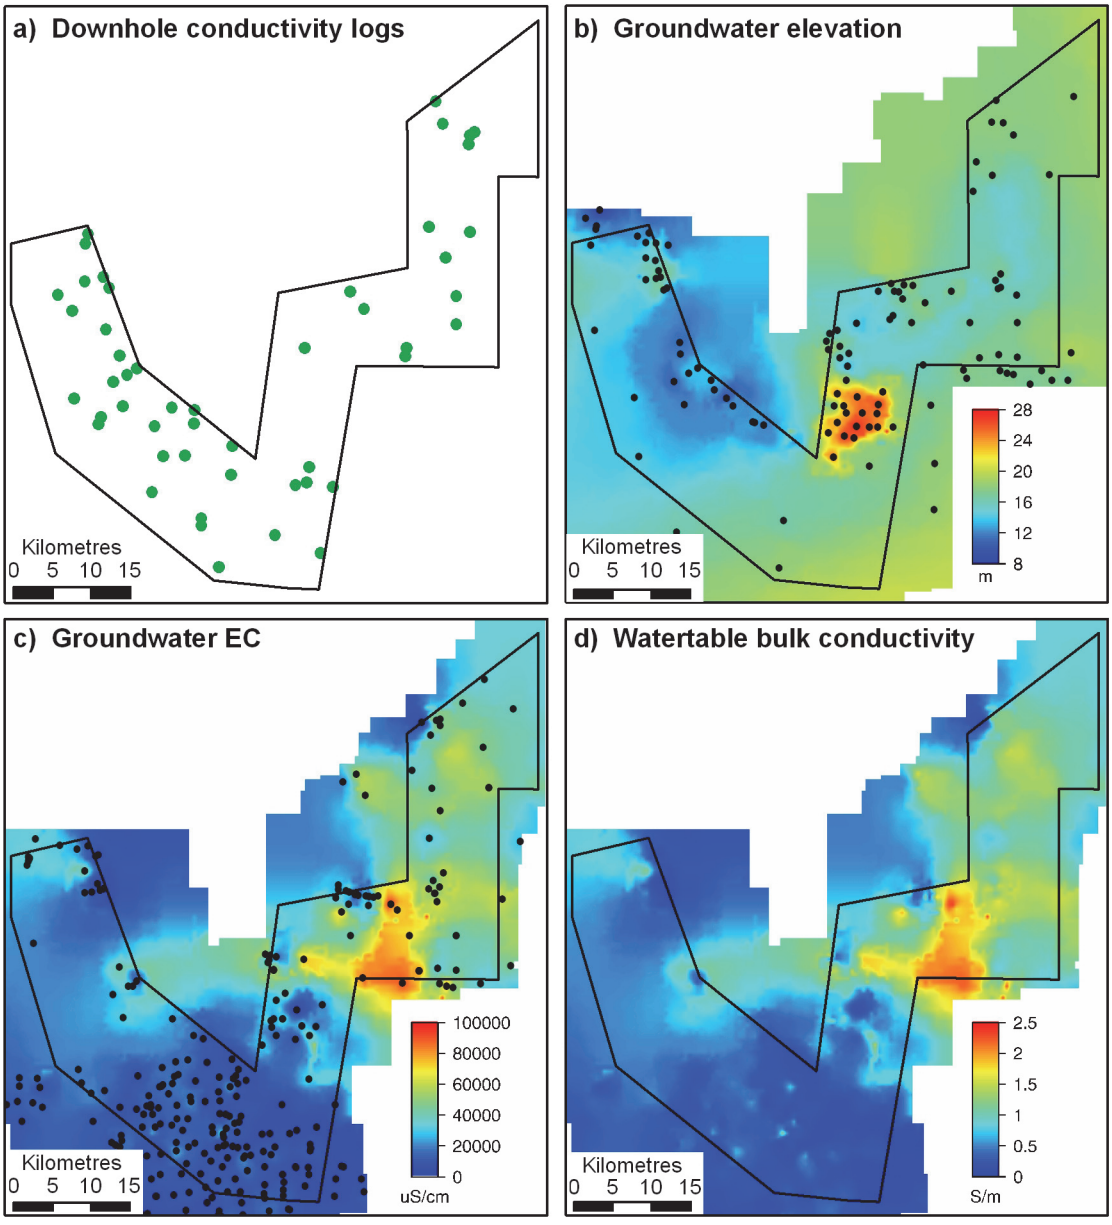

Figure 4.8 *Prior information included (a) downhole conductivity logs, (b) elevation of the watertable measured in bores, (c) EC of groundwater fluids measured in bores, and (d) bulk conductivity below the watertable estimated from conductivity logs and groundwater EC.*

In the few-layer holistic inversion example (Section 4.4) the downhole logs are used as explicit geoelectric data. In the multi-layer holistic inversion (Section 4.5), where a lack

of prior information is simulated, they are instead only used to evaluate the inversion results.

Additional prior information came from a South Australian Government database containing results from an extensive drilling program in the area. The database was used by Brodie et al. (2004b) to generate a gridded dataset of the elevation above sea level of the top of the saline groundwater watertable. The grid and the locations of boreholes used in its generation are shown in (Figure 4.8b). The grid was used in the previous study as a constraint on layer thicknesses, and is used here in the few layer holistic inversion (Section 4.4) as explicit interface-depth data.

Brodie et al. (2004b), used the same source to compile a grid of the apparent electrical conductivity of the pore water fluids (Figure 4.8c). They then converted this to a grid of the bulk conductivity of the subsurface below the watertable (Figure 4.8d) using the relationship,

$$\sigma = 2.5 \times 10^{-5} \times EC\,, \tag{4-2}$$

where σ is the bulk conductivity in units of S/m and EC is the pore water fluid conductivity in units of μS/cm. The relationship was empirically determined from linear regression of groundwater EC data against downhole log data, from the top 10 m of the watertable, if they were suitably close enough together. The resulting grid (Figure 4.8d) was used as a reference model constraint on the conductivity of the layers below the watertable, and is used in the same way in Section 4.4.

4.3 Inversion of synthetic data

4.3.1 Synthetic model and data

In this section the holistic inversion is applied to synthetically generated data. I demonstrate that the theory has been properly implemented, and that the method recovers the known model parameters correctly. The actual acquisition pattern, flight

path and measured bird height for a 10×10 km subset of the Riverland survey were used in synthesising the data. The area corresponded to 10 days of flying, 14 flights and 67 flight lines and 1 tie line. The acquisition pattern is depicted in Figure 4.9a and b using different colours to distinguish separate flying dates and flight numbers. I show later how these same patterns are reflected in diagrams of the synthetic calibration parameters (Figure 4.11 and Figure 4.12) and data (Figure 4.13). There was a distinctive variation in bird height over the area as shown Figure 4.9c.

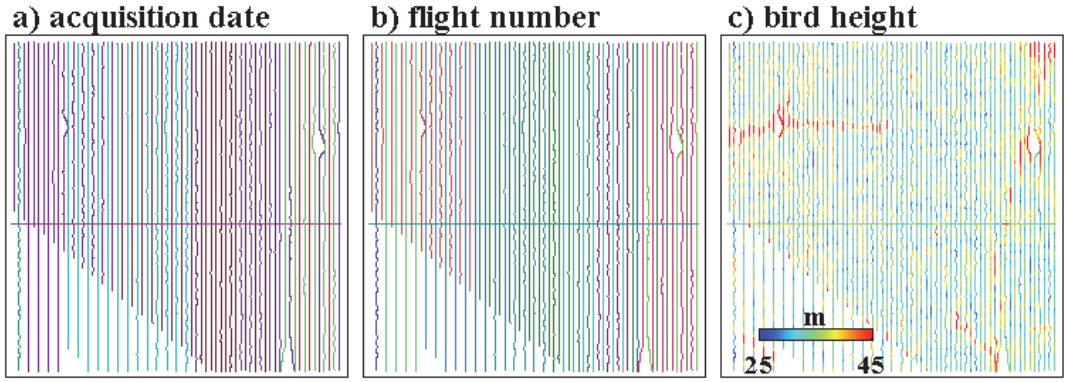

Figure 4.9 *The synthetic dataset showing the flight lines coloured according to (a) date, (b) flight number, and (c) bird height.*

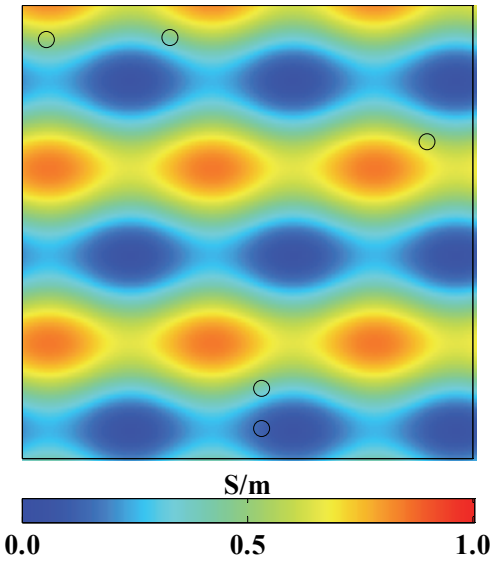

Figure 4.10 *The synthetic conductivity model and the location of the boreholes (black circles) for which synthetic geoelectric data were generated.*

The first step in generating the synthetic data over the chosen area was the definition of a conductivity model. The chosen model has a single layer whose conductivity varies

horizontally but not vertically. Cosine functions were used to synthesise the conductivity variations, which are shown in Figure 4.10. The next step was the definition of the calibration model to be used in the generation of the systematic errors included in the synthetic data.

Table 4.4 Synthetic gain and height error parameters.

Coilset	**1**	**2**	**3**	**4**	**5**	**6**
Frequency (Hz)	385	1,518	3,323	6,135	25,380	106,140
Gain	1.042	0.962	0.901	0.870	0.775	0.813
Height correction (m)	1.000					

The synthetic calibration model was parameterized using separate gains for each coilset, which were constant for the whole survey. The gain values were chosen to be equivalent to (i.e. the reciprocal of) the scaling factors derived for the Riverland data (Brodie et al., 2004a). The values of gain calibration parameters used are shown in Table 4.4.

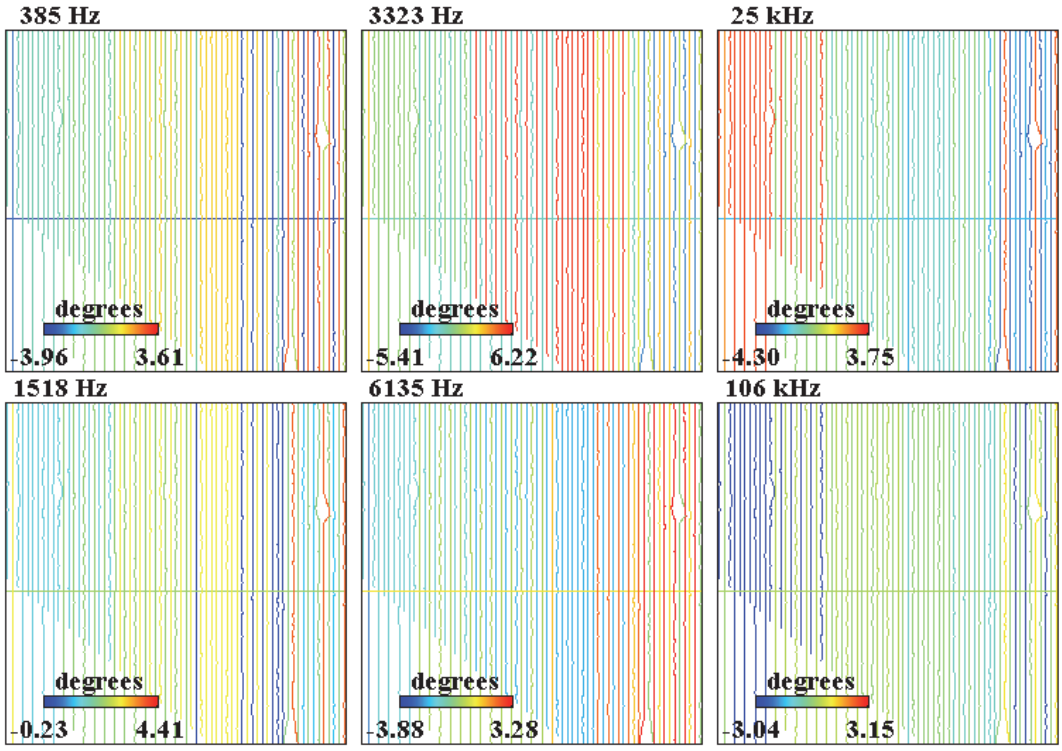

Figure 4.11 Flight lines coloured by the synthetically generated phase values.

The phase was synthesised to be constant for each day of flying and for each coilset. The phases were drawn from a Gaussian random distribution with a mean of 0.0º

degrees and standard deviation of 2.0º. The phase values are shown in Figure 4.11. It can be seen by comparing the pattern of flying days in (Figure 4.9a) with the phase values (Figure 4.12) that the phase is constant for each day of flying.

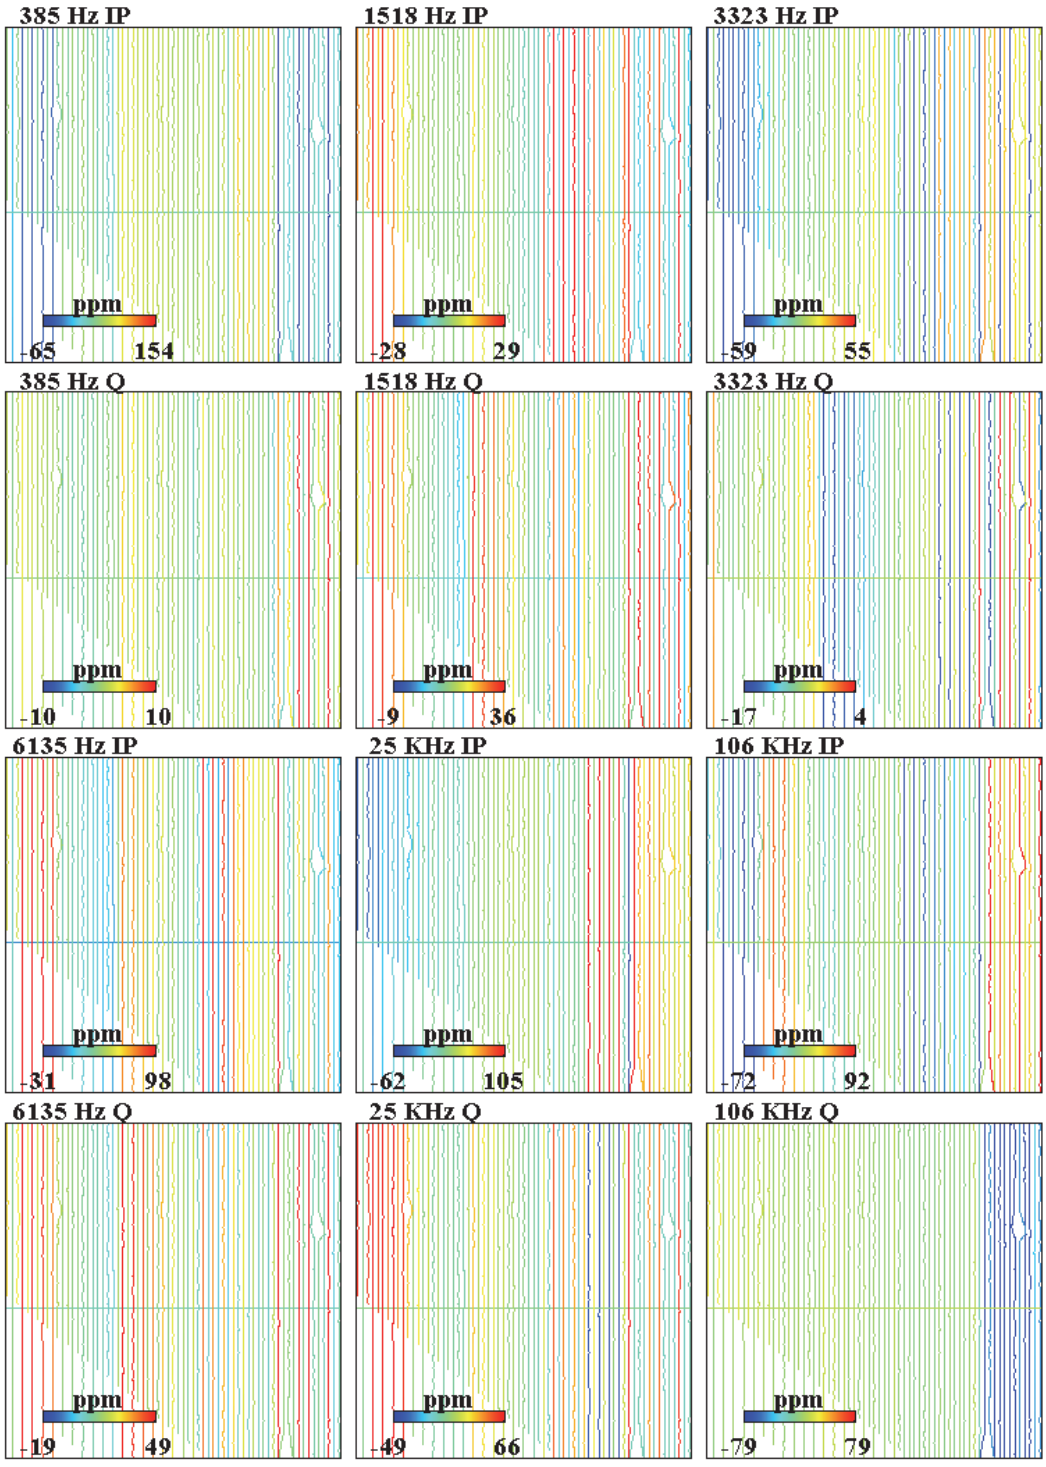

Figure 4.12 Flight lines coloured by the synthetically generated bias values.

The synthetic inphase and quadrature biases were taken to be a piecewise linear interpolation of the zero-level measured at high altitude during the actual survey. The

bias values are shown in Figure 4.12. Although difficult to see, because of the large range of values between flights, careful inspection of Figure 4.12 shows that the bias varies slowly throughout the flight.

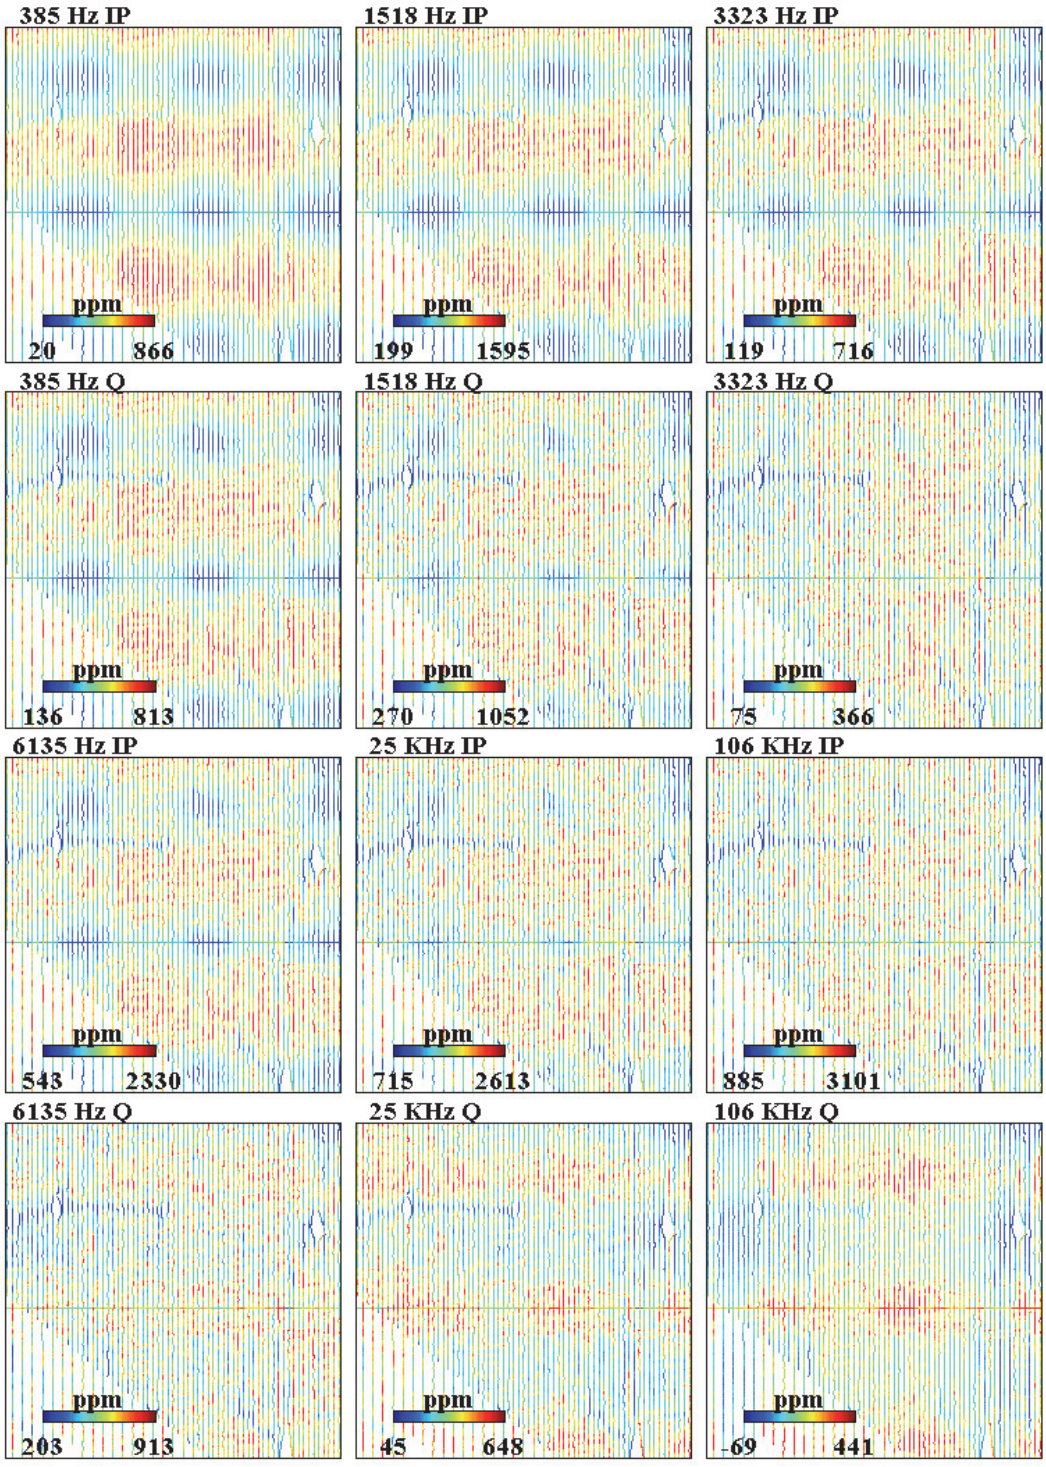

Figure 4.13 Flight lines coloured by the synthetically generated data, which are effected by calibration errors and height variations.

At the actual airborne sample locations local 1-D conductivity profiles were extracted and forward modelled (Equation 3-18) to generate airborne responses for a perfectly calibrated RESOLVE system with the bird at the altimeter measured height plus one metre (i.e. this simulated a the case of an altimeter measuring the bird height as 1 m less than the true height). Ideally full 3D forward modelling would have been used, but access to appropriate software was not available. The theoretical responses were transformed via the calibration model Equations 3-20 and 3-21 to synthesise uncalibrated responses. Gaussian random noise was computed from the noise model parameters (Table 4.3) and was added to the responses. A total of 236,976 airborne data were generated.

The synthetic data, including the systematic calibration errors and noise, are shown in Figure 4.13. In this figure, the effects that the calibration parameters have on the synthetic airborne data appear as artefacts elongate in the flight line direction with patterns that correspond to the flying date and flight number patterns depicted in Figure 4.9a and b. The effect of the variations in bird height (Figure 4.9c) are also readily apparent as the shorter spatial wavelength variations in Figure 4.13. For example, a prominent higher altitude (red) east–west trending feature in the top right of Figure 4.9c is reflected in as lower values (blue) in Figure 4.13, especially for the higher frequencies. The influence of the cosine function generated conductivity model is more apparent in the lower frequencies which are not so affected by height variations.

The area included five boreholes whose locations are shown on Figure 4.10. Downhole logs that were consistent with the conductivity model were generated at each of the borehole locations. The synthetic downhole logs were contaminated with 5% Gaussian random noise. A total of 129 geoelectric data input into the inversion since each downhole log was averaged over 2 m depth intervals. Synthetic interface depth data were not applicable in this single layer example.

4.3.2 Inversion

The 237,105 synthetic data were inverted to estimate 5,407 parameters. There were 4,620 spline node-coefficients on the single 150×150 m mesh which parameterized the conductivity and 787 calibration parameters. The calibration parameters were comprised of: 6 gain parameters, one for each of the 6 coilsets. There were 60 phase parameters, one per coilset for each of the 10 days of flying.

The bias was parameterised by separate piecewise linear function for each channel over each flight. The length of each piecewise interval was 12000 fiducials or 20 minutes. The method of 'start value and gradients' (Equation 3-4) was used so that unknowns to be solved for were the bias at the beginning of the flight and the gradients (bias drift rate) over each 20 minute interval. There were therefore 168 start of flight bias parameters, one for each of the 14 flights times the 12 channels of data. There were 552 drift rate parameters, because there were 12 channels and 46 sets of 20 minute intervals. The number of intervals depended of course on the duration of each flight. In addition there was a single height correction parameter to solve for.

A homogenous 0.35 S/m halfspace reference model was used in the inversion. This value was chosen as it was the average of all the synthetic downhole conductivity logs. This simulates what might typically be done in the inversion of real survey data when just a few downhole logs are available but there is not enough information available to generate a spatially variable reference model. The reference value for the gain parameters was unity and the reference values for phase, bias and height parameters were all zero. Therefore the prior expectation was that the airborne system was perfectly calibrated. The details of the inversion parameters and the assigned reference values and reference value uncertainties are summarised in Table 4.5.

Table 4.5 Summary of synthetic inversion model parameters.

Local conductivity profile & calibration values	Dependence	Inversion parameters	Number of unknowns	Reference model	Reference model standard deviation
σ_1	location	150×150 m mesh	4620	log(0.35) †	0.4 †
t_1	n/a – one layer	n/a	n/a	n/a	n/a
gain	coilset	gain	6	1.0	0.1
phase	day, coilset	phase	60	0.0º	2.0º
bias	flight, fiducial, coilset, channel	start bias	168	0.0 ppm	survey statistics*
		bias gradients	552		
height error	survey	Δ height	1	0.0 m	0.5 m

† units of log(S/m) * statistically calculated from survey data

This inversion was run using the first generation of code developed for a single processor computer (*cf.* Section 3.6.4). The convergence of the inversion's outer and inner loops (*cf.* Figure 3.8) is summarised in Figure 4.14. Its left hand panel shows that the normalized data misfit (i.e. data misfit divided by the total number of data *cf.* Equation 3-77) was reduced to unity, the level indicating that the data were acceptably fitted, within 8 outer loop iterations.

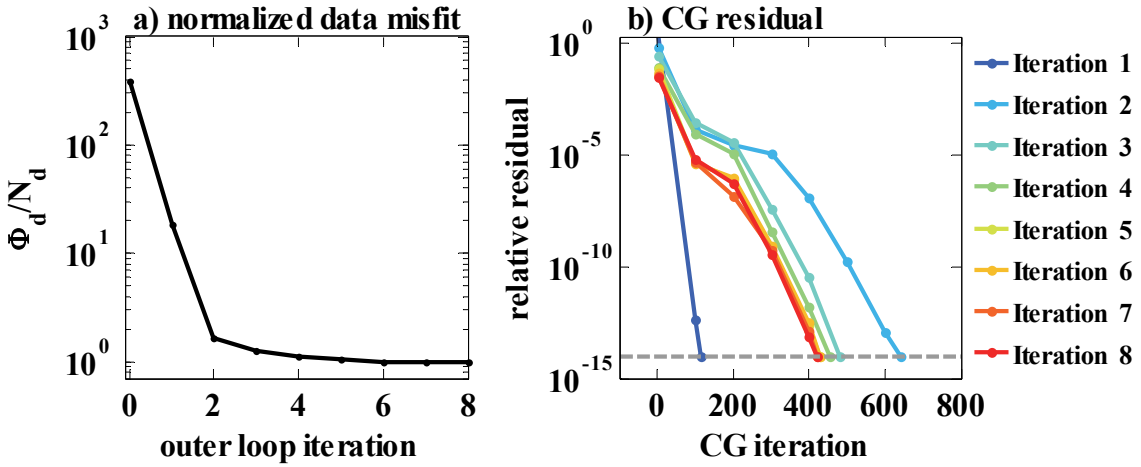

Figure 4.14 Summary of the synthetic inversion's convergence, showing, (a) the reduction of the normalized data misfit (Equation 3-77), and (b) the reduction of the relative residual of the conjugate gradient solution of the linearized system in each outer loop iteration (Equation 3-78).

Recall that inside each outer loop iteration a linearized system of equations, designed to estimate the next set of model parameters (Equation 3-71), is solved via the conjugate

gradient (CG) method. In Figure 4.14b the convergence of the iterative CG algorithm, which I call the inner loop, is shown for each of the outer loop iterations. Each coloured curve denotes the trajectory of the CG residual for one outer loop iteration, however note that the residuals were only logged every 100 CG iterations. The CG solver was set to return a solution when the relative residual (Equation 3-78) was reduced to a value $tol = 10^{-14}$. A solution was typically found within 400 to 500 iterations, and 640 iterations at most. This is significantly less than the 5,407 (N_m) dimension of the system, and thus indicates that the CG solver performed well.

Unlike for the larger examples presented later, for this relatively few parameter example it was possible to estimate the condition number of the linear system matrix **A**. Using MATLAB's *condest()* function, which is based on methods developed by Hager (1984) and Higram and Tisseur (2000), the condition number was estimated to be $\kappa = cond(\mathbf{A}) = 9.78 \times 10^{12}$. Since the reciprocal of κ is much larger than the machine precision (2.22×10^{-16}) the linear system would not be considered to be ill-conditioned (Parker, 1994; Aster et al., 2005).

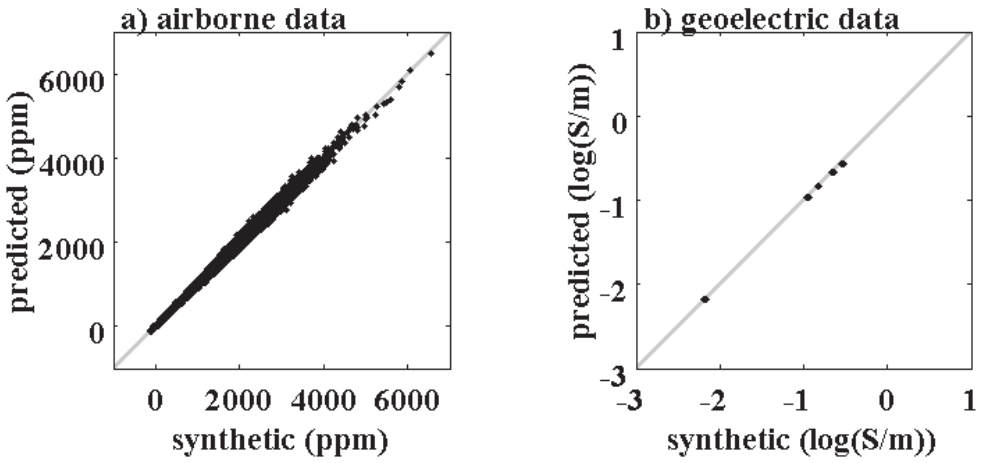

Figure 4.15 Synthetic versus predicted (a) airborne and (b) geoelectric data.

Figure 4.15 shows the fit between the synthetically generated data and the predicted data from the final inversion model. If the data fitted perfectly all points would lie on the diagonal line of each plot. The non-perfect fit, i.e. the scatter about the diagonal line, is due to the random noise that was included in the synthetic data.

There were 129 geoelectric data in the inversions, but it appears that there are only five data points on Figure 4.15b. This is because the geoelectric data were generated to be consistent with the one layer model from five synthetic downhole conductivity logs averaged over 2 m depth intervals. Hence the data from different depths in the one borehole are almost identical and plot coincidently. The simplicity of the model in this synthetic example contributes to the apparently low geoelectric data misfit. In reality however, downhole conductivity logs are always much more complex than any inversion model can accommodate, and thus the geoelectric data misfit will never be so low.

4.3.3 Results

In Figure 4.16 a comparison between the synthetic (a) and the estimated (b) conductivity models is presented. It is immediately obvious that the conductivity was recovered almost perfectly. Importantly, in the recovered model there are not artefacts elongate in the flight line direction that would be indicative of the systematic gain, phase and bias calibration errors that were in the synthetic data. This means that the holistic inversion's calibration model has fully accounted for the systematic errors. There is a small 'bull's eye' artefact along the bottom border of the estimated conductivity model indicated by the arrow. This was because of the gap in the flight path gap in that vicinity, which can be seen in Figure 4.9.

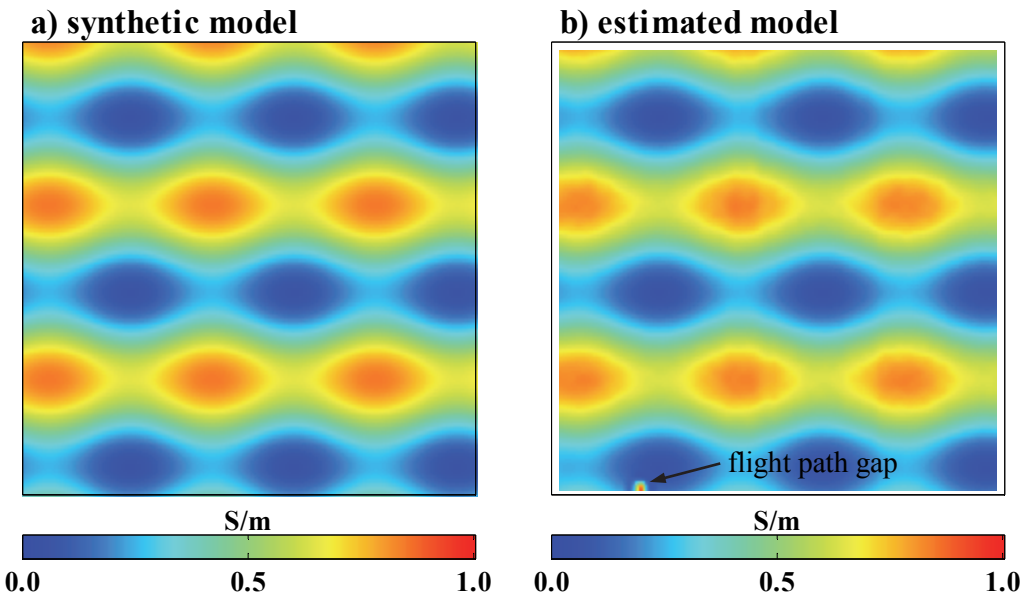

Figure 4.16 Comparison of (a) the synthetic conductivity model used to generate the synthetic airborne and geoelectric data, and (b) the conductivity model estimated by the holistic inversion.

The scatter plots shown in Figure 4.17 summarise the agreement between the synthetic and the estimated calibration errors. In each plot the diagonal grey line is the line along which all points would ideally lie if all parameters were estimated precisely. The gain (Figure 4.17a) and phase (Figure 4.17b) are estimated almost exactly. The inphase (Figure 4.17c) and quadrature (Figure 4.17d) bias scatter plots shows bias values plotted at 20 minute intervals throughout all flights. The biases are not recovered precisely but there is good correlation between the true and estimated values. The inversion gave a height calibration value of 1.008 m, which is effectively the same as the true value of 1.0 m.

The scatter of points about the diagonal in Figure 4.17 is due to the presence of noise. This synthetic example demonstrates that when a reasonable amount of Gaussian noise is added to the data, the recovered parameters of the conductivity and calibration models are acceptably close to the true values. This indicates that, at least in this idealized case where the parameterization is consistent with the synthetic models, the sensitivity of the estimated values to Gaussian noise is within an acceptable range. Ill-posedness, or

fundamental trade-off between the five classes of parameters, does not appear to be a serious problem in this case.

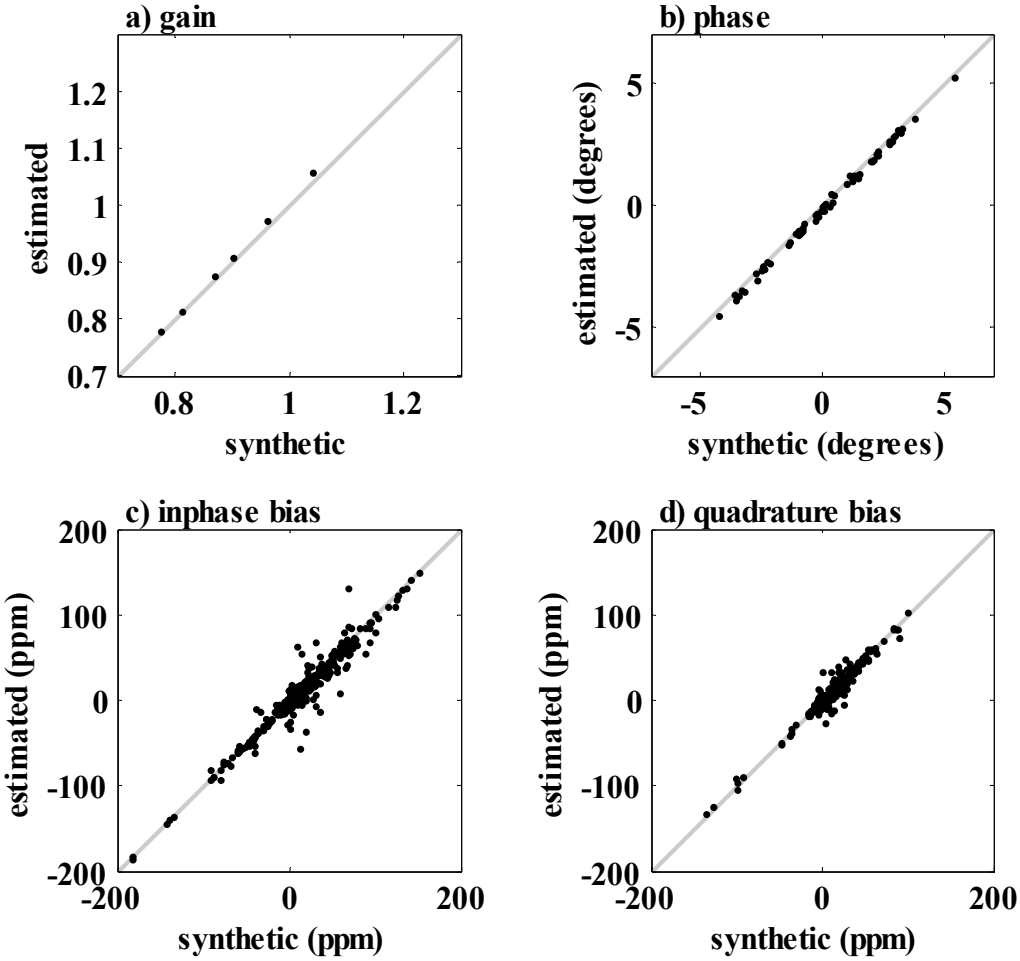

Figure 4.17 Synthetic versus estimated (a) gain, (b) phase, (c) inphase bias, and (d) quadrature bias.

When Gaussian noise is not included in the synthetic data, all of the parameters are estimated exactly. This confirms that the holistic inversion theory, i.e. the forward models, derivatives and minimization, have all been implemented correctly in code.

4.4 Few-layer inversion

4.4.1 Introduction

In this section the holistic inversion is demonstrated using real AEM survey data from the Riverland Airborne Electromagnetic Survey. The Riverland dataset has previously been demonstrated to be afflicted by systematic calibration errors, particularly from gain

or scaling errors (Brodie et al., 2004a; Ley-Cooper and Macnae, 2004). As was explained in Section 4.2.5, over one thousand separate gain, phase or bias adjustments were applied to the dataset during the conventional data processing. Nevertheless, Brodie et al. (2004b) had to rescale the data by up to 129% to successfully fit the data in their sample-by-sample inversion. The scaling factors had been derived by regressing forward models of downhole log data against observed airborne data (Brodie et al., 2004a).

It is not altogether surprising that data could not be fitted because many of the adjustments applied in the data processing were selected by inspection of apparent resistivity maps that were generated independently for each coilset frequency. While this may have made the apparent resistivity map for each frequency appear self-consistent or coherent, it did not guarantee mutual consistency between frequencies, which is required by the physics, and hence to fit the data. This holistic inversion example demonstrates how those somewhat subjective and time consuming data processing steps, designed to remove systematic calibration error, can be avoided by solving for calibration factors along with a conductivity model. In so doing, consistency between the data from every coilset, prior information, and the estimated conductivity model is guaranteed. Moreover, that it is achieved through a more objective and efficient procedure than the numerous data processing corrections.

In this holistic inversion example, the conceptual geological model and the prior information that were compiled by Brodie et al. (2004b) are used (*cf.* Section 4.2.7). However, neither data that have had the data processing corrections, nor the downhole log derived scaling corrections applied, are inverted. Therefore the data that are inverted are essentially raw and the systematic errors are instead accounted for by the gain, phase and bias calibration model parameters.

4.4.2 Inversion

Data

The subset of the Riverland Survey dataset chosen for the few-layer inversion of real data covers approximately one quarter (406 km^2) of the total survey area. The sub-area is shown on the locality diagram Figure 4.1. This is the largest portion of the dataset that could be inverted, using the chosen data sub-sampling and parameterization, on a computer with two gigabytes of random access memory. In this case, computer memory was the limiting factor because it was executed with the first generation code, which is a single processor program. A second generation parallel version of the code, designed for execution on a cluster computer, was later developed. It handles many more data and parameters, however that code is used in the multi-layer inversion in Section 4.5.

The sub-area was traversed by 2,590 line kilometres of airborne survey data. The data were acquired over 15 days of flying, 27 different flights and 171 flight lines, four of which are tie lines and 19 of which are daily repeat lines. The inverted data were raw, except for the aforementioned reversal of zero-level shifts that had been introduced during the high altitude calibrations (*cf.* Figure 4.7c), and the 9-point median and Hanning random noise reduction filters discussed in Section 4.2.5. The filtered airborne data were sub-sampled to every 17*th* sample (~56 m) along lines to reduce the number of data and hence memory requirements. Note however that the inverted samples were still closer together than the minimum footprint of the AEM system. A total of 47,934 samples were inverted, or 575,208 airborne datum given that there 6 coilsets and 12 channels per sample. The method of Green and Lane (2003) was employed to estimate noise in the airborne data using the parameters provided in Section 4.2.6.

The 16 downhole conductivity logs in the area were averaged over two meter depth intervals to provide 330 geoelectric data. These were assigned a noise level standard

deviation of 0.1 natural logarithm decades. The watertable elevation grid shown in Figure 4.8b was intersected to generate interface-depth data associated with the interface between the third and fourth layers of the conductivity model, which in the conceptual hydrogeological model (see next section) corresponded to the standing water level. The interface depth data were chosen to lie on a 60×60 m regular grid, so that the spacing was similar to the AEM system's minimum footprint. There were a total of 113,961 interface depth data which were assigned a noise level of 0.2 m.

Parameterization

The conceptual hydrogeological model shown in Figure 4.18, and used in the sample-by-sample inversion carried out by Brodie et al. (2004b), forms the basis for the parameterization in this holistic inversion. The conceptual model was based on sound prior knowledge of the geology of the survey area, groundwater depth and downhole log information. It consists of recently deposited dry and resistive sands that overlie a more conductive clay rich unit known as the Blanchetown Clay. Below this is the Loxton-Parilla Sands unit which tends to be resistive in the unsaturated zone. However the Loxton-Parilla Sands are conductive in the saturated zone due to the presence of saline groundwaters. Brodie et al. (2004b) also found that there tended to be elevated conductivities (see 'conductivity bulge' in Figure 4.18) in the top 10 m of the saturated zone immediately below the standing water level. This prompted the use of a five-layer geoelectric model consisting of the following layers: (i) Recent Sands; (ii) Blanchetown Clay; (iii) unsaturated Loxton-Parilla Sands; (iv) 10 m thick Conductive Bulge, and a (iv) Conductive Basement.

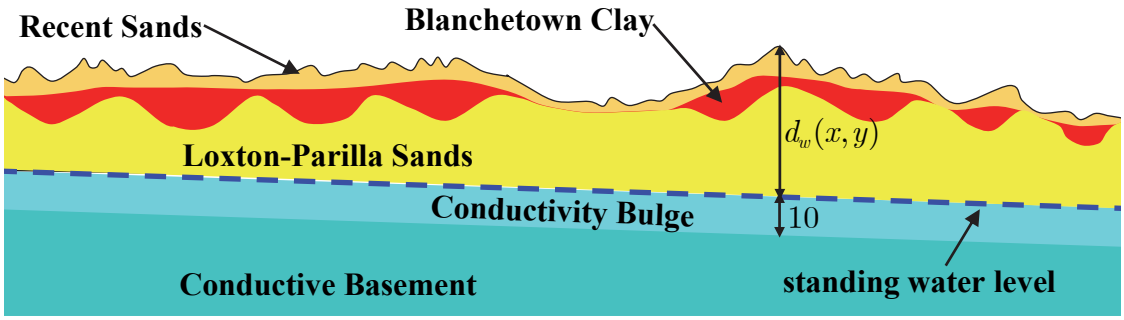

Figure 4.18 The conceptual hydrogeological model for the few-layer inversion.

Since the purpose of the survey was to map the distribution and thickness of the Blanchetown Clay, a few-layer variable-thickness formulation in which the clay is represented by the second layer is suitable. This is preferred over the potentially more stable multi-layer fixed-thickness Occam's style formulation (Constable et al., 1987) because the estimated second layer thickness is a proxy that can be directly used in generating a clay distribution map without further interpretation/model manipulations being required.

For the majority of the spline meshes the node spacing was set to 150 m in both the east–west and north–south directions. This distance was chosen because it was equal to the nominal survey line spacing and was found to be adequately fine to allow the data to be fitted. The Recent Sands unit in the conceptual model exists in part as east–west orientated surficial sand dunes that are visible on Figure 4.2b. The dunes create relatively rough high spatial frequency variations in the topography and hence in the thickness of the top layer. To accommodate this, the Layer 1 thickness spline mesh was set to have a smaller node spacing of only 75 m in the north–south direction, but was kept at 150 m in the east–west direction. The groundwater conductivity was known to vary slowly as the prior information in Figure 4.8d shows. It was expected then that the conductivity of the conductive basement would also vary slowly, and the spacing of the spline nodes for the Layer 5 conductivity was set at 1000×1000 m. A finer 300×300 m mesh was used for Layer 4 conductivity to allow for known more rapid changes in the Conductive Bulge's bulk conductivity thought to be associated with porosity and

permeability variations (Munday et al., 2004). The thickness of the Conductive Bulge (Layer 4 thickness) was kept fixed at 10 m.

Table 4.6 Summary of few-layer inversion model parameters.

Local conductivity profile & calibration values	Dependence	Inversion parameters	Number of unknowns	Reference model	Reference model standard deviation
σ_1	location, layer	150×150 m mesh	68,404	log(0.030)†	0.125†
σ_2		150×150 m mesh		log(0.240)†	0.250†
σ_3		150×150 m mesh		log(0.065)†	0.125†
σ_4		300×300 m mesh		$\log(\sigma_w(x,y))$†	0.625†
σ_5		1000×1000 m mesh		$\log(\sigma_w(x,y))$†	0.045†
t_1	location, layer	150×75 m mesh	78,170	log(3.0)‡	2.5‡
t_2		150×150 m mesh		log(3.0)‡	2.5‡
t_3		150×150 m mesh		$\log(d_w(x,y)-6)$‡	2.5‡
t_4	fixed	n/a	0	10 m	n/a
gain	survey, coilset	gain	6	1.0	0.1
phase	day, coilset	phase	90	0.0º	0.5º
bias	flight, fiducial, coilset, channel	start bias	324	0.0 ppm	survey stats*
		bias gradients	1,680		
height error	fixed	n/a	0	n/a	n/a

† units of log(S/m) ‡ units of log(m) * statistically calculated from survey data

The conductivity model reference values and their uncertainty standard deviations were chosen to the same as those used by Brodie et al. (2004b). The actual values are summarised in Table 4.6. Spatially variable reference values were used for the Layer 3 and 4 conductivity which were taken to be equal to the logarithm of the gridded value of the bulk conductivity below the watertable ($\log(\sigma_w(x,y))$), which had been compiled from prior measurements of groundwater EC and downhole conductivity log information (*cf.* Figure 4.8d). The reference value for Layer 3 thickness was also spatially variable and was the logarithm of depth to the standing water level less the cumulative thickness of the top two layers of the reference model (i.e. $\log(d_w(x,y) - 3 - 3)$). This placed the interface between the third and fourth layers of

the reference model at the standing water level as defined in the conceptual model. The groundwater depth $d_w(x, y)$ was derived by subtracting the gridded value of the elevation of the standing water level (*cf.* Figure 4.8b) from the gridded value of the elevation of the surface topography (*cf.* Figure 4.2b).

The gain was parameterized to have a constant value for each coilset over the entire survey, thus 6 parameters were required. This decision was made based on the observations made in the prior calibration study (Brodie et al., 2004a) that much of the apparent systematic error could be explained by applying survey-wide constant scaling factors to the data for each coilset.

The phase calibration was parameterized in piecewise constant fashion, with one value for every day of flying and for each coilset so that there were 90 (6×15) parameters. This parameterization was chosen to emulate the phase corrections made during ground calibrations that are typically performed on a daily basis.

The experience of airborne frequency-domain data acquisition service providers has determined that high altitude zero-level measurements need to be made at approximately 20 minute intervals to track zero-level drift. It is a typical survey contract specification and was the case for the Riverland survey. The bias was thus parameterised as piecewise linear variations over each flight where all piecewise intervals had the same length and were as close as possible to 20 minutes long. Separate functions were used for each channel. The first interval for a flight started at the first inverted sample in the flight and the final interval ended at the last inverted sample in the flight.. The number of intervals, of which there were 140 in total, depended of course on the duration of the flights. The method of 'start value and gradients' (Equation 3-4) was used, thus the unknowns to be solved for were the bias at the beginning of the flight and the gradients (bias drift rate) over each piecewise time

interval. There were therefore 324 (27 flights×12 channels) start of flight bias parameters, and 1,680 (140 piecewise intervals×12 channels).

Tests were carried out during the survey to check the output of the laser altimeter system mounted on the towed bird assembly. The calibrations, a series of passes along an aerodrome runway at different heights, confirmed linearity between the laser altimeter and altitude measurements made with a GPS receiver mounted on the bird (Cowey et al., 2003). The tests ruled out the possibility of systematic scaling errors in the altimeter data. The absolute error in the derived runway elevation was +0.49 m, which implies that the altimeter may have underestimated height by 0.49 m. However, this offset is probably smaller since the survey marker defining the aerodrome height was actually off to one side of the runway at a slightly lower elevation. There was therefore no strong motivation to include a height calibration parameter for this particular dataset.

One might infer altimeter error as the reason why Brodie et al. (2004a) had to apply scaling factors to the data in order for them to match forward models of downhole conductivity logs, and so that they could be fitted in inversions. However this is not a likely explanation because the derived scaling factors are not monotonically increasing or decreasing with frequency (*cf.* Table 4.4). A monotonic relationship would be expected if altimeter error was the cause. Furthermore, the height errors would have to be of the order of 3 m for the highest frequencies, which is not supported by the calibration tests described above.

In choosing the reference values for the calibration model, it was assumed that the system was perfectly calibrated. Therefore a value of unity was used for the gain parameters and zero was used for phase and bias calibration model parameters. The standard deviation uncertainty was set to 0.1 for the gain reference values and set to 0.5° for the phase reference values. The standard deviation uncertainty for the bias reference

values were computed statistically from the variations observed in high altitude zero-level data measured during the survey.

Solution

In total 689,499 data were inverted, consisting of 575,208 airborne data, 330 geoelectric data and 113,961 interface-depth data. There were 148,674 unknowns to be estimated, consisting of 146,574 were spline node coefficients and 2,100 were calibration model parameters.

Reference model constraints were applied ($\lambda_r = 5$) on all parameters using the reference model and standard deviation uncertainties summarised above in Table 4.6. Roughness constraints were applied in the east–west ($\lambda_x = 1$) and north–south ($\lambda_y = 1$) directions. The horizontal roughness was applied relative to layer number (*cf.* Equation 3-51) rather than relative to elevation. Since it was a few-layer model no vertical roughness constraints were applied ($\lambda_z = 0$). Bias roughness constraints were also applied ($\lambda_b = 3 \times 10^8$).

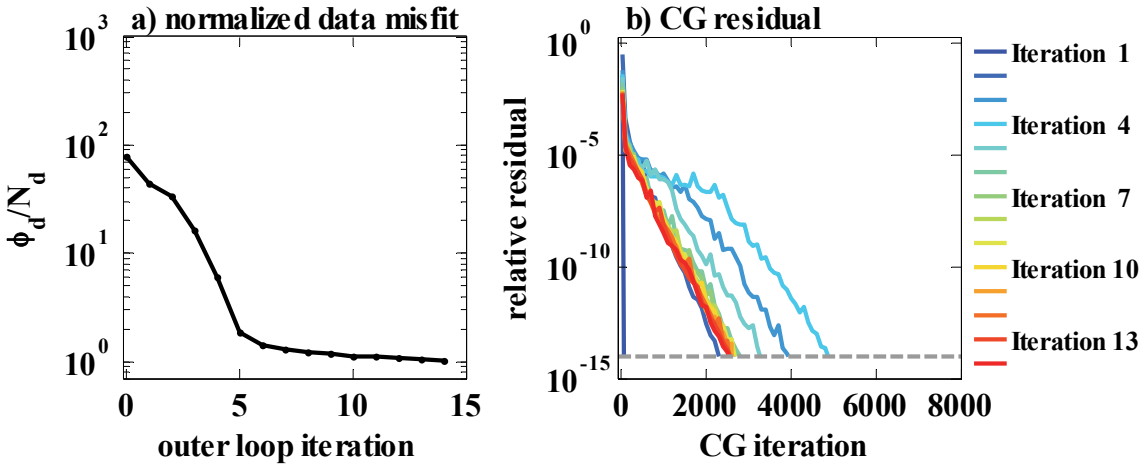

Figure 4.19 Summary of the few-layered inversion's convergence, showing, (a) the reduction of the normalized data misfit (Equation 3-77), and (b) the reduction of the relative residual of the conjugate gradient solution of the linearized system in each outer loop iteration (Equation 3-78).

The sparsity structure of the matrices involved in this particular inversion example was presented in Figure 3.10. The convergence of the inversion's outer and inner loops (*cf.*

Figure 3.8) is summarised in Figure 4.19. The left hand panel shows that the normalized data misfit (Equation 3-77) was reduced to unity, the level indicating that the data were acceptably fitted, within 14 outer loop iterations.

The conjugate gradient (CG) algorithm, which solves the linearized system inside each outer loop iteration, was set to return a solution when the relative residual (Equation 3-78) was reduced to a value $tol = 10^{-14}$. In Figure 4.19b each coloured curve denotes the trajectory of the relative residual (logged at every 100th CG iteration) for one outer loop iteration. It can be seen that the CG solution of the linearized system always returned a solution in less than 5,000 iterations which is much less that the 146,574 unknown model parameters. Due to the size of the linear system, its condition number could not be calculated like it was for the relatively few parameter synthetic case study presented earlier. However, the rapid and stable convergence of the CG solution shown in Figure 4.19b suggests that that the linear system probably was well conditioned.

4.4.3 Results

Calibration parameters

In Figure 4.20a the reciprocal of the gain estimated from the holistic inversion is plotted against the scaling factors derived by Brodie et al. (2004a). It is necessary to plot the reciprocal of gain for proper comparison because the scaling factors were designed to scale uncalibrated observed data to simulate perfectly calibrated data, whereas the holistic gain parameters did the reverse. The comparison shows that there is good agreement between the two sets of values. This provides some confidence in the credibility of the estimated gains. However, this is not a completely independent verification of these parameters since many of the same downhole logs were used in both methods.

There are no independent calibration studies to compare to the holistic inversion phase error estimates shown in Figure 4.20b. However for the 385 Hz coilset, the inversion returned an average phase error of -3.68º. This compares well, once the reversed transformation direction is taken into account, with the +3.5º phase rotation applied during conventional processing (Cowey et al., 2003). In the conventional processing phase rotations, ranging from +2.0º to -3.5º depending on the flight number, were also applied to the 106 kHz data. Phase rotations were not applied to other frequencies in the conventional processing.

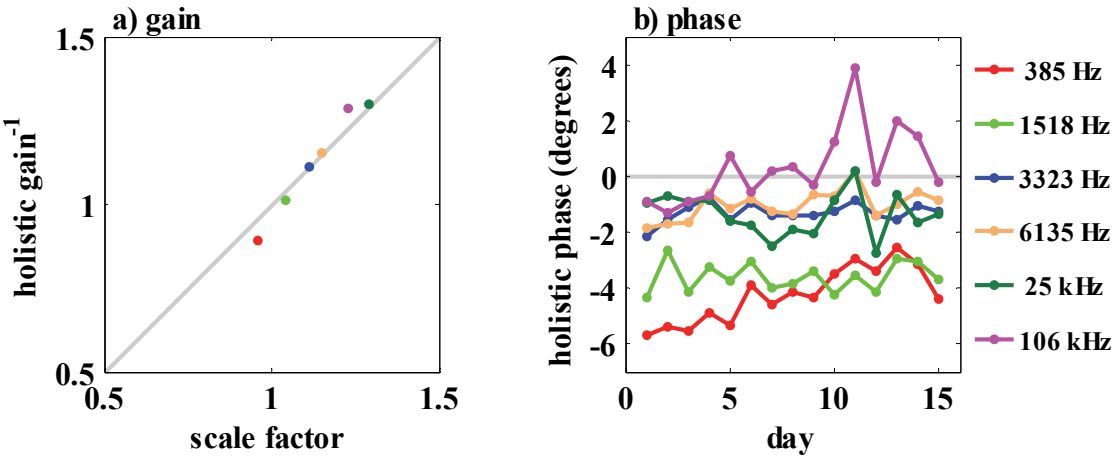

Figure 4.20 *Calibration parameters showing (a) scale factors derived by Brodie et al. (2004a) versus the reciprocal of the gain estimated from the holistic inversion, and (b) the phase estimates for each day of flying.*

Figure 4.21 shows comparisons between the zero-level values that were measured during the high altitude excursions and bias values predicted from the holistic inversion parameters. With the exception of the 106 kHz inphase data, there is a remarkably good correlation between observed and predicted values. Recall that the observed zero-level values were not used to create the reference or starting models, so the good correlations are indicative of the success of the algorithm rather than being artefacts of the setup.

In general it is desirable that if zero-level data are measured, as is usually the case, they should be used to create the reference model for the inversion. They were not used to set reference values in this inversion so that an independent assessment of the results could be made. The results demonstrate that, under the right conditions, it may actually

be feasible to eliminate high altitude zero-level observations altogether. This is of considerable practical benefit because it would reduce acquisition costs substantially.

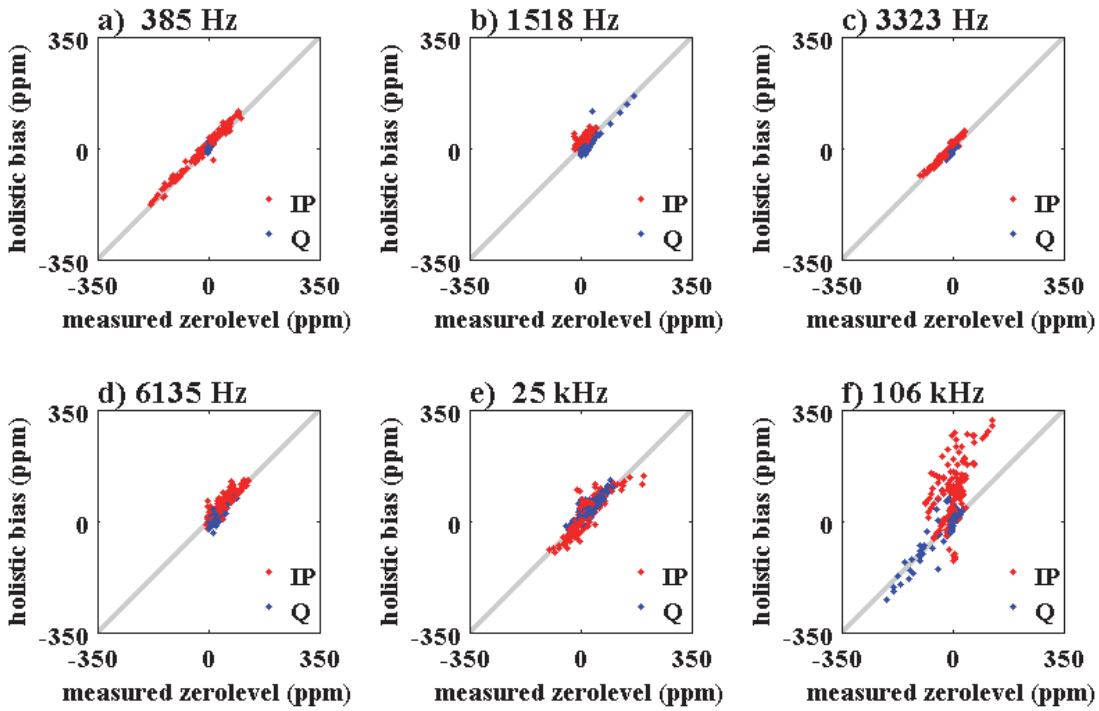

Figure 4.21 Comparison of the zero-level measured at high altitude during the survey and the bias estimated in the holistic inversion.

The reason for the poorer correlation in the 106 kHz inphase data (red markers on Figure 4.21f) is not fully understood. Later in this section it is demonstrated how the conductivity model obtained from the holistic inversion is superior to that obtained using the conventional high altitude zero-level estimates. Furthermore, according to Huang and Fraser (1999) the high altitude zero-level measurement procedure may not yield accurate results because of the differences in temperature between survey altitude (30 m) and calibration altitude (500 m). They note that a change in temperature with altitude can affect the transmitting-receiving coil separation or coil alignment and cause a response of 40 ppm for a 0.1 mm change in coil separation. Considering these points, I speculate that the holistic inversion may have provided a better estimate of bias than the high altitude zero-level measurements.

One possible explanation for the differences in measured zero-level and the holistic inversion estimated biases could be the 'zero-level resets'. It was shown in Figure 4.7b

how the response measured during the high altitude calibrations was sometimes reset to zero. As shown in Figure 4.7c any identified resets were reversed in the survey level data so that the bias could be parameterised as a continuous function. However, any resets not known about, for example those at the beginning of the flight, would not have been accounted for. This would lead to unknown constant differences between the true zero-level and the measured zero-level values shown in Figure 4.21. This explanation is speculative and cannot be confirmed with the available information.

Conventional processing and inversion conductivity models

So that the holistic inversion results can be fully appreciated, the same subset area was inverted using the conventional sample-by-sample method as described in Brodie et al. (2004b). However in this case the fully processed final data that was supplied by the survey contractor, which had been processed in conventional fashion, was inverted. The data were not scaled using factors derived by Brodie et al. (2004a) from a regression of the data against forward models of the downhole conductivity logs.

Images of the Layer 1 and Layer 2 conductivity and thickness grids generated by the conventional sample-by-sample inversion are shown in Figure 4.22. The properties for Layer 3, Layer 4 and Layer 5 are presented in Figure 4.23. Since Layer 4 thickness was fixed at 10 m and Layer 5 thickness is infinite, they have not been included in the figures. An overall north–south striping can be seen on the images of the conventional inversion results. Some specific north–south trending artefacts have been indicated by arrows on these Figure 4.22b and c, and many more are visible. The most prominent artefacts can be seen on other layer property images, for example the artefact running down the right hand (eastern) edge is visible on all Layer 1, 2 and 3 images.

The flight line orientated artefacts in the inversion of the conventionally processed data are symptomatic of calibration errors that have not been correctly removed from the

data. They could also be the effects of systematic errors that may have been unintentionally introduced into the dataset during the conventional processing.

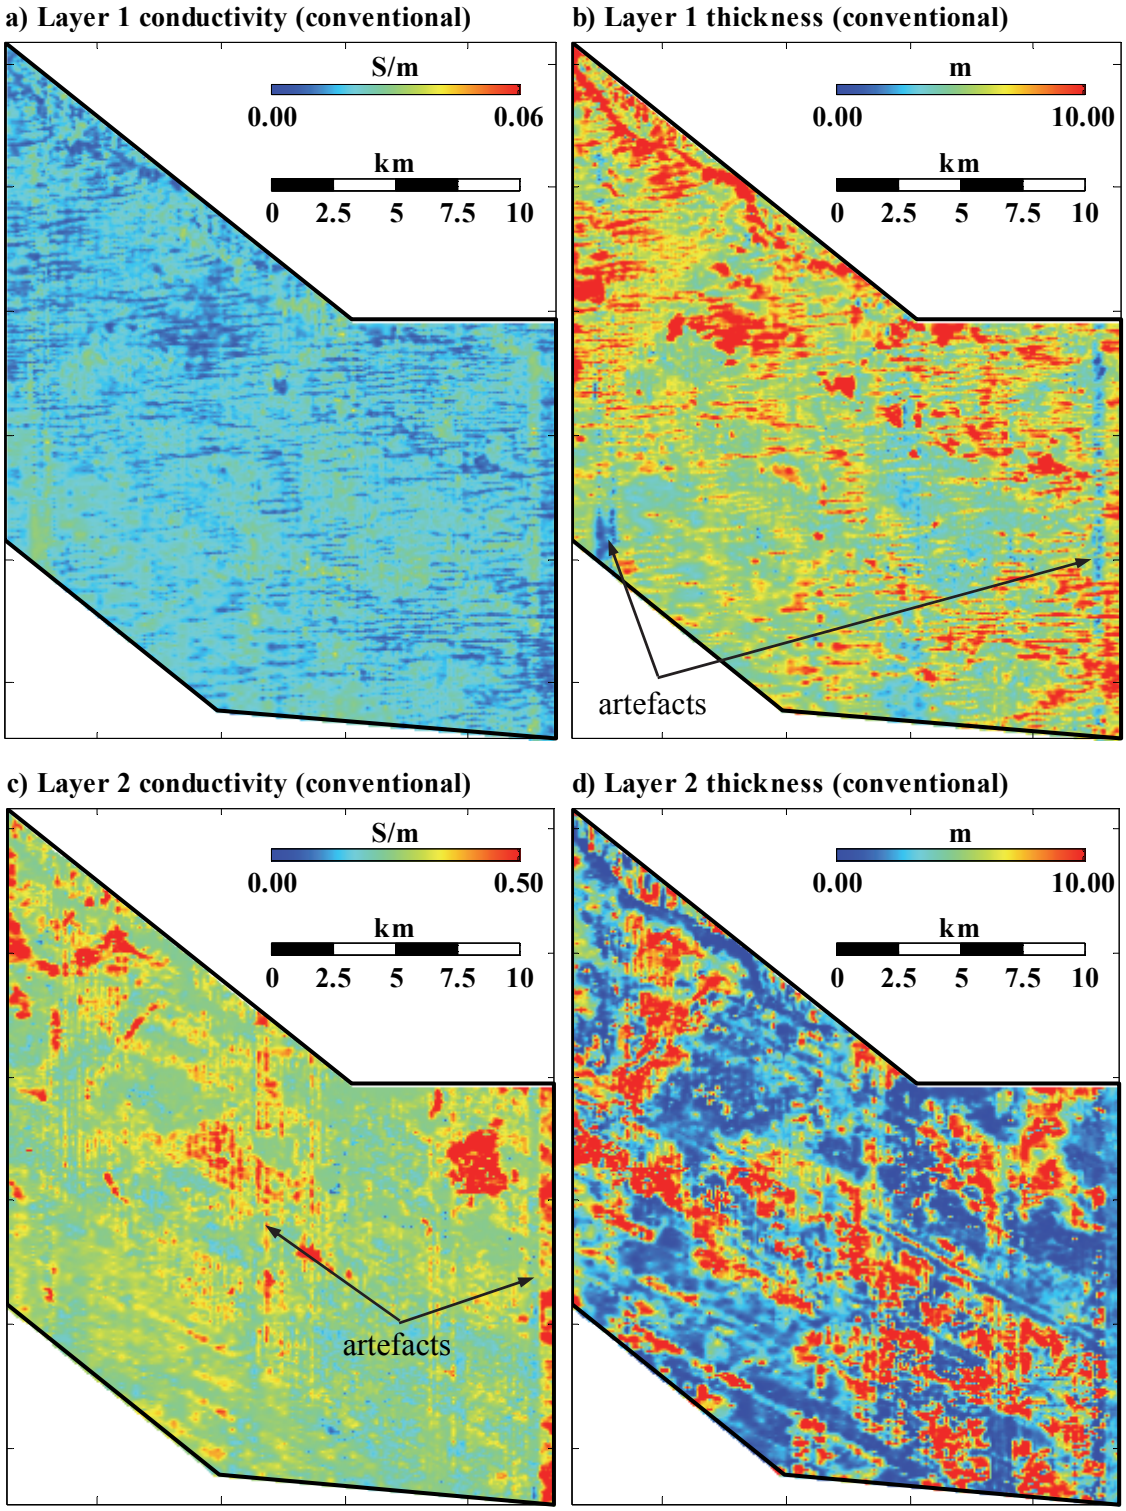

Figure 4.22 Images of the layer properties estimated from the sample-by-sample inversion of conventionally processed data; (a) Layer 1 conductivity, (b) Layer 1 thickness, (c) Layer 2 conductivity, and (d) Layer 2 thickness.

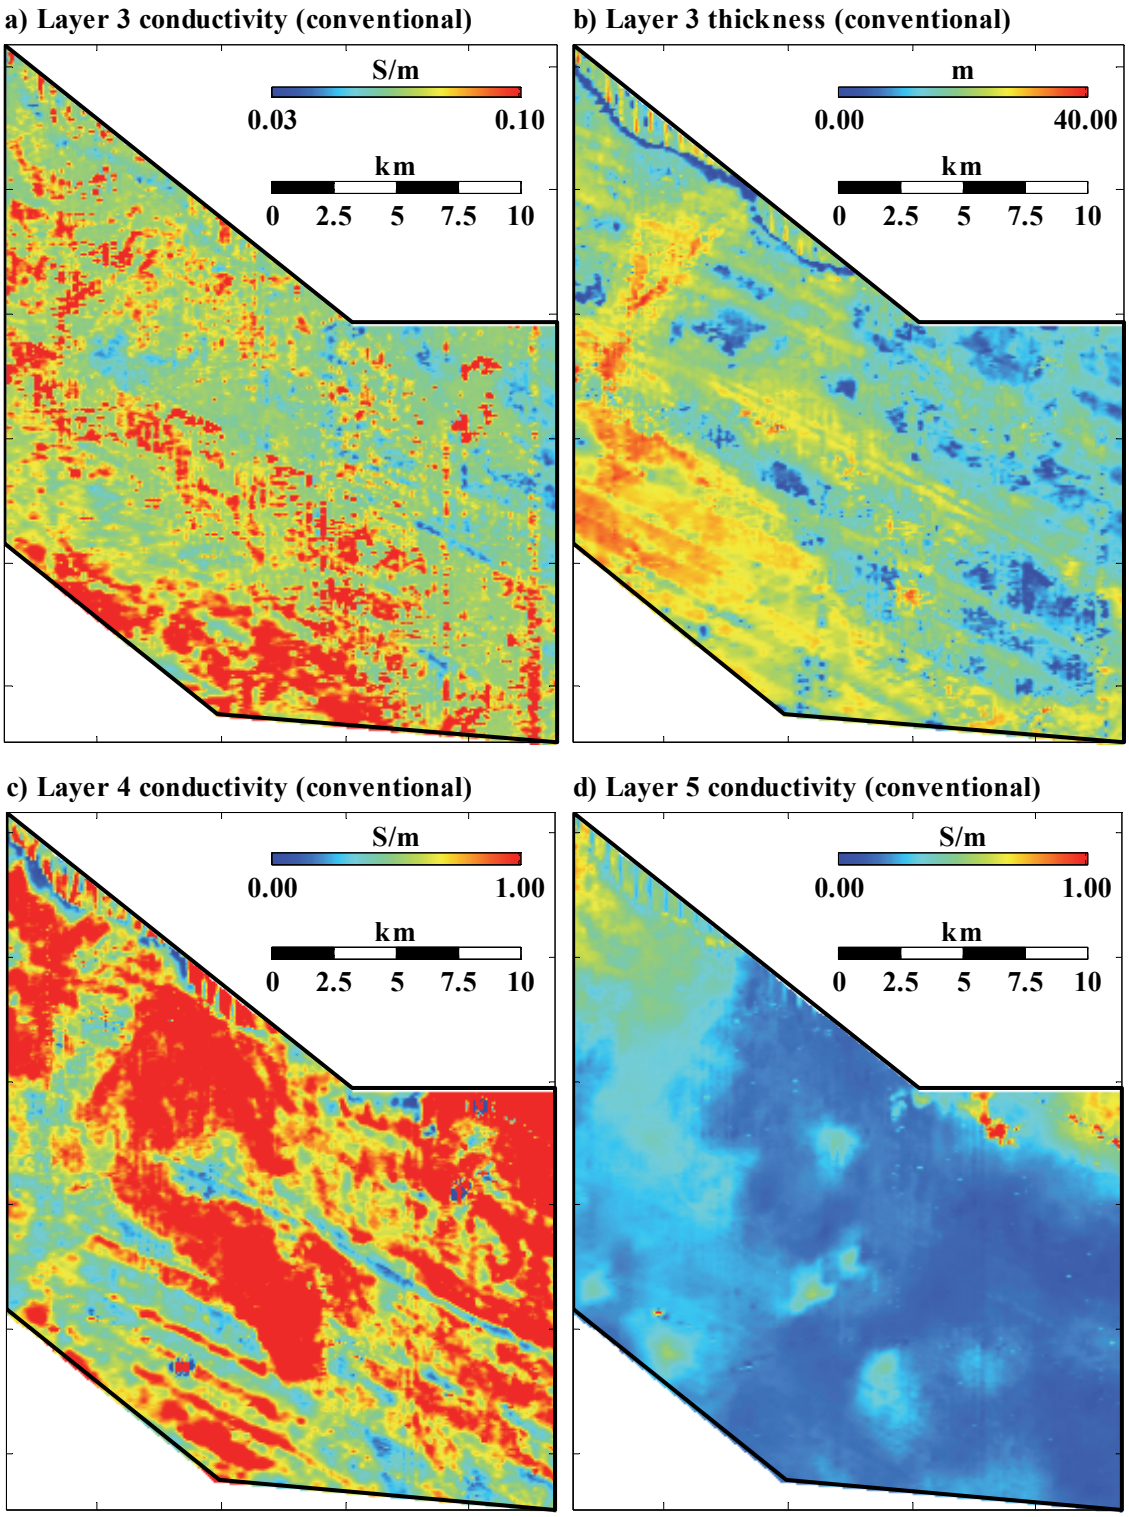

Figure 4.23 *Images of the layer properties estimated from the sample-by-sample inversion of conventionally processed data; (a) Layer 3 conductivity, (b) Layer 3 thickness, (c) Layer 4 conductivity, and (d) Layer 5 conductivity.*

Holistic inversion conductivity model parameters

At the completion of the holistic inversion 30×30 m grids were calculated for each layer's conductivity and thickness. This was achieved by using the final conductivity and thickness spline mesh coefficients parameters in the spline evaluation Equation 3-2, (i.e. where the x,y coordinate was the centre of each grid cell). The 30 m grid cell size was chosen to be one-fifth of the flight line spacing which is typically the cell size chosen for airborne geophysical survey data. The ability to create grids by evaluation of the splines at any position within the model is an appealing property of the holistic inversion. It means that there is no need for gridding, kriging or projection when generating grids and cross sections since conductivity and thickness values can be exactly evaluated from the analytic spline basis where the information is required.

Images of the Layer 1 and Layer 2 conductivity and thickness grids generated by the holistic inversion are shown in Figure 4.24. The properties for Layer 3, Layer 4 and Layer 5 are presented in Figure 4.25. Since Layer 4 thickness was fixed at 10 m and Layer 5 thickness is infinite, they have not been included in the figures.

The images are practically free of artefacts that are elongate in the north–south flight line direction. This is a significant improvement on the results for the sample-by-sample inversion of the conventionally processed data. It is also a satisfying outcome for the inversion of virtually raw airborne survey data. It shows that the holistic inversion has largely accounted for the systematic errors that conventionally would be removed with high altitude zero-level subtractions, manually determined phase and gain adjustments and manually determined levelling base-level shift and tilt corrections. There is an artefact, indicated by the arrow on the Layer 2 conductivity image (Figure 4.24c), which suggests that the bias may not have been accurately determined on that flight.

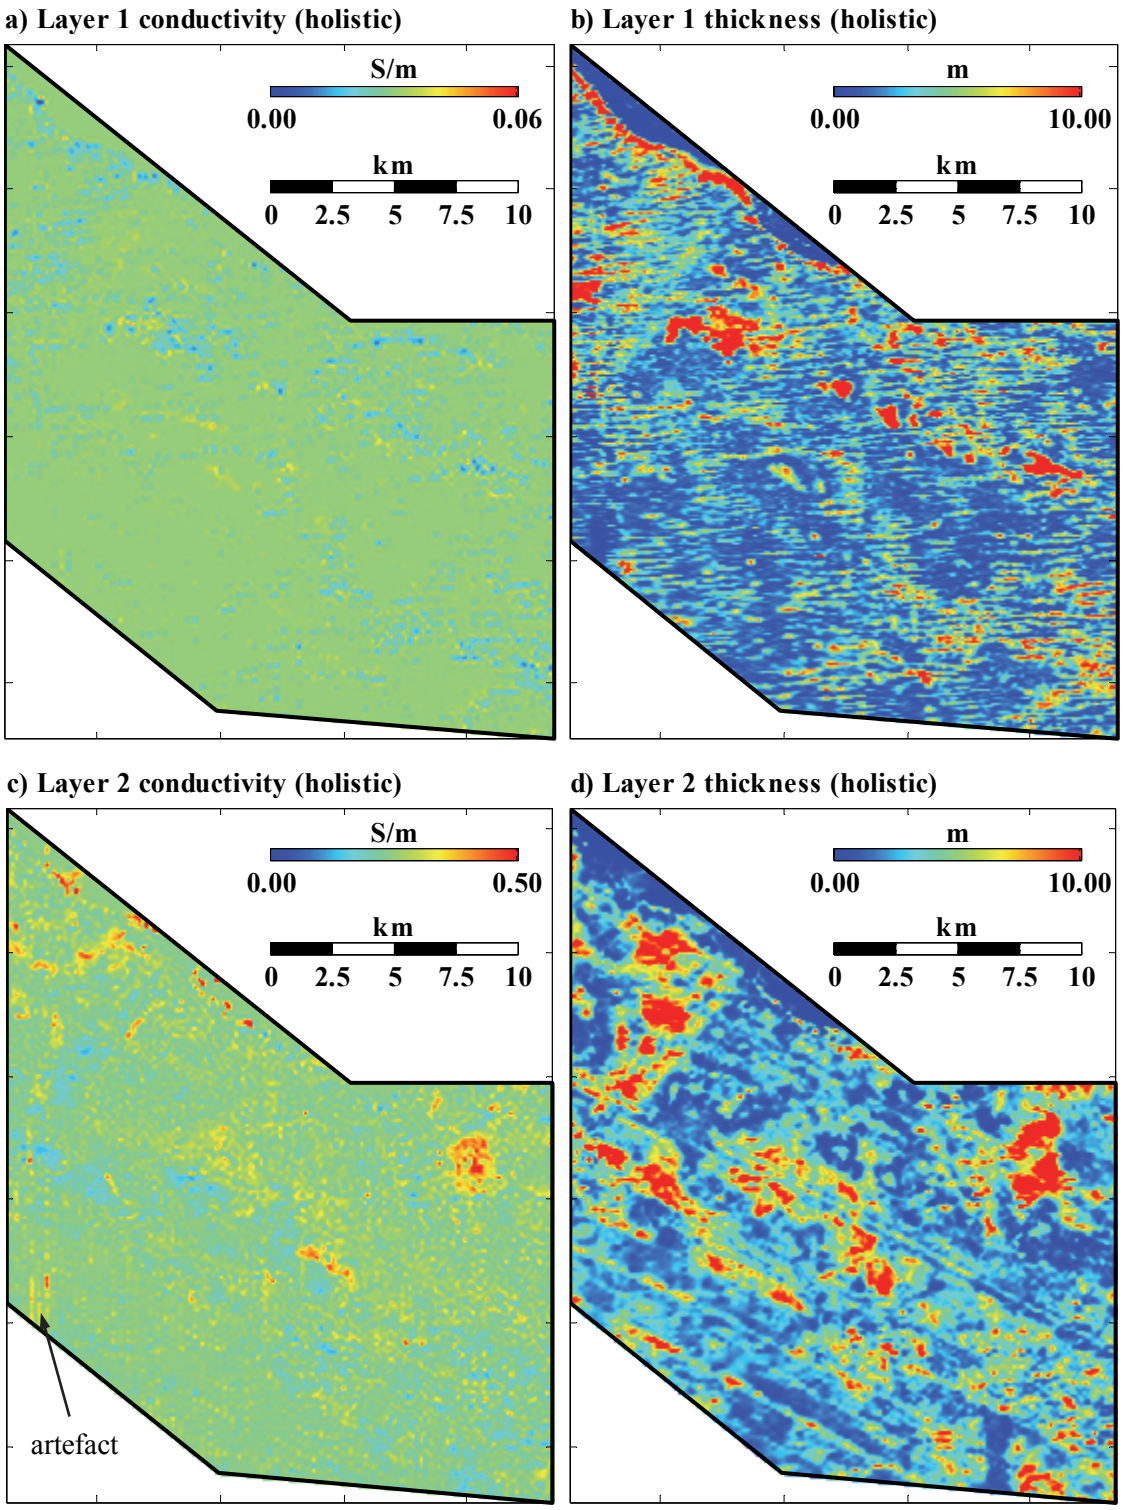

Figure 4.24 *Images of the layer properties estimated from the few-layer holistic inversion; (a) Layer 1 conductivity, (b) Layer 1 thickness, (c) Layer 2 conductivity, and (d) Layer 2 thickness.*

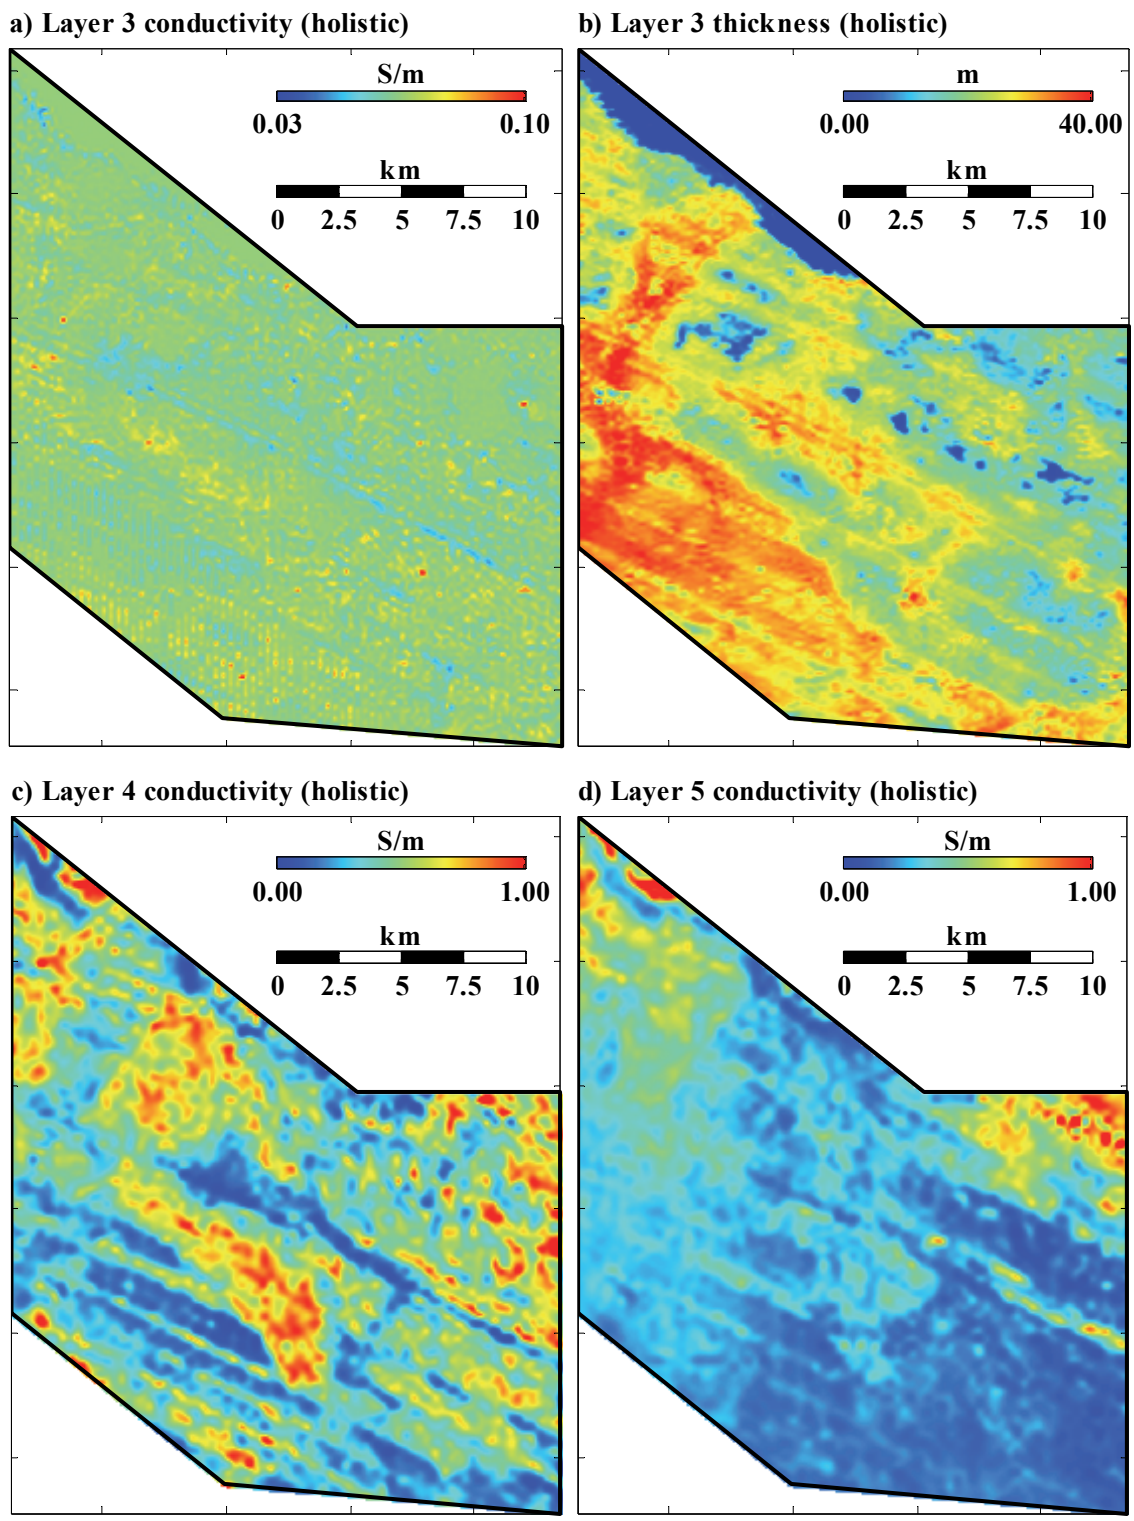

Figure 4.25 Images of the layer properties estimated from the few-layer holistic inversion; (a) Layer 3 conductivity, (b) Layer 3 thickness, (c) Layer 4 conductivity, and (d) Layer 5 conductivity.

Discussion of systematic errors

We can begin to understand how systematic calibration errors in the data migrate through to artefacts in the conventional inversion products by viewing the data, the

apparent conductivities and depths generated from the data, and the inversion residuals. The conventionally processed data from the 106 kHz coilset are displayed in Figure 4.26 as colour coded flight lines. Although masked to some degree by the response variation due to geological and bird height changes, the data contains several elongate north–south orientated anomalies that spatially correspond to the inversion artefacts visible on Figure 4.22 and Figure 4.23.

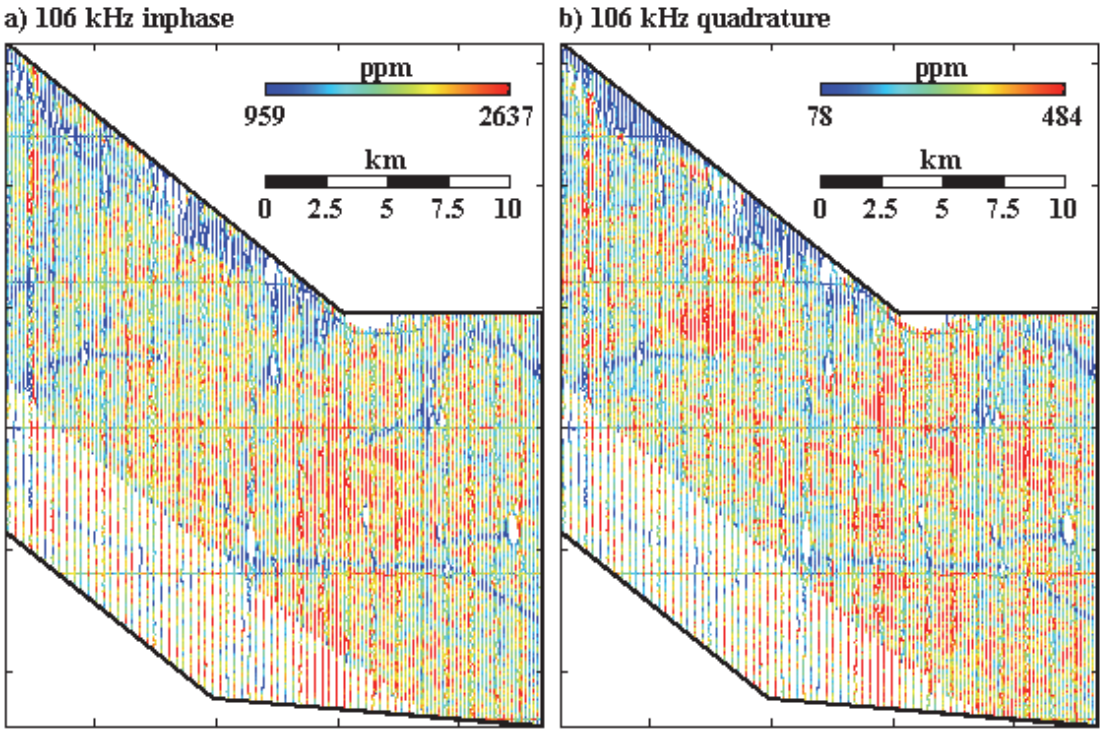

Figure 4.26 Conventionally processed 106 kHz data showing the north–south striping that is symptomatic of systematic errors; (a) inphase, and (b) quadrature.

Note however that the same anomalies are largely non-existent on the grids of apparent conductivity data (Figure 4.27a) that were generated from the conventionally processed data by the survey contractor. This is because in the conventional data processing one of the main objectives is to remove, from the apparent conductivity data, any flight line orientated anomalies that appear non-geological in origin and are expected to be caused by calibration error. Such a procedure, carried out independently for each coilset frequency, does not guarantee removal of calibration error. Instead some of the error may be 'absorbed' by the apparent depth data, which is not considered to be the primary product. This effect is illustrated by Figure 4.27b which shows north–south artefacts in

the apparent depth data that were generated in the same procedure as the artefact-free apparent conductivity data.

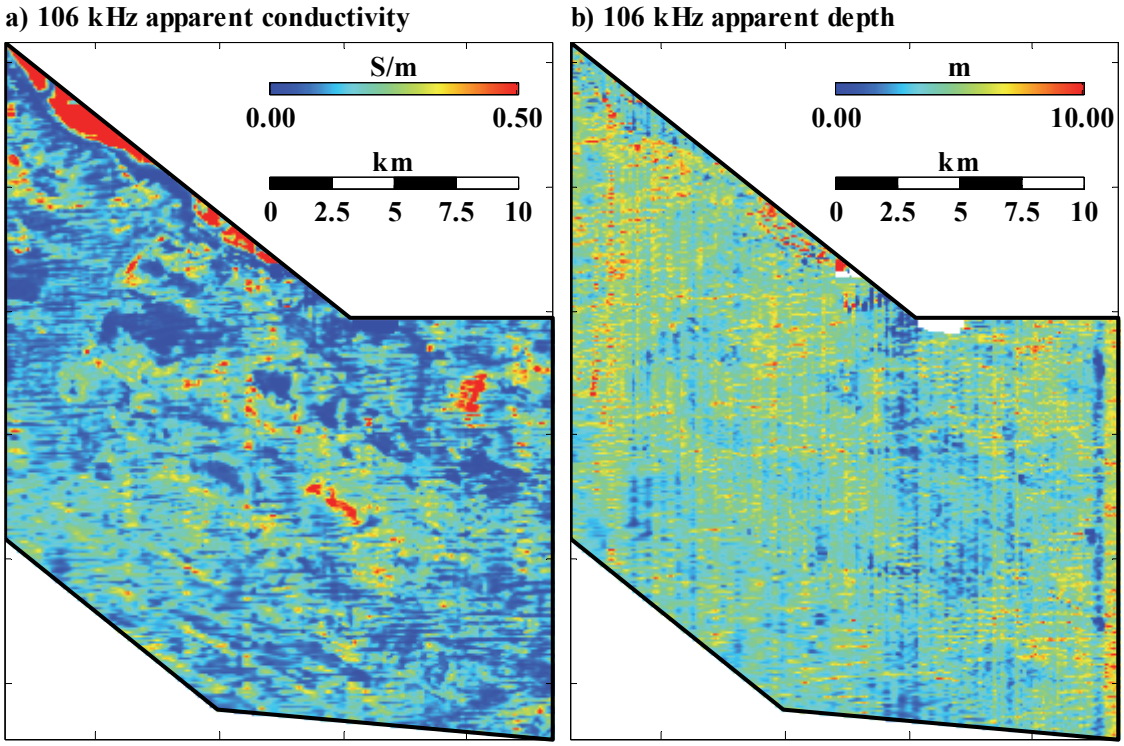

Figure 4.27 Images of the (a) apparent conductivity and, (b) apparent depth grids that were generated as part of the conventional data processing.

Further effects of the systematic errors is illustrated in Figure 4.28 which shows the 106 kHz coilset noise-normalized residuals (i.e. $\frac{predicted-observed}{noise}$) for the inversion of the conventionally processed data. It can be seen that the residuals are not randomly distributed about a mean value of zero. Instead, for the inphase data (Figure 4.28a) the patterns in the residuals are north–south orientated and clearly correspond to artefacts in the processed data (Figure 4.26) and the apparent depth data (Figure 4.27b). As expected, the systematic residual patterns migrate through to artefact in the layer property images already demonstrated in Figure 4.22 and Figure 4.23.

Additionally the residuals for the quadrature data (Figure 4.28b) are consistently positive and much greater than the expected average value of zero. This indicates that there was a systematic inconsistency, between the quadrature channel and other channels that did not allow the data to be fitted adequately.

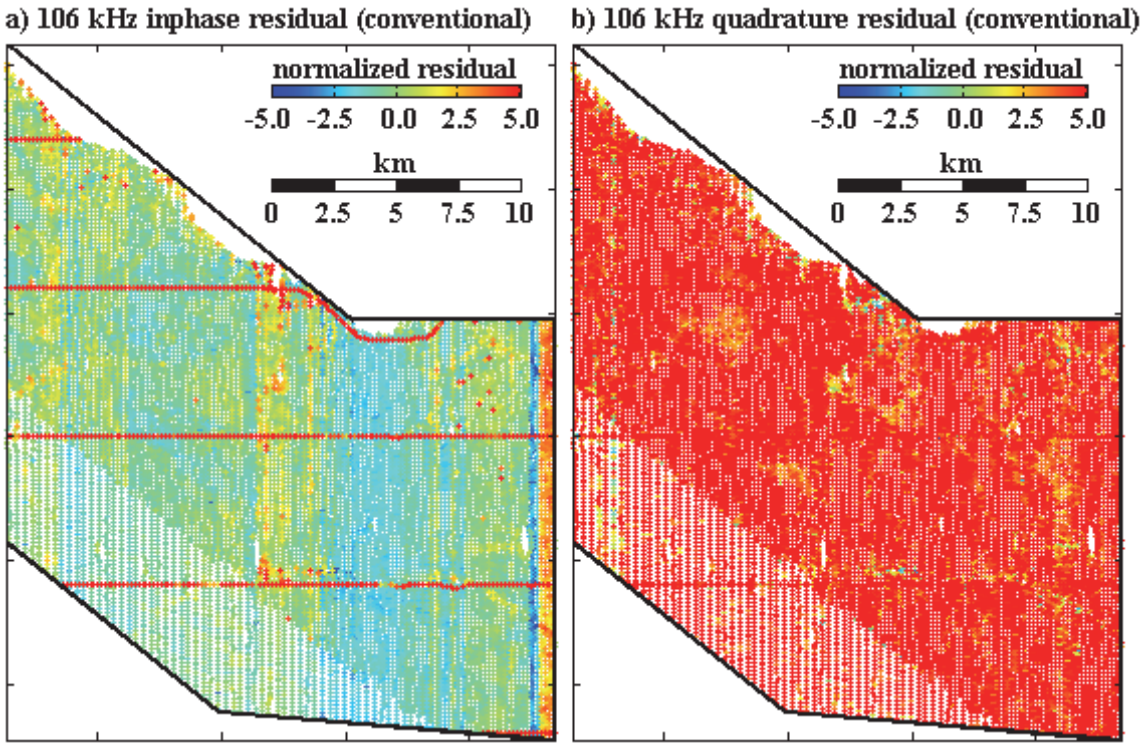

Figure 4.28 *Noise normalized residuals (data misfit) for the 106 kHz coilset from the sample-by-sample inversion of conventionally processed data; (a) inphase, and (b) quadrature.*

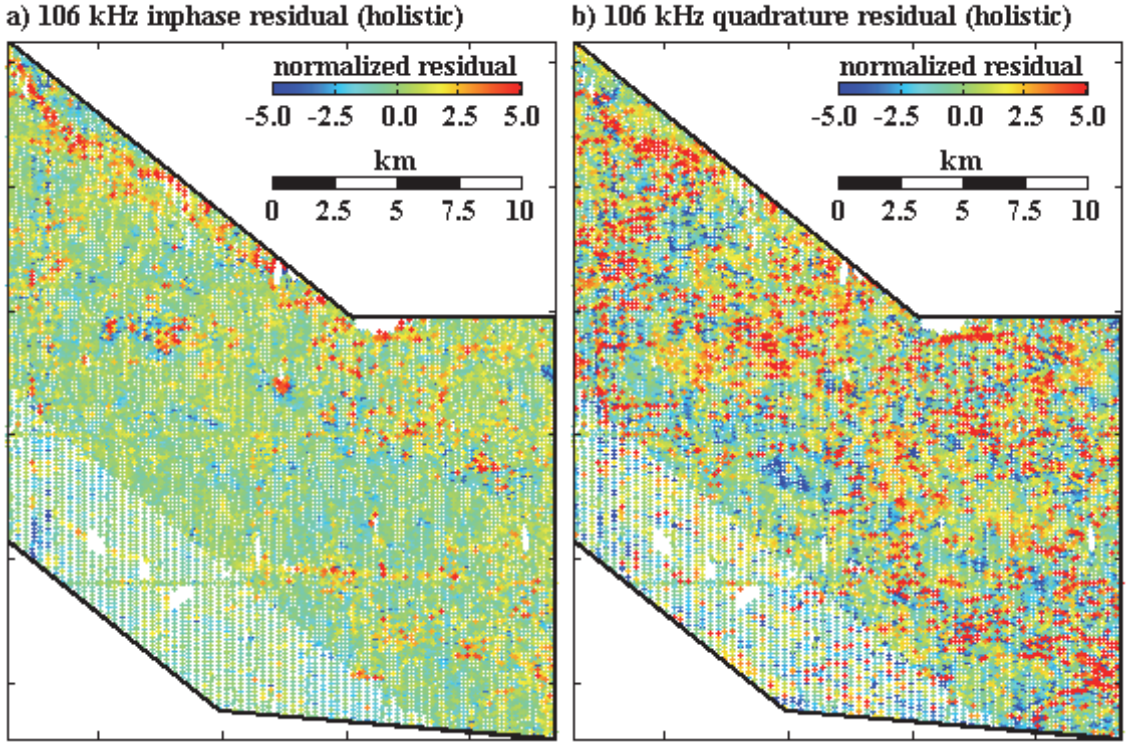

Figure 4.29 *Noise normalized residuals (data misfit) for the 106 kHz coilset from the few-layer holistic inversion; (a) inphase, and (b) quadrature.*

As a comparison to Figure 4.28, the equivalent 106 kHz residuals for the holistic inversion are shown in Figure 4.29. It can be seen that, in contrast to the conventional inversion residuals, that the holistic inversion residuals are generally randomly distributed about zero. This is manifested as spatially uncorrelated scatter of the high amplitude residuals in Figure 4.29 which contrasts to the spatially correlated (streaked) patterns in Figure 4.28. Furthermore there is a distinct lack of north–south striping patterns in the holistic inversion residuals. For the holistic inversion the residuals also have smaller amplitudes, indicating that the data were better fitted.

Considering together the observations made so far in this section, the north–south anomalies in the conventional fully processed data; the artefacts in the apparent conductivity and inversion results for the conventionally processed data; the distribution of inversion residuals for the conventionally processed data; and the general lack of such artefacts in the holistic inversion results; it is reasonable to conclude that the holistic inversion has generated substantially improved results over conventional approaches. A further appealing result is that the holistic inversion was carried out on practically raw data, rather than fully processed data, thereby saving much repetitive processing time by expert individuals.

Geological features of the holistic inversion conductivity model

We now turn our attention to some of the hydrogeological features shown in the holistic inversion results. The Layer 1 image Figure 4.24a and b reflect the dry and resistive sand sheets and the east–west orientated sand dunes that are present in the Riverland area. The strong northwest–southeast trends that are abundant in the Layer 2, 3 and 4 images, are consistent with the conceptual hydrogeological model (*cf.* Figure 4.18) in which the Blanchetown Clay was laid down over (or infills) the beach strandline dominated palaeo-topography that was left when the sea retreated from the Murray Basin in the early Pliocene. Brodie et al. (2004b) approximated the palaeo-topography

surface as the elevation of the bottom of the Blanchetown Clay layer (Layer 2). This surface was generated from the holistic inversion results by subtracting the Layer 1 and Layer 2 thicknesses from the modern day surface elevation. It is imaged in Figure 4.30 which shows a plausible and coherent mapping of the slightly curvilinear northwest–southeast trending strandlines. It is rather encouraging that such geologically revealing results are able to be produced from the holistic inversion of near raw survey data.

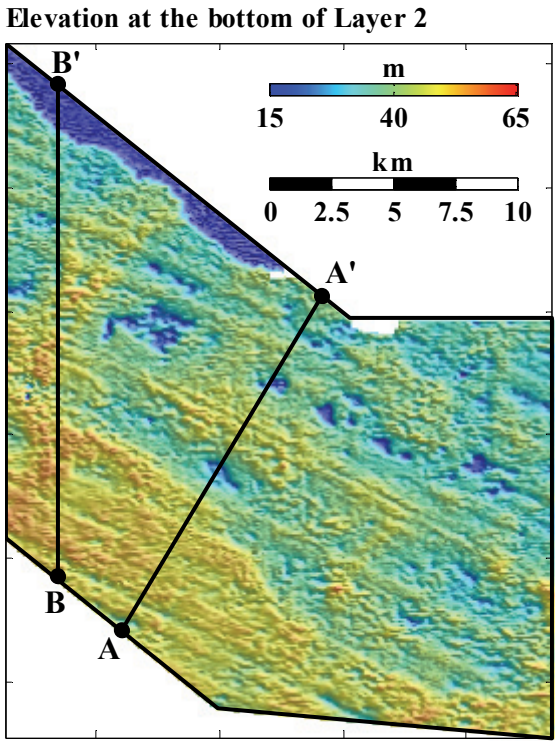

Figure 4.30 *Holistic inversion prediction of the elevation of the bottom of the Blanchetown Clay (bottom of Layer 2) which is a proxy for the Pliocene strandline dominated palaeo-topography. The image is artificially illuminated from the northeast and profiles AA' and BB' show the locations of the conductivity sections shown in Figure 4.31.*

The strong narrow banding that can be seen in the Layer 4 conductivity image (Figure 4.25c) were also detected by Brodie et al. (2004b) and were interpreted to be caused by systematic variation of the porosity in the Loxton-Parilla Sands between subsequent strandline deposits. In Layer 4, which is below the watertable, porosity variations created the strong conductivity banding because of the changes in the volumetric fraction of saline waters.

The purpose of the survey was to map the distribution and thickness of the Blanchetown Clay, which is summarised in the Layer 2 thickness image (Figure 4.24d). The image shows a geologically plausible distribution of Blanchetown Clay and how it is influenced by the interpreted strandlines. It also shows how this layer is locally dissected, or absent, due to the reworking of this material during the Quaternary. Also, we have already seen how this image is more interpretable because it is free of the numerous artefacts that are in the Layer 2 thickness image (Figure 4.22d) produced from the data that were conventionally processed and inverted.

In Figure 4.31a and b, holistic inversion conductivity depth cross-sections are presented for the profiles marked AA' and BB' on Figure 4.30. In contrast, Figure 4.31c is a section along profile BB' resulting from the conventional processing and inversion of the data. None of the three section has been smoothed. For the conventional inversion a section was not produced for the profile AA' because it runs at an angle to the flight lines, and hence inversion results would have to be projected from the nearest flight lines onto the profile in order to generate the section. This further emphasizes the utility of the spline parameterization of the holistic inversion in that information required for producing sections can be analytically and uniquely calculated at the necessary positions.

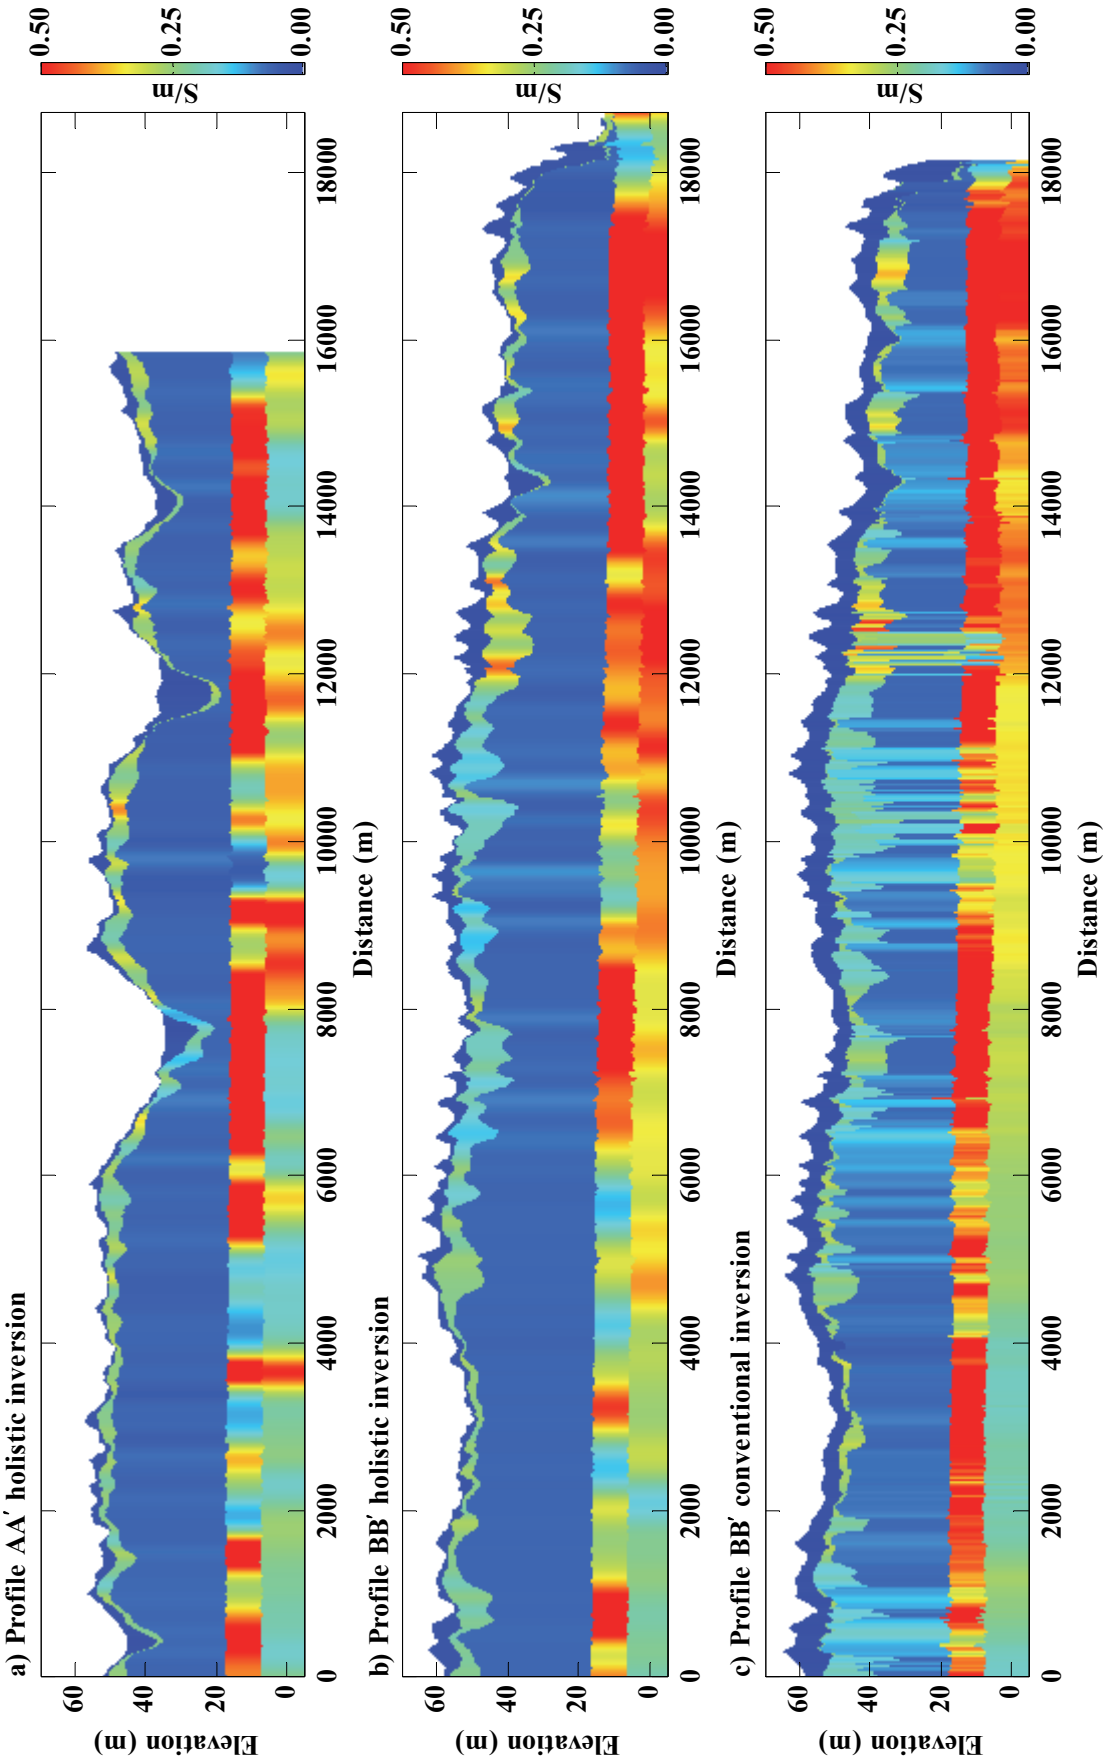

Figure 4.31 *Conductivity sections along the Profiles AA' and BB' shown on Figure 4.30. Panels a and b show results from the few-layer holistic inversion, and panel c shows results from the conventionally processed and inverted data.*

The holistic inversion sections (panels a and b) are more geologically realistic as they do not exhibit the many lateral discontinuities present, as vertical stripes, in the conventional sample-by-sample inversion section (panel c). Because of such discontinuities, sample-by-sample inversion results, like those shown on Figure 4.31c, often require some degree of smoothing before being compiled into more presentable sections for interpretation. The discontinuities are probably caused partly by fundamental ambiguity, partly by bird motion noise as shown in Fitterman and Yin (2004), and short period noise sources. The lateral continuity imposed by the spline parameterization of the conductivity model enables the holistic inversion to "fit through" the short period noise, which is not possible with a sample-by-sample inversion. This is an advantage shared by the laterally constrained inversions of Auken et al. (2005), who also claim that the lateral constraint improves the resolution of potential equivalences.

The relatively narrow bands visible on the sections in the Layer 4 conductivity may not look geologically realistic at first appearance. However it is clear from Figure 4.25c that the banding is spatially coherent, and as discussed earlier, is explained by porosity variations between alternate strandlines.

Figure 4.32a shows the fit between the observed geoelectric data (averages over each 2 m interval of the 16 downhole conductivity logs in the area) and the predicted geoelectric data computed from the few-layer holistic conductivity model. The largest discrepancies are for the data in the lower end of the predicted conductivity range, which correspond to intervals in Layers 1 and 3 where the conductivity was tightly constrained.

Figure 4.32b shows the observed interface-depth data that were compiled from the gridded dataset of groundwater elevations (*cf.* Figure 4.8b). They were successfully fitted by the predicted interface-depth data computed from the sum of the top three

layers of the holistic inversion conductivity model. The constraint that these interface-depth data provide is illustrated nicely in the conductivity cross sections in Figure 4.31, where it can be seen that the top of the watertable (interface between Layers 3 and 4) gently dips northward (to the right) as it approaches the River Murray.

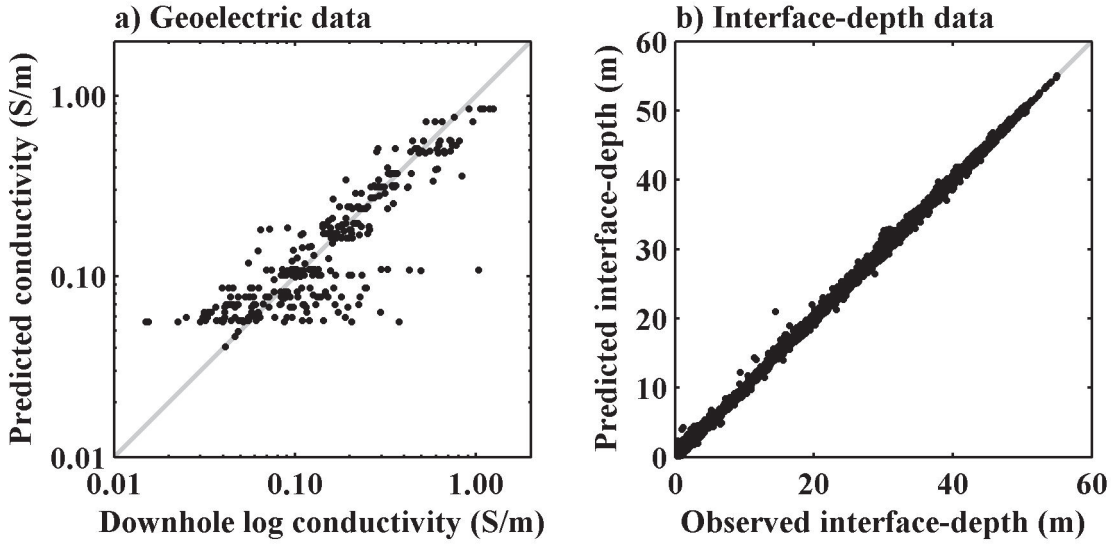

Figure 4.32 Geoelectric and interface-depth data; (a) observed downhole conductivity log data averaged over 2 m intervals versus predicted conductivity over the same intervals, and (b) observed water table depth versus predicted depth to the top of Layer 4.

A comparison between the 16 individual downhole conductivity logs and the holistic inversion model at their respective locations are displayed in Figure 4.33. It can be seen that the conductivity model recovers the main elements of each conductivity log well. But this is to be expected because the geoelectric data, which were compiled from the downhole logs, provide strong constraint in the vicinity of the boreholes. However, since the profile shape of the downhole conductivity logs is complex (i.e. the variations are gradational and sinuous) there is no five layer conductivity model that can precisely represent the conductivity log data. This factor also accounts for some of the misfit between the observed and predicted geoelectric data shown in Figure 4.32a. It will be seen in the following sections how the sinuous nature of conductivity logs can better represented by a multi-layer conductivity model.

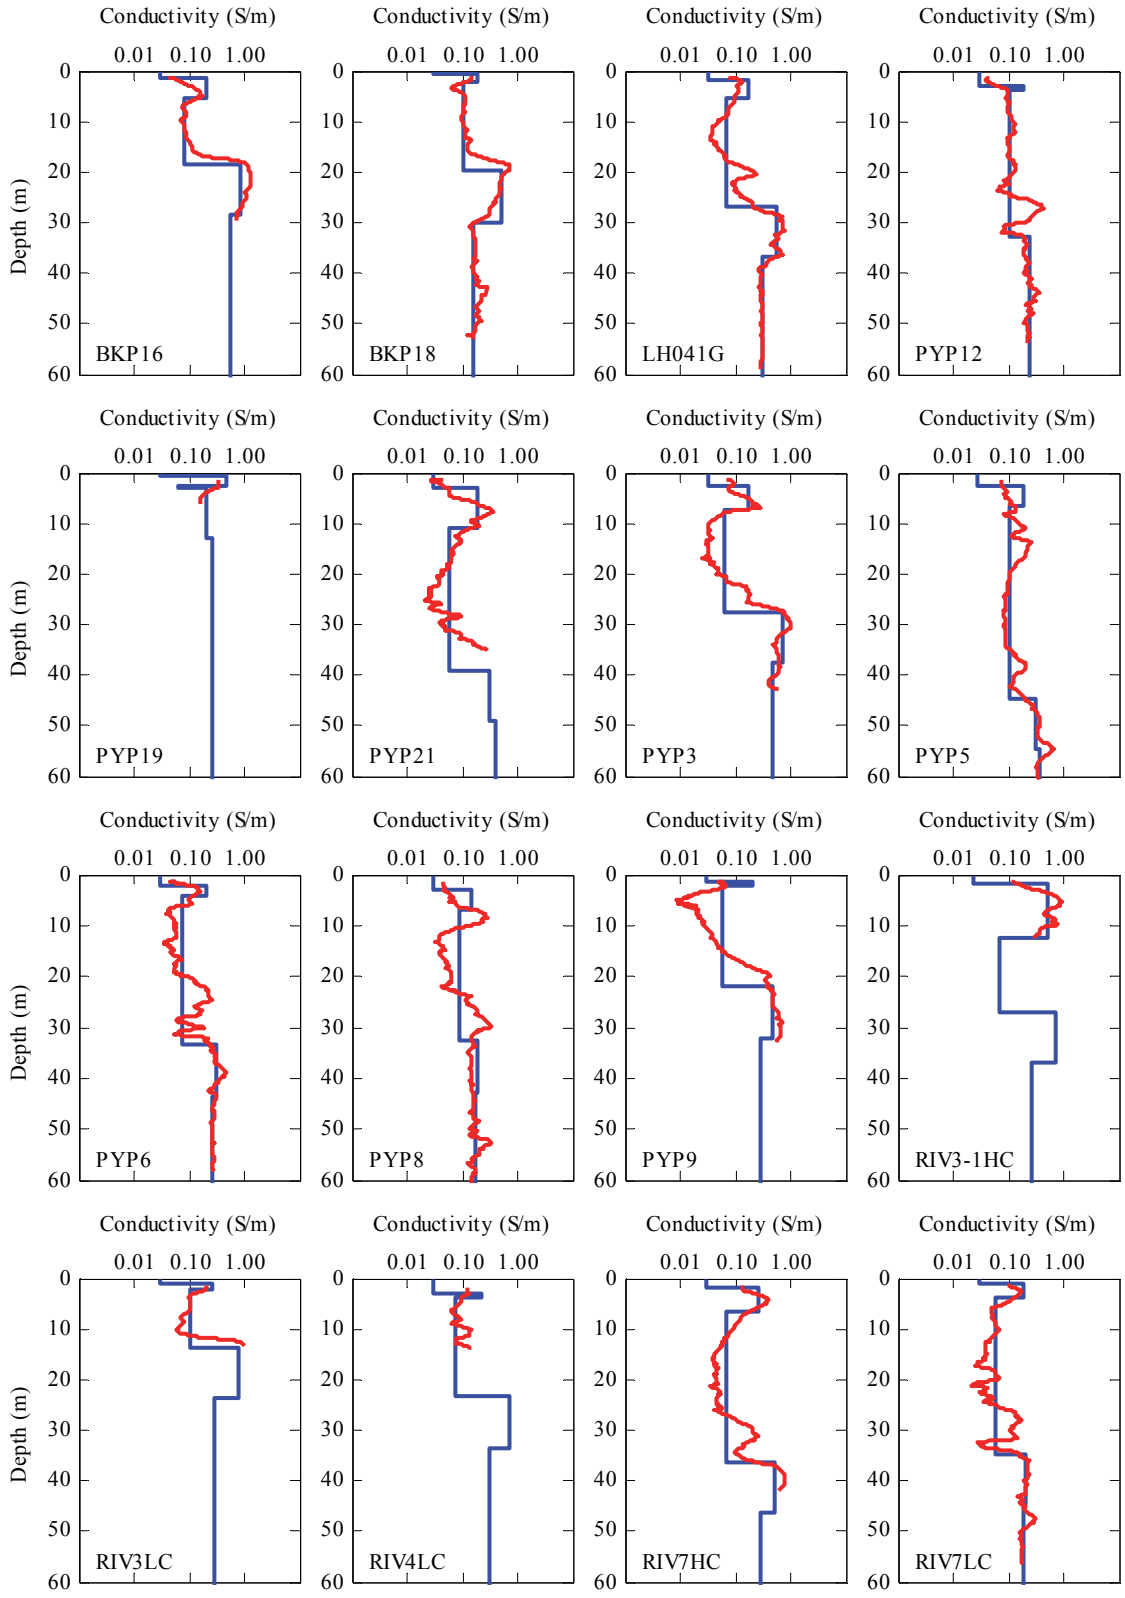

Figure 4.33 Comparison between the 16 downhole conductivity logs (red) in the area and the conductivity profile at each downhole log (blue) estimated by the few-layer holistic inversion.

4.5 Multi-layer inversion

4.5.1 Introduction

One question arising from the few-layer inversion example was whether or not its success could be attributed to the use of good prior information and well a developed knowledge of the survey area's geology. This prompted an investigation into how the method would perform if survey area specific prior information was not available. Prior information is often not available in the so called 'greenfields' areas where there has been little or no previous geological or geophysical mapping. In this multi-layer holistic inversion example the greenfields scenario is simulated, and hence no survey area specific conceptual conductivity model, geoelectric data or interface data are used.

A shortage of prior information limits the understanding of the geological layering, and hence how that layering should be translated in to a conceptual conductivity model. There is no way to know how to choose a generic few-layered model that can adequately represent the large, but unknown, range of vertical conductivity profile shapes that are possible. In such cases, a multi-layer parameterization is a suitable choice because a wide range of vertical profile shapes can be approximately represented by a large number of thin layers. In the current example, where a 20 layer fixed-thickness model is solved for, this reasoning and parameterization style is adopted.

A further consequence of not having prior knowledge of the geoelectric properties of the (hydro)geological units in greenfields areas, is that it is difficult to build an accurate reference model with associated uncertainties. The best that may be achievable is to use average reference values from a similar geological province where data may be available. In this multi-layer inversion example a homogeneous halfspace conductivity reference model was used to simulate the greenfields situation.

In the move from a few-layer parameterization to a multi-layer parameterization there are many more unknown parameters to be solved for. This increases the computer memory and computational requirements, or alternatively reduces the size of datasets that can be simultaneously inverted. However, part of the philosophy of the holistic inversion is to include as much data as possible so that spatial coherency of the conductivity and temporal coherency of systematic errors can be exploited. So ideally it is preferable to invert a complete dataset at once. Therefore to meet both goals, it was necessary to develop a second generation of the code that could be executed on a parallel computer. This parallelized code was used for the multi-layer inversion example, allowing the complete Riverland dataset (11,476 line km) to be simultaneously inverted to estimate a 20 layer conductivity model.

The example demonstrates that the method can be used when only minimal prior information is available. Without using explicit prior information, the recovered conductivity model is geologically plausible and consistent with downhole logs and groundwater depth data. Zero-levels predicted from the calibration model correlate well with zero-level measurements recorded at high altitude. Gain values are also similar to gain values that would be predicted from downhole log data.

4.5.2 Inversion

Data

In this multi-layer inversion example only the airborne data were used. Geoelectric and interface-depth data were not used because the aim is to simulate the greenfields mapping situation. The form of the airborne data inverted was unchanged from the few-layer inversion example (*cf.* Section 4.4.2.). That is, they were raw except for the reversal of zero-level shifts that had been introduced during the high altitude calibrations (*cf.* Figure 4.7c), and the 9-point median and Hanning random noise reduction filters discussed in Section 4.2.5.

Since a parallelized second generation of the code was used, a much greater volume of data could be inverted compared to the few-layer example. In this case the complete dataset inside the survey area boundary shown in (Figure 4.1), totalling 11,476 line kilometres, was able to be inverted. The amount of decimation (sub-sampling) of the airborne data, which was required to limit the memory resources needed, could be also reduced from every 17*th* to every 5*th* sample. Since the 5-sample or ~16.5 m along line interval is much smaller than the minimum system footprint (~60 m), it was expected that the sub-sampling would not cause significant information loss.

A small percentage (0.27%) of data was excised in zones where the filters did not adequately remove powerline and sferic noise. A total of 673,196 airborne samples were inverted, which translates to 8,078,352 data since each sample has 12 channels of data. There were 89 flights, and 597 lines in the airborne dataset acquired during 37 days of flying. The method of Green and Lane (2003) was employed to estimate noise in the airborne data using the parameters provided in Section 4.2.6.

Parameterization

So that a wide range of (unknown) vertical conductivity profiles could be represented, a 20 layer conductivity model was employed. The increased number of layers in a multi-layer inversion means that the ambiguity in the estimated vertical profiles will be increased and the inversion will become less well conditioned. One way to alleviate this situation is to fix the layer thicknesses and to regularise the vertical profile shape. This, so called 'Occam's inversion' (Constable et al., 1987) formulation, was used in the current example to improve inversion stability. In this type of parameterization, particular layers in the model do not correspond to particular geological features. Instead the features must be extracted by some post-inversion manipulation or interpretation.

Table 4.7 Layer thicknesses and depths for the multi-layer inversion.

Layer#	Thickness (m)	Top depth (m)	Bottom depth (m)	Layer# (cont'd)	Thickness (m)	Top depth (m)	Bottom depth (m)
1	1.50	0.00	1.50	**11**	4.08	24.50	28.58
2	1.66	1.50	3.16	**12**	4.51	28.58	33.09
3	1.83	3.16	4.99	**13**	4.98	33.09	38.07
4	2.02	4.99	7.01	**14**	5.50	38.07	43.57
5	2.24	7.01	9.25	**15**	6.08	43.57	49.65
6	2.47	9.25	11.72	**16**	6.72	49.65	56.37
7	2.73	11.72	14.45	**17**	7.43	56.37	63.80
8	3.02	14.45	17.47	**18**	8.21	63.80	72.01
9	3.34	17.47	20.81	**19**	9.07	72.01	81.08
10	3.69	20.81	24.50	**20**	∞	81.08	∞

The layer thicknesses and top and bottom depths are shown in Table 4.7. The topmost layer was 1.5 m thick, chosen to be approximately the minimum thickness layer that can be expected to be resolved by the airborne system. The layers increased in thickness with depth in logarithmic fashion to reflect the decreasing vertical resolution with depth. Twenty layers were considered sufficient because it meant the depth to the top of the last layer was beyond the expected maximum penetration depth (~80 m) of the airborne system under moderately conductive Australian conditions.

The node spacing for all layer conductivity spline meshes was set to be 100 m in both the east–west and north–south directions. Unlike for the few-layer inversion, the spacings were not varied in order to simulate the unknown geological variability of the greenfields situation. For the same reasoning all the conductivity model reference values for all layers were set to the constant 0.5 S/m. The uncertainty standard deviation for all conductivity reference values were set to the extremely large value of 200 natural logarithm units so that a wide range of conductivity values would be permitted. The details of the parameterization of both the conductivity and calibration model is provided in Table 4.8.

Table 4.8 Summary of multi-layer inversion model parameters.

Local conductivity profile & calibration values	Dependence	Inversion parameters	Number of unknowns	Reference model	Reference model standard deviation
σ_1	location, layer	100×100 m mesh	169,901×20 = 3,398,020	log(0.5)†	200†
σ_2		100×100 m mesh		log(0.5)†	200†
…		…		…	…
σ_{20}		100×100 m mesh		log(0.5)†	200†
t_1	fixed	n/a	0	1.50 m	n/a
t_2		n/a		1.66 m	n/a
…		….		…	…
t_{19}		n/a		9.07 m	n/a
gain	survey, coilset	gain	6	1.0	0.1
phase	day, coilset	phase	222	0.0º	3.0º
bias	flight, fiducial, coilset, channel	bias value	6,540	0.0 ppm	survey stats*
height error	fixed	n/a	0	n/a	n/a

† units of log(S/m) ‡ units of log(m) * statistically calculated from survey data

The systematic calibrations errors were parameterized in essentially the same manner, and with the same reasoning, as for the few-layer inversion example. There were of course more parameters since the complete dataset was inverted in this case. The gain was parameterized to have a constant value for each coilset over the entire survey, thus 6 parameters were required. The phase calibration was parameterized in piecewise constant fashion, with one value for every day of flying and for each coilset so that there were 222 (6×37) phase parameters in total.

As in the few-layer inversion example the bias was parameterised as piecewise linear variations over each flight. All piecewise intervals had the same length and were as close as possible to 20 minutes long. Separate functions were used for each inphase and quadrature channel. The first interval for a flight started at the first inverted sample in the flight and the final interval ended at the last inverted sample in the flight. In contrast to the few-layer inversion example, the method of 'parameterization by bias

value' (Equation 3-7) was trialled. Therefore the unknowns to be solved for were the bias values at set fiducials positioned at the start of each flight, at the times linking each piecewise interval, and at the end of each flight. There were 6,540 (545 set fiducials×12 channels). As explained in relation to the few-layer inversion, there was no motivation to include a height calibration parameter for the Riverland dataset.

In choosing the reference values for the calibration model, it was assumed that the system was perfectly calibrated. Therefore a value of unity was used for the gain parameters and zero was used for phase and bias calibration model parameters. The standard deviation uncertainty was set to 0.1 for the gain reference values and set to 3.0° for the phase reference values. The standard deviation uncertainty for the bias reference values were computed statistically from the variations observed in high altitude zero-level data measured during the survey.

Solution

In total 8,078,352 data were inverted, all of which were airborne data. There were 3,404,788 unknowns to be estimated, consisting of 3,398,020 were spline node coefficients and 6,768 were calibration model parameters. Reference model constraints were applied ($\lambda_r = 200000$) on all parameters using the reference model and standard deviation uncertainties summarised above in Table 4.8.

When a fixed-thickness formulation is used in conjunction with roughness constraints imposed relative to layer depth, there is a tendency for the conductivity variations to mimic or follow the topography variations. This may be desirable in certain geological settings where the conductivity is expected to be strongly related to depth, for instance, where it is related to the depth of weathering. However, in other settings it is expected to be more strongly dependent on elevation, for example, in depositional regimes or where groundwater influences conductivity. Since the latter was the case in the Riverland area, horizontal roughness (second derivative) constraints were imposed with

respect to elevation rather than depth. This was achieved by using horizontal roughening matrices $\mathbf{L}_x$ and $\mathbf{L}_y$ that were in the form shown in Equation 3-55. The horizontal roughness constraints were applied in both the east–west ($\lambda_x = 25$) and north–south ($\lambda_y = 25$) directions.

Because of the large number of layers, vertical roughness (second derivative) constraints were applied ($\lambda_z = 5$) to regularize the vertical profile shape. As explained earlier these constraints serve to prevent oscillations in the vertical profile shape that are not required by the data and improve the inversion stability. Bias roughness constraints were also applied ($\lambda_b = 1$) to encourage smooth variation of the biases through each flight. The degree of smoothing was chosen by initially setting the regularisation weights λ_x, λ_y, λ_z and λ_b to large values and then reducing them manually in trial runs until the data were able to be fitted.

The holistic inversion was run on 64 processors of the TerraWulf (http://rses.anu.edu.au/terrawulf). The TerraWulf is a Linux cluster computational facility based at the Research School of Earth Sciences, Australian National University. At the time of running this example, each processor had a clock speed of 2.4 GHz and had access to 1 GB of random access computer memory. The total elapsed execution time was 8.22 hours and approximately 51 GB of memory were required. The parallelized code was implemented via the Message Passing Interface (MPI) programming paradigm.

Since the 1D forward model and derivative calculations (Equation 3-18) for each airborne sample are independent of all other samples, these can be performed in parallel with linear speedup (i.e. doubling the number of processors halves the elapsed time). However the solution of the linear system (Equation 3-71) and other matrix operations require communications between the processors, and hence the speedup is not linear. The inter-processor communication is required because, to conserve memory, the

matrices are stored in row-distributed fashion (i.e. where each of the P processors stores $1/P$ rows of the matrices each).

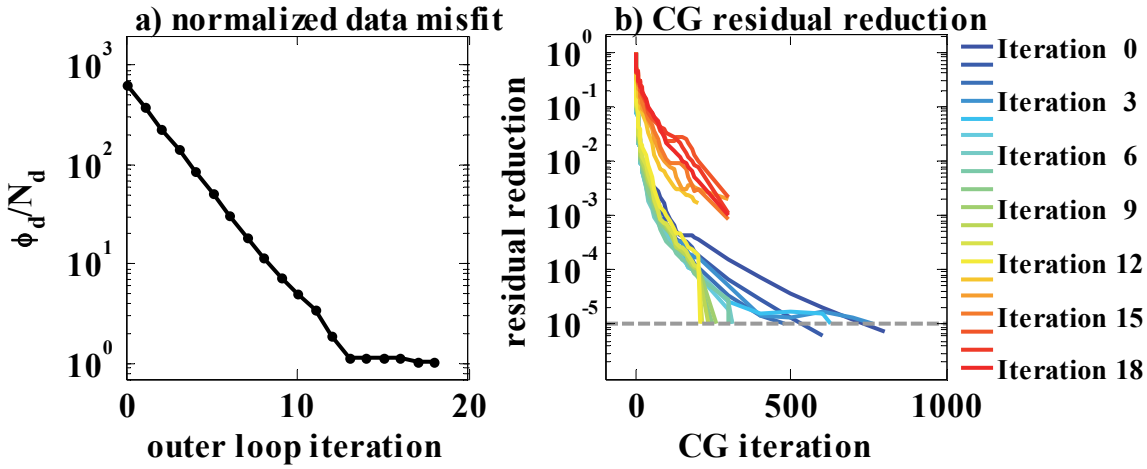

Figure 4.34 *Summary of the multi-layered inversion's convergence, showing, (a) the reduction of the normalized data misfit (Equation 3-77), and (b) the relative residual reduction of the conjugate gradient solution of the linearized system in each outer loop iteration (Equation 3-79).*

The convergence of the inversion's outer and inner loops (*cf.* Figure 3.8) is summarised in Figure 4.34. The left hand panel shows that the normalized data misfit (Equation 3-77) was reduced to unity, the level indicating that the data were acceptably fitted, within 18 outer loop iterations. For the first 13 outer loop iterations, the misfit was reduced to approximately 0.6 of its prior value in each iteration as designed by the minimization algorithm via step length damping (*cf.* Equation 3-75). However after iteration 13 the convergence was slow to achieve the remaining comparatively small reduction in data misfit.

Figure 4.34b shows the convergence of the conjugate gradient (CG) algorithm that solves the linearized system inside each outer loop iteration. Each coloured curve represents the trajectory of the reduction in the CG residual relative to the residual at the first CG iteration (i.e. $||\, Ax - b \,|| \,/\, ||\, Ax^0 - b \,||$). A different coloured curve is plotted for each outer loop iteration. In this parallel code the CG algorithm was set to return a solution when the residual was reduced to less than $rtol = 10^{-5}$ of its original value (*cf.* Equation 3-79). This was the case for the first 13 outer loop iterations as shown by the

cool coloured curves on Figure 4.34b. However, after the 13th iteration the CG algorithm returned when the relative residual was as high as 10^{-3} because a different termination condition was met. This condition occurred when the absolute residual was reduced to $|| Ax - b ||< 10^{-16}$, that is it approaching the machine precision. This indicates that the initial solution guess $x^0 = \mathbf{m}^n$ supplied to the CG algorithm was already close to the final solution. It also explains why the overall data misfit only converged slowly after iteration 13 as discussed above and shown on Figure 4.34a. It perhaps suggests that the inversion could have been safely terminated at iteration 13.

4.5.3 Results

Calibration parameters

In Figure 4.35a the reciprocal of the gains estimated from the holistic inversion are plotted against the scaling factors derived from downhole conductivity logs by Brodie et al. (2004a). Again, the reciprocal of the gains are plotted because they have the inverse effect to the scaling factors. The comparison shows excellent agreement between the two sets of independently derived values. This provides confidence in the credibility of the gain calibration values estimated from the holistic inversion.

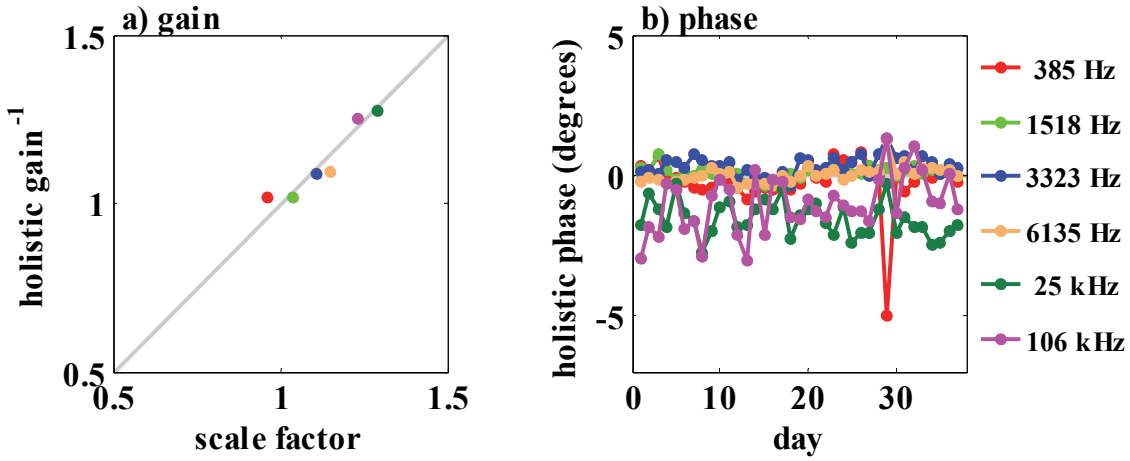

Figure 4.35 Calibration parameters showing (a) scale factors derived by Brodie et al. (2004a) versus the reciprocal of the gain estimated from the multi-layer holistic inversion, and (b) the phase estimates for each day of flying.

Inspection of Table 4.9 shows that the gain parameters derived for the multi- and few-layer inversions are not the same. If for the moment we accept that the scale factors derived from downhole conductivity logs by Brodie et al. (2004a) are correct, which is not necessarily the case, Table 4.9 suggests that the gains derived from the few-layer inversion are slightly better than for the multi-layer inversion. This might be expected since the downhole logs were used in the few-layer inversion. Another difference is that the multi-layer inversion was based on a super-set of the data inverted in the few-layer inversion.

Overall, the gains estimated from the multi-layer holistic inversion are within a few percent of those derived from both other methods. Since no downhole conductivity logs were present in this multi-layer inversion, it demonstrates that it is feasible to estimate gain factors using the holistic inversion, even for greenfields surveys where prior information is not generally available.

Table 4.9 Comparison of gain and scale factors.

Coilset	**1**	**2**	**3**	**4**	**5**	**6**
Frequency (Hz)	**385**	**1,518**	**3,323**	**6,135**	**25,380**	**106,140**
Brodie et al. (2004a) scale factor	0.96	1.04	1.11	1.15	1.29	1.23
Few-layer holistic gain^{-1}	0.89	1.01	1.11	1.15	1.30	1.29
Multi-layer holistic gain^{-1}	1.02	1.01	1.09	1.09	1.27	1.25

We turn now to the estimated phase errors parameters that were estimated by the multi-layer holistic inversion, which are shown in Figure 4.35b. There are no independent calibration studies to independently assess the accuracy of these error estimates. However, by comparing Figure 4.35b with Figure 4.20b it can be seen that both the multi-layer and the few-layer inversion estimated small negative phase errors. However the multi-layer inversion has estimated values closer to zero, but there is no way to assess which is more accurate.

Under normal circumstances the zero-level values recorded at high altitude would be used as the reference values for the bias calibration model. However in this work they were omitted to allow an assessment of whether or not the zero-level could be successfully estimated directly from the holistic inversion procedure. All reference values for bias parameters were set to zero. Figure 4.36 shows comparisons between the zero-level values that were measured during the high altitude excursions and bias values predicted from the multi-layer holistic inversion parameters. There is clearly a strong correlation between the two sets of independently derived results. Since the measured zero-level values were not used to create the reference or starting models, the good correlations are indicative of the success of the algorithm rather than being artefacts of the setup.

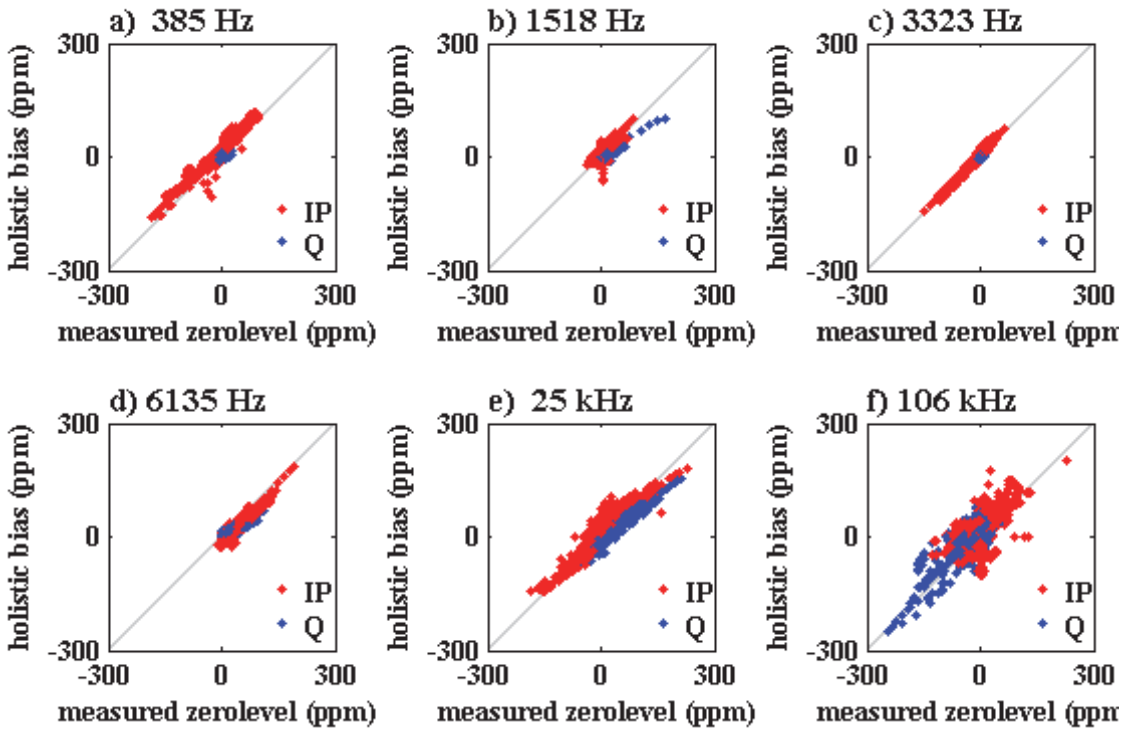

Figure 4.36 Comparison of the zero-level measured at high altitude during the survey and the bias estimated in the multi-layer holistic inversion

As was the case in the few-layer inversion example (*cf.* Figure 4.21) the greatest discrepancies are in the 106 kHz bias estimates. However, for the multi-layer inversion the discrepancies are substantially smaller. This may be a consequence of the greater

flexibility in the multi-layer conductivity model parameterization. It may also be because the complete dataset was used in this instance.

In Section 4.4.3 some possible reasons for discrepancies between the zero-level values measured at high altitude and bias values predicted from the holistic inversion were discussed. The assessment was made that the holistic inversion may have actually provided a better estimate of the true bias. This was concluded because the holistic conductivity model had less artefacts aligned with the flight lines than for the inversion of conventionally processed data. Similar conclusions can be drawn from the following section that discusses the multi-layer inversion conductivity model.

As for the few-layer case, the multi-layer inversion results demonstrate that, under the right conditions, it may be feasible to eliminate high altitude zero-level observations altogether. This would be of considerable practical benefit because of the likely substantial reduction in acquisition costs.

Conductivity model parameters

The images in Figure 4.37 through to Figure 4.41 show the conductivity of Layers 1, 5, 10, 15 and 20 estimated by the multi-layer holistic inversion. The most prominent artefact is marked on the Layer 20 conductivity image (Figure 4.41). Less prominent artefacts can be identified at the same location on the Layer 1 (Figure 4.37) and Layer 15 (Figure 4.40) conductivity images. These artefacts also happen to be related to the same flight, flown near the western edge of the few-layer inversion sub-area, that caused the artefacts which can be seen in the conventionally processed inversion (e.g. Figure 4.22c) and the few-layer holistic inversion (e.g. Figure 4.24b) conductivity images. This suggests that there may have been an especially serious calibration error associated with that particular flight that neither the holistic inversions nor the conventional processing has been able to fully account for.

Other than this particular instance the conductivity images are almost totally free of any artefacts elongate in the north–south flight line direction that would indicate the presence of systematic calibration errors. This shows that the calibration model has adequately accounted for systematic errors in the raw data that were not properly corrected by the conventional data processing, and which caused artefacts in the sample-by-sample inversion results shown in Figure 4.22 and Figure 4.23.

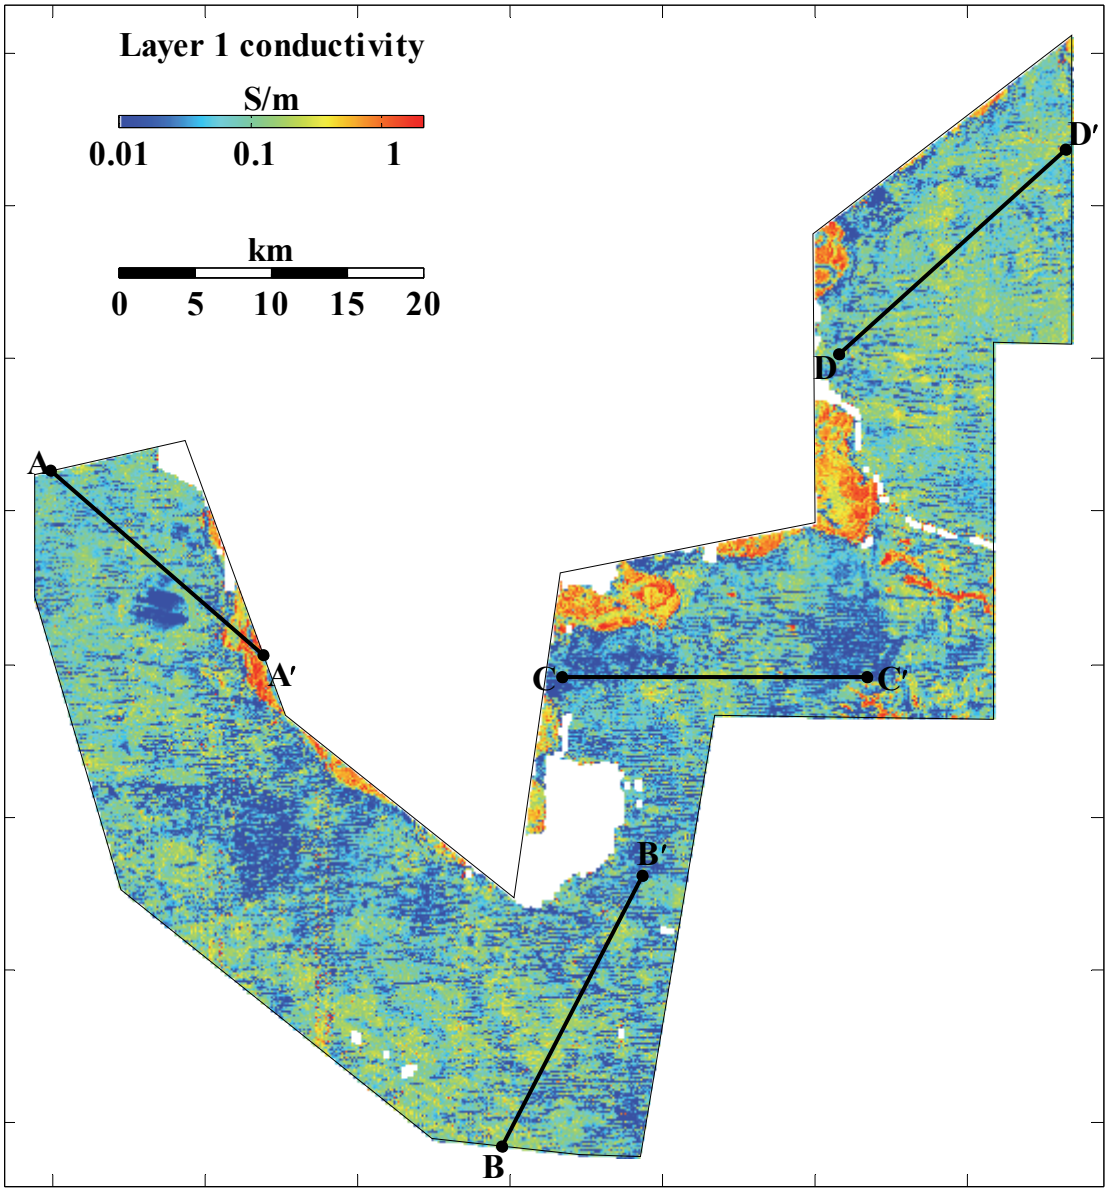

Figure 4.37 Image of the conductivity of Layer 1 estimated from the multi-layer holistic inversion. The profiles AA' through DD' show the locations of conductivity cross sections plotted in Figure 4.42 and Figure 4.43.

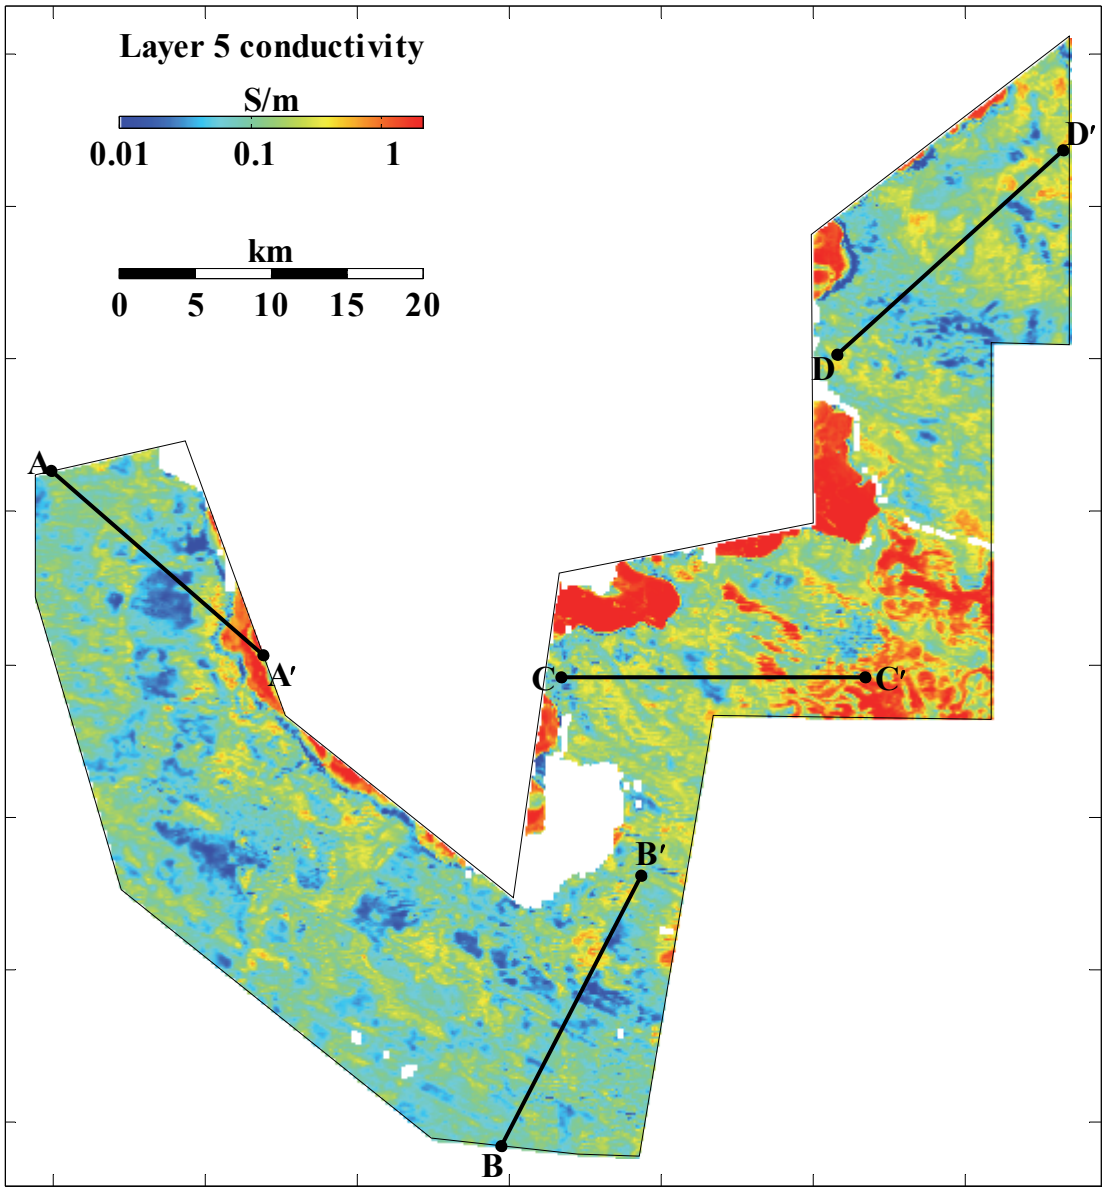

Figure 4.38 Image of the conductivity of Layer 5 estimated from the multi-layer holistic inversion. The profiles AA' through DD' show the locations of conductivity cross sections plotted in Figure 4.42 and Figure 4.43.

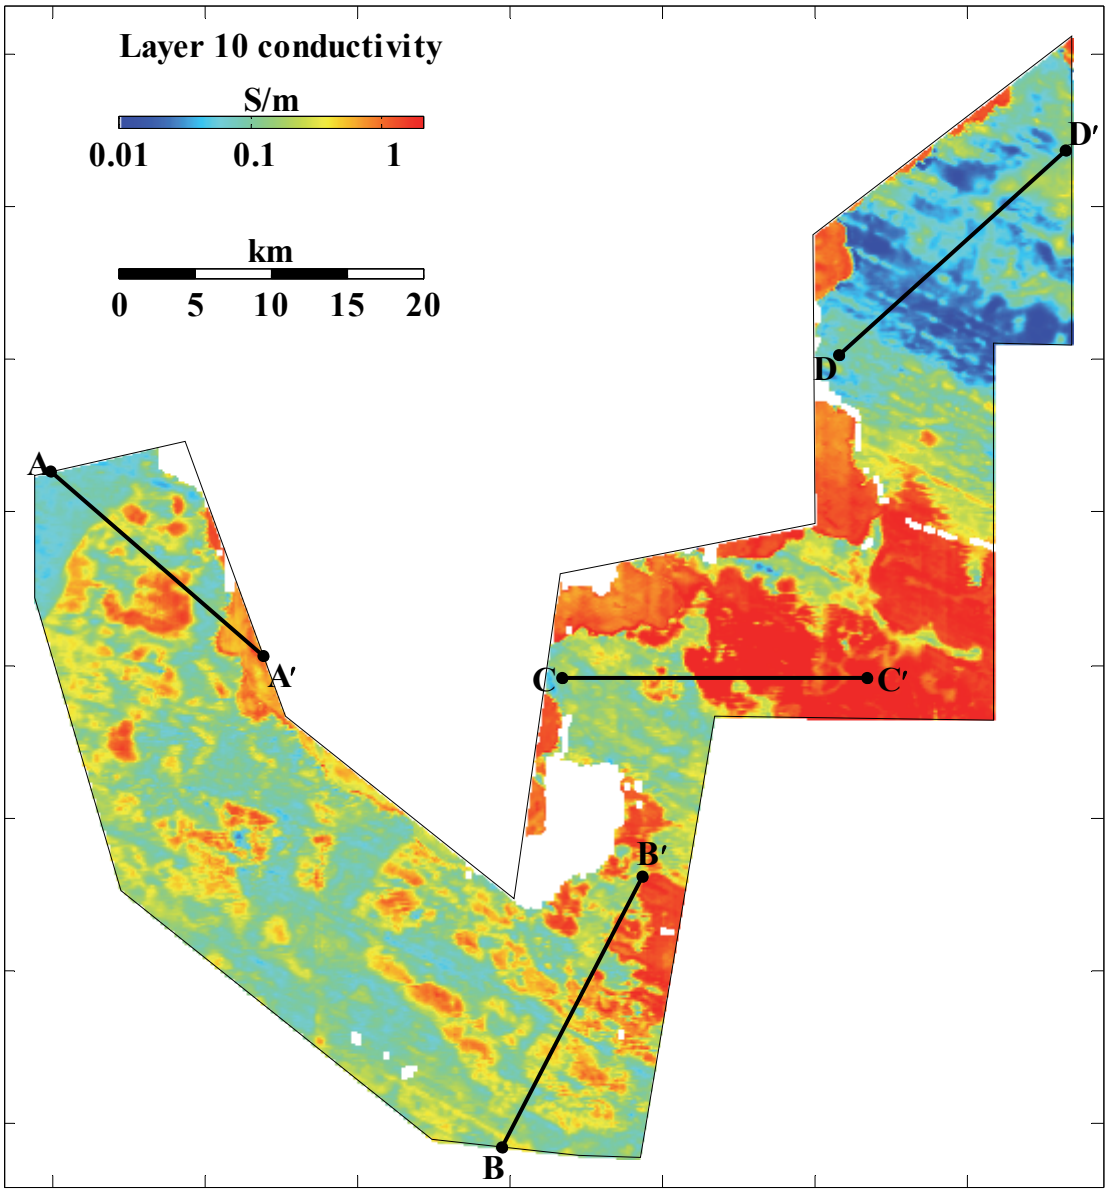

Figure 4.39 Image of the conductivity of Layer 10 estimated from the multi-layer holistic inversion. The profiles AA' through DD' show the locations of conductivity cross sections plotted in Figure 4.42 and Figure 4.43.

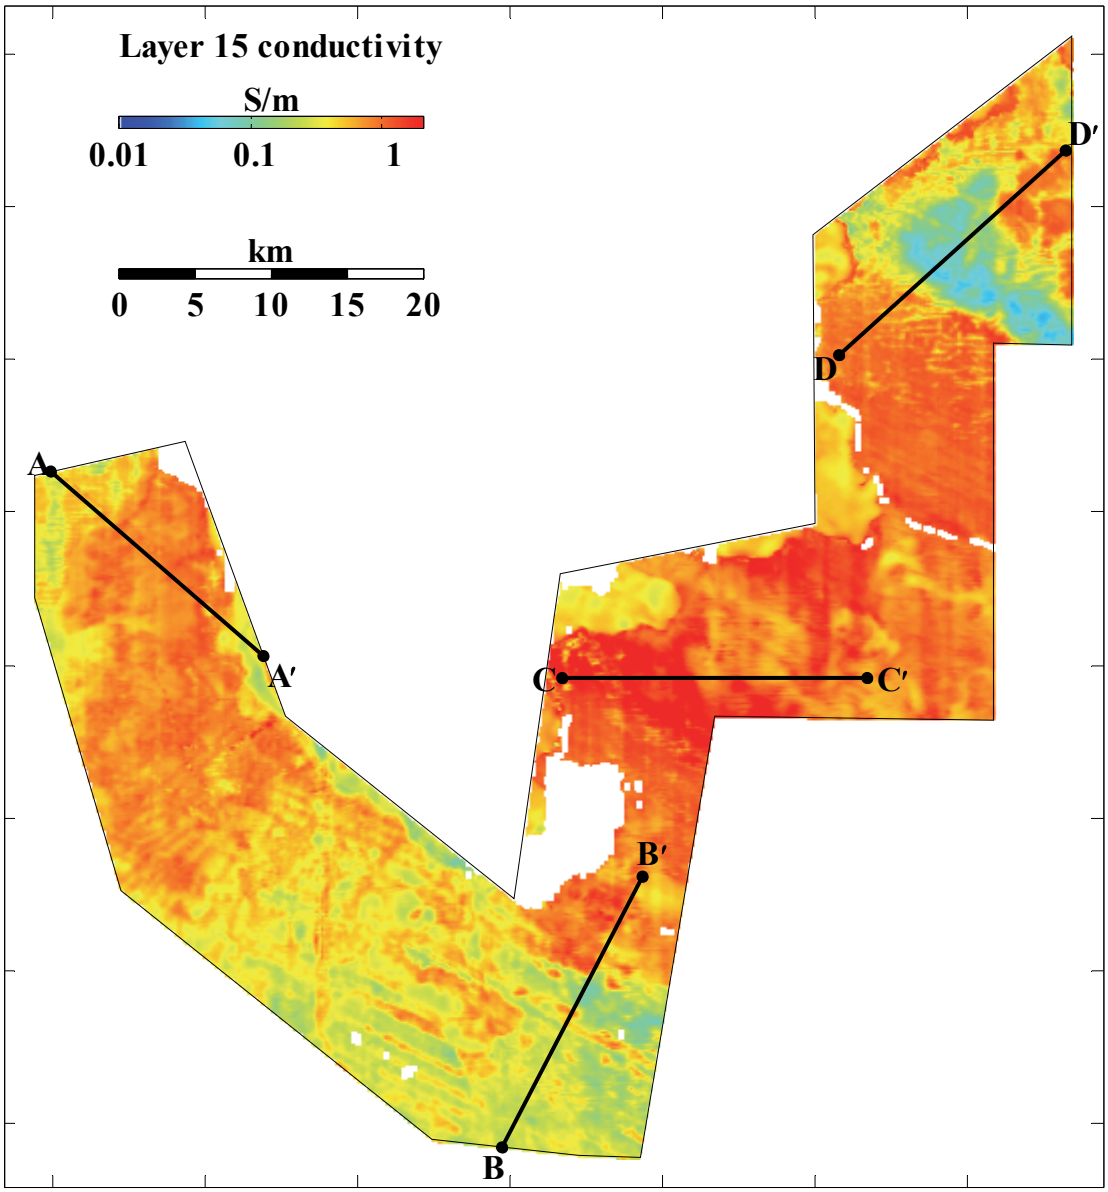

Figure 4.40 Image of the conductivity of Layer 15 estimated from the multi-layer holistic inversion. The profiles AA' through DD' show the locations of conductivity cross sections plotted in Figure 4.42 and Figure 4.43.

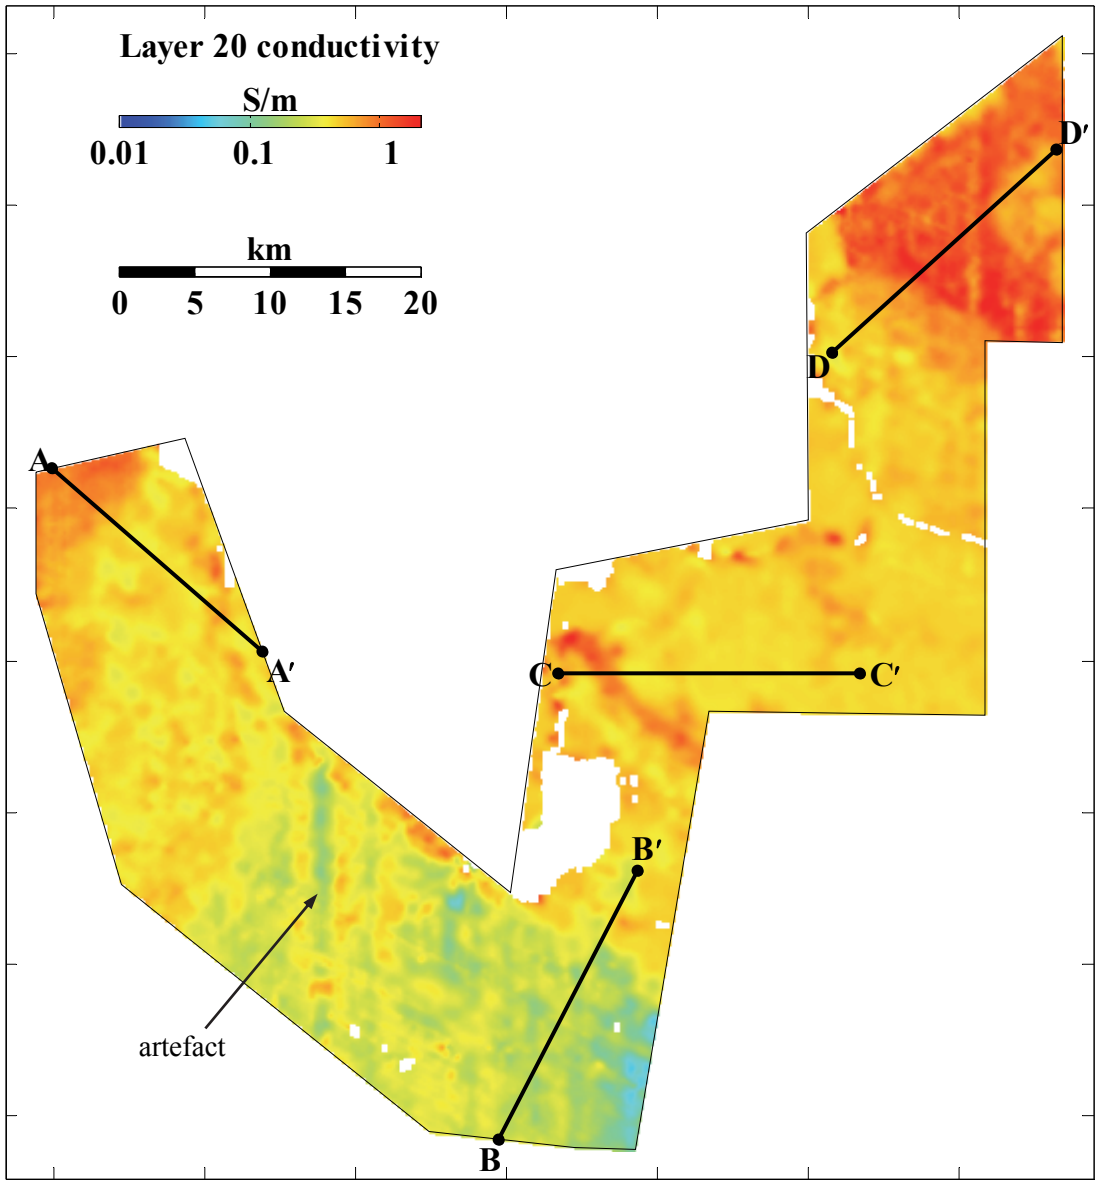

Figure 4.41 Image of the conductivity of Layer 20 estimated from the multi-layer holistic inversion. The profiles AA' through DD' show the locations of conductivity cross sections plotted in Figure 4.42 and Figure 4.43.

The conductivity images, along with the conductivity-depth sections that are shown in Figure 4.42 and Figure 4.43, are consistent with the conceptual geological model as well as prior information presented in Section 4.4.2 and in Brodie et al. (2004b). This is supported by the observations that follow. The Layer 1 image shows that the near surface is generally more resistive than the deeper layers and is largely influenced by the high spatial frequency east–west dry and resistive sand dunes. The high conductivities displayed in the Layer 15 and 20 images show that the saline groundwater is encountered at depth. The Layer 10 and Layer 15 conductivity images

show the strong curvilinear northwest–southeast grain that is consistent with the existence of Pliocene strandlines. The presence of a locally dissected near surface conductor, that is interpreted to be the Blanchetown Clay, is supported by the Layer 5 image. This conductor can be identified in the conductivity-depth sections (Figure 4.42 and Figure 4.43) at approximately 40 m elevation.

The groundwater elevation surface that is shown in Figure 4.8b was compiled by Brodie et al. (2004b) by gridding groundwater data from numerous boreholes. This surface was explicitly used as interface-depth data in the few-layer holistic inversion, but was not used at all in this multi-layer inversion example to simulate the greenfields scenario. However the surface has been plotted on the conductivity-depth sections and appears as a blue line at approximately 15-20 m elevation. It is noteworthy that the groundwater elevation surface correlates well with the section's pronounced rapid conductivity increase (yellow-orange transition) at depth. It is thus believed that the conductivity model provides an effective means of mapping the elevation of the saline watertable.

One might conclude that a highly conductive zone, evident well above the watertable surface on the Profile CC' between 16,000 and 19,000 m (Figure 4.43a), is suspicious or not consistent with the conceptual model. However, mineralogical analysis by Tan et al. (2004) of materials collected from borehole RIV9HC (shown on the section), which intersects the edge of the anomaly, indicated that this anomaly is real and due to a highly conductive kaolinitic clay between 4 and 16 m depth that sits above the water table which is at 22 m depth. This is also confirmed by the conductivity log for borehole RIV9HC shown on Figure 4.47.

Any downhole conductivity logs that are with 200 m of the conductivity-depth sections profiles have been plotted on the sections (Figure 4.42 and Figure 4.43) in a vertical column coloured using the same colour lookup table as the section. They downhole logs are consistent with the inversion model shown on the sections.

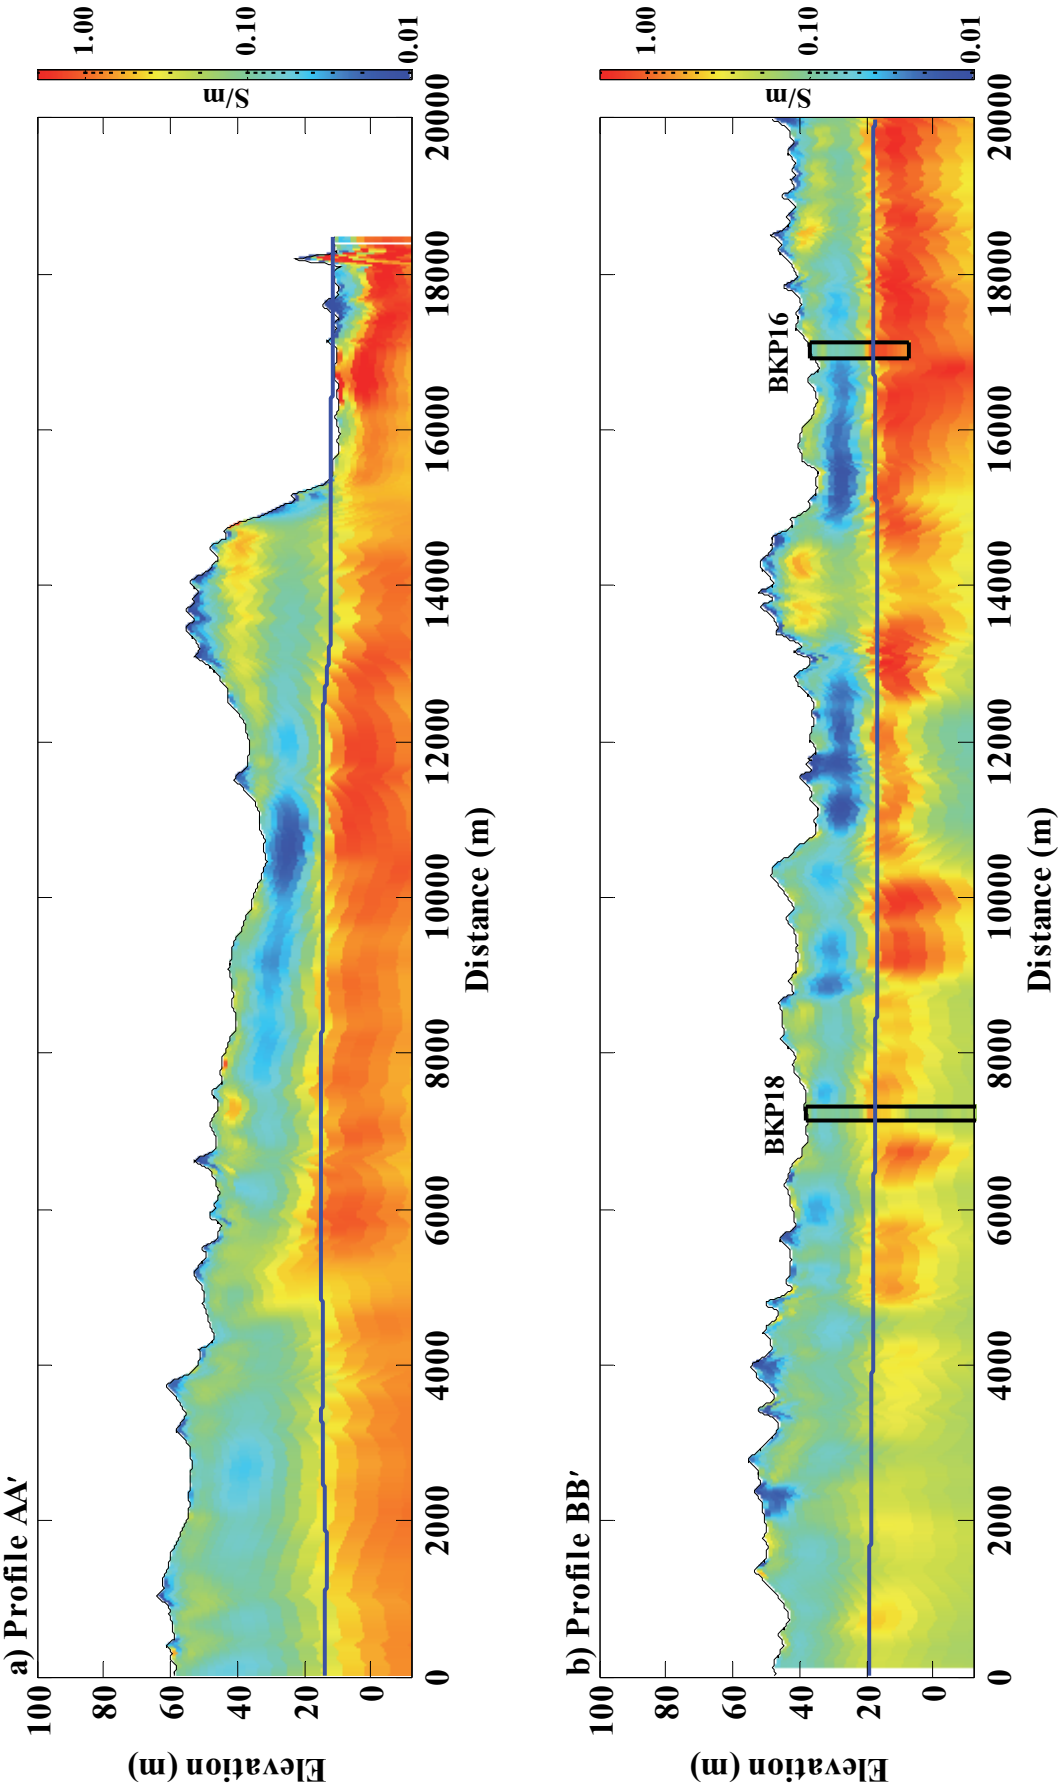

Figure 4.42 Conductivity sections along the profiles AA' and BB' shown on Figure 4.37. Also plotted, using the same colour mapping, are downhole conductivity logs (columns) and the watertable elevation (blue line).

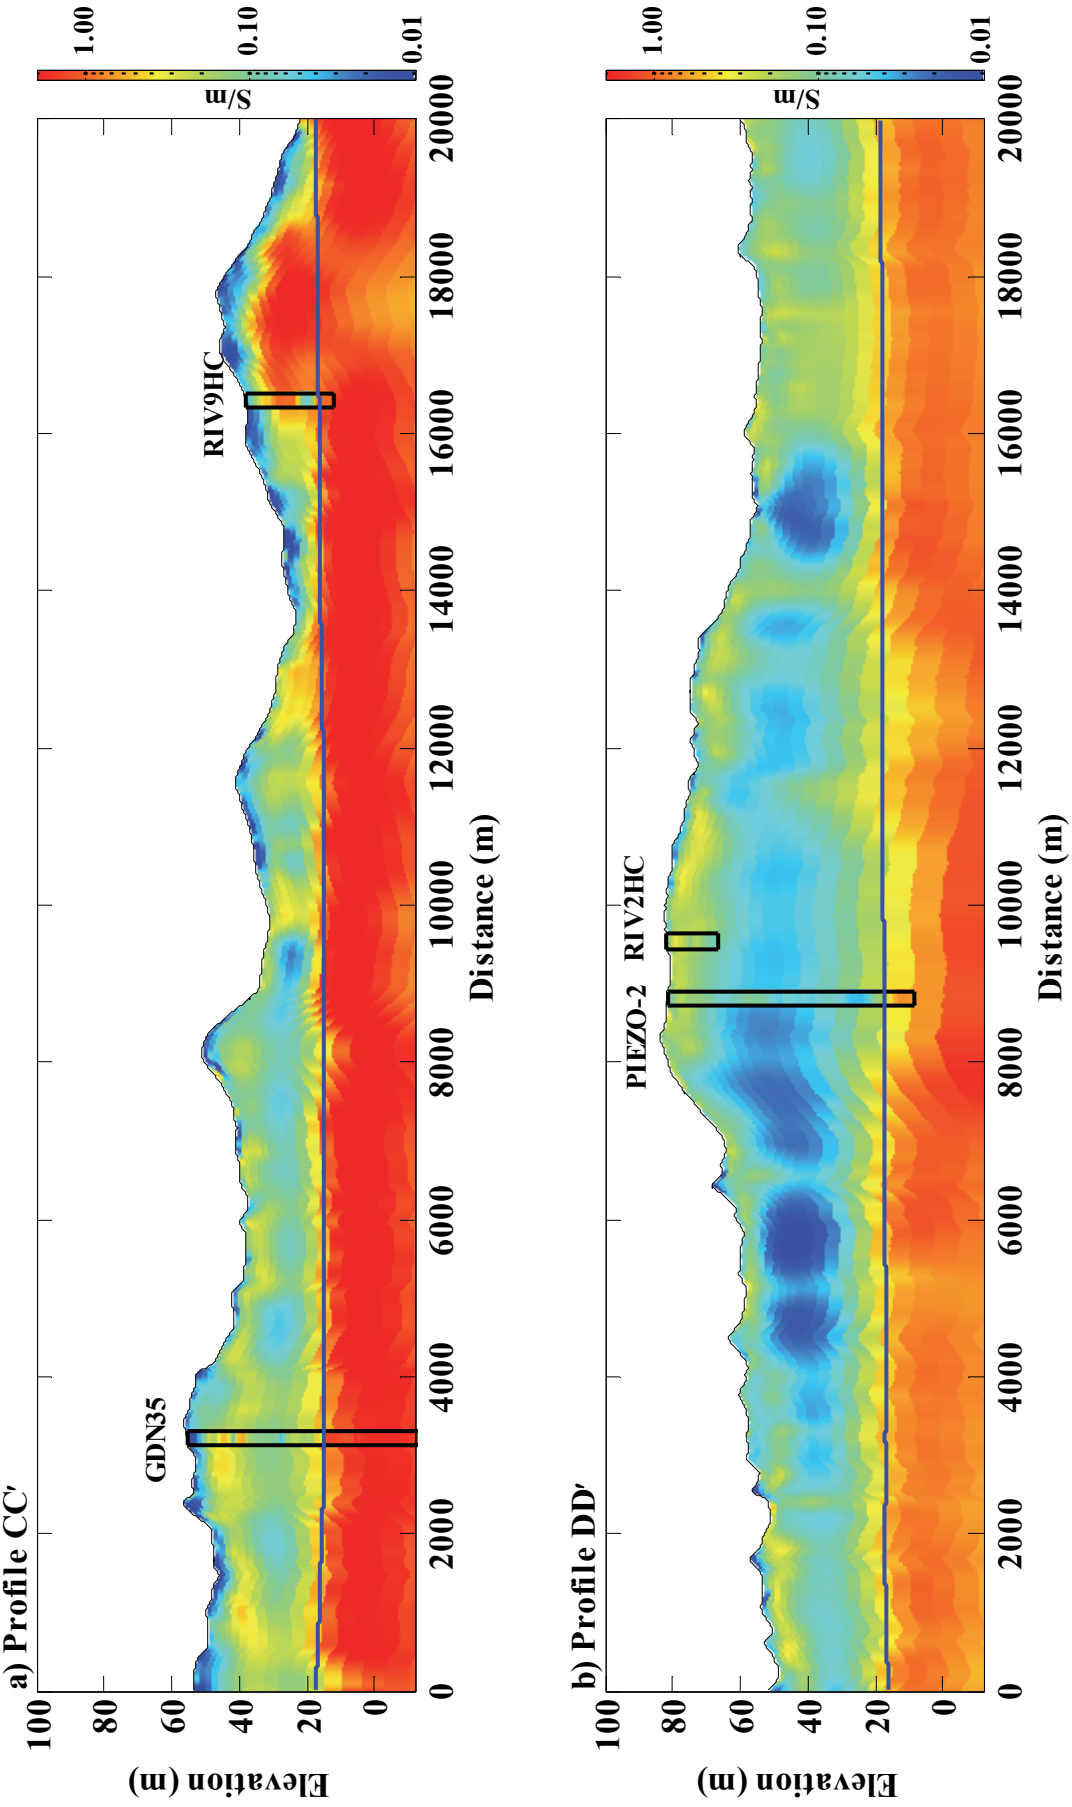

Figure 4.43 Conductivity sections along the profiles CC' and DD' shown on Figure 4.37. Also plotted, using the same colour mapping, are downhole conductivity logs (columns) and the watertable elevation (blue line).

A more general comparison of downhole conductivity log measurements and conductivities estimated by the multi-layer holistic inversion model is shown in Figure 4.44. The data relate to the average conductivity over every 5 m depth interval in all of the 44 boreholes that were in the defined portion of the conductivity model. Note that there were 49 downhole logs in the survey area, but 5 were in positions where airborne data could not be acquired due to built-up areas and the highway, and hence the conductivity model is not defined in these areas. The correlation is high (R^2=0.82), although the inversion estimates are slightly biased toward higher conductivities. Individual conductivity-depth plots for all 44 logs are shown in Figure 4.45, Figure 4.46, and Figure 4.47. The overall assessment is that the holistic inversion has been able to estimate the vertical conductivity structure well, which is an encouraging result given that the downhole log data were not used at all.

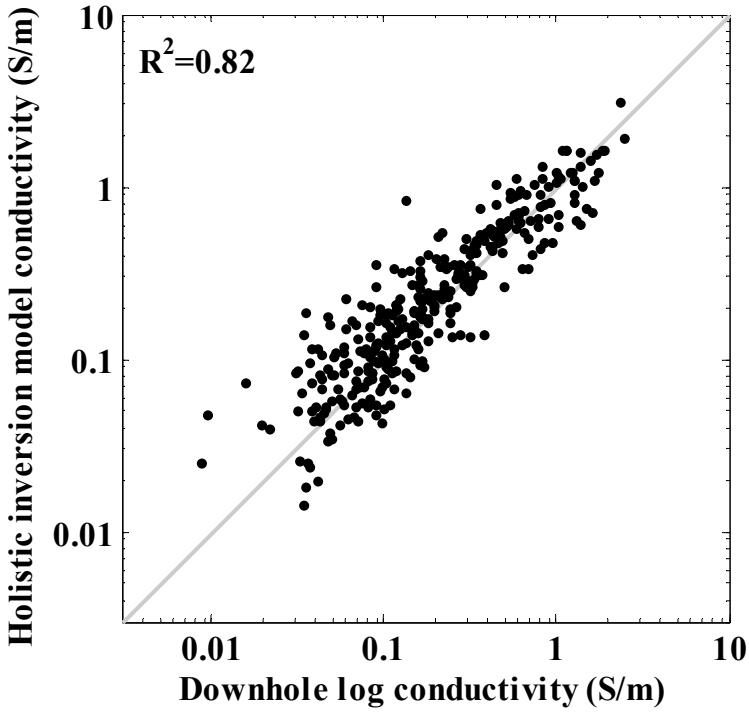

Figure 4.44 Comparison between downhole log conductivities measurements from all 44 boreholes and the multi-layer holistic inversion model conductivities. Each data point represents a 5 m depth interval in each borehole.

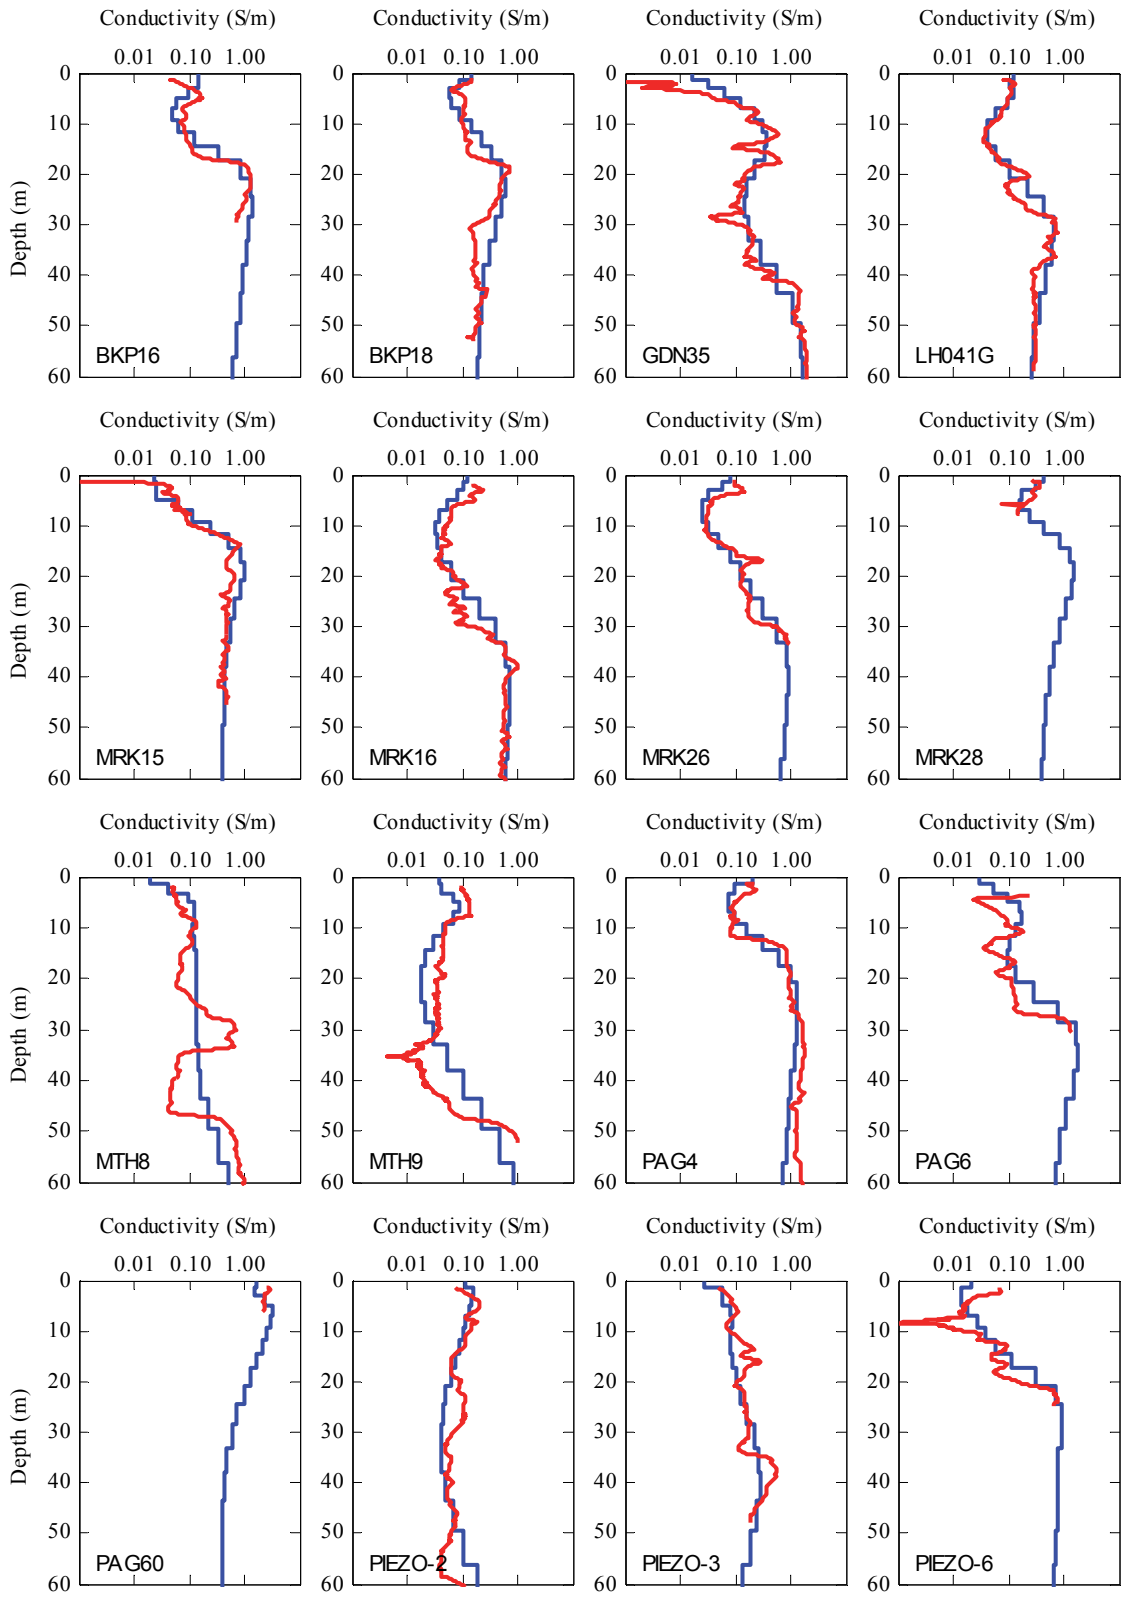

Figure 4.45 Comparison between downhole conductivity logs (red) and the conductivity profile at each downhole log estimated by the multi-layer holistic inversion (blue). Logs 1 to 16 of 44 - continued in Figure 4.46.

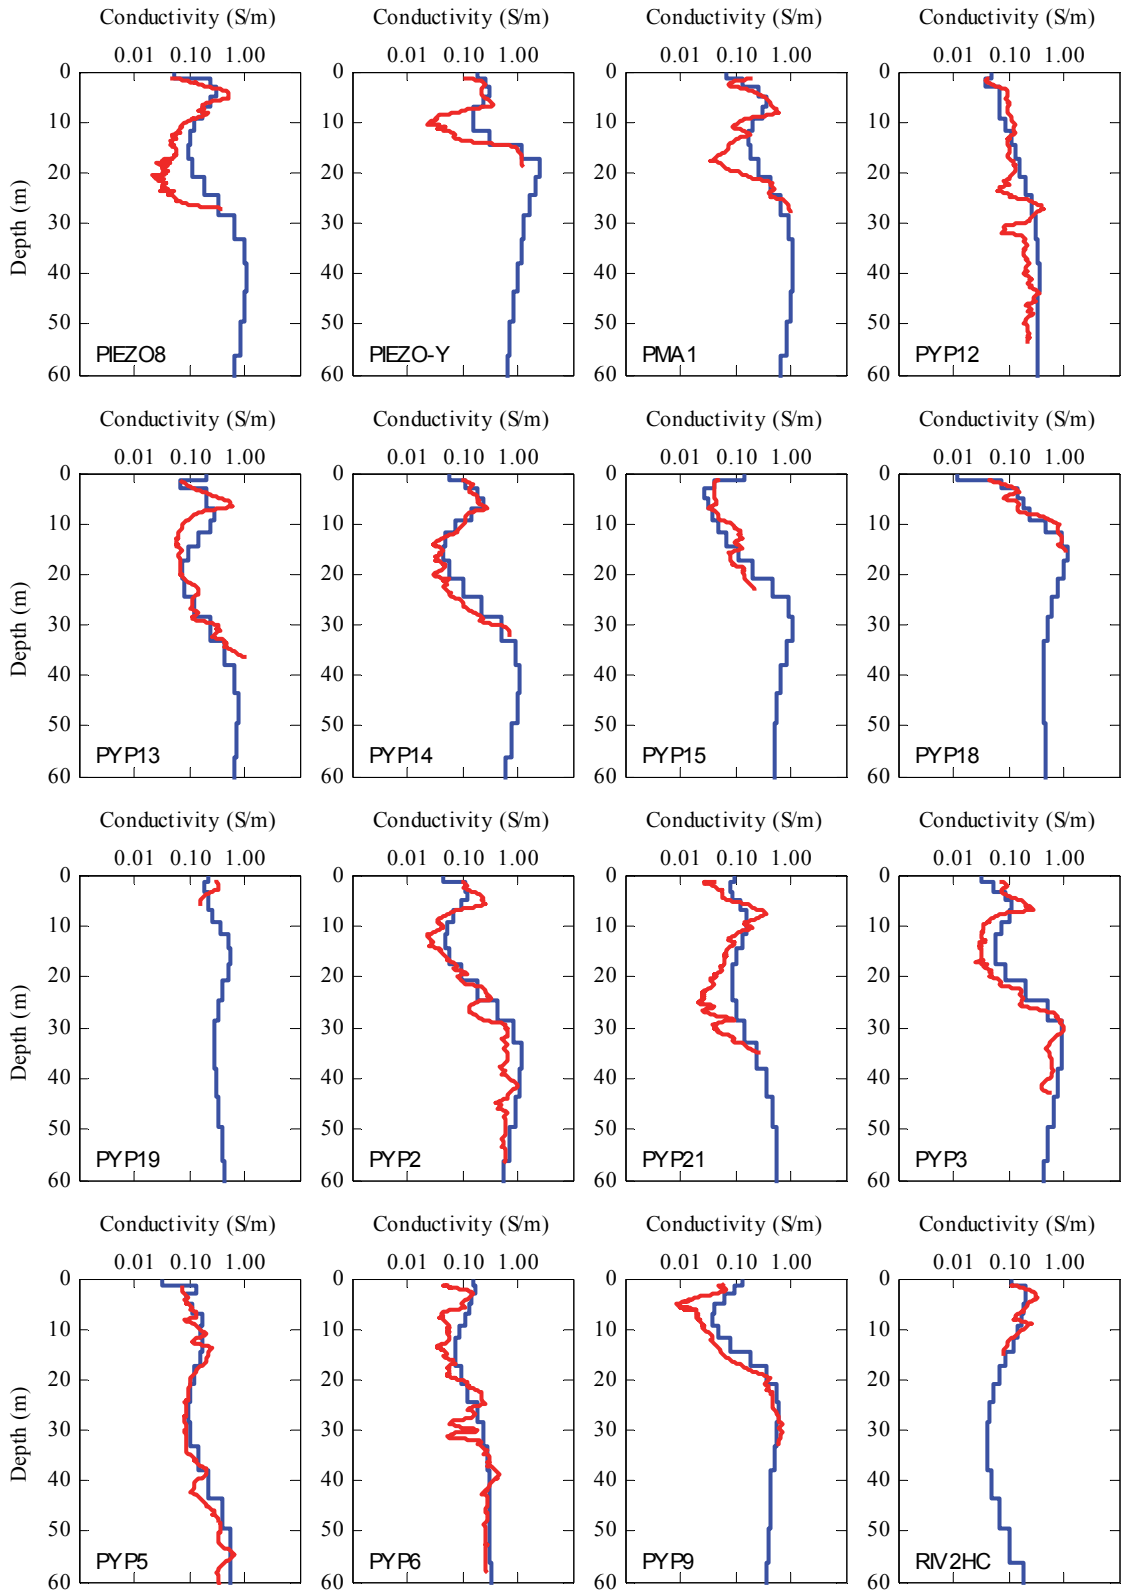

Figure 4.46 Comparison between downhole conductivity logs (red) and the conductivity profile at each downhole log estimated by the multi-layer holistic inversion (blue). Logs 17 to 32 of 44 - Continued in Figure 4.47.

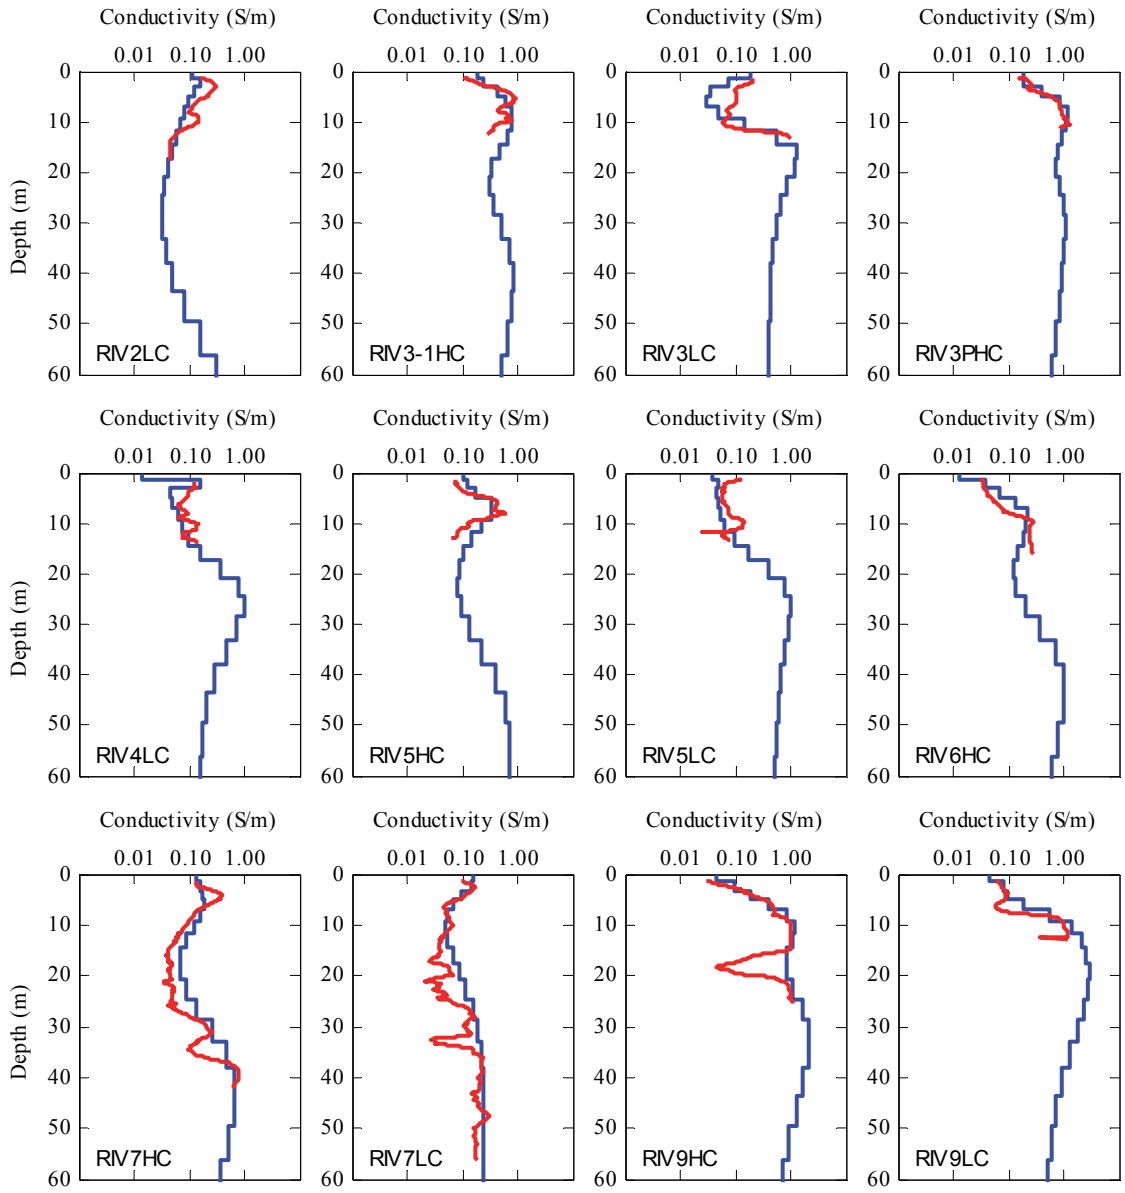

Figure 4.47 Comparison between downhole conductivity logs (red) and the conductivity profile at each downhole log estimated by the multi-layer holistic inversion (blue). Logs 33 to 44 of 44 - continued from Figure 4.45 and Figure 4.46.

In Figure 4.44 to Figure 4.47 it can be seen how, with multiple layers and vertical smoothness constraints, the multi-layer conductivity model is much better able to represent the typically gradational (curved) nature of the downhole logs than the few-layer inversion in (Figure 4.33) could. This would partly explain why the comparison between the measured downhole data and the multi-layer inversion conductivity model is superior to the comparison shown in Figure 4.32a between the observed and predicted geoelectric data for the few-layer inversion.

4.6 Conclusion

In this chapter different applications of the holistic inversion for frequency-domain electromagnetic data have been presented. In the first example data that had been synthetically generated from a conductivity and calibration model, plus added noise, were inverted. It demonstrated that even when a realistic amount of Gaussian noise is added to the data, the recovered parameters of the conductivity and calibration models were acceptably close to the known true values. It indicated that the sensitivity of the estimated values to Gaussian noise was within an acceptable range and that ill-posedness or fundamental trade-off between the five classes of parameters did not appear to be a serious problem, at least in the idealized case where the parameterization was consistent with the synthetic models. Since the parameters were exactly recovered when noise was not added to the synthetic data, the synthetic example also confirmed the theory had been correctly implemented in computer code.

The few-layer example demonstrated how the holistic inversion method can be successfully used to calibrate, process and invert practically raw airborne data to estimate a plausible calibration and conductivity models. The resulting conductivity model was superior to that produced from inversion of final, fully processed contractor delivered data using a conventional sample-by-sample algorithm. This was demonstrated by: (a) more coherent and artefact-free maps in plan and section form; (b) lower overall data misfits; and (c) the spatial non-correlation of inversion residuals.

The multi-layer inversion similarly demonstrated that the holistic inversion can be used to successfully calibrate, process and invert practically raw airborne data. However the multi-layer example additionally demonstrated that the method does not rely on having an advanced conceptual conductivity model, strong reference model constraints, downhole conductivity logs, or interface-depth data.

The few and multi-layer inversions of real data both estimated gain calibration parameters that were consistent with (reciprocal) scale factors derived by independent means. The bias estimates from the holistic inversions were strongly correlated with, but not identical to, the zero-level estimates that were measured in the conventional high-altitude method. It was concluded though, that because of its many fewer artefacts and lower data misfits, the holistic inversions may in fact have provided the better estimate of bias than the conventional technique.

Although downhole conductivity log data were not included in the multi-layer holistic inversion, so that the greenfields mapping scenario could be simulated, the estimated conductivity model was very well correlated with the downhole logs. It also demonstrated how the gradational or sinuous nature of downhole logs can be much better represented by a multi-layer model than by a few-layer model.

Although conventional sequential processing and inversion approaches are convenient and computationally less challenging to implement than the holistic approach, the examples shows that the holistic approach can be implemented. The few-layer example showed that the method is tractable on a standalone computer for a regular industry size (10×10 km) survey. The multi-layer example, which was run with parallelized code on a cluster computer, showed that the holistic inversion is also tractable for larger regional (>10,000 line km) surveys.

The holistic approach may in fact provide cost savings because it avoids the time consuming and therefore costly, iterative calibration-processing-recalibration-reprocessing carried out by expert individuals. Because superior results were obtained without specific use of zero-levels, it may be feasible to eliminate high altitude measurements altogether if the holistic approach is used. This would be of considerable practical benefit because of the likely substantial reduction in acquisition costs.

Chapter 5

Time-domain holistic inversion

5.1 Outline

In Chapter 3 the theory for the holistic inversion of frequency-domain airborne electromagnetic data was presented. The frequency-domain holistic inversion formulation is designed to invert raw data to estimate a bicubic B-spline based 3D conductivity model and the parameters of a calibration model that essentially account for systematic calibration errors in the observed data. The method was applied to synthetic and real survey data in Chapter 4, and the results demonstrated the method can achieve superior results when compared to sample-by-sample inversion of conventional fully processed data. In this chapter an analogous method for time-domain airborne electromagnetic data will be presented. The general philosophy in this case is the same as that for the frequency-domain method.

The challenges that are encountered in calibrating time-domain systems are different to those for frequency-domain systems. This means the calibration model for the time-domain method must represent a different set of systematic errors than the frequency-domain method. Furthermore, the most significant calibration issues for each class of time-domain system also differ. Therefore the development will concentrate on one particular set of calibrations problems, which is the determination of unmeasured elements of the system geometry in fixed-wing towed-bird systems.

Another difference between the frequency-domain holistic inversion and the time-domain method presented in this chapter is the parameterization of the conductivity

model. For the time-domain development the conductivity model will be parameterized as a continuous 2D rather than 3D model. The reason for adopting a 2D model is as follows. During the course of this research Geoscience Australia began a program of regional airborne electromagnetic surveys in which the flight line spacing is up to 5 km. A key feature of the frequency-domain holistic inversion is to take advantage of any between-line spatial geological correlation by using a continuous 3D model. However, with such widely spaced flight lines the between-line spatial correlation is unlikely to be of much value. Accordingly, a method that could take advantage of the along-line correlation, but that did not rely on between-line correlation, was required.

The overall idea then is to invert one complete flight line of data simultaneously, in other words it is a flight line-by-line inversion. The method therefore sits between the conventional sample-by-sample approach, in which individual samples are independently inverted; and the frequency-domain method presented in Chapter 3 in which complete datasets are simultaneously inverted. The obvious parameterization choice was to simplify the previous 3D model, based on **bicubic** B-spline meshes, to a 2D model based on (one dimensional) **cubic** B-splines.

For a 2D conductivity model parameterization geoelectric and interface-depth data are not as useful as when a 3D model is employed. This is because the boreholes, in which downhole conductivity log data are measured, or other ground-truth stations, are seldom located coincident with flight lines. Consequently, the vast majority of flight lines will not have any ground-truth stations within a suitable distance, and for those that do some interpolation of the data onto the flight line would be required. It is for this reason the geoelectric and interface-depth data have not been included in the 2D conductivity model time-domain formulation. This highlights one of the strengths of the 3D conductivity model used in the frequency-domain formulation, which allowed inclusion of any such data in the survey area.

The time-domain holistic inversion presented here is thus a simplification of the frequency-domain method. In principle, however, there is no reason why a 3D model parameterization could not be used in exactly the same way as it was for the frequency-domain, in which case it would be logical to include geoelectric and interface-depth data.

In the following sections a description of the method for fixed-wing towed-bird installations is provided. Since much of the development is similar to, or exactly the same as provided earlier in Chapter 3 for the frequency-domain case, the description is briefer with the emphasis on the differences. Possible methods of addressing calibrations errors that are more important in other types of time-domain systems will be discussed at relevant points, however they will not be fully developed. The code for implementing the method has so far not been fully developed. Currently the code only caters for inversions having the same spline node spacing for every layer, and a spatially constant reference model must be used. Nevertheless, an example of the method's application to real TEMPEST data is presented.

5.1.1 Nomenclature

Nomenclature that appears throughout this chapter is listed in the table below. Terms are described as they are encountered, but are provided here for the reader's convenience.

N_s	number of airborne data samples in the flight line to be inverted
N_c	number of receiver components recorded by the system
N_w	number of time windows per receiver component
$N_d = N_s \times N_c \times N_w$	number of airborne data to be inverted
N_L	number of layers in the earth model
N_e	number of earth model spline coefficients to be solved for
N_ψ	number of system geometry spline coefficients to be solved for
$N_m = N_e + N_\psi$	total number of model parameters to be solved for

N_r number of reference model values

N_h number of horizontal (along-line) roughness constraints

N_v number of vertical roughness constraints

d^X_{sw} observed X-component data for the wth window at the sth sample in the line

d^Y_{sw} observed Y-component data for the wth window at the sth sample in the line

d^Z_{sw} observed Z-component data for the wth window at the sth sample in the line

$\mathbf{d} = \left[d^X_{1,1} \ldots d^X_{1,Nw} \mid d^Y_{1,1} \ldots d^Y_{1,Nw} \mid \ldots\ldots \mid d^Z_{Ns,1} \ldots d^Z_{Ns,Nw} \right]$ observed data vector

ψth height of the transmitter loop above ground

ψtr roll of the transmitter loop

ψtp pitch of the transmitter loop

ψty yaw of the transmitter loop

$\psi \triangle x$ horizontal in-line separation between the TX-loop and the RX coils

$\psi \triangle y$ horizontal transverse separation between the TX-loop and the RX coils

$\psi \triangle z$ vertical separation between the TX-loop and the RX coils

ψrr roll of the receiver coils' assembly

ψrp pitch of the receiver coils' assembly

ψry yaw of the receiver coils' assembly

$\psi \triangle x^{est}$ estimate of $\psi \triangle x$ made during the data processing

$\psi \triangle z^{est}$ estimate of $\psi \triangle z$ made during the data processing

$\psi \upsilon$ shorthand for a geometry parameter (i.e. $\psi \upsilon$ may ψth, $\psi \triangle x$, ψrp etc.)

c^{σ}_{ik} coefficient of the ith node of the kth layer's conductivity spline

c^{t}_{ik} coefficient of the ith node of the kth layer's thickness spline

$c^{\psi \upsilon}_{i}$ coefficient of the ith node of the geometry parameter $\psi \upsilon$'s spline

x_s horizontal distance along the flight line of the sth sample

$w^{\sigma}_{ik}(x_s)$ weight of ith node coefficient used in evaluating $\sigma_k(x_s)$

$w^{t}_{ik}(x_s)$ weight of ith node coefficient used in evaluating $t_k(x_s)$

$w^{\psi \upsilon}_{i}(x_s)$ weight of ith node coefficient used in evaluating $\psi \upsilon(x_s)$

$\sigma_k(x_s)$ local 1D model's kth layer conductivity evaluated from its spline at x_s

$t_k(x_s)$ local 1D model's kth layer thickness evaluated from its spline at x_s

$\psi \upsilon(x_s)$ local geometry value for the parameter $\psi \upsilon$ evaluated from its spline at x_s

$\mathbf{m}^{ear}$ earth model spline coefficients to be solved for

$\mathbf{m}^{geom}$ system geometry model spline coefficients to be solved for

$\mathbf{m} = \left[\mathbf{m}^{ear} \mid \mathbf{m}^{geom} \right]^{\mathbf{T}}$ model parameter vector

$\boldsymbol{g}_s^{air}(\boldsymbol{\sigma}(x_s), \mathbf{t}(x_s), \boldsymbol{\psi}(x_s)) = \left[g_{s,1}^X \ldots g_{s,Nw}^X \mid g_{s,1}^Y \ldots g_{s,Nw}^Y \mid g_{s,1}^Z \ldots g_{s,Nw}^Z \right]$ sth sample forward model

$\boldsymbol{g}(\mathbf{m}) = \left[\boldsymbol{g}_1^{air}(\mathbf{m}) \mid \boldsymbol{g}_2^{air}(\mathbf{m}) \mid \ldots\ldots \mid \boldsymbol{g}_{Ns}^{air}(\mathbf{m}) \right]$ predicted data vector

$\mathbf{G}_\mathbf{m}$ Jacobian matrix where $\mathbf{G}_\mathbf{m}\big|_{ij} = \partial \boldsymbol{g}_i(\mathbf{m}) \big/ \partial \mathbf{m}_j$

5.2 Conductivity model

The subsurface is deemed to be comprised of discrete layers stacked in layer-cake fashion. The bottommost layer, or the basement halfspace, is infinitely thick. Figure 5.1 illustrates a possible three layer model, showing how both the conductivity (represented by colour) and thickness of each discrete layer, vary smoothly and continuously in the lateral sense. However at any one horizontal position the conductivity is constant from the top to the bottom of the layer.

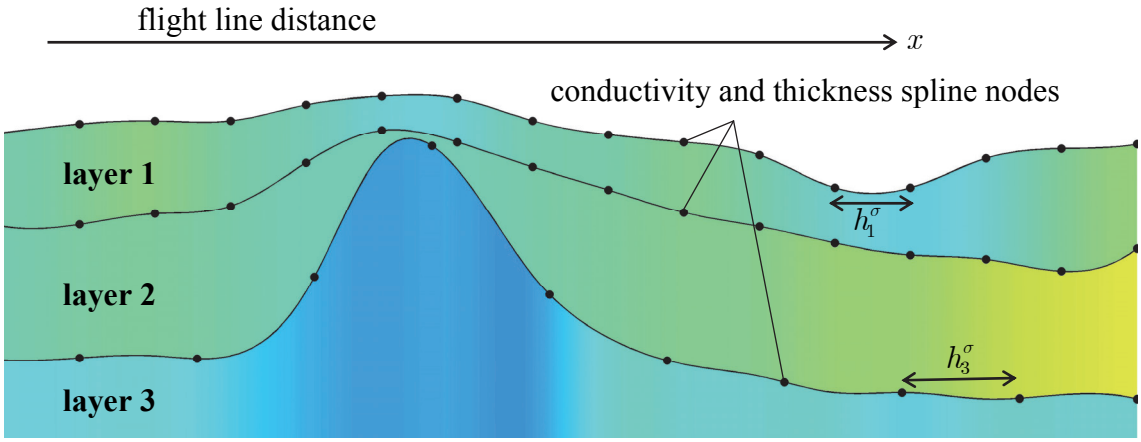

Figure 5.1 The layered and laterally continuous, one dimensional cubic B-spline based, conductivity model used for the time-domain holistic inversion.

The logarithm of conductivity and logarithm of thickness of each layer is parameterized by a uniform cubic B-spline. Cubic B-splines are a 1D simplification of the 2D bicubic B-splines which are fully described in Appendix C. The spline nodes are positioned at a regular spacing in distance along the flight line. The spline parameterization allows a node spacing to be chosen such that the subsurface conductivity can be adequately represented with as few model parameters as possible. The chosen node spacing is dependent on the minimum of, the expected scale length of lateral conductivity variations, and the airborne system's footprint. Each spline is completely independent

of the others. Different node spacings can be chosen for each layer and can also be different for the layer conductivity and thickness properties.

In general to parameterize a N_L layered-earth $2N_L$ - *1* cubic B-splines are required, N_L for the conductivities and N_L - *1* for the layer thicknesses. Since the bottom layer is infinitely thick, its thickness is not a variable. In a few-layer style inversion, where we choose to invert for both layer thickness and conductivity, all $2N_L$ - *1* splines are required. However for a multi-layer inversion, where the layer thicknesses are kept fixed, we may use splines to represent the fixed layer thicknesses, or alternatively the layer thicknesses can be constant. In the latter case only N_L splines are required in total. However for the ensuing development we will consider the more general case where both conductivities and thicknesses are allowed to vary.

The cubic B-spline representing the *k*th layer conductivity and thickness is denoted S_k^{σ} and S_k^{t} respectively. To avoid repetition, for the moment consider the spline S_k^{v} where v may represent either σ or t. The spline S_k^{v} is comprised of a topologically linear array, of N_k^{v} - *1* segments, which when joined together form a spline having N_k^{v} nodes that are spaced distance h_k^{v} horizontally apart. The *i*th node has a corresponding coefficient value c_{ik}^{v}. It is these earth model coefficients that are solved for in the inversion, of which in total there are,

$$N_e = \sum_{k=1}^{N_L} N_k^{\sigma} + \sum_{k=1}^{N_L-1} N_k^{t} \,. \tag{5-1}$$

At any given lateral distance x along the flight line that lies on the *i*th segment of the spline S_k^{v}, the conductivity or thickness parameter $v_k(x)$ is parameterized by,

$$l\,\mathrm{og}(v_k(x)) = S_{ik}^{v}(x) = \sum_{a=i-1}^{i+2} w_{ak}^{v}(x)\; c_{ak}^{v} \,, \tag{5-2}$$

where $w_{ak}^{v}(x)$ is the known weight that determines the contribution of the *a*th spline node coefficient, on the *k*th spline, c_{ak}^{v} to the spline surface value at point x. The

weights are a function only of the position x along the spline. For details of determining the weights see Appendix C. The positions of the spline nodes are set at the beginning of the inversion and remain fixed throughout, and therefore the weights remain fixed for each flight line distance x as well.

It can be seen then from Equation 5-2 that the spline evaluation is a simple weighted sum of four coefficients. Therefore the conductivity or thickness of any one layer at any one location is influenced by only the coefficients of the four immediately surrounding nodes for that layer. Accordingly the vast majority of the derivatives in the inversion's Jacobian matrix are zero, which allows use of sparse matrix computational methods in the inversion.

5.3 Calibration model

5.3.1 Background

In Chapter 3.3.1 it was discussed how errors in frequency-domain airborne electromagnetic systems can be sub-divided into 'non-systematic noise' and 'systematic noise' components. The situation is the same for time-domain systems. What is meant by non-systematic noise here are those errors caused by unpredictable events and that cannot readily be deterministically modelled. For time-domain systems some examples of non-systematic noise are: atmospheric sferic events, interference from powerlines, interference from VLF and radio transmitters, and high frequency mechanical vibration of the receiver coils. It is well recognised in the geophysical industry and has to some degree been the main focus of attempts to improve AEM systems over many years. This is because it can be difficult to discriminate between non-systematic noise and the responses from discrete geological conductors, which mineral explorers are most interested in, since it usually presents in data at the same high spatial frequencies as the geological conductor responses.

A variety of well-established techniques (e.g. stacking and filtering) are employed to successfully reduce the effect of non-systematic noise, nevertheless it remains an important contributor to the total error (Macnae et al., 1984). Since non-systematic errors cannot readily be deterministically modelled, they cannot be targeted by the holistic inversion. They are therefore treated stochastically, as they are in conventional electromagnetic inversion algorithms.

One aim of the holistic inversion is to deterministically model errors and thereby remove any detrimental effect that they may otherwise have on the conductivity estimates. Since systematic errors can be deterministically modelled they can be targeted by the holistic inversion. Each class of time-domain airborne electromagnetic system is affected by different types of systematic errors. The exact nature of the error will vary form system to system. Therefore, as noted at the beginning of the chapter, this thesis only gives a detailed development and examples for a particular set of calibration problems. That set, for which background information is given in Sections 5.3.2, 5.3.3, and 5.3.4, is caused by the fact that some elements of fixed-wing towed-bird systems geometry are not measured. In Section 5.8 some calibration errors that are common in other forms of time-domain system are discussed, and possible methods of parameterizing them are postulated.

5.3.2 Definition of fixed-wing towed-bird system geometry

Fixed-wing towed-bird time-domain installations include well known systems such as TEMPEST (Lane et al., 2000), SPECTREM (Klinkert et al., 1997) and GEOTEM (Annan and Lockwood, 1991). In these systems a wire(s) is suspended between the aircraft's wingtips, tail and nose to form the transmitter (TX) loop. The receiver (RX) coils are towed (~120 m behind and ~35 m below) the aircraft inside a shell known as the 'bird'. The transmitter loop's and receiver coils' height, relative positions and orientations, are in combination is known as the system geometry. Table 5.1 lists the

ten variables that define the full system geometry, which are also schematically illustrated in Figure 5.2.

Table 5.1 Elements of the fixed-wing towed-bird system geometry

Element	**Symbol**	**Measured**	**Inversion parameter**
TX loop height above ground	ψth	yes	possibly
TX loop's roll	ψtr	yes	no, usually measured
TX loop's pitch	ψtp	yes	no, usually measured
TX loop's yaw	ψty	yes	no, usually measured
TX loop to RX in-line distance	$\psi \triangle x$	no	yes
TX loop to RX transverse distance	$\psi \triangle y$	no	only if Y-component available
TX loop to RX vertical distance	$\psi \triangle z$	no	yes
RX coils' roll	ψrr	no	only if Y-component available
RX coils' pitch	ψrp	no	yes, if X and Z-components
RX coils' yaw	ψry	no	no

Figure 5.2 shows that the aircraft flies in the inertial coordinate systems x-axis (in-line) direction with the centre of the transmitter loop at height ψth above the ground. The orientation of the transmitter loop is defined by its roll ψtr, pitch ψtp, and yaw ψty. Roll, pitch and yaw are defined as counter-clockwise rotations about the origin of the inertial coordinate system's x-axis, y-axis and z-axis respectively, for an observer looking toward the origin (see Appendix A).

Relative to the transmitter loop, the receiver bird is at position $\psi \triangle x$ in the x-axis (in-line), $\psi \triangle y$ in the y-axis (transverse), and $\psi \triangle z$ in the z-axis (vertical) direction. The values of $\psi \triangle x$ and $\psi \triangle z$ will be negative because the bird always flies behind and below the aircraft. Inside the shell of the bird, three orthogonal receiver coils are mounted. They measure the time rate of change of the magnetic field in the direction of their axis. In general the receiver coils will not be aligned with the inertial coordinate system's axes as shown in the left hand inset of Figure 5.2. Instead the orientation of the receiver assembly is defined by its roll ψrr, pitch ψrp, and yaw ψry rotations about the origin. The rotations are defined relative to a reference orientation in which

the X-component coil is aligned with the x-axis, the Y-component coil is aligned with the y-axis, and the Z-component coil is aligned with the z-axis.

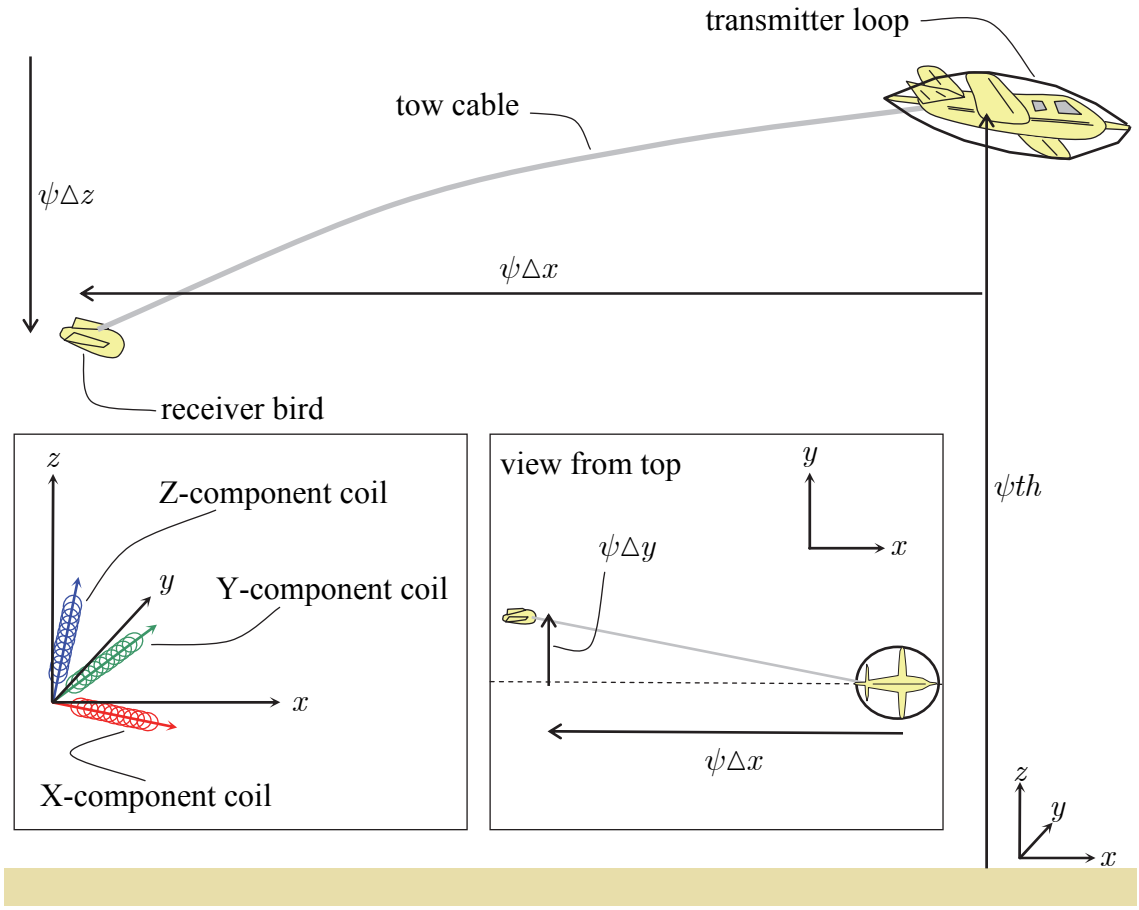

Figure 5.2 Schematic illustration of the fixed-wing towed-bird system geometry.

5.3.3 Measurement of system geometry

All ten of the system geometry elements are necessary inputs to calculate the observed response at the receiver, in the axial direction of the coils. However in general not all can be practically measured during routine production survey flying (Table 5.1).

The height of the aircraft, and hence transmitter loop, can be measured with a radar or laser altimeter. If calibrated correctly these instruments can provide measurements to approximately ±1 m accuracy. They are known to give incorrect results if the altimeter reflects off dense vegetation or infrastructure (Beamish, 2002) rather than the ground. Radar altimeters can also give incorrect results if the ground is rough at the scale of the radar's wavelength, for example over ploughed fields (Brodie and Lane, 2003). More

systematic altimeter calibration errors can occur for example if: the pointing direction of a laser altimeter is not known; the altimeter's voltage-to-height conversion factors are not correctly established; or the vertical offset between altimeter sensor and transmitter loop are not properly measured and accounted for. Therefore in some cases it may be necessary to invert for transmitter loop height, but it is probably unnecessary if the altimeters are sufficiently well calibrated, and the vegetation is not dense.

The roll and pitch of the transmitter loop can be measured by gyroscopes and inclinometer measurements mounted on the aircraft. Since these instruments are capable of better than 1° accuracy it is not considered necessary to invert for transmitter loop roll or pitch. The yaw is zero by definition of the coordinate system.

The problem of measuring the orientation and position of the receiver bird relative to the aircraft is an ongoing problem for fixed-wing tow-bird systems. Attempts have been made to track the bird's position with laser ranging devices (Smith, 2001b). This was successful at high altitude under calm conditions, but they found it to be impractical for routine production at survey altitude. It may also be possible to determine the bird's position and orientation using a GPS coupled to an inertial navigation unit attached to the bird (Vrbancich and Smith, 2005). This is an active area of research and development, however such systems are yet to be implemented because of the cost and difficulty of locating instruments close to the sensitive receiver coils without introducing noise.

5.3.4 Estimation of the primary field and system geometry

In the past, a nominal receiver bird position have been estimated using photography of the system in flight (Smith, 2001b). Photography has limitations because of distortions involved and the implicit assumption that the geometry at the instant the photograph is taken is representative of the nominal position. Another approach is to undertake numerical analysis of the aerodynamic properties of the bird combined with

photographic evidence (Annan, 1983). Since this requires information about prevailing atmospheric conditions (e.g. wind velocity, temperature, and pressure and humidity), which are not available on a dynamic basis throughout the survey area, the applicability of the method is limited.

More recent approaches have estimated the bird position from the electromagnetic data itself. The basic premise is that, if the primary field is known, then the bird position relative to the transmitter, can be solved for from the simple analytic expression for the magnetic fields due to a magnetic dipole in free space. The bird is usually assumed to flying straight and level (i.e. with zero roll, pitch and yaw). Knowledge of the primary field at the receiver only became feasible with the advent of the more sophisticated broadband systems that measure the response during both the on- and off-time of the transmitted waveform, and that have two or three component receivers. These systems observe the total, primary plus secondary, fields. So the primary field, or some estimate of it, must first be isolated from the total response before the bird position can be solved for. Some variations on these methods are now summarized.

A method first suggested by Annan (1984) and implemented by Smith (2001a) uses a least squares inversion to solve for the bird position using the electromagnetic data recorded at high altitude. If the data are recorded at sufficient altitude, the secondary field, originating from currents induced in the ground, is negligible and the recorded data can be considered to be pure primary field. This allows the receiver position to be solved for as explained above. The high altitude position estimates could then be averaged to produce a nominal bird position to be subsequently used in the interpretation of survey altitude data.

Smith (2001b) introduced a method to first estimate the primary field, and hence bird position, not only from high altitude data, but from survey altitude data as well. It has typically been the method used in the processing of GEOTEM and MEGATEM system

data. The method decomposes the recorded total-field response into what the author calls the 'time-domain inphase and quadrature responses'. The decomposition is performed for every sample (transient) along the survey altitude flight line. The time-domain inphase response is determined through a procedure that minimises, in a least squares sense, the difference between the observed total field and a scaled version of the total field reference transient which is recorded at high altitude. The scaled version of the reference transient becomes the time-domain inphase response and the remainder become the quadrature response. Using the time-domain inphase response as an estimate of the primary field, the bird position is then solved for. So essentially the method estimates that the primary field to be any part of the total field response that has the same shape and phase as the high altitude receiver reference waveform. As the author points out, the procedure is therefore not a method for estimating the primary field, but more correctly a procedure for estimating the total inphase response, which is actually comprised of the primary plus secondary response. It therefore, assumes that there is no secondary field, due to currents in the ground, in the time-domain inphase response.

The method used for SPECTREM is to equate the last time window of the total field transient to the primary field (Leggatt et al., 2000). This in essence assumes that any secondary fields have decayed away to zero by the end of the transient. An alternate strategy is to fit the last three windows of the total field transient to a decaying exponential and the primary field is estimated to be the asymptote of the exponential as time approaches infinity.

For TEMPEST data processing the primary field is estimated from the deconvolved total field frequency-domain spectrum. An unknown constant, across all frequencies in the spectrum, must be determined and removed (Lane et al., 2000). The constant to be removed is determined by making assumptions about the secondary fields, and hence ground conductivity, at either late time or low frequency. The assumptions depend on

the selection of a preferred model of the ground response (e.g., layered-earth response, sum of exponential decays).

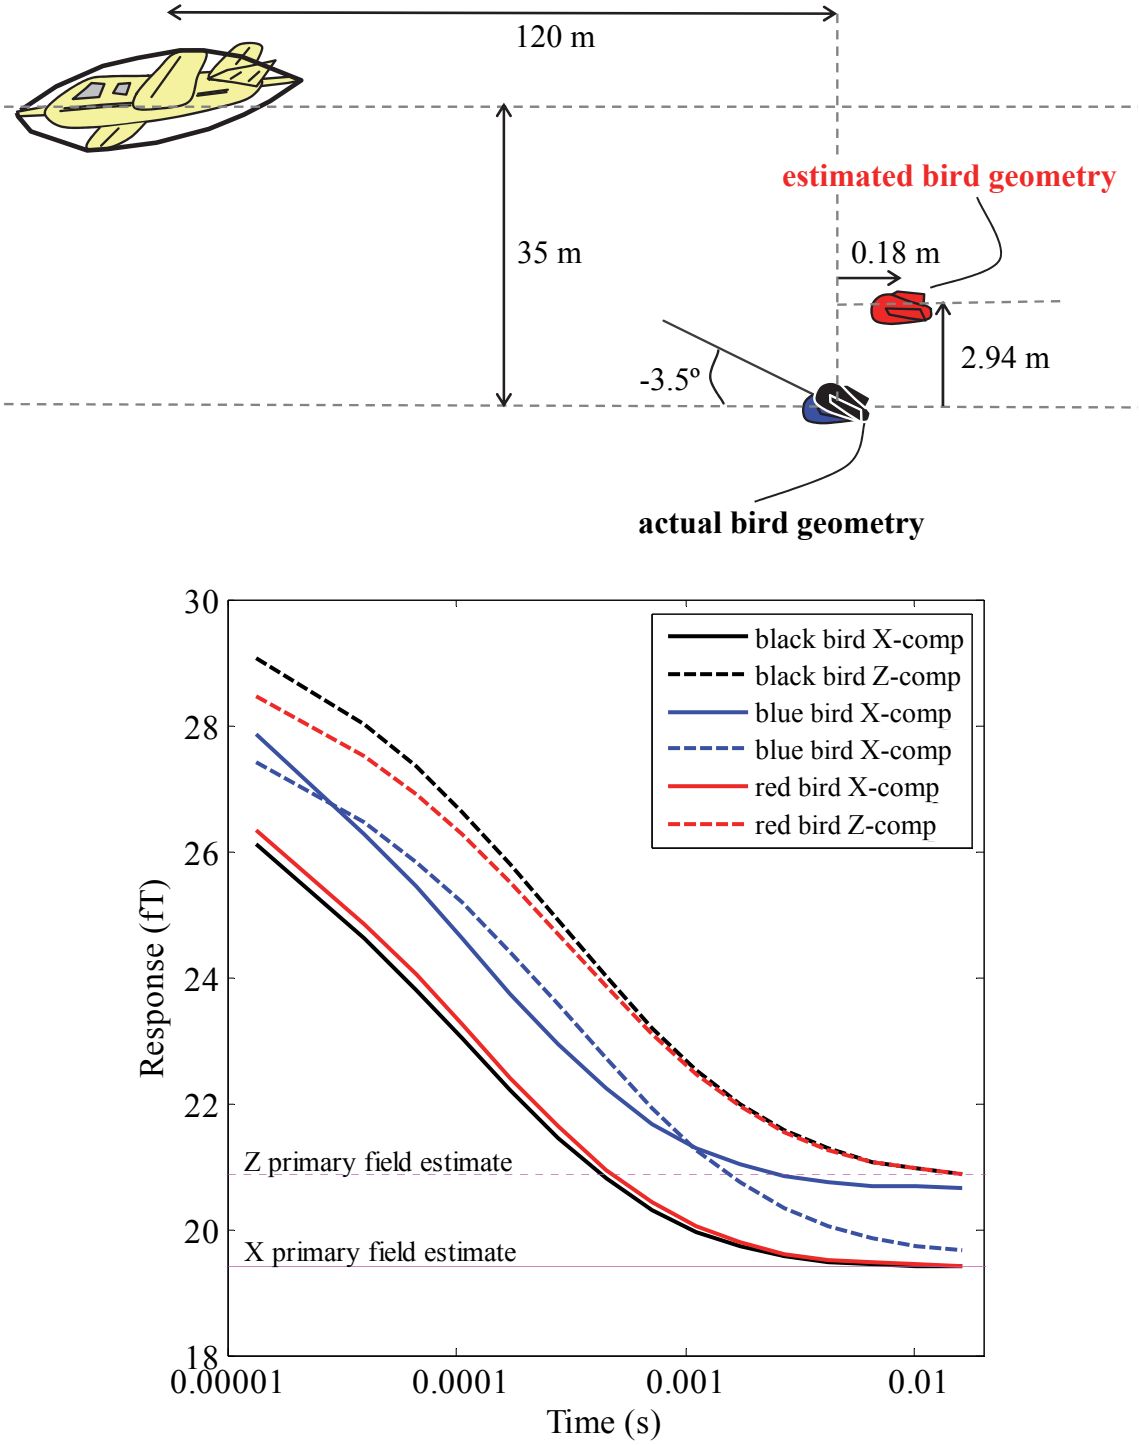

Figure 5.3 *An illustration of the primary field and receiver geometry estimation problem. The curves on the graph show the total field TEMPEST response for a 0.1 S/m halfspace for three different bird geometries, which are indicated by the same colour bird above the graph. Note that the negative Z-component data have been reversed in sign to keep all values positive.*

To help understand the effect of bird position and orientation and illustrative example is presented in Figure 5.3. The curves on the graph show the total field TEMPEST response for a 0.1 S/m halfspace conductivity model for three different bird geometries. The colour of each response corresponds to the colour of the bird shown above the graph. The black bird and the blue bird are located at the same position, but the black bird is pitched nose-up by 3.5°. It can be seen on the graph, by comparing the black with the blue responses, how the relatively small rotation changes the response markedly.

Consider now that the black bird is in the true position and orientation of the receiver. If the primary field was estimated by assuming the late time response from the ground was small, then the dashed lines marked X and Z primary field on the graph, would represent the constant primary field to be removed. The subsequent estimate of the bird position, under the assumption that the bird was flying straight and level, would be that of the red bird, i.e. 2.94 m higher and 0.18 m further back from the transmitter than its true position.

In summary, all of the methods for estimating bird position and orientation in contemporary fixed-wing towed-bird systems, first require separation of the primary and secondary fields from the total field response which is actually measured. Then the bird position can be solved for from the primary field estimate if a particular bird orientation is assumed. However, the primary field separation requires assumptions to be made about the secondary field response, and hence the ground conductivity. These assumptions mostly favour a resistive ground conductivity model, especially at depth. This can be a reasonable assumption in many cases, however it becomes less suitable as the conductivity of the ground increases and as the high-frequency content of the transmitted waveform increases. In the conductive regolith environments of Australia (Lane et al., 2004a), or over seawater (Sattel et al., 2004), this can be an especially poor assumption. In essence a resistive ground model is pre-imposed onto the processed data

during the processing, and the geometry estimates can be poor. In contrast, the holistic inversion formulation does not impose any assumptions about the earth conductivity until the final stage.

5.3.5 Parameterization of system geometry

In the previous section we have seen how in conventional data processing the primary field is estimated, from which the bird position is derived. The estimated primary field is then removed from the total field, to yield the processed secondary field data that are conventionally interpreted. When the processed secondary field data are quantitatively interpreted, using the estimated bird position and the assumed (zero roll, pitch and yaw) orientation, two problems can occur. Firstly, the recovered model will be biased toward the conductivity model that was pre-imposed by the respective primary field estimation routines. Secondly, the data may not be able to be fitted to any plausible conductivity model, or any model at all. For example, it is sometimes found with TEMPEST data that the X-component and the Z-component data can be inverted independently, but that each component produces substantially different conductivity models. Then, if inverted together no model can be found that fits both components simultaneously.

This situation is not dissimilar to the predicament faced in the inversion of frequency-domain airborne electromagnetic data. Conventionally, in the frequency-domain case, removal of calibration errors is attempted by processing the inphase and quadrature data with the goal of being able to generate spatially coherent and artefact free apparent conductivity maps from the resultant data,. Although this can produce spatially coherent apparent conductivity maps, the associated apparent depth maps may not be artefact free. This indicates that the calibration errors have not been removed, but have simply been disguised by, or absorbed into, the apparent depth maps. Also, since the processing is done independently for each coilset frequency, based on totally different pseudo-layer halfspace models, the mutual consistency between each coilset's data is

not necessarily maintained. In other words the data from each coilset are not guaranteed to be consistent with one common geological model. It is often found that when frequency-domain data are finally inverted, with all frequencies together, they cannot be fitted because of the inter-frequency inconsistencies.

Therefore, for both frequency and time-domain data, calibration and processing errors may not become apparent until all components of the data are considered together in an inversion to solve for a common geological model. The frequency-domain holistic inversion overcame this problem by replacing the sequentially applied data calibrations, processing and then inversion steps by a simultaneous procedure. Similarly, the idea of the time-domain holistic inversion is to invert the less processed total field data, and solve for the bird's position and orientation (or equivalently the primary field) as calibration parameters simultaneously with the conductivity model. Since it is not calibration errors in the electromagnetic data that are solved for per se, but rather unknown aspects of the system, 'calibration parameters' may not entirely be the best terminology. Nevertheless the term is used for analogy with the frequency-domain holistic inversion formulation.

The concept of solving for system geometry or primary field parameters during the inversion has been successfully employed before with conventional sample-by-sample inversion. Owers et al. (2001) reported on the inversion of a single component of TEMPEST data where transmitter loop height and a primary field correction were solved for. Lane et al. (2004a), also reported in Lane et al. (2004b), inverted total field TEMPEST data and solved for receiver position and pitch. Sattel et al. (2004) introduced an algorithm for three component GEOTEM data in which the transmitter pitch, receiver position and orientation, and primary field corrections were all solved for. Wolfgram and Vrbancich (2006) also inverted GEOTEM data for transmitter loop height and pitch, bird position, and primary field corrections.

This work employs a modified form of the system geometry parameterization introduced by Lane et al. (2004a) for sample-by-sample inversion, and which has since been further developed by Brodie and Fisher (2008). The parameterization is modified to be compatible with the holistic inversion philosophy, specifically to take advantage of the along-line continuity of system geometry variations.

As the aircraft moves along the flight line the system geometry varies in slow and continuous manner relative to the base period (e.g. 0.04 s) of the transmitted waveform and the stacked data output rate (e.g. 0.2 s). Annan (1984) reported that position of the bird relative to the transmitter tends to vary at periods of 10 to 100 s, and that its orientation varies at periods of 0.5 to a few seconds. Because the bird moves with pendulum like motion (Smith, 2001b) the variations are sinusoidal in nature. It seems then that one obvious way to parameterise each element of the system geometry is to use 1D cubic B-splines in an exactly analogous way as the conductivity model has been represented.

Similarly to the various sample-by-sample methods, the system geometry could be parameterized with separate geometry parameters associated with each airborne sample that is inverted. However, the along-line spline parameterization is a natural way of coupling all the airborne measurements along the flight line so that we can take advantage of the continuity that we know exists. One advantage of using a spline representation is that the true continuous form of the variations can be captured in a natural way.

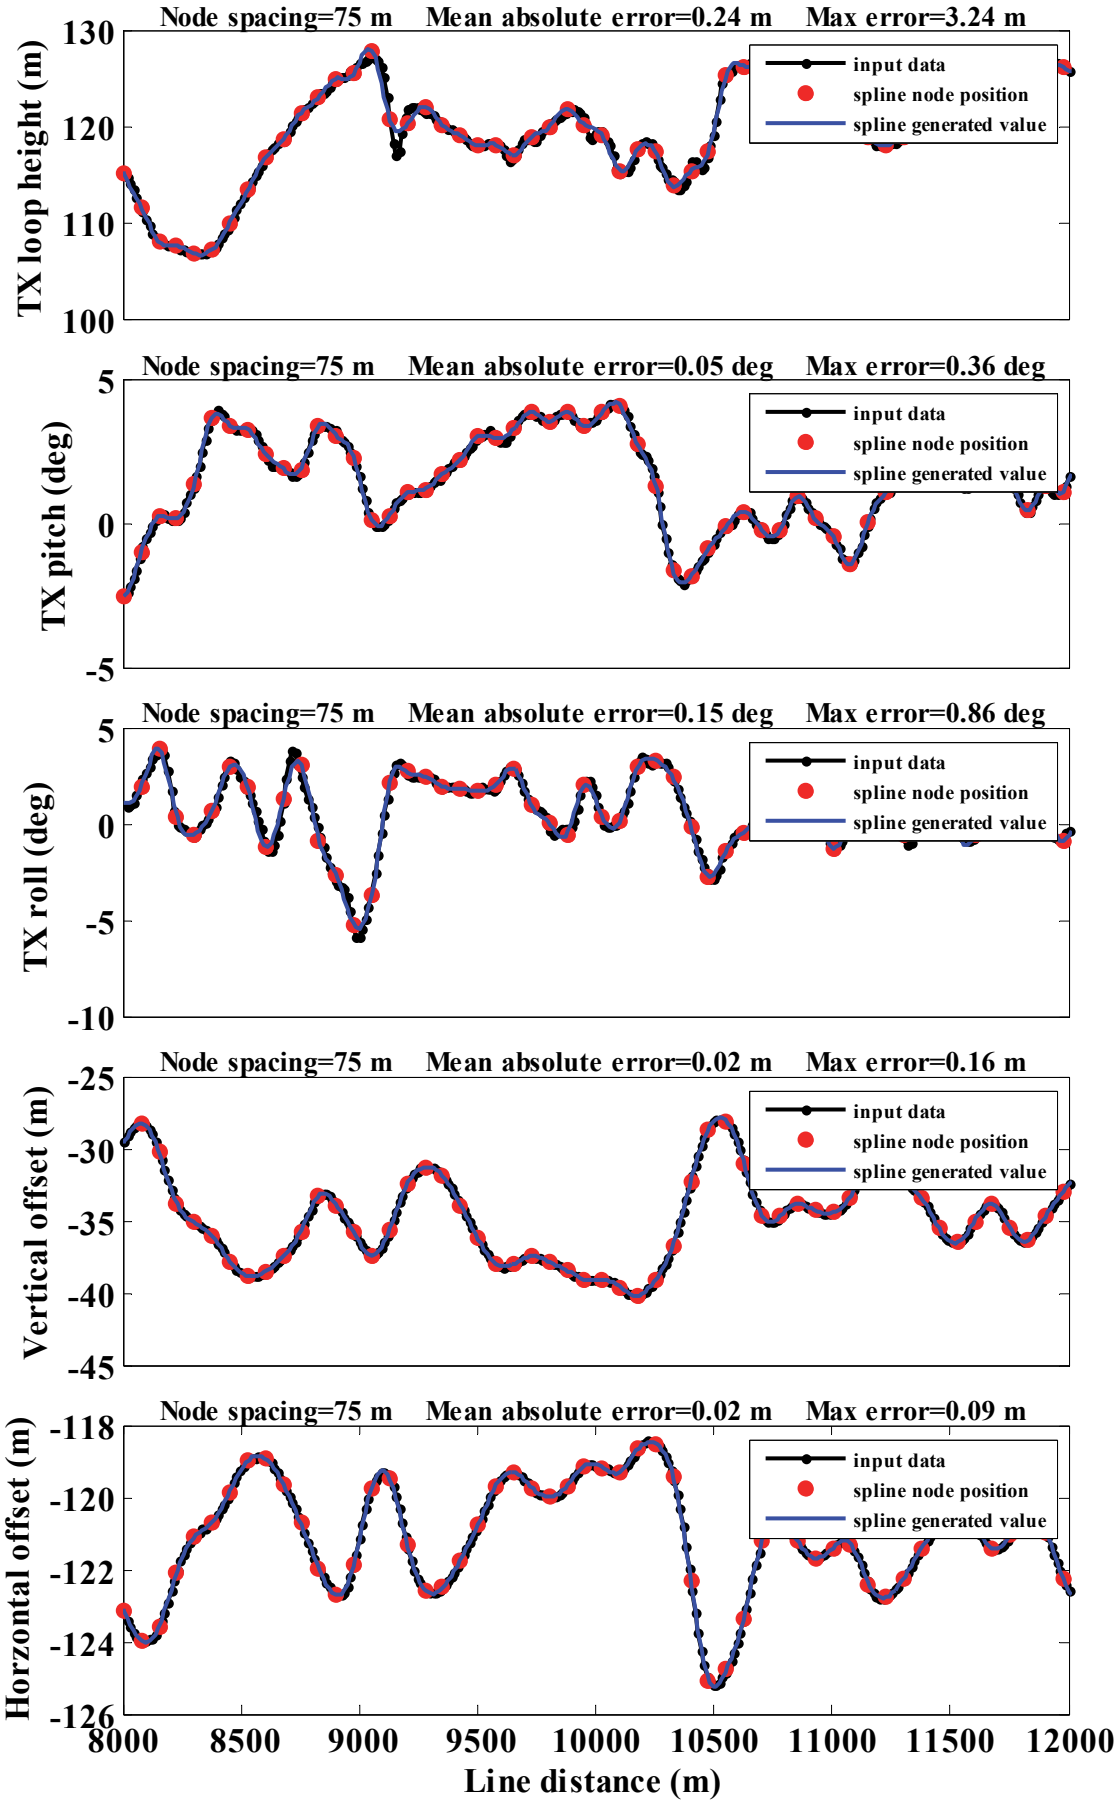

Figure 5.4 *Part of a TEMPEST survey flight line showing the geometry parameters which have been accurately represented by 1D cubic B-splines whose nodes are 75 m apart.*

To demonstrate this, in Figure 5.4 the system geometry parameters are shown for a 4 km portion of a 20 km long TEMPEST survey flight line. The figure shows the measured transmitter height, pitch and roll, and the conventional processing estimates of the receiver's horizontal and vertical offsets. The airborne sample measurements of the geometry (black dots) are approximately 12 m apart. The input data have been fitted with 1D cubic B-splines (blue line) over the whole length of the line. The spline nodes were spaced 75 m apart (red dots). Above each panel the mean and the maximum absolute differences between the input values and the spline generated values are shown.

The differences show that on average the splines are capable of representing the form of the geometry variations with sufficient accuracy, even with a node spacing at six times the airborne sample spacing. This underlines an additional advantage of a spline parameterization. Since the spline nodes can be placed much further apart than the airborne samples, the number of parameters is much reduced compared to having a set of geometry parameters for every sample, which should improve the inversion stability.

For every element of the system geometry that we wish to invert for, a separate cubic B-spline is required. They shall be denoted by $S^{\psi h}$, $S^{\psi \Delta x}$, $S^{\psi \Delta z}$ and $S^{\psi rp}$ etcetera according to the notation given Table 5.1. To avoid repetition, for the moment consider the spline $S^{\psi v}$ where v may represent one of the geometry elements to be solved for.. The spline $S^{\psi v}$ is comprised of a topologically linear array, of $N^{\psi v} - 1$ segments, which when joined together form a spline having $N^{\psi v}$ nodes that are spaced distance $h^{\psi v}$ horizontally apart. The ith node has a corresponding coefficient value $c_i^{\psi v}$. It is these coefficients that are solved for in the inversion, of which in total there are,

$$N_\psi = \sum^{\psi v} N\psi_v \quad (\psi v = \psi th, \psi tr, \psi tp, \psi ty, \psi \Delta x, \psi \Delta y, \psi \Delta z, \psi rr, \psi rp, \psi ry). \quad (5\text{-}3)$$

At any given lateral distance x along the flight line that lies on the ith segment of the spline $S^{\psi v}$, the system geometry parameter $\psi v(x)$ is parameterized by,

$$\psi v(x) = S_i^{\psi v}(x) = \sum_{a=i-1}^{i+2} w_a^{\psi v}(x)\, c_a^{\psi v}\,, \tag{5-4}$$

where $w_a^{\psi v}(x)$ is the known weight that determines the contribution of the ath spline node coefficient, $c_a^{\psi v}$ toward the spline surface value at point x. The weights are a function only of the position x along the spline.

Therefore the system geometry parameterization is exactly analogous to that of the layer conductivities and thicknesses, except that the logarithm of the variable is not used in this case. The various aspects of the system geometry parameterization are shown in Figure 5.5 for the transmitter loop to receiver coil vertical separation $\psi v(x)$. The same parameterization is repeated for each of the elements of the system geometry that are solved for, but with possibly different spline node spacings. Note that the spline nodes are not coincident with the airborne sample points. It will be explained later that when the forward model is calculated, for a particular airborne sample, a local set of system geometry values are evaluated from each spline (Equation 5-4), and then are input into the forward model function (Equation 5-12).

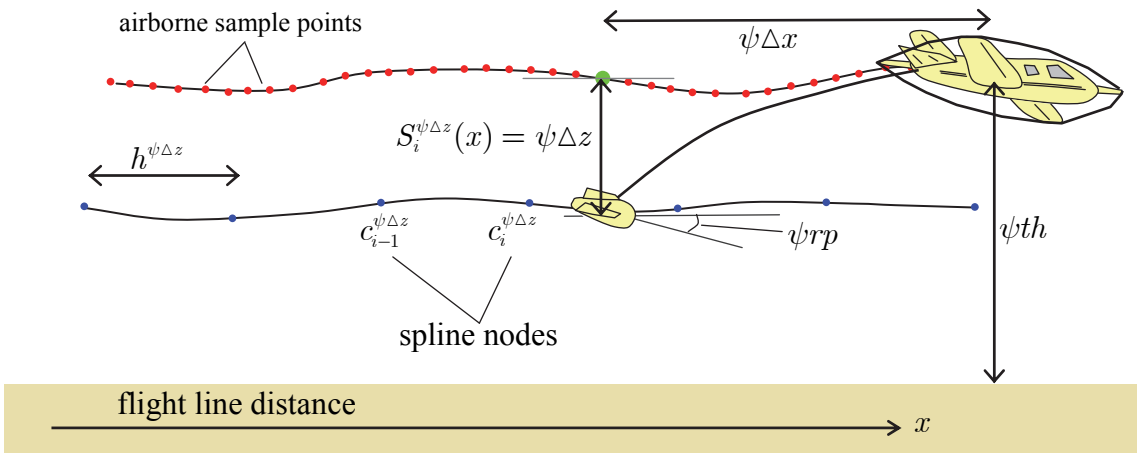

Figure 5.5 Schematic illustration of the parameterization of the transmitter-loop to receiver coils vertical offset. This style of parameterization is repeated for each element of the system geometry that is solved for.

In general there would be no need to solve for all of elements of the system geometry. In contemporary systems the transmitter height, roll, pitch and yaw are all likely to be measured to sufficient enough accuracy that there would be no advantage in solving for them. In fact, solving for them will increase the number of free parameters and may introduce irresolvable trade-offs. Furthermore when only X- and Z-component data are available, for example TEMPEST data, it will not be generally possible to resolve transmitter receiver transverse offset, receiver coils' roll or yaw because it is the Y-component data that are highly sensitive to these parameters. Therefore when inverting TEMPEST data it is only the transmitter to receiver in-line and vertical offsets and the receiver coils' pitch that are solved for. The transmitter receiver transverse offset, receiver coils' roll and yaw are all assumed to be zero.

5.4 Observed data

In the time-domain holistic inversion so far only airborne data have been included in the formulation. The reason for this is that the algorithm only inverts one flight line of data at a time to recover a 2D conductivity model. The boreholes, in which downhole conductivity log data are measured, and other ground-truth stations, are seldom coincident with flight lines. Consequently, the vast majority of flight lines will not have any ground-truth stations within a useful distance, and for those that do some interpolation of the data onto the flight line would be required. If we were inverting multiple lines to estimate a 3D model ground-truth data in between flight lines are much more useful. In that case the geoelectric and interface-depth data could be included in precisely the same manner as in the frequency-domain formulation (*cf.* Sections 3.4.2 and 3.4.3).

The form of the observed data for different time-domain platforms will necessarily change depending on the system geometry and which elements of it are measured, as well as the components of data that are recorded. Here we consider the formulation for

contemporary fixed-wing towed-bird systems that employ a three orthogonal coil receiver arrangement as shown in Figure 5.2. The data that are measured by each receiver coil, whose axis is nominally aligned with the x- (in-line), y- (transverse) and z- (vertical) coordinate system axes are called the X-, Y- and Z-component data respectively.

During the processing of the recorded data, many data reduction steps could be carried out. However, typically several half-periods of recorded data from each receiver component are stacked and binned into, so called, windowed data transients. Each component's transient consists of N_w data, one for each time window. A window being the time interval, after for example the transmitter loop current pulse is switched off, over which the electromagnetic response at the receiver coil are averaged or integrated. I will call the three components of the electromagnetic window data, and the associated auxiliary (e.g. GPS position, altimeter, gyroscope) data an airborne sample. Typically an airborne sample is drawn from the data stream every 0.1 or 0.2 seconds of flying time, which corresponds to approximately 7 to 15 m along the flight line.

As explained in Section 5.3.2, since the intention is to avoid the problems associated with the conventional processing methods of separating the measured total field into primary and secondary fields, it is the total field data that are inverted. Some contractors only supply secondary field window data to clients. However the total field window data can usually be reconstituted by calculating the primary field and then adding it back to the supplied secondary field to yield total field window data. The primary field is easily calculated via the elementary dipole formulae (Equation 2-27) using the supplied receiver position and orientation estimates.

We consider a general flight line in which the N_s airborne samples in the flight line are regularly spaced in time but irregularly spaced in distance along the flight path. The ith sample is located at distance x_i. The total field X-, Y- and Z-component

electromagnetic data for the wth window of the ith airborne sample on the line are denoted X_{iw}, Y_{iw} and Z_{iw} respectively. The complete inversion data vector has length $N_d = N_s \times N_c \times N_w$. Therefore if all three components of are available, the complete observed data vector would be written as,

$$\mathbf{d} = \left[X_{1,1} \ldots X_{1,Nw} \mid Y_{1,1} \ldots Y_{1,Nw} \mid Z_{1,1} \ldots Z_{1,Nw} \mid \ldots\ldots \mid Z_{Ns,1} \ldots Z_{Ns,Nw} \right]^{\mathbf{T}}. \tag{5-5}$$

In some systems only X and Z component data are available and accordingly the Y component would be omitted from the observed data vector $\mathbf{d}$. To get an appreciation of the size of the data vector, consider a 20 km long flight line with the airborne samples spaced on average 12.5 metres apart. The 1600 samples would result in a 48,000 length data vector for the 15 window 2 component TEMPEST system.

5.5 Forward model and partial derivative calculations

5.5.1 Outline

In this section the mathematical simulation, or prediction, of a set of data that would be expected to be observed for a given set of model parameters is described. The time-domain holistic inversion presented here employs a 2D representation of subsurface conductivity, but for computational reasons 1D layered-earth electromagnetic forward modelling theory is used to approximate the airborne response. Because of this and as for the frequency-domain formulation, the procedure for forward modelling is divided into two stages.

The first stage is the extraction of the ***local 1D conductivity profile*** from the 2D model at the location of the airborne sample. This step also involves extraction of its ***local geometry values*** for the sample. The procedure and governing equations are discussed in Sections 5.5.2 and 5.5.3. The second stage takes the local 1D conductivity profile and the local geometry values and uses them to simulate the observed airborne electromagnetic data. Section 5.5.4 details the forward modelling.

The predicted data for a given set of model parameters $\mathbf{m}$ is the combined result of the vector-valued nonlinear forward model functions for each airborne sample in the line. The predicted data vector is denoted as,

$$\boldsymbol{g}(\mathbf{m}) = \left[\boldsymbol{g}_1^{air}(\mathbf{m}) \mid \boldsymbol{g}_2^{air}(\mathbf{m}) \mid \ldots \mid \boldsymbol{g}_{Ns}^{air}(\mathbf{m}) \right]^{\mathbf{T}}, \tag{5-6}$$

where $\boldsymbol{g}_s^{air}(\mathbf{m})$ is the 1D electromagnetic forward model for the sth airborne sample on the flight line.

Since a gradient based scheme is used to iteratively minimize an objective function in the inversion, a Jacobian matrix must be constructed. It is therefore necessary to determine the partial derivatives of each datum with respect to each inversion model parameter that is to be solved for. Computation of the derivatives firstly requires the derivatives of the data with respect to the local 1D profile layer properties and local geometry values. Then the partial derivatives of the local layer properties and local calibration values with respect to the underlying conductivity and calibration model parameters are computed. These must then be combined, using the chain rule, to attain the derivatives of the data with respect to the actual model parameters that are being solved for. Analytic expressions have been derived for all required derivative combinations, and are given in Section 5.5.4.

5.5.2 Extraction of local 1D conductivity profile

Here we consider the extraction of the 1D local conductivity profile that is used in the forward modelling of a given airborne sample. The 1D conductivity profile can be thought of as the profile that would be encountered if you drill vertically down through the model at the airborne sample's position. The concepts are schematically illustrated in Figure 5.6 for the forward modelling of the airborne sample that is coloured green.

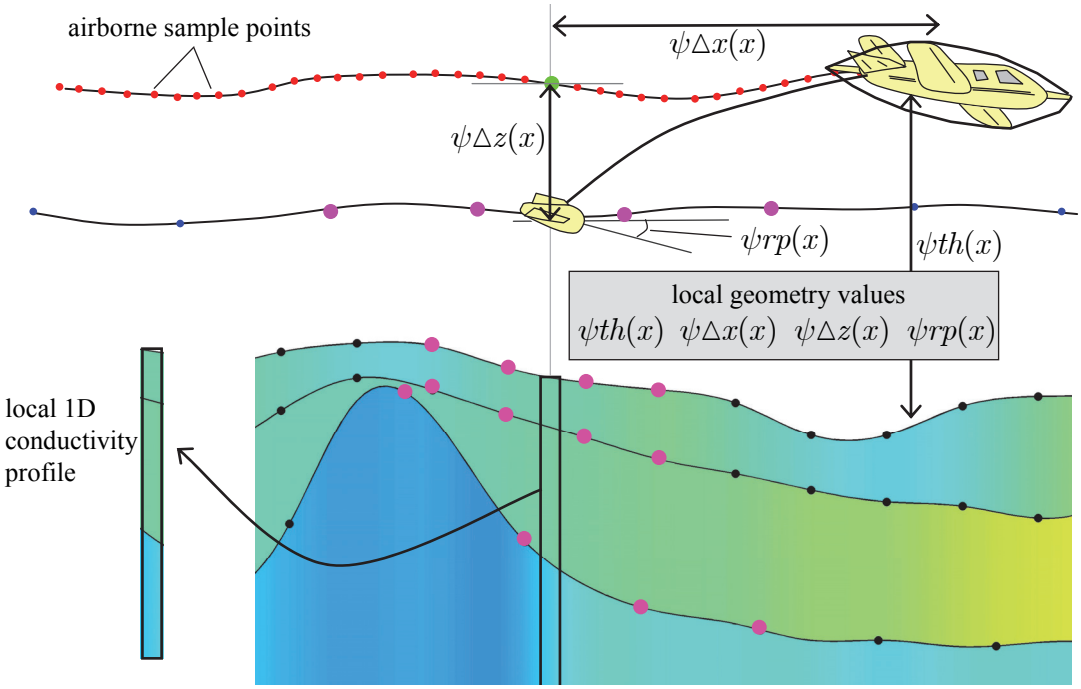

Figure 5.6 *Schematic illustration of the extraction of the 'local 1D conductivity profile' and 'local geometry values' from the model. These are subsequently used in the forward modelling of the data for the green coloured airborne sample. Only the magenta coloured spline nodes contribute to the forward modelled data for the green airborne sample.*

If we say the airborne datum is located at lateral position x, then the local 1D profile layer conductivities and thicknesses are denoted $\boldsymbol{\sigma}(x) = [\sigma_1(x), \sigma_2(x), \cdots, \sigma_{N_L}(x)]^{\mathrm{T}}$ and $\mathbf{t}(x) = [t_1(x), t_2(x), \cdots, t_{N_L-1}(x)]^{\mathrm{T}}$. Consider the case where the airborne sample position x lies between the i and $i+1$th nodes of the cubic spline that parameterizes the kth layer conductivity or thickness in the 2D model. Then, the kth layer conductivity is calculated by evaluating the cubic spline formula Equation 5-2, thus

$$\sigma_k(x) = \exp\left\{\sum_{a=i-1}^{i+2} w_{ak}^{\sigma}(x)\, c_{ak}^{\sigma}\right\}, \tag{5-7}$$

where c_{ak}^{σ} are the spline node coefficients and $w_{ak}^{\sigma}(x)$ are the known corresponding weights for the airborne sample. Similarly, the thickness $t_k(x)$ can be written as,

$$t_k(x) = \exp\left\{\sum_{a=i-1}^{i+2} w_{ak}^{t}(x)\, c_{ak}^{t}\right\}. \tag{5-8}$$

The spline node coefficients c_{ak}^{σ} and c_{ak}^{t} are elements of the model parameter vector which are solved for in the inversion. The exponentials are required because the

conductivities and thicknesses are parameterized by their logarithms. All datum that share a common lateral position, for example all components and time windows of data for an airborne sample, will share the same local 1D conductivity profile (i.e. it only has to be computed once per airborne sample).

In the following sections partial derivatives of the local layer conductivities and thicknesses, with respect to the underlying spline node coefficients, will be required. These are respectively given by,

$$\frac{\partial \sigma_k(x)}{\partial c^{\sigma}_{uk}} = \exp\left\{\sum_{a=i-1}^{i+2} w^{\sigma}_{ak}(x)\; c^{\sigma}_{ak}\right\} \times w^{\sigma}_{uk}(x) = \sigma_k(x) w^{\sigma}_{uk}(x) \tag{5-9}$$

and,

$$\frac{\partial t_k(x)}{\partial c^{t}_{uk}} = \exp\left\{\sum_{a=i-1}^{i+2} w^{t}_{ak}(x)\; c^{t}_{ak}\right\} \times w^{t}_{uk}(x) = t_k(x) w^{t}_{uk}(x)\,. \tag{5-10}$$

The weights $w^{\sigma}_{uk}(x)$ and $w^{t}_{uk}(x)$, and hence the derivatives, are nearly always zero because of the local-support of the cubic spline basis functions. The weights are non-zero only when $i-1 \le u \le i+2$, which for the green coloured airborne sample in Figure 5.6 are the magenta coloured spline nodes only. Additionally derivatives with respect to coefficients from any layer other than the *k*th layer are also always zero.

5.5.3 Extraction of local geometry values

The local geometry values corresponding to a given airborne sample are just the geometry spline values at the lateral position of the sample. We consider an airborne sample at lateral position x that lies between the i and $i+1$th nodes of the cubic spline that parameterizes a particular geometry variable. Then the local geometry value $\psi\upsilon$ is evaluated from Equation 5-4, so that,

$$\psi\upsilon(x) = \sum_{a=i-1}^{i+2} w^{\psi\upsilon}_{a}(x)\; c^{\psi\upsilon}_{a}\,, \tag{5-11}$$

where ψv represents the particular local geometry value that is required(e.g. ψth, $\psi \triangle x$, $\psi \triangle z$ or ψrp), and $c_a^{\psi v}$ are the spline node coefficients to be solved for and $w_a^{\psi v}(x)$ are the known corresponding weights for that airborne sample.

5.5.4 Forward model and partial derivatives

Consider the task of simulating the X-, Y- and Z-component data for the sth airborne sample on the flight line located at distance x_s along the flight line. Firstly the local 1D conductivity profile ($\boldsymbol{\sigma}_s = \boldsymbol{\sigma}(x_s)$ and $\mathbf{t}_s = \mathbf{t}(x_s)$) and the local geometry values are computed as described in Sections 5.5.2 and 5.5.3 respectively. Then the local conductivities, thicknesses and geometry values are input into the layered-earth forward modelling routine (Chapter 2) to calculate the theoretical forward model. The three component total field vector response for the wth window is written as,

$$\boldsymbol{g}_s^{air}(\boldsymbol{\sigma}(x_s), \mathbf{t}(x_s), \boldsymbol{\psi}(x_s)) = \left[g_{s,1}^X \ldots g_{s,Nw}^X \mid g_{s,1}^Y \ldots g_{s,Nw}^Y \mid g_{s,1}^Z \ldots g_{s,Nw}^Z \right], \tag{5-12}$$

where $\boldsymbol{\psi}(x_s)$ denotes the complete set of ten geometry variables, which include the local geometry values (e.g. $\psi th(x_s)$, $\psi \triangle x(x_s)$, $\psi \triangle z(x_s)$, and $\psi rp(x_s)$) and the remainder of geometry variables that are not being solved for.

The forward modelling procedure used to simulate the system response (Equation 5-12) is fully described in Section 2.6 of Chapter 2. The function must take into account the particular specifications of the airborne system being modelled. For example it must account for, the transmitter current waveform or receiver voltage waveform, the window positions and windowing methods, any applied normalizations for transmitter loop and receiver coil area and number of turns, and any unit conversions. The procedure also accounts for non-zero roll, pitch and yaw of the transmitter loop and the receiver coil assembly.

The predicted airborne data vector $\boldsymbol{g}(\mathbf{m})$ is then built up by repeating the procedure for all samples on the flight line, thus,

$$\boldsymbol{g}(\mathbf{m}) = \left[\boldsymbol{g}_1^{air} \mid \boldsymbol{g}_2^{air}, \ldots \mid \boldsymbol{g}_{Ns}^{air} \right]^{\mathbf{T}}. \tag{5-13}$$

In the inversion's minimization procedure the partial derivatives of each forward modelled airborne response component with respect to the inversion model parameters are required. To attain the derivative with respect to the uth node coefficient c_{uk}^{σ} of the kth layer conductivity spline, the chain rule is used in conjunction with Equation 5-9 to yield,

$$\frac{\partial g_{sw}^{*}}{\partial c_{uk}^{\sigma}} = \frac{\partial g_{sw}^{*}}{\partial \sigma_k(x_s)} \times \frac{\partial \sigma_k(x_s)}{\partial c_{uk}^{\sigma}} = \frac{\partial g_{sw}^{*}}{\partial \sigma_k(x_s)} \times \sigma_k(x_s) \times w_{uk}^{\sigma}(x_s), \tag{5-14}$$

where g_{sw}^{*} can be either g_{sw}^{X}, g_{sw}^{Y} or g_{sw}^{Z}. Using the same logic, the derivative with respect to the thickness spline nodes coefficient c_{uk}^{t} is,

$$\frac{\partial g_{sw}^{*}}{\partial c_{uk}^{t}} = \frac{\partial g_{sw}^{*}}{\partial t_k(x_s)} \times \frac{\partial t_k(x_s)}{\partial c_{uk}^{t}} = \frac{\partial g_{sw}^{*}}{\partial t_k(x_s)} \times t_k(x_s) \times w_{uk}^{t}(x_s). \tag{5-15}$$

Since the system geometry variables, unlike layer conductivity and thicknesses, are not parameterized in terms of their logarithms, the partial derivatives are simpler for the geometry parameters. The partial derivative with respect to the uth node coefficient $c_u^{\psi\upsilon}$ of the spline that parameterizes the system geometry variable $\psi\upsilon$ is attained by using the chain rule in conjunction with Equation 5-11, which gives,

$$\frac{\partial g_{sw}^{*}}{\partial c_u^{\psi\upsilon}} = \frac{\partial g_{sw}^{*}}{\partial \psi\upsilon(x_s)} \times \frac{\partial \psi\upsilon(x_s)}{\partial c_u^{\psi\upsilon}} = \frac{\partial g_{sw}^{*}}{\partial \psi\upsilon(x)} \times w_u^{\psi\upsilon}(x_s). \tag{5-16}$$

Again g_{sw}^{*} can be either g_{sw}^{X}, g_{sw}^{Y} or g_{sw}^{Z}, and $\psi\upsilon(x_s)$ can be any of the system geometry variables that are being solved for (e.g. $\psi th(x_s)$, $\psi\triangle x(x_s)$, $\psi\triangle z(x_s)$, and $\psi rp(x_s)$).

The partial derivatives $\partial g_{sw}^{*} / \partial \sigma_k(x_s)$, $\partial g_{sw}^{*} / \partial t_k(x_s)$, and $\partial g_{sw}^{*} / \partial \psi\upsilon(x_s)$, of the 1D electromagnetic forward model functions that are required for evaluation of Equations 5-14, 5-15, and 5-16 can all be calculated using the analytic partial derivative expressions that are detailed in Appendix B.

5.6 Inversion scheme

5.6.1 Outline

The inversion scheme used for the time-domain holistic inversion is similar to that developed for the frequency-domain method which was extensively discussed in Section 3.6 of Chapter 3. Therefore, in this chapter, where possible the frequency-domain development will not be repeated. Instead emphasis will be placed on the differences.

As an overview, in the time-domain inversion an initial set of model parameters is iteratively updated until there is an acceptable level of agreement between data predicted from the model and the observed data. The scheme is designed to minimize an objective function that is comprised of data misfit and model regularization terms. The data misfit term guides the inversion toward the principal goal of matching the observed and predicted data to within expected noise levels. The model regularization term is required in the objective function because there are usually many alternate models that would allow the data to be adequately fitted. In these ambiguous cases, the solution is guided toward the models that are closer to an a priori reference model and are spatially the smoothest. Lagrange multipliers are used to weight the relative influence of the model regularization terms.

The iterative procedure begins with an initial set of inversion model parameters (spline coefficients) that map to a set of reference values (layer properties and system geometry values). Like the frequency-domain method, the procedure consists of two logical loops (*cf.* Figure 3.8). The outer loop iteratively updates the both the conductivity and system geometry model parameters. It constructs a linearized system of equations that are designed such that, when solved, they yield a new set of model parameters whose corresponding objective function value is smaller than that corresponding to the original

parameters at the start of the loop. It is in the inner loop that the linear system is actually solved via an iterative preconditioned conjugate gradient algorithm.

In contrast to the frequency-domain formulation, in the time-domain method each of the individual penalty terms that make up the objective function (i.e. data misfit, deviation from the reference model, and roughness) are normalized to account for the number of data, parameters, or constraints involved. In the frequency-domain method multiple fixed model regularization parameters ($\lambda_r, \lambda_x, \lambda_y, \lambda_z$, and λ_b) were used to weight the relative influence of the data misfit and model regularization penalty terms in the objective function. In the time-domain method a single, but variable from iteration to iteration, parameter λ weights the relative influence of the data misfit and the composite model regularization penalty term. The value of λ for each iteration is chosen, via an automated line search, in such a way that the data misfit is only gradually reduced at each iteration. The line search employs both bisection and golden search methods. These changes do not alter the general philosophy of the inversion, but have been made to simplify and automate the choice of regularization parameters at a practical level.

5.6.2 Objective function definition

The inversion scheme seeks to minimize an objective function of the form,

$$\Phi(\mathbf{m}) = \Phi_d(\mathbf{m}) + \lambda \left[\alpha_r \Phi_r(\mathbf{m}) + \alpha_h \Phi_h(\mathbf{m}) + \alpha_v \Phi_v(\mathbf{m}) \right], \tag{5-17}$$

where $\Phi_d(\mathbf{m})$ is a data misfit term, $\Phi_r(\mathbf{m})$ is a reference-model misfit term, $\Phi_h(\mathbf{m})$ and $\Phi_v(\mathbf{m})$ are horizontal (along-line) and vertical conductivity model roughness terms respectively. The model regularization parameter λ weights the relative importance of the data misfit term and the combined model regularization term ($\alpha_r \Phi_r(\mathbf{m}) + \alpha_h \Phi_h(\mathbf{m}) + \alpha_v \Phi_v(\mathbf{m})$). The α's weight the relative influence of the three individual components of the model regularization term. The exact definition of each term is detailed in the following sub sections.

In Equation 5-17's frequency-domain counterpart, Equation 3-45, each individual model regularization term is weighted by different λ values rather than by a common λ and individual α values. Therefore, the λ_r in the frequency-domain formulation is analogous to the $\lambda\alpha_r$ here, and similarly for the other model regularization terms (i.e. $\lambda_x \approx \lambda\alpha_x$ etc.). This change was made to simplify and automate the choice of model regularization terms. During the inversion the relative influence of the three model regularization terms is kept fixed by keeping α_r, α_v and α_h constant. On the other hand, the relative influence of data misfit and model regularization terms can be varied by changing the single parameter λ. This allows an automated method devised by Constable et al. (1987) of selecting λ in each iteration to be implemented. As explained later, the automated method involves a line search (single parameter search) on λ to find its value that incrementally reduces the data misfit to a predefined large fraction of its previous value in each iteration. If the objective function was expressed with no coupling between the model regularization terms, as it was in the frequency-domain formulation, an automated technique would necessarily involve a multi-parameter search on several λ's. A multi-parameter search is unlikely to be computationally feasible, and would not provide the practitioner with control over the relative influence of the three model regularization terms.

The model regularization objective function terms are a subset of the broader class of L_2 model-structure measures, which are widely used in geophysical inversion to constrain solutions (e.g. Constable et al., 1987; Menke, 1989; Farquharson and Oldenburg, 1998). As explained in relation to the frequency-domain formulation, a nuance in the way they are used in the holistic inversion, is that the constraints are not imposed directly on the inversion model parameters. The justification for this is that it is not particularly physically appealing to directly place constraints on the somewhat abstract spline coefficient parameters. Instead the decision has been made to impose the constraints on the more physically meaningful layer properties that the coefficients parameterize.

In principle, it would be possible to include a further model regularization term that penalizes along-line roughness of the system geometry variables. However so far this has not been found to be necessary in practice and has not been implemented. In the frequency-domain method a bias roughness penalty term (Φ_b) was included to encourage temporal smoothness of the bias calibration which was inverted for. A similar concept could in the time-domain method if it were extended, as is proposed in Section 5.8, to include zero-level drift calibration parameters.

Data misfit

The data misfit Φ_d is defined in the usual noise weighted least-squares sense,

$$\Phi_d(\mathbf{m}) = \frac{1}{N_d}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m})\right]^{\mathbf{T}} \mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m})\right]. \tag{5-18}$$

Here $\mathbf{d}$ is the N_d length vector of observed total field airborne data. The vector $\mathbf{m} = (\mathbf{m}^{ear} \mid \mathbf{m}^{geom})^{\mathbf{T}}$ is the N_m length vector of unknown conductivity model and system geometry spline coefficient parameters to be solved for. The function $\boldsymbol{g}(\mathbf{m})$ is the nonlinear vector-valued forward model function that maps the model parameters to predictions of data (i.e. Equations 5-13). The $N_d \times N_d$ matrix $\mathbf{C}_d$ expresses covariance of the errors expected in the observed data. If the errors are assumed to be uncorrelated and normally distributed then $\mathbf{C}_d\Big|_{ii}$ is the variance of the error associated with the ith datum, and all off-diagonal entries of $\mathbf{C}_d$ are zero.

Conceptually Equation 5-18 is the same as its frequency-domain method counterpart, Equation 3-46, except that in this case it is normalized by the number of data. It will also be seen in the following sections that the model regularization terms are similarly normalized by the number of reference values and constraint equations involved in each term. The normalizations do not represent a fundamental change in the philosophy of the inversion because numerically they can be offset by changes in the λ or α values. However, operationally they provide a more convenient formulation because they allow

users to alter the relative number of data, parameters and constraints (e.g. by choosing a different data sub-sampling rate, changing the spline node spacing or changing the number of layers), without unduly effecting the relative influence of the objective function terms.

Reference model misfit

The reference model misfit part of the objective function Φ_r allows constraint to be imposed on the solution via prior expectation of the conductivity structure that is likely to be encountered in the survey area, and the likely system geometry. I refer to these as probabilistic, or soft, constraints because they do not set hard bounds on model parameters. Nor are not they intended to mimic precise measurements.

The intention of the reference model regularisation is to encourage the solution toward a model that, subject to fitting the data, is as close as possible to an a priori conductivity and system geometry reference model. An appropriate penalty function to be minimised is a least squares difference measure, weighted by the prior uncertainty on the reference values, which can be is written as,

$$\Phi_r = \frac{1}{N_r}\left[\mathbf{r}_0 - \mathbf{Sm}\right]^{\mathrm{T}} \mathbf{C}_r^{-1}\left[\mathbf{r}_0 - \mathbf{Sm}\right]. \tag{5-19}$$

Here $\mathbf{r}^0$ is the vector of N_r reference values selected from the prior expectations. As explained in detail in relation to Equation 3-48, the reference values are chosen to represent the physically meaningful logarithms of layer conductivities, logarithms of layer thicknesses, and system geometry values, instead of the spline coefficients. This is because it is more intuitive and meaningful to express reference values and uncertainties in terms of real physical quantities rather than the somewhat abstract spline coefficient values. Each reference value has an associated position, which for simplicity is chosen to coincide with the spline node locations (i.e. $N_r = N_m$). A prior level of uncertainty can be placed on the reference values via a $N_r \times N_r$ covariance

matrix $\mathbf{C}_r$. This would typically be a diagonal matrix in which $\mathbf{C}_r\big|_{ii}$ is the variance of the uncertainty on the ith reference value.

It can also be seen from Equation 5-19 that this time the reference model objective function is normalized by the number of reference values, or equivalently the number of constraint equations. The matrix $\mathbf{S}$ is a $N_r \times N_m$ linear operator that maps the underlying model parameters $\mathbf{m}^0$ to the elements of $\mathbf{r}^0$, such that,

$$\mathbf{S}\mathbf{m}^0 = \mathbf{r}^0. \tag{5-20}$$

The operator $\mathbf{S}$ is sparse and known by definition of the model. In other words, the i,jth entry in $\mathbf{S}$ is the weight that would be assigned to the jth spline coefficient model parameter when evaluating the spline corresponding to the ith reference value.

The prior expected reference values ($\mathbf{r}^0$) might be chosen, for example, from statistical analysis of downhole conductivity log data from the survey area or from a geological environment known to be similar. The altimeter measured transmitter loop height is likely to be used as a reference if ψh is solved for. The continuously varying transmitter-receiver offsets estimated from survey altitude data during routine processing could be used as reference values for $\psi \triangle x$ and $\psi \triangle z$. Since estimates of transmitter-receiver offset derived from high altitude flight data are not affected by incorrect separation of the primary and secondary fields, a better approach may be to use the average offsets derived from the high altitude data.

If receiver coil pitch (ψrp) is solved for it is likely that its reference values would be set to zero to reflect an expectation that on average the receiver bird will fly level (i.e. with zero pitch). However it should be noted that even if the receiver bird shell on average flies level, the coil assembly inside the shell is not necessarily aligned precisely with the shell itself. Inversion work carried out by (Lane et al., 2004a) suggests that this may have been the case for at least one TEMPEST survey. If the bird is not expected to fly

level, or the coils are not aligned with the shell, non-zero reference values would be appropriate if the expected deviations from zero were known.

Conductivity model lateral roughness

Although the 2D conductivity model is locally smooth in the along line direction by virtue of the cubic B-spline parameterization, a greater degree of smoothness can be enforced by including conductivity model lateral roughness regularization. This regularization is implemented in the same manner as in the frequency-domain method. However, since we are only dealing with a 2D model the two lateral roughness measures (Φ_x and Φ_y) terms are replaced by a single along-line roughness measure Φ_h. Also, so far in the time-domain method development only the lateral roughness measure defined relative to layer number, rather than elevation above sea level, has been developed and implemented. A further difference is in the normalization of the roughness term as will be shown below.

The lateral roughness term Φ_h is a measure of the cumulative lateral roughness over the whole model. Again, because it is the layer properties and not the coefficients that we wish to be smooth, the roughness measure is based on the logarithms of the layer properties that the model coefficients represent rather than the coefficients themselves. The roughness is approximated by using a second finite difference roughening matrix. Consider now the second derivative approximation for the triplet centred on the ith node of the kth layer property spline (either conductivity or thickness). The logarithm of the layer property value at that intersection is just one of the elements of the vector $\mathbf{r} = \mathbf{Sm}$, which is denoted $r_{i,k}$. Then the second vertical finite difference approximation is,

$$l_{ik}^{h} = \frac{r_{i+1,k} - r_{i,k}}{\triangle h_k} - \frac{r_{i,k} - r_{i-1,k}}{\triangle h_k} = \frac{1}{\triangle h_k}[r_{i-1,k} - 2r_{i,k} + r_{i+1,k}], \tag{5-21}$$

where $\triangle h_k$ is the spline node spacing for the *k*th layer property. Therefore a vector of derivatives $\mathbf{l}_h$ can be built up by the product,

$$\mathbf{l}_h = \mathbf{L}_h\mathbf{r} = \mathbf{L}_h\mathbf{S}\mathbf{m}\,, \tag{5-22}$$

where $\mathbf{L}_h$ is a $N_h \times N_m$ roughening matrix of the form,

$$\mathbf{L}_h = \begin{bmatrix} \ddots & & & & & & & & & & \\ & \frac{1}{\triangle h_k} & \cdots & \frac{-2}{\triangle h_k} & \cdots & \frac{1}{\triangle h_k} & & & & & \\ & & & \frac{1}{\triangle h_k} & \cdots & \frac{-2}{\triangle h_k} & \cdots & \frac{1}{\triangle h_k} & & & \\ & & & & \cdots & \frac{1}{\triangle h_{k+1}} & \cdots & \frac{-2}{\triangle h_{k+1}} & \cdots & \frac{1}{\triangle h_{k+1}} & \\ & & & & & & & & & & \ddots \end{bmatrix}. \tag{5-23}$$

Each row of $\mathbf{L}_h$ corresponds to a triplet of adjacent spline nodes in the flight line direction. The number of rows is equal to the total number of spline nodes less those at the beginning and end of the line, because differences cannot be applied on end nodes. If there is no motivation to penalise roughness of the layer thicknesses, then the corresponding rows are simply omitted from the roughening matrix. The precise column positioning of the nonzero $\left[1/\triangle h_k \quad \cdots \quad -2/\triangle h_k \quad \cdots \quad 1/\triangle h_k\right]$ entries in each row depends on the ordering of the coefficients in the model parameter vector $\mathbf{m}$. The overall model roughness in the lateral direction to be minimised is,

$$\Phi_h = \frac{1}{N_h}\mathbf{l}_h{}^{\mathrm{T}}\mathbf{l}_h = \frac{1}{N_h}\mathbf{m}^{\mathrm{T}}\mathbf{S}^{\mathrm{T}}\mathbf{L}_h^{\mathrm{T}}\mathbf{L}_h\mathbf{S}\mathbf{m}\,. \tag{5-24}$$

By comparing Equation 5-24 to Equation 3-52, it can be seen that the lateral roughness term is analogous to the frequency-domain method's equivalent. However here the term is normalized by the number of constraints (i.e. number of rows of $\mathbf{L}_h$). As explained in Chapter 3, there are a number of alternative measures of model-structure

that could be used if it was deemed appropriate for a particular survey area. For example, rather than promote smoothness measured by second derivatives, it may be more appropriate to promote flatness defined by first derivatives (*cf.* Menke, 1989, p. 53). In some circumstances it might be considered appropriate to encourage smoothness or flatness of the model's deviation from the reference model, in which case an expressions in the form $\Phi_h = (1/N_h)(\mathbf{m}-\mathbf{m}^0)^{\mathrm{T}}\mathbf{S}^{\mathrm{T}}\mathbf{L}_h^{\mathrm{T}}\mathbf{L}_h\mathbf{S}(\mathbf{m}-\mathbf{m}^0)$ would be used in Equation 5-24 instead.

Conductivity model vertical roughness

The vertical conductivity profile at any position along the 2D model is not implicitly smooth because each layer conductivity and thickness is independently parameterized. When inverting for a few-layer (e.g. 3 to 5 layers) variable thickness conductivity model, where both the layer conductivities and thicknesses are solved for, it is typically acceptable and/or necessary to allow a vertically rough model so that the data can be adequately fitted. However, if a multi-layer fixed thickness conductivity model is inverted for, it is usually necessary to impose constraints on the roughness of the vertical conductivity profile. Since it does not make sense to apply vertical roughness constraints to layer thicknesses, this discussion only applies to the vertical roughness of the model conductivity.

The vertical roughness term Φ_v is a measure of the cumulative vertical roughness over the whole model which is approximated by using a second finite difference roughening matrix. Consider now the second vertical derivative calculation for the triplet centred on the ith node of the kth layer conductivity spline. The logarithm of the layer conductivity value at that intersection is just one of the elements of the vector $\mathbf{r} = \mathbf{Sm}$, which is denoted $r_{i,k}$. Then the second vertical derivative finite difference approximation is,

$$
\begin{aligned}
l_{ik}^{v} &= \frac{r_{i,k+1} - r_{i,k}}{0.5(t_{i,k} + t_{i,k+1})} - \frac{r_{i,k} - r_{i,k-1}}{0.5(t_{i,k-1} + t_{i,k})} \\
&= \frac{2r_{i,k-1}}{(t_{i,k-1} + t_{i,k})} - 2r_{i,k}\left[\frac{1}{(t_{i,k-1} + t_{i,k})} + \frac{1}{(t_{i,k} + t_{i,k+1})}\right] + \frac{2r_{i,k+1}}{(t_{i,k} + t_{i,k+1})}
\end{aligned}, \tag{5-25}
$$

where $t_{i,k}$ is the thickness of the kth layer at the ith spline node. A vector of derivatives $\mathbf{l}_v$ can be formed by the product,

$$
\mathbf{l}_v = \mathbf{L}_v\mathbf{r} = \mathbf{L}_v\mathbf{S}\mathbf{m}, \tag{5-26}
$$

where $\mathbf{L}_v$ is a $N_v \times N_m$ roughening matrix of the form,

$$
\mathbf{L}_v = \begin{bmatrix}
\ddots & & \ddots & & \ddots & \\
& \dfrac{2}{(t_{i,k-1} + t_{i,k})} & \cdots & \left[\dfrac{-2}{(t_{i,k-1} + t_{i,k})} + \dfrac{-2}{(t_{i,k} + t_{i,k+1})}\right] & \cdots & \dfrac{2}{(t_{i,k} + t_{i,k+1})} \\
& & \ddots & & \ddots & & \ddots
\end{bmatrix}. \tag{5-27}
$$

Each row of $\mathbf{L}_v$ corresponds to a triplet of vertically adjacent spline nodes. The number of rows is equal to the total number of conductivity model spline nodes less those on the top and bottom layers, because finite differences cannot be applied on end nodes. Again the precise column positioning of the nonzero entries in each row depends on the ordering of the coefficients in the model parameter vector $\mathbf{m}$.

The overall model roughness in the vertical direction to be minimised is,

$$
\Phi_v = \frac{1}{N_v}\mathbf{l}_v^{\mathrm{T}}\mathbf{l}_v = \frac{1}{N_v}\mathbf{m}^{\mathrm{T}}\mathbf{S}^{\mathrm{T}}\mathbf{L}_v^{\mathrm{T}}\mathbf{L}_v\mathbf{S}\mathbf{m}. \tag{5-28}
$$

As in the case of the horizontal roughness, this term is analogous to the frequency-domain method's equivalent, except here the term is normalized by the number of constraints (i.e. number of rows of $\mathbf{L}_v$). Also in this case the layer thickness has been included into the vertical difference calculations, which they were not in the frequency-domain method.

Since the roughness has been calculated by the second finite differences across vertically adjacent spline node triplets, it imposes a limitation on the current formulation. It means that when vertical roughness constraints are to be imposed, the node spacing must be the same for every layer conductivity spline. This limitation could be avoided by choosing to calculate the second differences at positions other than at spline nodes positions, for example at some regular spacing.

5.6.3 Objective function minimization

Since the relationship between the model parameters $\mathbf{m}$ and the observed data $\mathbf{d}$ is nonlinear, the forward model function $\boldsymbol{g}(\mathbf{m})$ and $\Phi(\mathbf{m})$ are also nonlinear with respect to $\mathbf{m}$. An iterative method must be used to minimize the objective function. The scheme adopted here is based on a gradient-based minimization method devised by Constable et al. (1987) and that has also been used, for example, by Farquharson and Oldenburg (1993) for time-domain electromagnetic inversion. The scheme is designed to minimize the objective function $\Phi(\mathbf{m})$, while also ensuring that the data misfit $\Phi_d(\mathbf{m})$ is gradually reduced throughout the inversion. The rationale for reducing the data misfit gradually is to prevent unnecessary model structure (i.e. reference model misfit and roughness) from entering into the solution and also to avoid instability caused by the nonlinearity of the optimized function.

In this scheme the single parameter λ that weights the relative influence of the data misfit and the combined model regularization penalty terms is varied from iteration to iteration. It is selected in each iteration by an automated line search so as to ensure the desired gradual reduction in the data misfit. The values of the three parameters that weight the relative influence of the individual components of the combined model regularization term (α_r, α_h, and α_v) are selected at the beginning of the inversion and they are all left fixed throughout the inversion. The α's are selected manually by the practitioner and must be made by assessing the resulting models.

The iterative scheme proceeds as follows. An initial estimate of the model parameters $\mathbf{m}^0$ is made by solving the linear system,

$$\mathbf{S}\mathbf{m}^0 = \mathbf{r}^0, \tag{5-29}$$

using the conjugate gradient algorithm (*cf.* Section 3.6.4). In essence this sets the initial model parameters so that they would reproduce the a priori reference values.

In the $n+1$th iteration the current model parameters $\mathbf{m}^n$ are updated to a new set of model parameters $\mathbf{m}^{n+1}$. The goal is to reduce current overall objective function value $\Phi(\mathbf{m}^n)$, while also ensuring that the data misfit is also reduced to a target data misfit Φ_d^{target} that is a fraction κ of its previous value, i.e.,

$$\Phi_d(\mathbf{m}^{n+1}) \approx \Phi_d^{target} = \max[\kappa\,\Phi_d(\mathbf{m}^n)\,,\,\Phi_d^{\min}]. \tag{5-30}$$

where $\Phi_d^{\min}$ represents and acceptable level of data misfit. The value of κ is set to a large fraction, for example $\kappa = 0.7$ was used in the example shown in Section 5.7. This is so that a gradual reduction in data misfit, rather than rapid reduction which risks unnecessary model structure being introduced into the solution, is promoted.

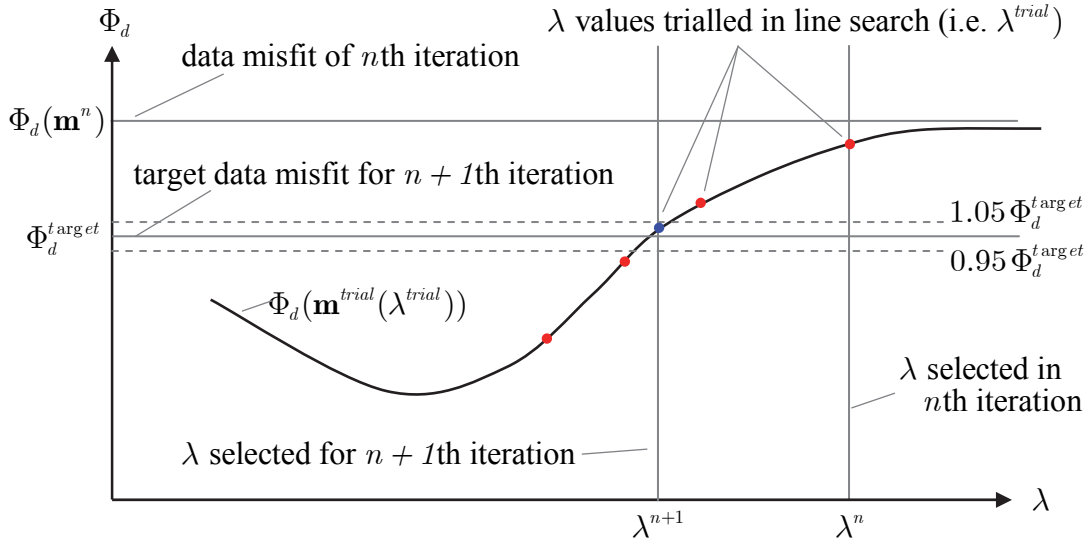

Figure 5.7 *Schematic illustration of the line search for an appropriate value of λ in the n+1th iteration. The line search progressively trials different values of λ (red and blue dots), each time solving for new model parameters and checking the corresponding data misfit, until a value suitably close to the target data misfit is found (blue dot).*

A particular value $\lambda = \lambda^{n+1}$ must be chosen for the model regularization parameter to be used for the current iteration. It is determined by an automated line search routine that progressively trials different values λ. The procedure is schematically illustrated in Figure 5.7. In each trial a value λ^{trial} is chosen and a corresponding set of trial model parameters $\mathbf{m}^{trial}$ are determined by solving the linear system (Equation 5-39) derived in Section 5.6.4. The method of solving for $\mathbf{m}^{trial}$, for a particular λ^{trial}, is conceptually identical to that used in the frequency-domain method (*cf.* Section 3.6.3).

Once a particular value of λ^{trial} is selected and the trial model parameters $\mathbf{m}^{trial}$ have been solved for, the corresponding objective function $\Phi(\mathbf{m}^{trial})$ and data misfit $\Phi_d(\mathbf{m}^{trial})$ are computed. Different values of λ^{trial} are trialled until the overall objective function is reduced **and** the data misfit has been reduced to a value sufficiently close to the target data misfit, i.e. until,

$$\begin{cases} \Phi(\mathbf{m}^{trial}, \lambda^{trial}) < \Phi(\mathbf{m}^{n}), \text{ and} \\ 0.95\,\Phi_d^{target} < \Phi_d(\mathbf{m}^{trial}, \lambda^{trial}) < 1.05\,\Phi_d^{target} \end{cases}. \tag{5-31}$$

Once a λ^{trial} and corresponding $\mathbf{m}^{trial}$ that satisfy the conditions set out in Equation 5-31 have been found, then $\lambda^{n+1} = \lambda^{trial}$ is adopted as the model regularization parameter for the $n+1$th iteration and the updated model parameters become $\mathbf{m}^{n+1} = \mathbf{m}^{trial}$.

The search routine uses a combination of bracketing, bisection search and golden search algorithms (Press et al., 2002) to methodically isolate a suitable value of λ^{trial}. It begins by trialling the value of λ from the previous iteration, or a large value if it the first iteration (e.g. 1×10^8). If this is not suitable, than the target data misfit Φ_d^{target} is enclosed using the bracketing routine to find values of λ^{low} and λ^{high} such that $\Phi_d(\mathbf{m}^{low}, \lambda^{low}) < \Phi_d^{target} < \Phi_d(\mathbf{m}^{high}, \lambda^{high})$. Once bracketed, a bisection search, starting with the values λ^{low} and λ^{high} is performed to find the value of λ^{bis} such that $0.95\,\Phi_d^{target} < \Phi_d(\mathbf{m}^{bis}, \lambda^{bis}) < 1.05\,\Phi_d^{target}$. In some iterations the data misfit may be

greater than the target data misfit Φ_d^{target} for all values of λ and thus the target cannot be bracketed. In that case a golden section search is carried out to find a value λ^{gold} that results in the minimum achievable data misfit for the iteration, so that $\Phi_d(\mathbf{m}^{gold}, \lambda^{gold}) = \min[\Phi_d(\mathbf{m}, \lambda)]$.

The iterative procedure continues until an acceptable level of misfit $\Phi_d^{\min}$ between the observed and predicted data is achieved. An acceptable level of misfit occurs when the normalized chi-squared data misfit is unity, i.e. the inversion terminates when,

$$\Phi_d(\mathbf{m}^{n+1}) \leq \Phi_d^{\min} = \chi^2 = \frac{1}{N_d}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^{n+1})\right]^{\mathbf{T}} \mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^{n+1})\right] = 1. \tag{5-32}$$

A further termination condition is triggered when the improvement in the data misfit between successive iterations is less than a specified percentage minimum $p^{\min}$, that is,

$$100 \times \left[\frac{\Phi_d(\mathbf{m}^n) - \Phi_d(\mathbf{m}^{n+1})}{\Phi_d(\mathbf{m}^n)}\right] < p^{\min}.$$

Typically a value of $p^{\min} = 1\%$ is used.

Before moving on to an example application of the time-domain holistic inversion, the following section gives the details of the derivation of the linear system used to solve for the trial model parameters $\mathbf{m}^{trial}$ that correspond to a given value λ^{trial} trialled by the line search routine.

5.6.4 Solving for the trial model parameters

The method of determining the trial model parameters $\mathbf{m}^{trial}$, for a given trial value of the model regularization parameter λ^{trial}, is similar to that used in the frequency-domain algorithm described Chapter 3. Because of the similarities, the necessary formulae can be derived by inspection of the corresponding equations in Section 3.6.3. This section will summarize the minimization scheme but the reader should refer to Section 3.6.3 for

a step by step derivation. It should be noted however that in this section $\mathbf{m}^{n+1}$ is replaced by $\mathbf{m}^{trial}$.

To make an appropriate choice of the unknown trial model parameters we make use of the knowledge that the derivative of the objective function with respect to the new model parameters $\mathbf{m}^{trial}$ will be zero at the minimum. That condition is,

$$\frac{\partial \Phi}{\partial \mathbf{m}^{trial}} = \frac{\partial \Phi_d}{\partial \mathbf{m}^{trial}} + \lambda^{trial} \left[\alpha_r \frac{\partial \Phi_r}{\partial \mathbf{m}^{trial}} + \alpha_h \frac{\partial \Phi_h}{\partial \mathbf{m}^{trial}} + \alpha_v \frac{\partial \Phi_v}{\partial \mathbf{m}^{trial}} \right] = 0 . \tag{5-33}$$

To proceed, the individual parts of Equation 5-33 must be determined. By analogy with the derivation in Chapter 3 (*cf.* Equations 3-59 to 3-63) the derivative of the first term is,

$$\frac{\partial \Phi_d}{\partial \mathbf{m}^{trial}} = -\frac{2}{N_d} \mathbf{G}_n^{\mathrm{T}} \mathbf{C}_d^{-1} \left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^n) - \mathbf{G}_n (\mathbf{m}^{n+1} - \mathbf{m}^n) \right], \tag{5-34}$$

where,

$$\mathbf{G}_n \Big|_{ij} = \frac{\partial g_i}{\partial m_j^n} = \frac{\partial \boldsymbol{g}_i(\mathbf{m}^n)}{\partial m_j^n}, \tag{5-35}$$

is the Jacobian matrix whose entry at the *i*th row and *j*th column of is the partial derivative of the *i*th predicted datum with respect to the *j*th model parameter evaluated at point $\mathbf{m}^n$ in model space. Similarly, the remaining terms in Equation 5-33 can be written as,

$$\frac{\partial \Phi_r}{\partial \mathbf{m}^{trial}} = -\frac{2}{N_r} \mathbf{S}^{\mathrm{T}} \mathbf{C}_r^{-1} \left[\mathbf{r}^0 - \mathbf{S}\mathbf{m}^{n+1} \right], \tag{5-36}$$

$$\frac{\partial \Phi_h}{\partial \mathbf{m}^{trial}} = \frac{2}{N_h} \mathbf{S}^{\mathrm{T}} \mathbf{L}_h^{\mathrm{T}} \mathbf{L}_h \mathbf{S}\mathbf{m}^{n+1}, \tag{5-37}$$

$$\frac{\partial \Phi_v}{\partial \mathbf{m}^{trial}} = \frac{2}{N_v} \mathbf{S}^{\mathrm{T}} \mathbf{L}_v^{\mathrm{T}} \mathbf{L}_v \mathbf{S}\mathbf{m}^{n+1}, \tag{5-38}$$

Collecting the results from Equations 5-34 to 5-38, substituting them into Equation 5-33 and collecting the terms in the unknown vector $\mathbf{m}^{trial}$ yields the system of linear equations in the form,

$$\mathbf{A}\mathbf{m}^{trial} = \mathbf{b}\,, \tag{5-39}$$

where,

$$\mathbf{A} = \left[\frac{1}{N_d}\mathbf{G}_n^{\mathrm{T}}\mathbf{C}_d^{-1}\mathbf{G}_n + \lambda^{trial}\left(\frac{\alpha_r}{N_r}\mathbf{S}^{\mathrm{T}}\mathbf{C}_r^{-1}\mathbf{S} + \frac{\alpha_h}{N_h}\mathbf{S}^{\mathrm{T}}\mathbf{L}_h^{\mathrm{T}}\mathbf{L}_h\mathbf{S} + \frac{\alpha_v}{N_v}\mathbf{S}^{\mathrm{T}}\mathbf{L}_v^{\mathrm{T}}\mathbf{L}_v\mathbf{S}\right)\right]\,, \tag{5-40}$$

and,

$$\mathbf{b} = \frac{1}{N_d}\mathbf{G}_n^{\mathrm{T}}\mathbf{C}_d^{-1}\left[\mathbf{d} - \boldsymbol{g}(\mathbf{m}^n) + \mathbf{G}_n\mathbf{m}^n\right] + \frac{\lambda^{trial}\alpha_r}{N_r}\mathbf{S}^{\mathrm{T}}\mathbf{C}_r^{-1}\mathbf{r}^0\,. \tag{5-41}$$

The new set of model parameters $\mathbf{m}^{trial}$ can then be found by solving the linearized system of equations (Equation 5-39) using the preconditioned conjugate gradient method as described in Section 3.6.4. By comparing Equations 5-40 and 5-41 with their frequency-domain method counterparts, Equations 3-72 and 3-73, it can be seen that the form of the equations are nearly identical. As explained earlier, the only differences are in the normalizations and weightings of the objective function terms.

5.7 Application

5.7.1 Introduction

In this section the time-domain holistic inversion is demonstrated with an example of its application to real survey data. As explained in the introduction to this chapter, the code for implementing the method has so far not been fully developed. Currently the code only caters for inversions having the same spline node spacing for every layer, and a spatially constant reference model must be used. The method has not been deployed to invert complete surveys at this point in time. However it has been used to invert

selected lines from a fixed wing time-domain survey, the results of which will be presented here.

In the example X and Z component TEMPEST survey data are simultaneously inverted to solve for a 25 layer fixed-thickness conductivity model. Also solved for are three unmeasured system geometry variables, the transmitter-receiver horizontal ($\psi\triangle x$) and vertical ($\psi\triangle z$) separations and the receiver bird's pitch (ψrp). Reference model constraints are applied to the conductivity and system geometry models, and both horizontal and vertical smoothness constraints are applied to the conductivity model.

The inversion produces a conductivity model that is geologically credible and consistent with downhole conductivity logs. The results are also compared to a conventional sample-by-sample inversion in which the system geometry is not solved for, and another where the system geometry is solved for. The comparison demonstrates how solving for system geometry allows the data to be adequately fitted to within the expected noise levels, which is not achievable otherwise. It also shows how the holistic inversion's smooth and continuous parameterization results in improved resolution and interpretability of the conductivity model in comparison to the sample-by-sample inversions.

5.7.2 Data

The data that are inverted in this example were acquired as part of the Lower Macquarie River airborne electromagnetic survey which was flown northwest of Dubbo in the Central West region of New South Wales, Australia. The survey was funded by the Australian Government under the National Action Plan for Salinity and Water Quality. It was commissioned and project managed by the Department of Agriculture, Fisheries and Forestry's Bureau of Rural Sciences as part of the broader Community Stream Sampling and Salinity Mapping Project. The data have previously been inverted using conventional sample-by-sample methods (Brodie and Fisher, 2008).

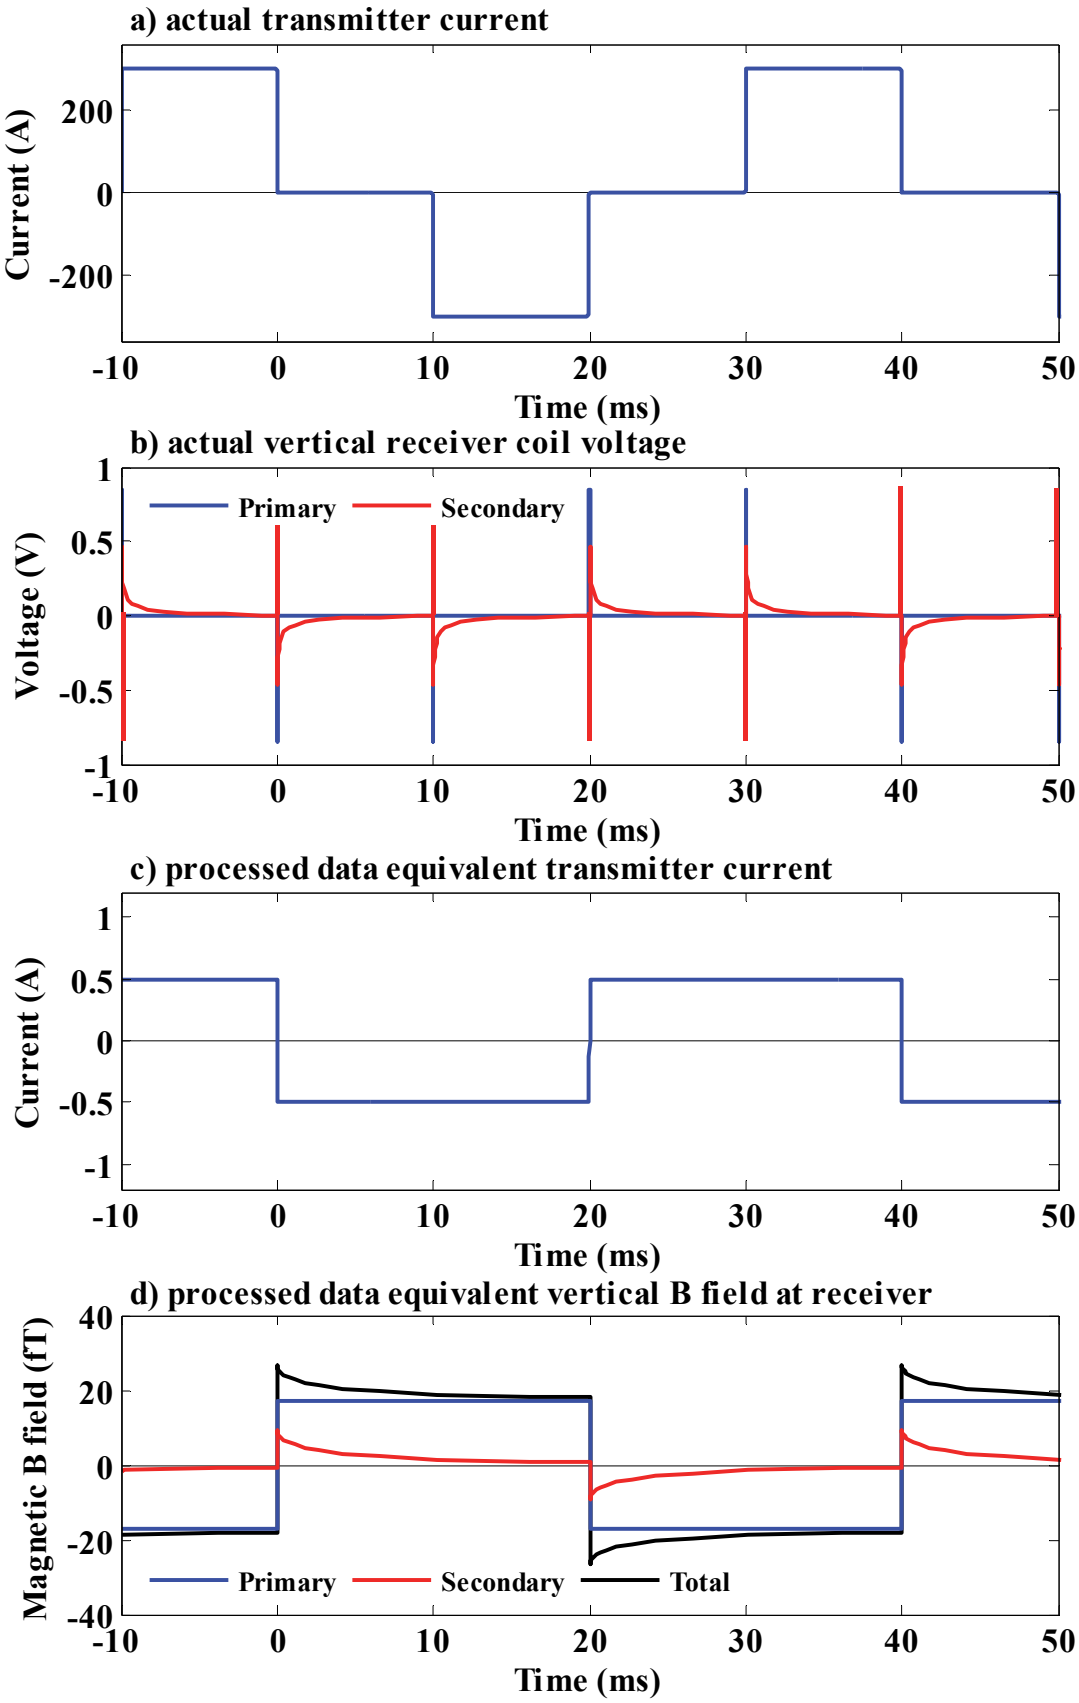

Figure 5.8 *Illustration of TEMPEST waveforms showing the (a) actual transmitter current waveform, (b) actual voltages measured in the vertical component receiver coil, (c) the processed data equivalent transmitter current waveform, and (d) the vertical component of the magnetic B field at the receiver for the processed data (i.e. for the equivalent transmitter current waveform).*

The data were acquired and processed in 2007 by Fugro Airborne Surveys (FAS) using its TEMPEST airborne electromagnetic system (Lane et al., 2000). The system nominally flew at 115 m above ground level and the receiver coils were towed 120 m behind and 40 m below the aircraft approximately. The single turn transmitter loop, which was slung between the slung around the extremities of the aircraft, had an area of 221 m^2. The loop transmitted a bipolar 50% duty cycle, approximately-square, waveform with a base frequency of 25 Hz (40 ms period) and a peak current of 280 A as shown in Figure 5.8a. The voltages in the orthogonal receiver coils are induced by both (primary) currents flowing in the transmitter loop and (secondary) currents flowing in the ground (Figure 5.8b).

The receiver coil voltages and transmitter current are sampled at 75 kHz to produce a continuous time series of streamed data (i.e. having 1500 samples per 20 ms half cycle) that are inputs to the data processing. While voltages in the X-, Y-, and Z-component coils are all recorded, the Y-component data are not to the same standard and are not processed and delivered to clients. The recording of the continuous time series of streamed data allows the measured TEMPEST data to be processed with a range of sophisticated signal processing techniques.

The details of the data acquisition and processing is reported in Noteboom and Stenning (2008). The steps involved in the data processing can be summarized as follows;

a) The streamed data are filtering via wavelet transforms to remove sferic related noise.

b) Application of a stacking algorithm that rejects 50 Hz powerline noise via synchronous signal detection and improves the signal to noise ratio. The stacking filter was cosine-tapered in shape and was 3.04 s (i.e. 152 20 ms half-periods) wide, and output was drawn at 0.2 second intervals.

c) Frequency-domain filtering to suppress VLF and coil-motion noise.

d) A deconvolution algorithm is applied to remove the effect of the system's transfer function that is determined from pre- and post-flight high altitude calibrations. This step removes the system self response and accounts for any slow variations in the transmitted current waveform's amplitude and shape that may occur during a flight.

e) The resultant transmitter waveform and transfer function independent frequency-domain spectrum data are convolved with the spectrum of a 100% duty cycle square-wave waveform (Figure 5.8c) to yield an equivalent square-wave B-field response (Figure 5.8d). Therefore, while the system actually records a voltage ($\propto \partial B/\partial t$) induced by a 50% duty cycle waveform with, the processed data are presented as if the system had measured the magnetic B field response induced by a 100% duty cycle waveform.

f) Normalizations are applied for the transmitter loop's current, area and number of turns, and the receiver coil's effective area so that the B-field response is equivalent to that which would be obtained by a 1 m^2 transmitter loop carrying a 1 A peak to peak current. The responses are converted to units of femtotesla.

g) After conversion back to the time-domain the processed streamed data are windowed using a box-car shaped filter (i.e. the 13.3$\dot{3}$ μs samples that lie in or on each window boundary are averaged) to the 15 window times shown in Table 5.2.

h) The primary B-field is estimated from the total B-field response and is subsequently subtracted from the total response to leave an estimated secondary field or ground response. Details of this procedure are discussed in Section 5.3.4.

i) Using the primary field estimated in the previous step, an estimate of the transmitter-receiver horizontal $\psi\triangle x^{est}$ and vertical $\psi\triangle z^{est}$ offsets can be determined. This assumes that the transmitter-receiver transverse offset, receiver roll, pitch and yaw are zero ($\psi\triangle y^{est} = \psi rp^{est} = \psi rr^{est} = \psi ry^{est} = 0$).

j) The data can be further reduced to a common datum, i.e. as if they were measured with constant system geometry. This is known as a height, pitch, roll and geometry correction (Green, 1998b).

Table 5.2 TEMPEST window delay times and estimated noise levels.

Window number	**Window times (seconds)**		**Estimated additive noise (fT)**	
	Start time†	**End time**†	**X-component**	**Z-component**
1	0.0000066667	0.0000200000	0.017	0.011
2	0.0000333333	0.0000466667	0.014	0.008
3	0.0000600000	0.0000733333	0.011	0.007
4	0.0000866667	0.0001266667	0.008	0.006
5	0.0001400000	0.0002066667	0.007	0.005
6	0.0002200000	0.0003400000	0.007	0.005
7	0.0003533333	0.0005533333	0.006	0.005
8	0.0005666667	0.0008733333	0.006	0.005
9	0.0008866667	0.0013533333	0.006	0.004
10	0.0013666667	0.0021000000	0.005	0.004
11	0.0021133333	0.0032733333	0.005	0.004
12	0.0032866667	0.0051133333	0.004	0.003
13	0.0051266667	0.0079933333	0.003	0.003
14	0.0080066667	0.0123933333	0.002	0.002
15	0.0124066667	0.0199933333	0.003	0.002
			Estimated multiplicative noise	
			1.3%	1.7%

† Time *t=0* is defined as when the "processed data equivalent transmitter current" (Figure 5.8c) switches polarity.

The data were acquired and processed in 2007 by Fugro Airborne Surveys (FAS) using its TEMPEST airborne electromagnetic system (Lane et al., 2000). The system nominally flew at 115 m above ground level and the receiver coils were towed 120 m behind and 40 m below the aircraft approximately. The single turn transmitter loop, which was slung between the slung around the extremities of the aircraft, had an area of 221 m^2. The loop transmitted a bipolar 50% duty cycle, approximately-square, waveform with a base frequency of 25 Hz (40 ms period) and a peak current of 280 A as shown in Figure 5.8a. The voltages in the orthogonal receiver coils are induced by both (primary) currents flowing in the transmitter loop and (secondary) currents flowing in the ground (Figure 5.8b).

The receiver coil voltages and transmitter current are sampled at 75 kHz to produce a continuous time series of streamed data (i.e. having 1500 samples per 20 ms half cycle) that are inputs to the data processing. While voltages in the X-, Y-, and Z-component coils are all recorded, the Y-component data are not to the same standard and are not processed and delivered to clients. The recording of the continuous time series of streamed data allows the measured TEMPEST data to be processed with a range of sophisticated signal processing techniques.

The input data to the inversion were then the 30 (15 X- and 15 Z-component) equivalent total B-field data associated with the stacked and windowed 0.2 s samples acquired over the entire length of the flight line. In the example that is discussed a 40 km long flight line of data was inverted. The flight line had 2,981 samples spaced on average at 13.5 m intervals. Every second sample was inverted, so that a total of 1,491 samples and 44,730 (1491 samples × 15 windows × 2 components) data were included in the observed data vector. The decision to invert every second sample, rather than every sample, was simply to minimize the inversion time. Because the footprint of the TEMPEST system is of the order of 250 m and the data are stacked over 3.04 s

(~205 m) it is considered that there would be no loss of information by inverting every second sample.

Other inputs to the inversion routine were the auxiliary information, which included the transmitter loop height above ground that was measured by laser altimeter, and the transmitter loop roll and pitch that were measured by gyroscopes. Noise levels for the Lower Macquarie survey data have previously been estimated by Brodie and Fisher (2008) from analysis of repeat line data acquired during the survey. They used the additive/multiplicative noise model method (Green and Lane, 2003) that is described in Section 4.2.6. These same noise level estimates, which are shown in Table 5.2, were used for this inversion example.

5.7.3 Parameterization

Conductivity model

A 25-layer fixed-thickness conductivity model parameterization was used. The spline node spacing for every layer in the conductivity model was 100 m. The choice of spline node spacing was guided by using the lateral variation in the early-time window data as a proxy for the lateral variation in the subsurface conductivity. Early-time data from the survey were fitting with splines of various node spacings. This ascertained what node spacing was required to adequately represent the data (and hence by proxy the conductivity) in a similar manner as was discussed for the system geometry parameters and demonstrated in Figure 5.4. It was also found that the 100 m node spacing was adequate to allow the data to be fitted in the inversion. Over the 40 km line length there were 401 layer conductivity spline nodes, which resulted in a total of 10,025 (401 nodes per layer × 25 layers) conductivity model parameters to be solved for.

The implementation of the code so far only allows a laterally constant reference and starting model to be used. Because of prior knowledge, from downhole logs, of high

conductivities that occur in the survey area, a homogenous 0.1 S/m conductivity reference and starting model was used. The prior uncertainty standard deviation for all conductivity reference values were set to the 3.0 logarithm decades so that a range of conductivity values would be permitted.

The conductivity model had 25 layers, where the top layer was 2.0 m thick and each layer got progressively thicker by 10%. The second last layer was 17.9 m thick and its base was at 177.0 m depth. All layer thicknesses and depth ranges are summarized in Table 5.3.

Table 5.3 Layer thicknesses and depths for the inversion of TEMPEST data.

Layer#	Thickness (m)	Top depth (m)	Bottom depth (m)	Layer# (cont'd)	Thickness (m)	Top depth (m)	Bottom depth (m)
1	2.00	0.00	2.00	**14**	6.90	49.06	55.96
2	2.20	2.00	4.20	**15**	7.59	55.96	63.55
3	2.42	4.20	6.62	**16**	8.35	63.55	71.90
4	2.66	6.62	9.28	**17**	9.19	71.90	81.09
5	2.93	9.28	12.21	**18**	10.11	81.09	91.20
6	3.22	12.21	15.43	**19**	11.12	91.20	102.32
7	3.54	15.43	18.97	**20**	12.23	102.32	114.55
8	3.90	18.97	22.87	**21**	13.45	114.55	128.00
9	4.29	22.87	27.16	**22**	14.80	128.00	142.80
10	4.72	27.16	31.88	**23**	16.28	142.80	159.08
11	5.19	31.88	37.07	**24**	17.91	159.08	176.99
12	5.71	37.07	42.78	**25**	∞	176.99	∞
13	6.28	42.78	49.06				

System geometry model

In this particular inversion example, three unmeasured elements of the system geometry were solved for. These were the transmitter-receiver horizontal in-line ($\psi\triangle x$) and vertical ($\psi\triangle z$) separations and the receiver coils' assembly pitch (ψrp). Each element was parameterized using a spline node spacing of 75 m. Similarly, to the conductivity

model, this node spacing was chosen by fitting various measured and data processing estimates of the system geometry to determine the largest node spacing that could adequately represent the along-line system geometry variations. As was demonstrated in Figure 5.4 it was found that 75 m was a suitable choice. Over the 40 km line length 534 spline nodes were required to parameterise the geometry parameters, resulting in a total of 1,602 (534 nodes per element × 3 elements) system geometry model parameters to be solved for.

The reference and starting model values for the transmitter receiver separations ($\psi\triangle x$ and $\psi\triangle z$) were set to be equal to the values made during the standard TEMPEST data processing (i.e. $\psi\triangle x^{est}$ and $\psi\triangle z^{est}$). The reference and starting model values for the receiver coils' pitch (ψrp) was set to be the assumed data processing value of zero degrees. The prior uncertainty standard deviations for the $\psi\triangle x$, $\psi\triangle z$ and ψrp reference values were set to 2 m, 4 m and 4° respectively.

Since there was no evidence to suggest that the measured transmitter loop height, pitch, roll and yaw were incorrect, those geometry parameters were not solved for. However, their measured values were directly used in the inversion for computation of the forward responses. The remaining three unmeasured elements of the system geometry, the transmitter-receiver horizontal transverse separation, and the receiver coils' assembly roll and yaw were not solved for because these parameters are not resolvable unless Y-component data are available. Their assumed value of zero was used in the inversion for computation of the forward responses.

5.7.4 Solution

In total airborne 44,730 data were inverted and 11,627 (10,025 conductivity and 1,602 system geometry) parameters were solved for. The model regularization parameter λ after initially being set to a large value (1×10^8), was automatically determined in each iteration using the line search scheme described in Section 5.6.3. Reference model

constraints ($\alpha_r = 1$) were applied according to the reference model and prior uncertainties described above and in Table 5.4. Roughness constraints were applied in both the along-line ($\alpha_h = 10^6$) and vertical directions ($\alpha_v = 10^3$) directions to the layer conductivities.

Table 5.4 Summary of model parameters for the inversion of TEMPEST data.

Local conductivity profile & geometry values	**Dependence**	**Inversion parameters**	**Number of unknowns**	**Reference model**	**Reference model standard deviation**
σ_1	line distance, layer	100 m spline	401 ×25 = 10,025	log(0.1)†	3†
σ_2		100 m spline		log(0.1)†	3†
…		…		…	…
σ_{25}		100 m spline		log(0.1)†	3†
t_1	fixed	n/a	0	2.00 m	n/a
t_2		n/a		2.20 m	n/a
…		….		…	…
t_{24}		n/a		17.91 m	n/a
$\psi\triangle x$	line distance	75 m spline	534	data processing estimate $\psi\triangle x^{est}$	2 m
$\psi\triangle z$	line distance	75 m spline	534	data processing estimate $\psi\triangle z^{est}$	4 m
ψrp	line distance	75 m spline	534	0.0	4 m
ψth , ψtr , ψtp , ψty	fixed at measured values	n/a	0	measured values	n/a
$\psi\triangle y$, ψrr , ψry	fixed at the assumed value of zero	n/a	0	0.0	n/a

† units of log(S/m)

The inversion converged to the desired data misfit (Equation 5-18) $\Phi_d = 1$ in 19 iterations. Figure 5.9a shows the regularization parameter λ that was automatically determined in each iteration by the line search algorithm. It can be seen in Figure 5.9b how, by searching for an appropriate value λ in each iteration with the goal of reducing

Φ_d to 0.7 of its previous value, the data misfit converged predictably to the desired value.

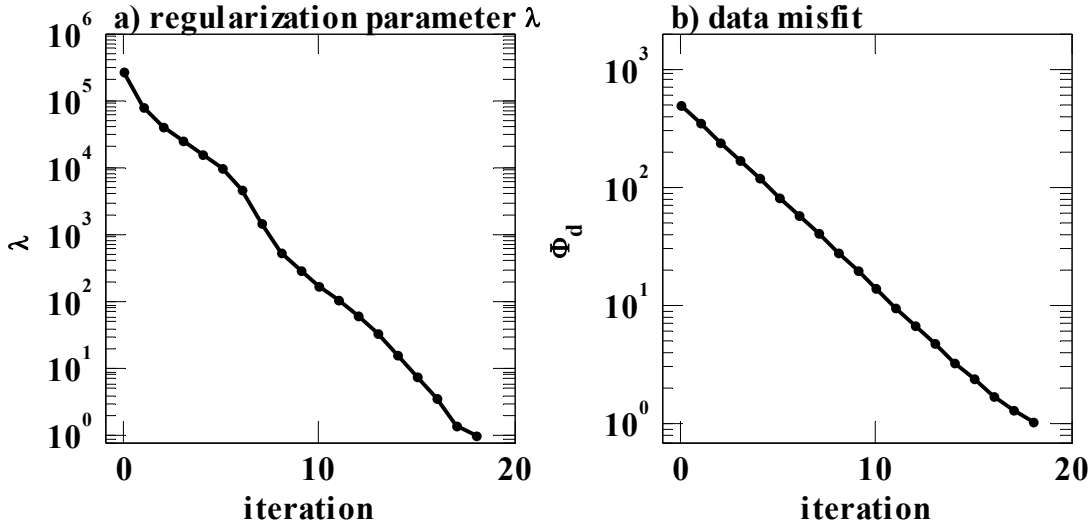

Figure 5.9 Summary of the inversion's convergence, showing, (a) the regularization parameter λ that was automatically determined in each iteration, and (b) the data misfit Φ_d (Equation 5-18) at each iteration.

5.7.5 Results

Sample-by-sample inversion results

So that the holistic inversion results could be compared to conventional results, two different sample-by-sample inversions were run on the data. The results of these two inversions will now be presented. The first was a conventional sample-by-sample inversion where no system geometry parameters were solved for. In the second inversion the same three system geometry parameters, as in the holistic inversion, were solved for. For convenience these inversion will respectively be denoted the SBS and the SBSG inversions.

To allow rational comparison, the same settings were used for both the SBS and SBSG as were used in the holistic inversion where applicable. Therefore, the same number of layers and the same layer thicknesses were used. The same homogenous 0.1 S/m conductivity starting and reference values and prior certainties were used. The same system geometry starting and reference and certainty values were used for the SBSG

(not applicable in the SBS where geometry was not solved for). Vertical smoothness constraints were applied in both the SBS and SBSG as they were in the holistic inversion.

The stitched conductivity section (i.e. individual inversion models beneath each inverted sample are compiled into a section) resulting from the SBS inversion is shown in the bottom panel of Figure 5.10. The central 14 km portion of the section has been displayed because it contains two boreholes named GW800232 and LMQ03 that were close to the flight line (79 m and 18 m respectively). To enable a comparison, the downhole conductivity logs that were measured in the boreholes have been plotted (i.e. inside the black columns) on the conductivity section with the same colour lookup scheme. In the panels above the conductivity section the various measured and estimated components of the system geometry are shown. The topmost panel shows the final data misfit (Φ_d) that was achieved by the inversion.

It can be seen in the topmost panel that the data has not been fitted to the within the expected noise levels because the misfit ($\Phi_d \cong 40$) was not reduced to a value approaching unity. This is indicative that there is be some inconsistency in the dataset because, with a layered-earth model, we would expect to be able to fit the data quite easily in this broad scale sedimentary environment.

It is also noticeable in that the conductivity in the lower part of the SBS section is relatively resistive. On the other hand, both the downhole conductivity logs suggest that the conductive material extends to greater depth. One possible reason why the inversion section is resistive at depth is because the primary field (and hence the system geometry) estimation procedure in the data processing mostly favours a resistive basement assumption (*cf.* Section 5.3.4).

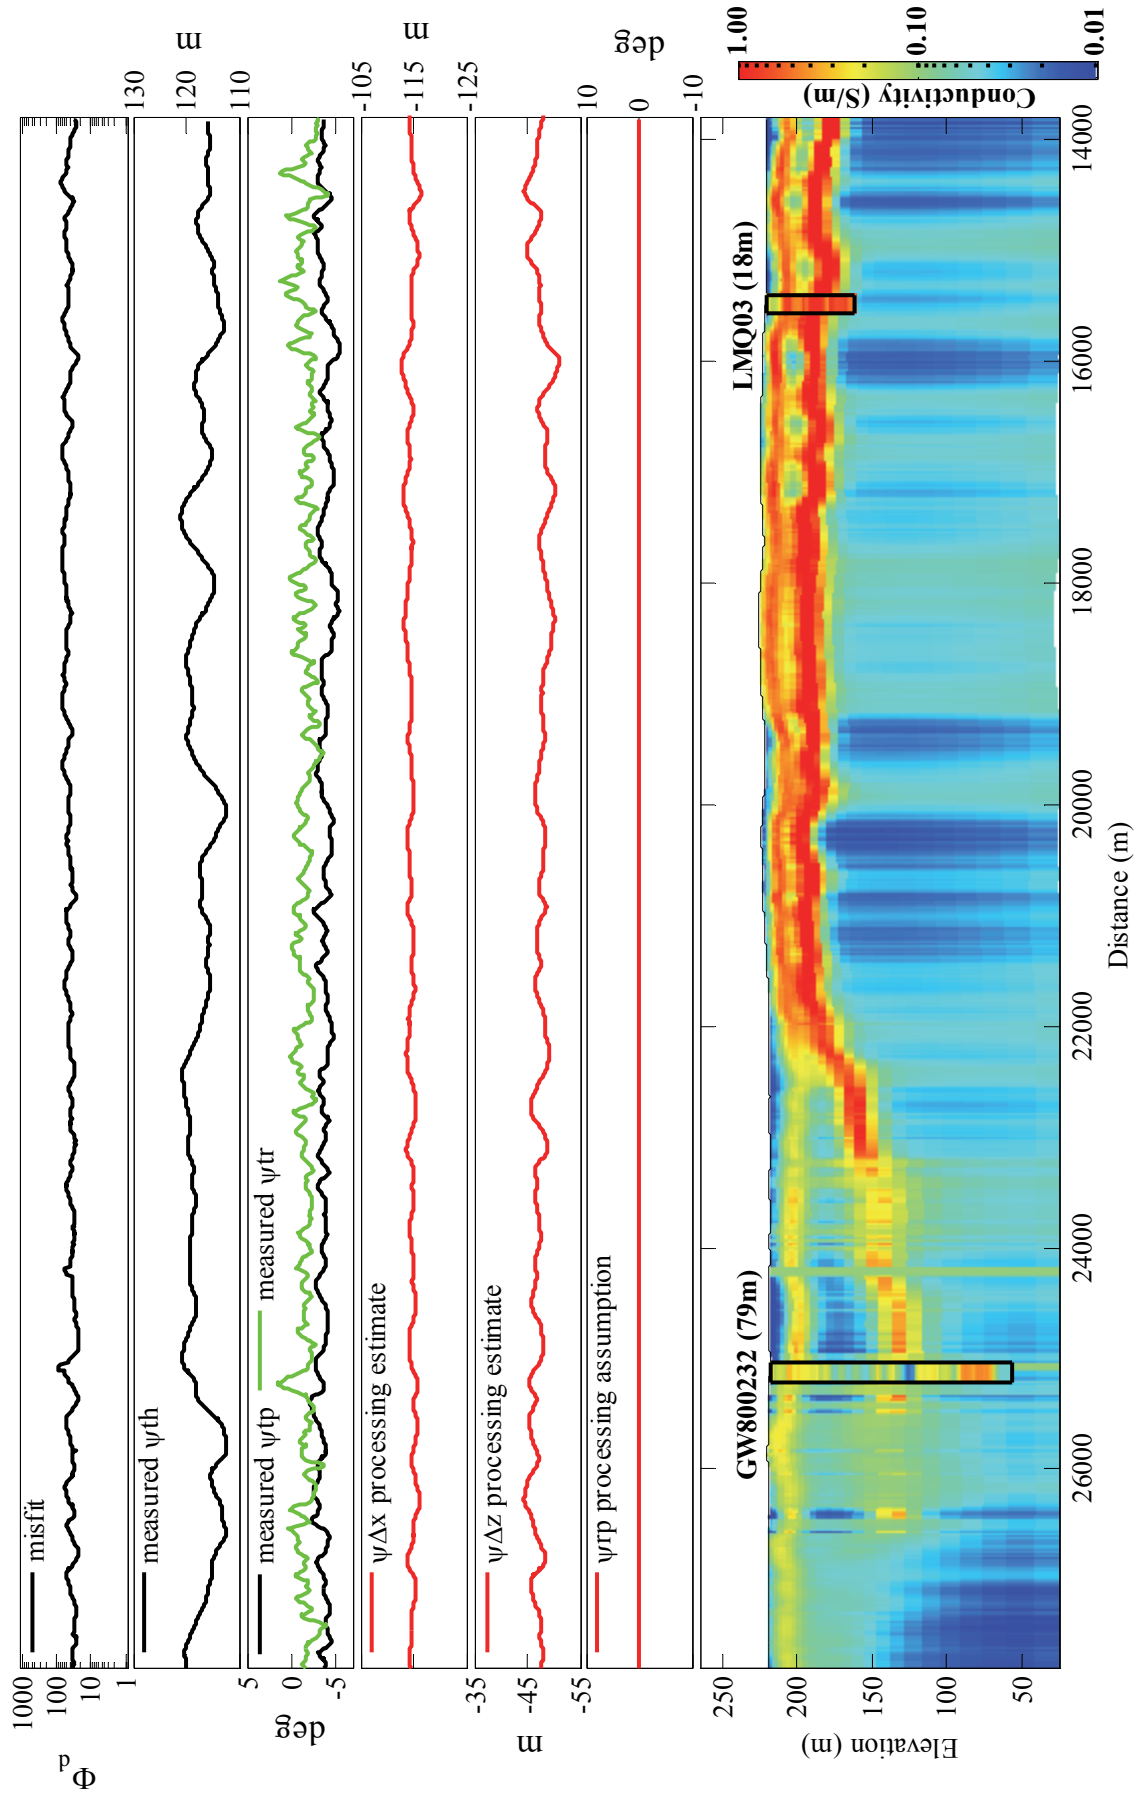

Figure 5.10 **SBS:** *Conductivity section resulting from the* conventional sample-by-sample inversion where the system geometry **was not** solved for. Note the many vertical artefacts and that the data was not fitted ($\Phi_d >> 1$).

The SBS conductivity section also contains several artefacts (i.e. vertical stripes in the section) and has a generally discontinuous or broken-up appearance which is not geologically credible. This is likely to be caused by changes in the system geometry that have not been correctly resolved by the data processing estimates. Similarly caused artefacts have been noted, for example, by Smith (2001a) and Fitterman and Yin (2004) in conductivity sections derived from data acquired by fixed-wing time-domain and helicopter frequency-domain systems respectively.

The discussion will now switch to the results of the SBSG inversion, which are shown in Figure 5.11. The general layout of this figure is the same as for the SBS inversion results (Figure 5.10). However in this case the three elements of the system geometry that were solved for are shown (as blue profiles) in the three panels immediately above the conductivity section.

It can be seen in the top panel of the SBSG inversion results the data has been fitted to within the expected noise levels ($\Phi_d \cong 1$). It can be seen that receiver pitch ψrp estimated by the SBSG inversion (blue profile) is systematically less than the data processing estimate (red profile). The same is the case for the transmitter-receiver horizontal offset $\psi \triangle x$. Note also that the conductivity section is much less conductive towards the bottom than was the case for the SBS inversion. Because of this the SBSG inversion matches the lower part of the downhole conductivity logs better. There are also many fewer of the vertical striping artefacts that were noted in the SBS inversion. The SBSG section is therefore somewhat more geologically credible. Combined, these observations suggest that the apparent incompatibility between the data, the system geometry estimates, and the geology that was apparent in the SBS inversion results have largely been resolved by solving for the three unmeasured elements of the system geometry.

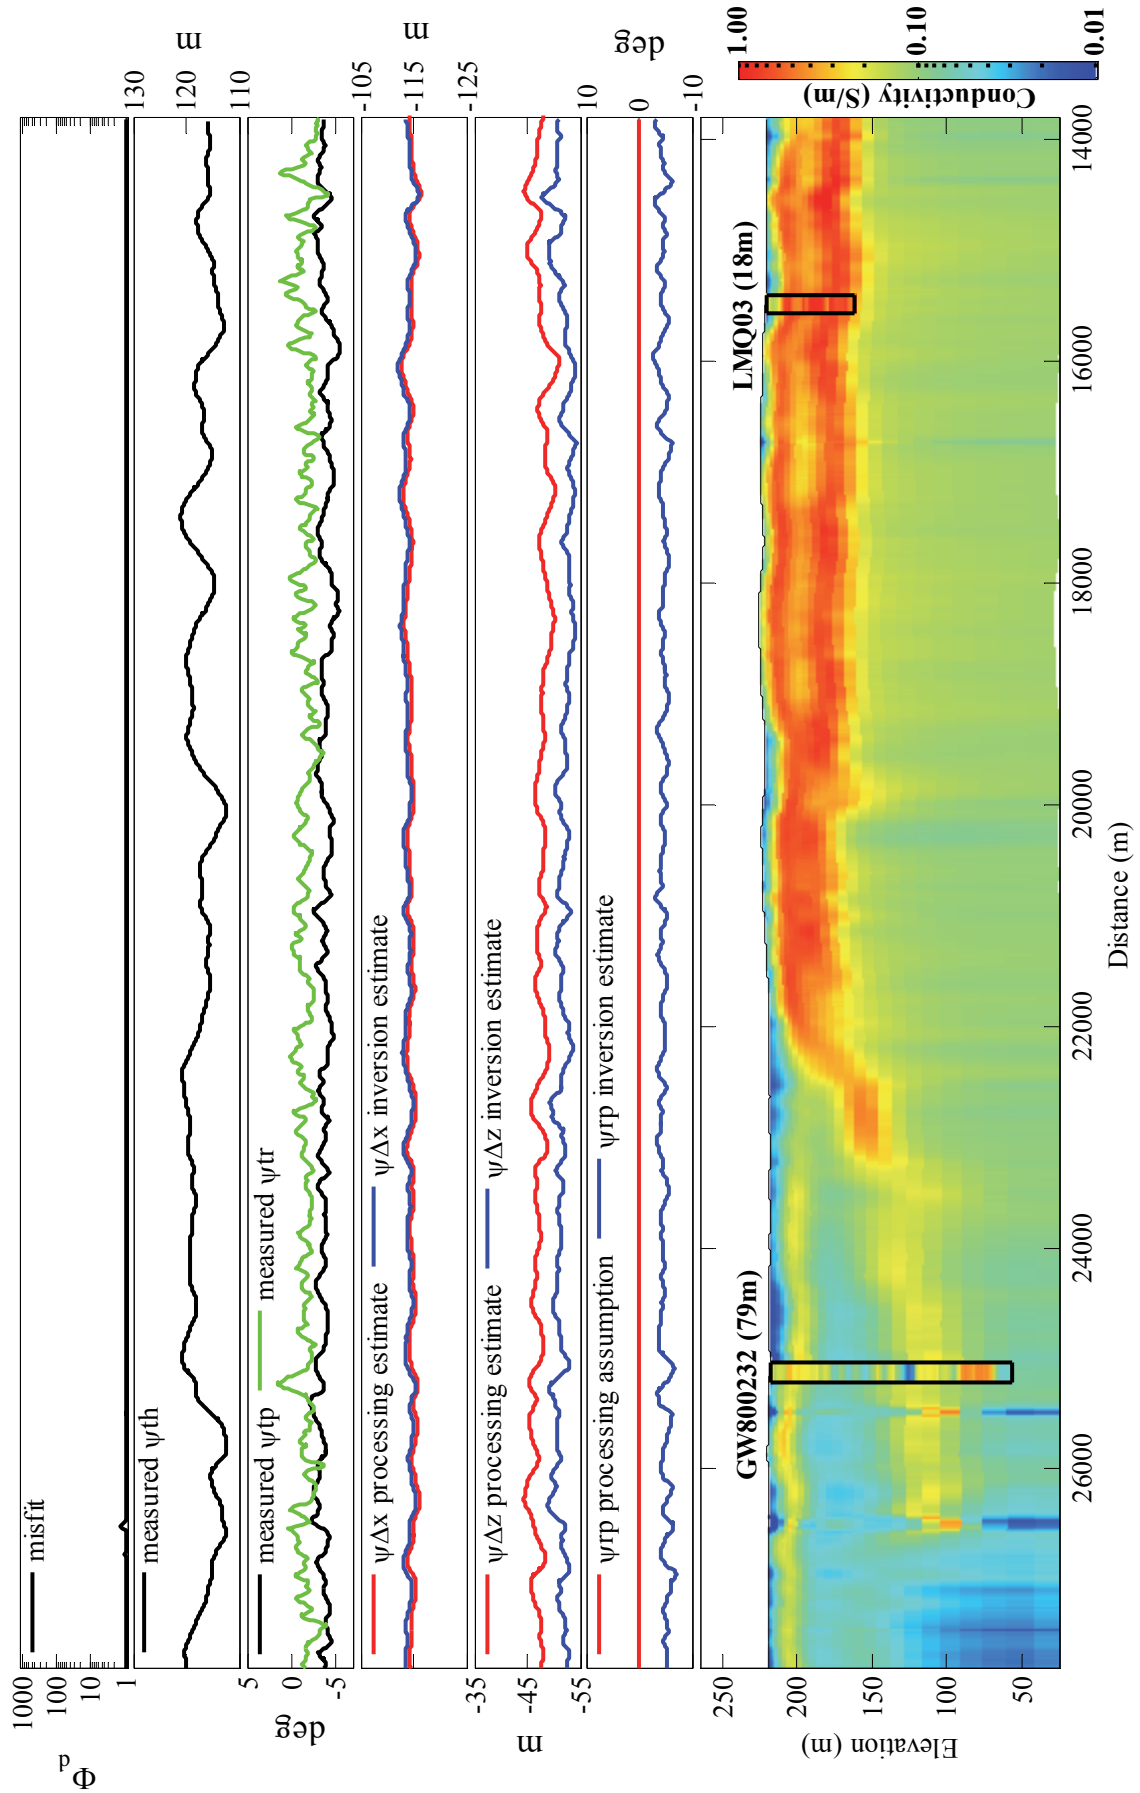

Figure 5.11 **SBSG:** *Conductivity section and system geometry estimates resulting from the* sample-by-sample inversion where three unmeasured elements of the system geometry were solved for.

5.7.6 Holistic inversion results

The results of the time domain holistic inversion are shown in Figure 5.12. The layout of this figure is the same as for the SBSG results (Figure 5.11), but with some additional information that is only relevant to the holistic inversion. To give the reader some appreciation of the model parameterization, the horizontal position of the spline nodes used in the model have been plotted toward the bottom of the conductivity section (magenta dots). Similarly the position of the spline nodes used for the system geometry parameterization have been plotted on the panel immediately above the section (green dots).

In the topmost panel two data misfit profiles have been plotted. The black profile shows the data misfit ($\Phi_d = 0.99$) for the inversion of the whole line. This is the quantity (Equation 5-18) that was successfully minimized in the inversion algorithm to the desired level. The green profile shows the local misfit which is calculated as a check on the data misfit for every individual sample in the flight line. In other words, it is equivalent to the data misfit that would be calculated in a sample-by-sample inversion for each inverted sample (i.e. as shown in the top panel of Figure 5.10 and Figure 5.11).

The local misfit is generally less than unity as desired. However, in places it rises to higher values but never to values near 40 as was the case in the SBS inversion in which the geometry was not estimated (*cf.* Figure 5.10). In the few places where the local misfit is higher the reason may be that not enough lateral variability was allowed in the conductivity model. Alternatively it could be that the holistic inversion is fitting through some local noise. This second argument would be supported by the presence of two artefacts that appear on either side of the profile distance 26,000 m, on the SBSG inversion section. The SBSG inversion, which does not limit lateral variability, fitted the data at these points at the expense of introducing an artefact into the conductivity section.

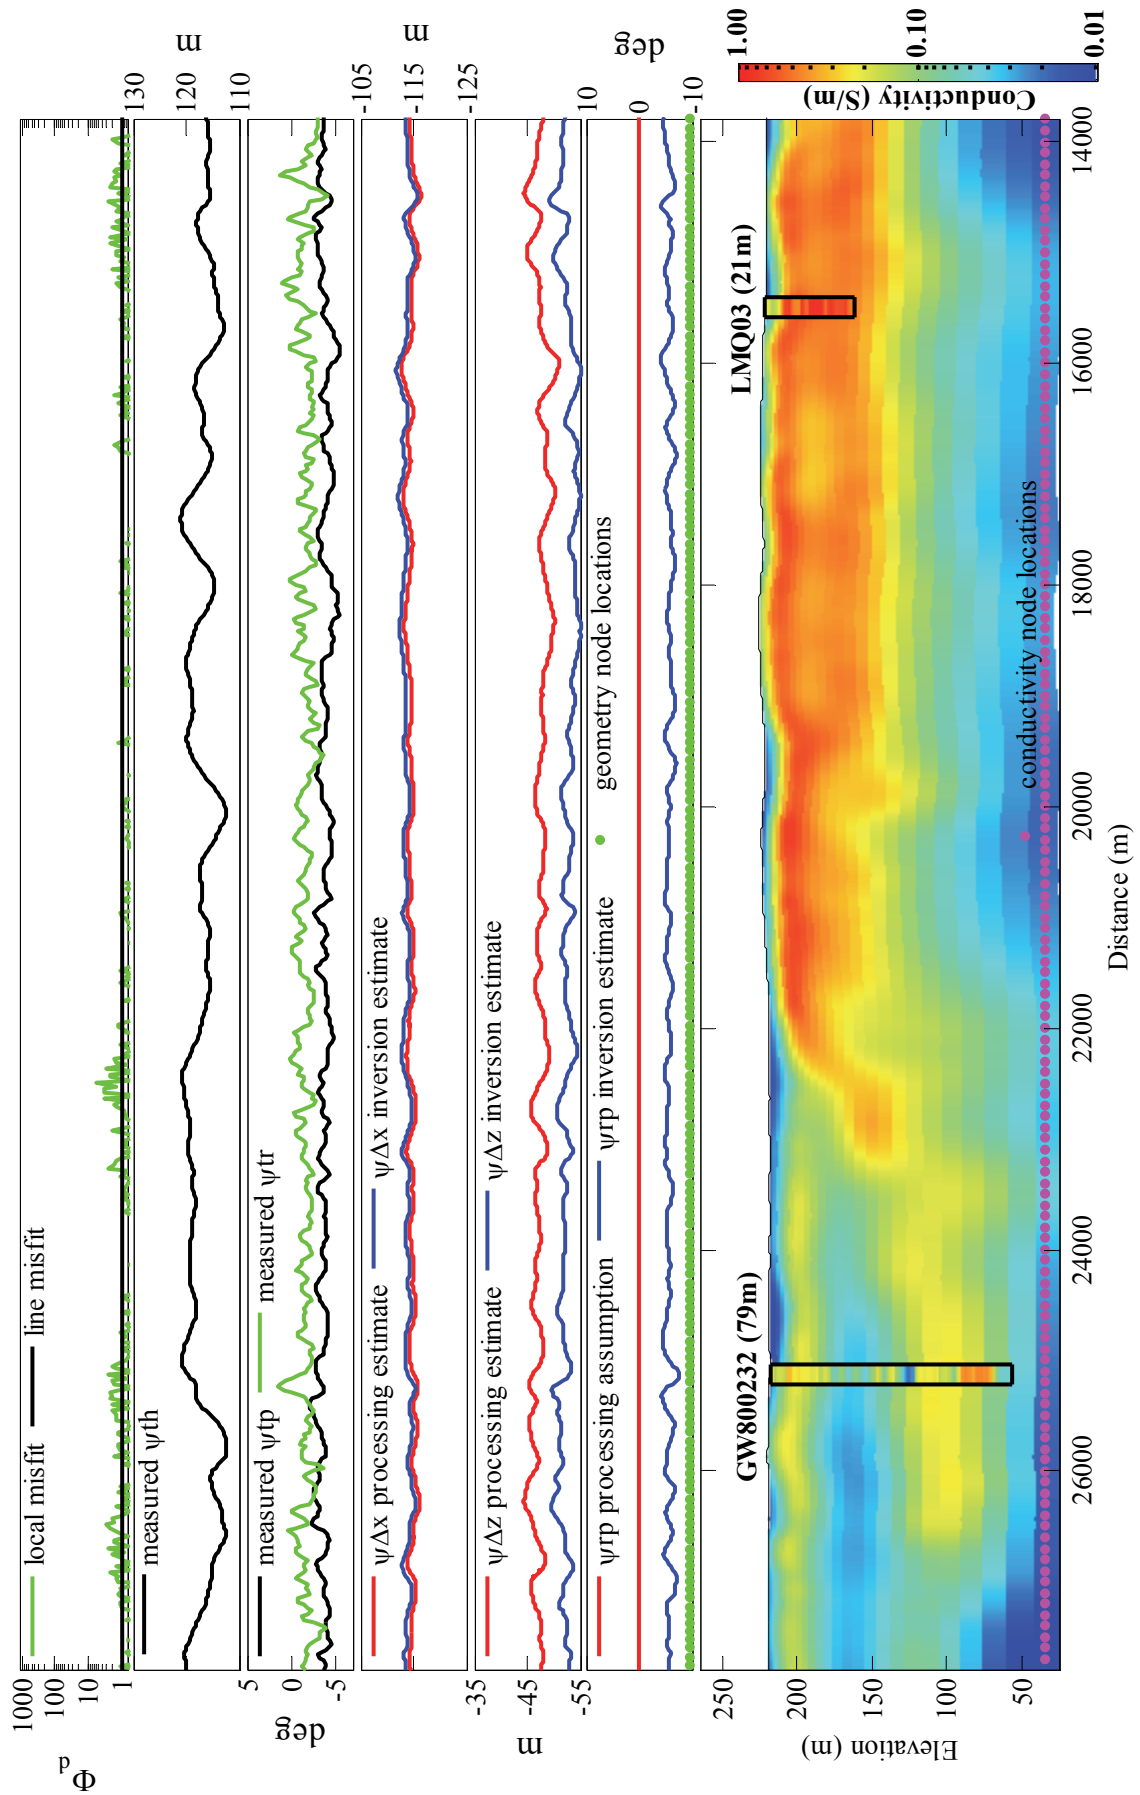

Figure 5.12 ***Holistic:*** *Conductivity section and system geometry estimates resulting from the* sample-by-sample holistic inversion where three unmeasured elements of the system geometry were solved for.

The system geometry estimates resulting from the holistic inversion are similar to those from the SBSG inversion. Figure 5.13 shows a detailed comparison between the two downhole conductivity logs and the nearest conductivity models from each inversion. It can be clearly seen that both the SBSG and the holistic inversions represent the downhole conductivity logs better than the SBS inversion does. On the basis of Figure 5.13, the holistic and SBSG inversions would have to be judged as equally good fits to the downhole logs. This suggests that the main improvement is gained by solving for the system geometry estimates rather than simply assuming that the data processing estimates are correct.

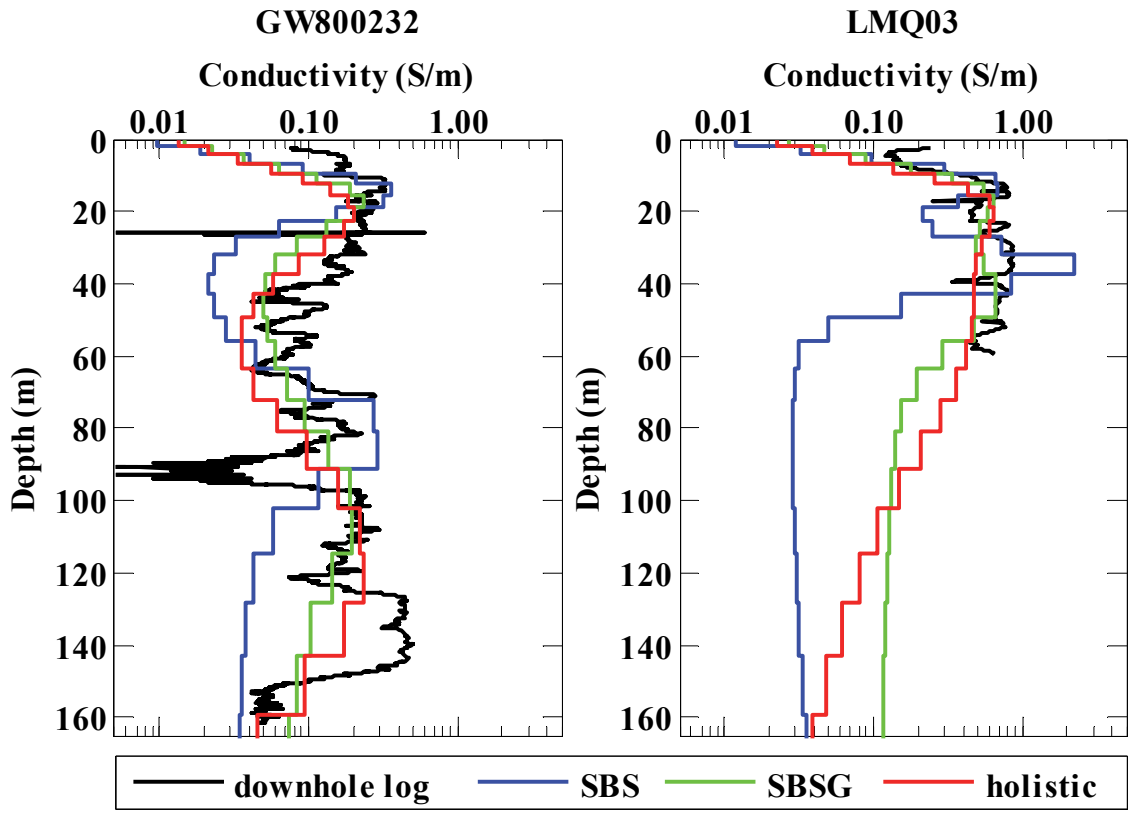

Figure 5.13 *A detailed profile comparison of the downhole conductivity logs, that are also shown on the sections, and the SBS, SBSG and holistic inversion models at the nearest airborne sample to the borehole.*

However the holistic inversion conductivity section is more geologically realistic than the SBSG inversion. This assessment is made because the holistic inversion section does not have any of the artefacts that appear on the SBSG (and SBS) section. The holistic inversion does show improved resolution, continuity and interpretability in

comparison to the SBSG inversion. The holistic inversion conductivity section is geologically credible and is also consistent with both the downhole conductivity logs and the electromagnetic data. The improvement over the SBS inversion is quite substantial, however in this particular case the improvement on the SBSG is incremental. While yet to be demonstrated, the advantage of the holistic inversion is expected to be greater in survey areas where the topography, and hence the system geometry, is more variable than in the example presented here.

5.8 Extension to other time-domain configurations

The time-domain holistic inversion has so far only been developed to address one particular set of calibration issues, that being the estimation of unmeasured elements of the system geometry in fixed-wing systems. The ideas could potentially be extended to address other types of calibration problems. There would be no need to change the conductivity model parameterization. Instead it would involve the addition of different types of calibration error to the existing calibration model.

The most obvious example is that the existing formulation could be used to estimate the unmeasured elements of system geometry in helicopter mounted systems. While altimeters and tilt meters are mounted on the transmitter/receiver assembly in some helicopter systems, some systems still do not measure height above the ground or orientation of the assembly. It would be preferable if these quantities were actually measured, however if they are not it is distinctly feasible that they could be solved for. The existing spline parameterization would be directly applicable without change because the transmitter assembly exhibits pendulum-like motion (Davis et al., 2006), which being relatively smooth and continuous could be naturally represented by 1D splines.

There are other forms of calibrations problems that are similar to those addressed in the frequency-domain formulation developed in Chapter 3. For example Vrbancich and

Fullagar (2007) have noted the existence of amplitude errors in the helicopter time-domain HoistEM system. They determined, using ground truth data, that there was a systematic amplitude (scaling) error in each window of their dataset. It is likely that this type of error could be addressed in a holistic inversion by including gain parameters into the calibration model in the same way as in Chapter 3.

The helicopter time-domain systems, which have a small separation between transmitter and receiver, are also affected by small residual currents that flow in the transmitter loop even after the current pulse is switched off. These currents introduce primary field contamination into the measured response. This effect can be largely reduced by positioning the receiver so that it is minimally coupled with the primary field (e.g. SkyTEM). However in other systems (e.g. VTEM, AEROTEM and HoistEM) the contamination must be estimated by making regular zero-level measurements at high altitude. These high altitude zero-level estimates of the primary field contamination are later subtracted out of the data recorded at survey altitude. This problem is obviously comparable to the zero-level bias calibration error in frequency-domain surveys. It is therefore likely that a piecewise linear parameterization of the bias errors, similar to that shown in Figure 3.4, could be added to the time-domain holistic inversion formulation to address this calibration problem.

An additional area in which the time-domain holistic inversion could be beneficial is in the inversion of data acquired by the SkyTEM helicopter system. This system is novel in that it employs a dual moment configuration not used by any other system (Sorensen and Auken, 2004). The system transmits two different types of pulses. One pulse is a low transmitter moment (LM) pulse having a ~4 μs turn-off time that is designed for shallow investigations. The other is a high transmitter moment (HM) pulse having a ~40 μs turn-off time that is designed for deeper sounding. Therefore in effect it is a dual transmitter system.

A group of LM and then a group of HM pulses are transmitted alternately as the helicopter flies along the flight line. Each group of transients are stacked into alternate LM and HM samples. Therefore the data from each moment are spatially separated. The separation depends on the number of transient pulses that are stacked into each output sample, but it might typically be of the order of 20 m. In conventional sample-by-sample methods, the LCI (Auken et al., 2005) or the SCI (Viezzoli et al., 2008) the model parameterizations use many separate 1D conductivity models that are 'placed beneath' airborne samples. Under this paradigm, when inverting SkyTEM data, a choice must be made as to whether a different 1D conductivity model is associated with adjacent LM and HM samples or if they will share a common model. The first option would double the number of parameters to be solved for. The second option would involve interpolation of the spatially separated HM and LM data to a common model location, or simply accepting the inconsistency of the two sets of data not being exactly co-located with the model. Neither option is particularly desirable.

The continuously defined parameterization used in the holistic inversion offers a convenient remedy to this problem. The spatially separated LM and HM data samples would simply be forward modelled at their true location, by extracting the local 1D model from the continuously defined spline in the usual way. There would be no change in the number of model parameters, no interpolation of data, and no inconsistency between model and data locations.

Chapter 6

Summary, conclusions and future work

6.1 Summary and conclusions

Since its inception over sixty years ago, AEM has been a widely and successfully used geophysical technique in the exploration for massively conductive mineral ore bodies. In recent years the method has been increasingly applied to environmental and (hydro)geological mapping tasks. The new stakeholders have called for more accurate estimates of subsurface conductivity, which has required more quantitative analysis of AEM data. It was soon recognised that AEM data were not sufficiently accurate to allow the desired level of quantitative interpretation for the new applications. Despite the industry having responded to the challenge by developing systems with better bandwidth, monitoring and calibration, even data from the most advanced systems still contain significant calibration error.

In the introductory chapter, it was concluded that: a) non-uniqueness; b) systematic calibrations errors; and c) the limitations of how systematic calibration errors are dealt with in data processing are three significant factors that negatively impact on the accuracy of AEM inversions. Non-uniqueness cannot be entirely eliminated, however the range of possible solutions can be narrowed by reducing data errors, adding independent data, or restricting the solutions to those that are geologically plausible by adding regularizing constraints. It was explained how the conventional processing of frequency-domain data, based on levelling apparent conductivity maps a single frequency at a time: is subjective; cannot resolve inter-channel inconsistencies; and

showed that it can be (wrongly) achieved by having the associated apparent depth map 'absorb' the error instead. I showed examples of how systematic calibration errors lead to systematic biases in inversions results, and indeed how they can prevent the data from being fitted to any plausible geological model at all.

Motivated by the need to reduce the influence of these three negative factors, this research re-examines the problem from the ground up. I have proposed, formulated and put into practice an entirely new way of thinking about AEM practice. To change from the conventional approach, where that data are sequentially calibrated, processed and inverted, after which we often find that the inversion results are inconsistent with independent information. To switch to a new framework, where all of the available information is considered together in a formal inversion problem, so that consistency is guaranteed.

In Chapter 2 the established background theory that forms the basis for all of the 1D electromagnetic modelling carried out as part of the research was presented. The forward modelling algorithm was written in original C++ language source code. The code caters for general AEM system transmitter and receiver geometries, and includes the computationally efficient propagation matrix method of analytically computing derivatives of the forward response with respect to the layer properties. I have derived, and presented in the appendices, new analytic derivatives for all of the system geometry parameters, which were required for efficient implementation of the holistic inversion.

In the main chapter of the thesis, Chapter 3, the formal description of the holistic approach for frequency-domain data was set out. The method is, in essence, designed to simultaneously calibrate, process and invert all data channels at all spatial locations. The problem is formulated in such a way that it ensures: a) processing assumptions, conductivity estimates, and independent prior information are all kept mutually consistent; b) spatial coherency of the geology and temporal coherency of systematic

errors are maximally capitalized upon; c) inter-channel inconsistencies are resolvable; d) errors are not propagated from one sequential step to the next; e) the expensive and time consuming circularly iterative paradigm of calibration→processing→inversion→validation→recalibration→etcetera is avoided; and f) the whole procedure is more objective.

Much inspiration was drawn from previously successful work, namely, the use of a model for calibration error (Deszcz-Pan et al., 1998), and exploitation of the along-line lateral continuity of the geology (Auken et al., 2005). These ideas have been adopted and made more valuable by inserting them into a broader framework. Because all data is inverted at once, the calibration model now benefits from the temporal coherency of systematic error throughout a whole flight and between flights. For the same reason, not only along-line continuity, but spatial continuity in all horizontal directions is now exploited.

The continuous 3D conductivity model introduces a fundamentally new way of parameterizing the subsurface in 1D AEM inversion. It conceptualises the conductivity structure in precisely the way one envisages a layered-earth to be. In my assessment this alone is an advance because it is a more natural parameterization than conceptualising the 3D Earth as a collection of 1D Earths that must somehow be stitched together. Its other advantages are that: a) it is an implicit continuity and smoothness constraint; b) it reduces the number of inversion parameters required per datum, c) its continuous definition provides a natural mechanism for inclusion of independent geoelectric and interface-depth data that do not lie directly on a flight line: and d) the conductivity estimates do not need to be gridded or stitched post inversion.

In Chapter 4 three applications of the holistic inversion to frequency-domain AEM data were presented. The synthetic example demonstrated that even when a realistic amount of Gaussian noise was present, the recovered parameters of the conductivity and

calibration models were acceptably close to the known true values. It indicated that the sensitivity to noise was within an acceptable range and that ill-posedness or fundamental trade-off between the five classes of parameters was not a serious problem. The synthetic example also confirmed the theory had been correctly implemented in computer code.

The second, a few-layer inversion example, demonstrated how the holistic inversion approach can be successfully used to calibrate, process and invert practically raw AEM data. It produced a plausible calibration model and a conductivity model that was consistent with prior information. The resulting conductivity model was superior to that produced from inversion of final, fully processed contractor delivered data using a conventional sample-by-sample algorithm. The same conclusions were drawn from the third example, a multi-layer inversion. In addition, this application demonstrated that the method is not reliant on having an advanced conceptual conductivity model, strong reference model constraints, downhole conductivity logs, or interface-depth data. Despite, downhole conductivity log data not being included in the multi-layer inversion, the estimated conductivity model was very well correlated with the downhole logs.

A major achievement in this chapter was the parallelization of the code to run on a cluster computer, showing that the holistic inversion is applicable to large regional (>10,000 line km) surveys. Although conventional sequential processing and inversion approaches are computationally less challenging to implement, it was demonstrated that the holistic approach is tractable. Furthermore, it can in fact generate cost savings because it avoids the time consuming, and therefore costly, calibration→processing→ recalibration→reprocessing that, with conventional approaches, is so often necessary.

Based on the superior results, it was concluded that the holistic inversions may in fact have provided the better estimate of bias than the conventional high-altitude zero-level estimates. This leads to the possibility that, if the holistic approach is used, it may be

feasible to eliminate high altitude zero-level measurements altogether. This would be of immense practical benefit to the industry because it would result in a substantial reduction in acquisition costs.

The research presented in Chapter 5 was originally only intended to examine the feasibility of applying the same general philosophy, as in frequency-domain holistic inversion, to time-domain AEM data. However, what eventuated was an operational line-by-line inversion method formulated specifically for regional surveys where the line spacing is wide and the inter-flight line coherency cannot be relied upon. The current development deals only with fixed-winged towed-bird AEM data. It also deals only with specific calibration issues faced by those systems, that being, primary field and system geometry estimation.

It was found that, by inverting total field data and solving for three unmeasured elements of the system geometry, the data could be satisfactorily fitted and that the resulting conductivity model was consistent with independent prior information. This was a vast improvement over the conventional approach, in which the data could not be satisfactorily fitted, nor was the conductivity model consistent with prior information. It was further established that by using the holistic inversion spline parameterization, the resulting conductivity model was more continuous and interpretable than if the conventional style discrete parameterization was used. The main advance though was gained by solving for the system geometry calibration parameters.

My overall assessment is that the research presented herein represents a fundamental advance over the conventional way frequency-domain AEM data are calibrated, processed and inverted. Also that the approach may results in cost savings for industry contractors and more accurate conductivity estimates for stakeholders. It can be argued, convincingly in my view, that even though the time-domain holistic inversion development is less comprehensive, it also delivers better outcomes for stakeholders.

6.2 Future work

The holistic approach has introduced greater objectivity and reproducibility through use of a mathematical calibration model and a formal inversion procedure constrained by prior independent information. However, some subjective elements remain. The choice of the regularization parameters (λs) in the objective functions is manual. The λs are chosen to allow convergence to an acceptable data misfit while yielding feasible calibration and conductivity models. Choosing the λs is somewhat subjective and the least robust part of the method. This situation may potentially be improved by using one of the automated methods described by Farquharson and Oldenburg (2004).

An obvious future development for the time-domain method would be its extension to the use of a continuous 3D conductivity model and to simultaneously invert a whole dataset rather than a single flight line. This would allow the across-line spatial coherency of the geology to be exploited in addition to the along-line coherency. This would be straightforward since the frequency-domain formulation could be used without alteration.

Its applicability to different types of time-domain system calibration problems could also be investigated. For example, it is likely that the zero-level drift that is experienced by helicopter time-domain systems could be addressed by a calibration model that parameterizes the drift in piecewise linear fashion in much the same way as it was in the frequency-domain case. Similarly amplitude scaling errors (e.g. Vrbancich and Fullagar, 2007) may be able to be resolvable through the use of a gain parameter in the calibration model.

It is likely that the holistic approach could be successfully employed in other geophysical methods in which the knowledge of the measurement system is incomplete and multiple data channels are coupled via the common underlying geology. Airborne gamma-ray spectrometry is one obvious candidate.

Bibliography

Ackman, T.E. 2003, An introduction to the use of airborne technologies for watershed characterization in mined areas: Mine Water and the Environment, **22**, p62–68.

Anderson, H.A., Duncan, A.C. and Lynch, S.M. 1993, Geological mapping capabilities of the QUESTEM airborne electromagnetic system for mineral exploration — Mt. Isa Inlier, Queensland: Exploration Geophysics, **24**, p333–340.

Anderson, W.L. 1982, Fast Hankel transforms using related and lagged convolutions . ACM Transactions on Mathematical Software, **8**, p344–368.

Annan, A.P. 1983, Effect of differential transmitter/receiver motion on airborne transient EM interpretation: SEG Technical Program Expanded Abstracts, **2**, no. 1, p622-623.

Annan, A.P. 1984, Compensation of towed bird AEM system data for differential transmitter receiver motion: SEG Technical Program Expanded Abstracts, **3**, no. 1, p80-81.

Annan, A.P. and Lockwood, R. 1991, An application of airborne GEOTEM in Australian conditions: Exploration Geophysics, **22**, no. 1, p5-11.

Annetts, D., Sugeng, F. and Raiche, A., 2003, Modelling and matching the airborne EM response of Harmony and Maggie Hays: 16th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Archie, G.E. 1942, The electrical resistivity log as an aid in determining some reservoir characteristics: Transactions of the American Institute of Mining and Metallurgy and Petroleum Engineers, **46**, p54–62.

Aster, R., Borchers, B. and Thurber, C. 2005, Parameter Estimation and Inverse Problems Elsevier Academic Press.

Auken, E. and Christiansen, A.V. 2004, Layered and laterally constrained 2D inversion of resistivity data: Geophysics, **69**, no. 3, p752-761.

Auken, E., Christiansen, A.V., Foged, N. and Sørensen, K., 2008, Accuracy of airborne TEM data for quasilayered modelling: AEM2008: 5th International Conference on Airborne Electromagnetics, Helsinki University of Technology, Conference Abstracts.

Auken, E., Christiansen, A.V., Jacobsen, B.H. and Foged, N. 2005, Piecewise 1D laterally constrained inversion of resistivity data: Geophysical Prospecting, **53**, p497-506.

Auken, E., Westergaard, J., Christiansen, A.V. and Sorensen, K., 2007, Processing and inversion of SkyTEM data for high resolution hydrogeophysical surveys: 19th International Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Backus, G.E. and Gilbert, J.F. 1967, Numerical applications of a formalism for geophysical inverse problems: Geophysical Journal of the Royal Astronomical Society, **13**, p247–276.

Balay, S., Buschelman, K., Eijkhout, V., Gropp, W.D., Kaushik, D., Knepley, M.G., McInnes, L.C., Smith, B.F. and Zhang, H. 2005, PETSc Users Manual: ANL-95/11 - Revision 2.3.0, Argonne National Laboratory, (http://www.mcs.anl.gov/petsc/petsc-as/snapshots/petsc-current/docs/manual.pdf, accessed February 11, 2005).

Barrett, R., Berry, M., Chan, T.F., Demmel, J., Donato, J., Dongarra, J., Eijkhout, V., Pozo, R., Romine, C. and Van der Vorst, H. 1994, Templates for the solution of linear systems: building blocks for iterative methods: SIAM Press, (http://www.netlib.org/templates/templates.pdf, accessed 26 October, 2004).

Barsky, B.A. 1982, End conditions and boundary conditions for uniform B-spline curve and surface representations: Computers in Industry, **3**, p17-29.

Bartels, R.H., Beatty, J.C. and Barsky, B.A. 1987, An introduction to splines for use in computer graphics and geometric modeling: Morgan Kaufmann.

Beamish, D. 2002, The canopy effect in airborne EM: Geophysics, **67**, no. 6, p1720-1728.

Brodie, R. and Sambridge, M., 2004, Holistically calibrating, processing and inverting frequency-domain AEM surveys: 17th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Brodie, R.C. and Fisher, A. 2008, Inversion of TEMPEST AEM survey data, Lower Macquarie River, New South Wales: Report by Geoscience Australia to the Bureau of Rural Sciences.

Brodie, R.C., Green, A.A. and Munday, T.J. 2004a, Calibration of RESOLVE airborne electromagnetic data - Riverland and East Tintinara, South Austrtalia: Open file report 173, Cooperative Research Centre for Landscapes, Environment and Mineral Exploration, (http://crcleme.org.au/Pubs/OPEN%20FILE%20REPORTS/OFR%20171-180/OFR173.pdf, accessed 8 September 2005).

Brodie, R.C., Green, A.A. and Munday, T.J. 2004b, Constrained inversion of RESOLVE electromagnetic data - Riverland, South Australia: Open file report 175, Cooperative Research Centre for Landscapes, Environment and Mineral Exploration, (http://crcleme.org.au/Pubs/OPEN%20FILE%20REPORTS/OFR%20171-180/OFR175.pdf, accessed 8 September 2005).

Brodie, R.C. and Lane, R. 2003, The importance of accurate altimetry in AEM surveys for land management: Exploration Geophysics, **34**, p77-81.

Brodie, R.C. and Sambridge, M. 2006, A holistic approach to inversion of frequency-domain airborne EM data: Geophysics, **71**, no. 6, pG301-G312.

Carter, S., Baron-Hay, S., Combrinck, M. and Moreton, M. 2009, Survey and Logistics Report on a helicopter borne Versatile Time Domain Electromagnetic (VTEM) survey: Pine Creek (Kombolgie) Project Area, Northern Territory, Geotech Airborne Pty. Ltd. report to Geoscience Australia.

Chave, A.D. 1983, Numerical integration of related Hankel transforms by quadrature and continued fraction expansion: Geophysics, **48**, no. 12, p1671-1686.

Christensen, N., Boie 1990, Optimized fast Hankel transform filters: Geophysical Prospecting, **38**, no. 5, p545-568.

Constable, S.C., Parker, R.L. and Constable, C.G. 1987, Occam's inversion; a practical algorithm for generating smooth models from electromagnetic sounding data: Geophysics, **52**, no. 3, p289-300.

Cowey, D., Garrie, D. and Tovey, A. 2003, Riverland and Tintinara, South Australia, RESOLVE Geophysical Survey, Acquisition and Processing Report, Fugro Airborne Surveys report to the Bureau of Rural Sciences, available from Geoscience Australia.

Cunion, E., 2009, Comparison of ground TEM and VTEM responses over kimberlites in the Kalahari of Botswana: 20th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Davis, A.C., Macnae, J.C. and Robb, T. 2006, Pendulum motion in airborne HEM systems: Exploration Geophysics, **37**, p355–362.

Deszcz-Pan, M., Fitterman, D.V. and Labson, V.F. 1998, Reduction of inversion errors in helicopter EM data using auxiliary information: Exploration Geophysics, **29**, no. 1&2, p142–146.

Dongarra, J., Lumsdain, A., Pozo, R. and Remington, K.A. 1996, IML++ v. 1.2: Iterative methods library, reference guide, National Institute of Standards and Technology, U.S.A, (ftp://math.nist.gov/pub/pozo/docs/iml.ps.gz, accessed October 26, 2004).

Ellis, R.G. 1995, Airborne electromagnetic 3d modelling and inversion: Exploration Geophysics, **26**, no. 2 & 3, p138-143.

Ellis, R.G. 1998, Inversion of airborne electromagnetic data: Exploration Geophysics, **29**, no. 1&2, p121-127.

Farquharson, C.G. 2000, Background for Program "EM1DFM", University of British Columbia, Geophysical Inversion Facility, (http://www.eos.ubc.ca/ubcgif/iag/sftwrdocs/em1dfm/bg.pdf, accessed.

Farquharson, C.G. and Oldenburg, D.W. 1993, Inversion of time-domain electromagnetic data for a horizontally layered Earth: Geophysical Journal International **114**, p433-442.

Farquharson, C.G. and Oldenburg, D.W. 1998, Non-linear inversion using general measures of data misfit and model structure: Geophysical Journal International, **134**, no. 1, p213-227.

Farquharson, C.G. and Oldenburg, D.W. 2004, A comparison of automatic techniques for estimating the regularization parameter in non-linear inverse problems: Geophysical Journal International, **156**, no. 3, p411-425.

Farquharson, C.G., Oldenburg, D.W. and Routh, P.S. 2003, Simultaneous 1D inversion of loop-loop electromagnetic data for magnetic susceptibility and electrical conductivity: Geophysics, **68**, no. 6, p1857-1869.

Fitterman, D.V. 1997, Analysis of errors in HEM bird calibration: OF 97-0742, U.S.G.S., (http://pubs.er.usgs.gov/usgspubs/ofr/ofr97742, accessed.

Fitterman, D.V. 1998, Sources of calibration errors in helicopter EM data: Exploration Geophysics, **29**, no. 1&2, p65–70.

Fitterman, D.V. and Deszcz-Pan, M. 1998, Helicopter EM mapping of saltwater intrusion in Everglades National Park, Florida: Exploration Geophysics, **29**, no. 1&2, p240-243.

Fitterman, D.V. and Yin, C. 2004, Effect of bird maneuver on frequency-domain helicopter EM response: Geophysics, **69**, no. 5, p1203-1215.

Fitzpatrick, A. and Munday, T., 2007, The application of airborne geophysical data as a means of better understanding the efficacy of disposal basins along the Murray River; an example from Stockyard Plains, South Australia: 19th International Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Fountain, D. 1998, Airborne electromagnetics — 50 years of development: Exploration Geophysics, **29**, no. 2, p1-11.

Fraser, D.C. 1978, Resistivity mapping with an airborne multicoil electromagnetic system: Geophysics, **43**, no. 1, p144-172.

Fraser, D.C. 1990, Electromagnetic anomaly recognition within geological and cultural noise: in Fitterman, D.V. ed., Developments and applications of modern airborne electromagnetic surveys: U.S. Geological Survey Bulletin, **1925**, p53–64.

Frischknecht, F.C. 1967, Fields about an oscillating magnetic dipole over a two layered earth: Colorado School of Mines Quarterly, **62**, no. 1.

Fugro Airborne Surveys, 2009a, RESOLVE - The World's Most Advanced HEM System, (http://www.fugroairborne.com/service/resolve.php, accessed 23 September 2009).

Fugro Airborne Surveys, 2009b, Calibration of the RESOLVE Resistivity Mapping System, (http://www.fugroairborne.com/resources/technical_notes/heli_em/resolve_1.html, accessed 23 September 2009).

Gamey, T.J., Doll, W.E., Duffy, A. and Millhouse, D.S., 2000, Evaluation of improved airborne techniques for detection of UXO: Proceedings of the Symposium on the Application of Geophysics to Engineering and Environmental Problems, Environmental and Engineering Geophysical Society. Wheat Ridge, CO, United States, Abstracts, p57–66.

George, R.J., Beasley, R., Gordon, I., Heislers, D., Speed, R., Brodie, R., McConnell, C. and Woodgate, P.W. 1998, National Report: Evaluation of Airborne Geophysics for Catchment Management: Technical report to the Department of Agriculture, Fisheries and Forestry Australia and the National Dryland Salinity Program.

Ghosh, D., P 1971, The application of linear filter theory to the direct interpretation of geoelectrical resistivity sounding: Geophysical Prospecting, **19**, no. 2, p192-217.

Green, A. 1998a, The use of multivariate statistical techniques for the analysis and display of AEM data: Exploration Geophysics, **29**, no. 1 & 2, p77-82.

Green, A. 1998b, Altitude correction of time domain AEM data for image display and geological mapping using the Apparent Dipole Depth (ADD) method: Exploration Geophysics, **29**, no. 2, p87-91.

Green, A., 2003, Correcting drift errors in HEM data: 16th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Green, A. and Lane, R., 2003, Estimating noise levels in AEM data: 16th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Green, A.A., Brodie, R.C. and Munday, T.J., 2004, Interpretation of helicopter AEM data from the Riverland area: 17th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Guptasarma, D. and Singh, B. 1997, New digital linear filters for Hankel J0 and J1 transforms: Geophysical Prospecting, **45**, no. 5, p745-762.

Gyulai, A. and Ormos, T. 1999, A new procedure for the interpretation of VES data; 1.5-D simultaneous inversion method: Journal of Applied Geophysics, **41**, no. 1, p1-17.

Hager, W. 1984, Condition estimates: SIAM Journal on Scientific and Statistical Computing, **5**, no. 2, p311-316.

Higram, N.J. and Tisseur, F. 2000, A block algorithm for matrix 1-norm estimation, with an application to 1-Norm pseudospectra: SIAM Journal on Matrix Analysis and Applications, **21**, no. 4, p1185-1201.

Hodges, G., 1999, A world of applications for helicopter electromagnetics to environmental and engineering problems: Proceedings of the symposium on the Application of geophysics to engineering and environmental problems, Environmental and Engineering Geophysical Society, Wheat Ridge, CO, United States, p899-907.

Hodges, G., 2004, Mapping conductivity, magnetic susceptibility, and dielectric permittivity with helicopter electromagnetic data: 74th Annual Meeting, Society of Exploration Geophysicists, Expanded Abstracts.

Hodges, G., Munday, T., Heydorn, A. and Fitzpatrick, A. 2007, Improved Mapping of Conductive Clays and Groundwater Salinity Using Attitude-Corrected Helicopter-Borne EM: Symposium on the Application of Geophysics to Engineering and Environmental Problems, **20**, no. 1, p418-425.

Huang, H. 2008, Airborne geophysical data leveling based on line-to-line correlations: Geophysics, **73**, no. 3, pF83-F89.

Huang, H. and Fraser, D.C. 1999, Airborne resistivity data leveling: Geophysics, **64**, no. 2, p378-385.

Huang, H. and Fraser, D.C. 2001, Mapping of the resistivity, susceptibility, and permittivity of the Earth using a helicopter-borne electromagnetic system: Geophysics, **66**, no. 1, p148-157.

Huang, H. and Fraser, D.C. 2003, Inversion of helicopter electromagnetic data to a magnetic conductive layered earth: Geophysics, **68**, no. 4, p1211-1223.

Johansen, H.K. and Soerensen, K. 1979, Fast Hankel transforms: Geophysical Prospecting, **27**, no. 4, p876-901.

Jones, G. and Spring, J. 2003, Report on downhole geophysical logging of bores in the Riverland Region, SA for the SA-SMMSP, Bureau of Rural Sciences, Canberra.

Keller, G.V. 1988, Rock and mineral properties, *in* Nabighian, M.N. ed., Electromagnetic methods in applied geophysics, Volume 1, Theory; Investigations in Geophysics No. 3: Society of Exploration Geophysicists, p13-51.

Klinkert, P.S., Leggatt, P.B. and Hage, T.B., 1997, The Spectrem airborne electromagnetic system – latest developments and field examples: Fourth International Conference on Mineral Exploration, Proceedings of Exploration 97, p557-563.

Koefoed, O. 1972, A note on the linear filter method of interpreting resistivity sounding data: Geophysical Prospecting, **20**, no. 2, p403-405.

Konishi, N. 1998, Landslide surveys in Tertiary soft rock areas using HEM: Exploration Geophysics, **29**, p234–239.

Kovacs, A. and Valleau, N.C. 1990, Airborne electromagnetic measurement of sea ice thickness and sub-ice bathymetry: in Fitterman, D.V. ed., Developments and applications of modern airborne electromagnetic surveys: U.S. Geological Survey Bulletin, **1925**, p165–196.

Lane, R. 2002, Ground and airborne electromagnetic methods, *in* Papp, E. ed., Geophysical and remote sensing methods for regolith exploration: Cooperative Research Centre for Landscapes, Environment and Mineral Exploration, p53-79. (http://crcleme.org.au/Pubs/OPEN%20FILE%20REPORTS/OFR%20144/08Electromagnetics.pdf, accessed 10 October, 2009).

Lane, R., Brodie, R.C. and Fitzpatrick, A. 2004a, Constrained inversion of AEM data from the Lower Balonne area, Southern Queensland, Australia: Open file report 163, Cooperative Research Centre for Landscapes, Environment and Mineral Exploration, (http://crcleme.org.au/Pubs/OPEN%20FILE%20REPORTS/OFR%20161-162-163-164-166-167/OFR163.pdf, accessed 14 December 2005).

Lane, R., Brodie, R.C. and Fitzpatrick, A., 2004b, A revised inversion model parameter formulation for fixed wing transmitter loop – towed bird receiver coil time-domain airborne electromagnetic data: 17th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Lane, R., Green, A., Golding, C., Owers, M., Pik, P., Plunkett, C., Sattel, D. and Thorn, B. 2000, An example of 3D conductivity mapping using the TEMPEST airborne electromagnetic system: Exploration Geophysics, **31**, p162-172.

Lawrie, K.C., Munday, T.J., Dent, D.L., Gibson, D.L., Brodie, R.C., Wilford, J., Reilly, N.S., Chan, R.N. and Baker, P. 2000a, A geological systems approach to understanding the processes involved in land and water salinisation; the Gilmore project area, central-west New South Wales: AGSO Research Newsletter, **32**, p13-15.

Lawrie, K.C., Munday, T.J., Gibson, D.L., Chan, R., Brodie, R., Wilford, J., Reilly, N. and Foster, K. 2000b, Mapping porphyry Au-Cu and epithermal Au mineral systems under complex regolith cover, Lachlan fold belt, N.S.W: 15th Australian geological convention. Sydney, N.S.W. Australia, Geological Society of Australia Sydney N.S.W. Australia.

Lawson, C.L. and Hanson, D.J. 1974, Solving Least Squares Problems: Englewood Cliffs.

Leggatt, P.B., Klinkert, P.S. and Hage, T.B. 2000, The Spectrem airborne electromagnetic system—further developments: Geophysics, **65**, no. 6, p1976-1982.

Ley-Cooper, Y. and Macnae, J. 2004, Model-consistent rescaling to correct amplitude calibration problems in HEM data: Exploration Geophysics, **35**, no. 4, p277-282.

Ley-Cooper, Y., Macnae, J., Robb, T. and Vrbancich, J. 2006, Identification of calibration errors in helicopter electromagnetic (HEM) data through transform to the altitude-corrected phase-amplitude domain: Geophysics, **71**, no. 2, pG27-G34.

Lines, L., R and Treitel, S. 1984, Tutorial: A review of least-squares inversion and its application to geophysical problems: Geophysical Prospecting, **32**, p159-186.

Liu, G. and Becker, A. 1990, Two-dimensional mapping of sea-ice keels with airborne electromagnetics: Geophysics, **55**, no. 2, p239-248.

Luyendyk, A.P.J. 1997, Processing of airborne magnetic data: AGSO Journal of Australian Geology and Geophysics, **17**, no. 2, p31-38.

Macnae, J.C. 1995, Esoteric and mundane geophysics for diamondiferous pipe exploration: Exploration Geophysics, **26**, no. 2 & 3, p131-137.

Macnae, J.C., King, A., Stolz, N., Osmakoff, A. and Blaha, A. 1998, Fast AEM data processing and inversion: Exploration Geophysics, **29**, p163–169.

Macnae, J.C. and Lamontagne, Y. 1987, Imaging quasi-layered conductive structures by simple processing of transient electromagnetic data: Geophysics, **52**, no. 4, p545-554.

Macnae, J.C., Lamontagne, Y. and West, G.F. 1984, Noise processing techniques for time-domain EM systems: Geophysics, **49**, no. 7, p934-948.

Marquardt, D., W 1963, An algorithm for least squares estimation of non-linear parameters: Journal of the Society of Industrial and Applied Mathematics, **11**, p431-441.

Maxwell, J.C. 1892, A treatise on electricity and magnetism: Clarendon Press. (An unabridged, slightly altered, republication of the two third edition volumes bound as one: Dover Publications, 1954)

McNeill, J.D. 1980, Electrical conductivity of soils and rocks, Geonics Ltd Technical Note TN-5.

Menke, W. 1989, Geophysical data analysis : discrete inverse theory: Academic Press.

Minty, B.R.S. 1991, Simple micro-leveling for aeromagnetic data: Exploration Geophysics, **22**, p591–592.

Munday, T.J., Hill, A.J., Wilson, T., Hopkins, B., Telfer, A.L., White, G. and Green, A.A. 2004, Combining geology and geophysics to develop a hydrogeologic framework for salt interception in the Loxton Sands Aquifer, Central Murray Basin, Australia: Open file report 180, Cooperative Research Centre for Landscapes, Environment and Mineral Exploration, (http://crcleme.org.au/Pubs/OPEN%20FILE%20REPORTS/OFR%20171-180/OFR180.pdf, accessed 8 September 2005).

Nabighian, M.N. 1979, Quasi-static transient response of a conducting half-space - An approximate representation: Geophysics, **44**, no. 10, p1700-1705.

Nabighian, M.N. and Macnae, J.C. 1991, Time domain electromagnetic prospecting methods, *in* Nabighian, M.N. ed., Electromagnetic methods in applied geophysics, Volume 2, Application, Parts A and B; Investigations in Geophysics No. 3: Society of Exploration Geophysicists, p427-520.

Noteboom, M. and Stenning, L. 2008, Lower Macquarie River Airborne Electromagnetic Mapping Survey Acquisition and Processing Report, Fugro Airborne Surveys report to the Bureau of Rural Sciences, available from Geoscience Australia.

Owers, M., Sattel, D. and Stenning, L. 2001, Billabong Creek, NSW, TEMPEST Geophysical Survey, Acquisition and Processing Report, Fugro Airborne Surveys report to the Bureau of Rural Sciences, available from Geoscience Australia.

Palacky, G.J. 1990, Airborne electromagnetics, geological rnapping, and prospecting for nontraditional targets: in Fitterman, D.V. ed., Developments and applications of modern airborne electromagnetic surveys: U.S. Geological Survey Bulletin, **1925**, p89—101.

Palacky, G.J. 1993, Use of airborne electromagnetic methods for resource mapping: Advances in Space Research, **13**, no. 11, p5-13.

Palacky, G.J. and West, G.F. 1991, Airborne electromagnetic methods, *in* Nabighian, M.N. ed., Electromagnetic methods in applied geophysics, Volume 2, Application, Parts A and B; Investigations in Geophysics No. 3: Society of Exploration Geophysicists, p811-879.

Parker, R.L. 1994, Geophysical inverse theory: Princeton University Press.

Pellerin, L. 2002, Applications of electrical and electromagnetic methods for environmental and geotechnical investigations: Surveys in Geophysics, **23**, p101–132.

Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P. 2002, Numerical recipes in C++: The art of scientific computing, Second edition: Cambridge University Press.

Raiche, A. 1998, Modelling the time-domain response of AEM systems: Exploration Geophysics, **29**, no. 1&2, p103-106.

Rawlinson, N. 2000, Inversion of seismic data for layered crustal structure: Ph.D. thesis, Department of Earth Sciences, Monash University.

Reid, J. and Viezzoli, A., 2007, High-resolution near surface airborne electromagnetics; SkyTEM survey for uranium exploration at Pells Range, WA: 19th International Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Reid, J.E., Pfaffling, A. and Vrbancich, J. 2006, Airborne electromagnetic footprints in 1D earths: Geophysics, **71**, no. 2, pG63-G72.

Reid, J.E. and Vrbancich, J. 2004, A comparison of the inductive-limit footprints of airborne electromagnetic configurations: Geophysics, **69**, no. 5, p1229-1239.

Reid, J.E., Vrbancich, J. and Worby, A.P. 2003, A comparison of shipborne and airborne electromagnetic methods for Antarctic sea ice thickness measurements: Exploration Geophysics, **34**, p46–50.

Sambridge, M. 1999, Geophysical inversion with a Neighbourhood algorithm, I, Searching a parameter space: Geophysical Journal International, **138**, p479–494.

Sambridge, M. and Mosegaard, K. 2002, Monte Carlo methods in Geophysical inverse problems: Reviews in Geophysics, **40**, no. 3, p1-29.

Sambridge, M., Rickwood, P., Rawlinson, N. and Sommacal, S. 2007, Automatic differentiation in geophysical inverse problems: Geophysical Journal International, **170**, p1-8.

Sasaki, S., 2004, Inversion of airborne EM data accounting for terrain and inaccurate flight height: 74th Annual Meeting, Society of Exploration Geophysicists, Expanded Abstracts.

Sattel, D. 1998, Conductivity information in three dimensions: Exploration Geophysics, **29**, no. 1&2, p157-162.

Sattel, D. and Kgotlhang, L. 2004, Groundwater Exploration with AEM in the Boteti area, Botswana: Exploration Geophysics, **35**, p147-156.

Sattel, D., Lane, R. and Pears, G., 2004, Novel ways to process and model GEOTEM data: 17th Geophysical Conference and Exhibition., Australian Society of Exploration Geophysicists, Extended Abstracts.

Scales, J.A., Docherty, P. and Gersztenkorn, A. 1990, Regularisation of nonlinear inverse problems: imaging the near-surface weathering layer: Inverse Problems, **6**, no. 1, p115-131.

Schaefer, M.J., Gingerich, J. and Lemieux, J. 1998, Case study: The evolution of airborne time domain electromagnetic applications for geologic mapping; a Noranda perspective: Exploration Geophysics, **29**, p204–210.

Sengpiel, K.P. 1983, Resistivity/depth mapping with airborne eletromagnetic survey data: Geophysics, **48**, no. 2, p181-196.

Sengpiel, K.P. and Siemon, B. 1998, Examples of 1-D inversion of multifrequency HEM data from 3-D resistivity distributions: Exploration Geophysics, **29**, no. 1&2, p133-141.

Siemon, B. 2009, Levelling of helicopter-borne frequency-domain electromagnetic data: Airborne Geophysics, **67**, no. 3, p206-218.

Smith, R., Fountain, D. and Allard, M. 2003, The MEGATEM fixed-wing transient EM system applied to mineral exploration: a discovery case history: First Break, **21**, no. 7, p73-77.

Smith, R.S. 2001a, On removing the primary field from fixed–wing time–domain airborne electromagnetic data: some consequences for quantitative modelling, estimating bird position and detecting perfect conductors: Geophysical Prospecting, **49**, p405–416.

Smith, R.S. 2001b, Tracking the transmitting–receiving offset in fixed–wing transient EM systems: methodology and application: Exploration Geophysics, **32**, p14–19.

Sorensen, C., Fisher, A. and Costelloe, M., 2009, Airborne electromagnetic survey results from the Paterson Province, WA: 20th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Sorensen, K., I. and Auken, E. 2004, SkyTEM; a new high-resolution helicopter transient electromagnetic system: Exploration Geophysics, **35**, no. 3, p194-202.

Spies, B., Fitterman, D., Holladay, S. and Liu, G. 1998, Summaries, discussion and future trends: Exploration Geophysics, **29**, p263–271.

Street, G.J., Pracilio, G., Nallan-Chakravartula, P., Nash, C., Sattel, D., Owers, M., Triggs, D. and R., L. 1998, National Dryland Salinity Program Airborne geophysical surveys to assist planning for salinity control; 1. Willaura SALTMAP Survey Interpretation Report: National Airborne Geophysics Project, World Geoscience Corporation, Perth, WA.

Sugeng, F. and Raiche, A.P., 2004, Modelling the electromagnetic response in complex geological structures using the 3D finite-element method based on the hexahedral and the tetrahedral edge-element technique: 17th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Tan, K.P., Munday, T.J. and Leaney, F.W. 2004, The validation of RESOLVE helicopter EM data: Mineralogical and petrophysical results from field investigations in the Riverland area, South Australia: Open file report 179, Cooperative Research Centre for Landscapes, Environment and Mineral Exploration, (http://crcleme.org.au/Pubs/OPEN%20FILE%20REPORTS/OFR%20171-180/OFR179.pdf, accessed 8 September 2005).

Tikhonov, A.N. and Arsenin, V.Y. 1977, Solutions of ill-posed problems: John Wiley and Sons, Inc.

Valleau, N.C. 2000, HEM data processing — a practical overview: Exploration Geophysics, **31**, p584-594.

Viezzoli, A., Christiansen, A.V., Auken, E. and Sorensen, K. 2008, Quasi-3D modeling of airborne TEM data by spatially constrained inversion: Geophysics, **73**, no. 3, pF105-F113.

Vrbancich, J. and Fullagar, P.K. 2007, Improved seawater depth determination using corrected helicopter time-domain electromagnetic data: Geophysical Prospecting, **55**, p407–420.

Vrbancich, J., Hallett, M. and Hodges, G. 2000, Airborne electromagnetic bathymetry of Sydney Harbour: Exploration Geophysics, **31**, no. 1&2, p179-186.

Vrbancich, J. and Smith, R. 2005, Limitations of two approximate methods for determining the AEM bird position in a conductive environment: Exploration Geophysics, **36**, p365-373.

Wait, J.R. 1982, Geo-electromagnetism: Academic Press Inc.

Walker, S. and Rudd, J., 2008, Advanced processing of helicopter TEM data: AEM2008: 5th International Conference on Airborne Electromagnetics, Helsinki University of Technology, Conference Abstracts.

Ward, S.H. 1967, Electromagnetic theory for geophysical applications, *in*, Mining Geophysics: Society of Exploration Geophysicists, p10-196.

Ward, S.H. and Hohmann, G.W. 1988, Electromagnetic theory for geophysical applications, *in* Nabighian, M.N. ed., Electromagnetic methods in applied geophysics, Volume 1, Theory; Investigations in Geophysics No. 3: Society of Exploration Geophysicists, p131-311.

West, G.F. and Macnae, J.C. 1991, The physics of the electromagnetic induction exploration method, *in* Nabighian, M.N. ed., Electromagnetic methods in applied geophysics, Volume 2, Application, Parts A and B; Investigations in Geophysics No. 3: Society of Exploration Geophysicists, p5-45.

Witherly, K. and Irvine, R., 2006, The VTEM heli-time domain EM system-Four Case Studies: 18th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Witherly, K., Irvine, R. and Morrison, E., 2004, The Geotech VTEM time-domain helicopter EM system: 74th Annual International Meeting, Society of Exploration Geophysicists, Expanded Abstracts, p1217-1220.

Wolfgram, P. and Golden, H. 2001, Airborne EM Applied to Sulphide Nickel - Examples and Analysis: Exploration Geophysics, **32**, p136 - 140.

Wolfgram, P. and Vrbancich, J., 2006, Layered-earth inversions of AEM bathymetry data incorporating aircraft attitude and bird offset – a case study of Torres Strait: 17th Geophysical Conference and Exhibition, Australian Society of Exploration Geophysicists, Extended Abstracts.

Worrall, L., Munday, T. and Green, A. 1998, Beyond bump finding — airborne electromagnetics for mineral exploration in regolith dominated terrains: Exploration Geophysics, **29**, p199–203.

Zollinger, R., Morrison, H.F., Lazenby, P.F. and Becker, A. 1987, Airborne electromagnetic bathymetry: Geophysics, **52**, p1127–1137.

Appendix A

Roll, pitch and yaw

Roll, pitch and yaw are defined as counter-clockwise rotations about the origin of the inertial coordinate system's x-axis, y-axis and z-axis respectively, for an observer looking toward the origin. Individually these can be defined by a roll matrix $\mathbf{R}_R$, a pitch matrix $\mathbf{R}_P$, and yaw matrix $\mathbf{R}_Y$ as follows,

$$\mathbf{R}_R = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\psi_R & \sin\psi_R \\ 0 & -\sin\psi_R & \cos\psi_R \end{bmatrix}, \tag{A-1}$$

$$\mathbf{R}_P = \begin{bmatrix} \cos\psi_P & 0 & -\sin\psi_P \\ 0 & 1 & 0 \\ \sin\psi_P & 0 & \cos\psi_P \end{bmatrix}, \text{ and,} \tag{A-2}$$

$$\mathbf{R}_Y = \begin{bmatrix} \cos\psi_Y & \sin\psi_Y & 0 \\ -\sin\psi_Y & \cos\psi_Y & 0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{A-3}$$

where ψ_R is the roll, ψ_P is the pitch and ψ_Y is the yaw rotation angle.

After a loop or coil with reference orientation $\mathbf{v}_0$ undergoes successive yaw then pitch then roll manoeuvres it will have new orientation $\mathbf{v}$. The two vectors $\mathbf{v}_0$ and $\mathbf{v}$ are related through the orthogonal rotation matrix equations,

$$\mathbf{v}_0 = \mathbf{R}\mathbf{v}, \text{ and} \tag{A-4}$$

$$\mathbf{v} = \mathbf{R}^T\mathbf{v}_0 \tag{A-5}$$

where,

$$\mathbf{R} = \mathbf{R}_R \mathbf{R}_P \mathbf{R}_Y \tag{A-6}$$

is the composite or compound rotation matrix. Since rotation matrices are orthogonal the following identities hold true, $\mathbf{R}_R^{-1} = \mathbf{R}_R^T$, $\mathbf{R}_P^{-1} = \mathbf{R}_P^T$, $\mathbf{R}_Y^{-1} = \mathbf{R}_Y^T$ and $\mathbf{R}^{-1} = \mathbf{R}^T$. Using the shorthand $\cos\psi = \mathrm{c}\,\psi$ and $\sin\psi = s\psi$ the composite rotation matrix is,

$$\mathbf{R} = \begin{bmatrix} \mathrm{c}\,\psi_P\,\mathrm{c}\,\psi_Y & \mathrm{c}\,\psi_P\,\mathrm{s}\,\psi_Y & -\mathrm{s}\,\psi_P \\ \mathrm{s}\,\psi_R\,\mathrm{s}\,\psi_P\,\mathrm{c}\,\psi_Y - \mathrm{c}\,\psi_R\,\mathrm{s}\,\psi_Y & \mathrm{s}\,\psi_R\,\mathrm{s}\,\psi_P\,\mathrm{s}\,\psi_Y + \mathrm{c}\,\psi_R\,\mathrm{c}\,\psi_Y & \mathrm{s}\,\psi_R\,\mathrm{c}\,\psi_P \\ \mathrm{c}\,\psi_R\,\mathrm{s}\,\psi_P\,\mathrm{c}\,\psi_Y + \mathrm{s}\,\psi_R\,\mathrm{s}\,\psi_Y & \mathrm{c}\,\psi_R\,\mathrm{s}\,\psi_P\,\mathrm{s}\,\psi_Y - \mathrm{s}\,\psi_R\,\mathrm{c}\,\psi_Y & \mathrm{c}\,\psi_R\,\mathrm{c}\,\psi_P \end{bmatrix}. \tag{A-7}$$

Note that, as usual, the order of multiplication of the rotation matrices is important. In some AEM systems instruments that measure roll, pitch and yaw may use a convention of ordering of the roll, pitch and yaw. It is therefore necessary to ensure that the order of the roll, pitch and yaw matrices in Equation A-6 are arranged so that they are consistent with the convention of the airborne system being modelled.

In the inversion routines we may wish to solve for the orientation of a loop or coil. In which case the partial derivative of the rotation matrix with respect to the rotation angles are required. They are given by,

$$\frac{\partial \mathbf{R}}{\partial \psi_R} = \begin{bmatrix} 0 & 0 & 0 \\ \mathrm{c}\psi_R\mathrm{s}\psi_P\mathrm{c}\psi_Y + \mathrm{s}\psi_R\mathrm{s}\psi_Y & \mathrm{c}\psi_R\mathrm{s}\psi_P\mathrm{s}\psi_Y - \mathrm{s}\psi_R\mathrm{c}\psi_Y & \mathrm{c}\psi_R\mathrm{c}\psi_P \\ -\mathrm{s}\psi_R\mathrm{s}\psi_P\mathrm{c}\psi_Y + \mathrm{c}\psi_R\mathrm{s}\psi_Y & -\mathrm{s}\psi_R\mathrm{s}\psi_P\mathrm{s}\psi_Y - \mathrm{c}\psi_R\mathrm{c}\psi_Y & -\mathrm{s}\psi_R\mathrm{c}\psi_P \end{bmatrix}. \tag{A-8}$$

$$\frac{\partial \mathbf{R}}{\partial \psi_P} = \begin{bmatrix} -\mathrm{s}\psi_P\mathrm{c}\psi_Y & -\mathrm{s}\psi_P\mathrm{s}\psi_Y & -\mathrm{c}\psi_P \\ \mathrm{s}\psi_R\mathrm{c}\psi_P\mathrm{c}\psi_Y & \mathrm{s}\psi_R\mathrm{c}\psi_P\mathrm{s}\psi_Y & -\mathrm{s}\psi_R\mathrm{s}\psi_P \\ \mathrm{c}\psi_R\mathrm{c}\psi_P\mathrm{c}\psi_Y & \mathrm{c}\psi_R\mathrm{c}\psi_P\mathrm{s}\psi_Y & -\mathrm{c}\psi_R\mathrm{s}\psi_P \end{bmatrix}. \tag{A-9}$$

$$\frac{\partial \mathbf{R}}{\partial \psi_Y} = \begin{bmatrix} -\mathrm{c}\psi_P\psi_Y & \mathrm{c}\psi_P\mathrm{c}\psi_Y & 0 \\ -\mathrm{s}\psi_R\mathrm{s}\psi_P\psi_Y - \mathrm{c}\psi_R\mathrm{c}\psi_Y & \mathrm{s}\psi_R\mathrm{s}\psi_P\mathrm{c}\psi_Y - \mathrm{c}\psi_R\mathrm{s}\psi_Y & 0 \\ -\mathrm{c}\psi_R\mathrm{s}\psi_P\mathrm{s}\psi_Y + \mathrm{s}\psi_R\mathrm{c}\psi_Y & \mathrm{c}\psi_R\mathrm{s}\psi_P\mathrm{c}\psi_Y + \mathrm{s}\psi_R\mathrm{s}\psi_Y & 0 \end{bmatrix}. \tag{A-10}$$

Appendix B

Layered-earth model partial derivatives

B.1 Primary field tensor partial derivatives

The primary field Green's tensor $\mathbf{G}^P$ is shown below, followed by its partial derivatives with respect to the geometric parameters $x,\ y,\ z,\ h$ and the conductivity of the kth layer σ_k, and thickness of the kth layer t_k.

$$\mathbf{G}^P = \frac{1}{4\pi}\begin{bmatrix} \frac{3x^2-R^2}{R^5} & \frac{3xy}{R^5} & \frac{3x(z-h)}{R^5} \\ \frac{3xy}{R^5} & \frac{3y^2-R^2}{R^5} & \frac{3y(z-h)}{R^5} \\ \frac{3x(z-h)}{R^5} & \frac{3y(z-h)}{R^5} & \frac{3(z-h)^2-R^2}{R^5} \end{bmatrix} \tag{B-1}$$

$$\frac{\partial \mathbf{G}^P}{\partial x} = \frac{1}{4\pi}\begin{bmatrix} -15\frac{x^3}{R^7}+9\frac{x}{R^5} & -3\frac{y\left(5x^2-R^2\right)}{R^7} & -3\frac{\left(z-h\right)\left(5x^2-R^2\right)}{R^7} \\ -3\frac{y\left(5x^2-R^2\right)}{R^7} & 3\frac{x\left(R^2-5y^2\right)}{R^7} & -15\frac{y\left(z-h\right)x}{R^7} \\ -3\frac{\left(z-h\right)\left(5x^2-R^2\right)}{R^7} & -15\frac{y\left(z-h\right)x}{R^7} & 3\frac{x\left(R^2-5\left(z-h\right)^2\right)}{R^7} \end{bmatrix} \tag{B-2}$$

$$\frac{\partial \mathbf{G}^P}{\partial y} = \frac{1}{4\pi}\begin{bmatrix} -3\frac{y\left(5x^2-R^2\right)}{R^7} & 3\frac{x\left(R^2-5y^2\right)}{R^7} & -15\frac{y\left(z-h\right)x}{R^7} \\ 3\frac{x\left(R^2-5y^2\right)}{R^7} & -15\frac{y^3}{R^7}+9\frac{y}{R^5} & 3\frac{\left(z-h\right)\left(R^2-5y^2\right)}{R^7} \\ -15\frac{y\left(z-h\right)x}{R^7} & 3\frac{\left(z-h\right)\left(R^2-5y^2\right)}{R^7} & 3\frac{y\left(R^2-5\left(z-h\right)^2\right)}{R^7} \end{bmatrix} \tag{B-3}$$

$$\frac{\partial \mathbf{G}^P}{\partial z} = -\frac{\partial \mathbf{G}^P}{\partial h} = \frac{1}{4\pi}\begin{bmatrix} -3\dfrac{(z-h)(5x^2-R^2)}{R^7} & -15\dfrac{xy(z-h)}{R^7} & 3\dfrac{x\left(R^2-5(z-h)^2\right)}{R^7} \\ -15\dfrac{xy(z-h)}{R^7} & 3\dfrac{(z-h)(R^2-5y^2)}{R^7} & 3\dfrac{y\left(R^2-5(z-h)^2\right)}{R^7} \\ 3\dfrac{x\left(R^2-5(z-h)^2\right)}{R^7} & 3\dfrac{y\left(R^2-5(z-h)^2\right)}{R^7} & -15\dfrac{(z-h)^3}{R^7}+9\dfrac{z-h}{R^5} \end{bmatrix} \tag{B-4}$$

$$\frac{\partial \mathbf{G}^P}{\partial \sigma_k} = \frac{\partial \mathbf{G}^P}{\partial t_k} = 0 \tag{B-5}$$

B.2 Secondary field tensor partial derivatives

The secondary field Green's tensor $\mathbf{G}^S$ is shown below, followed by its partial derivatives with respect to the geometric parameters x, y, z, h and the conductivity of the kth layer σ_k, and thickness of the kth layer t_k. Due to lack of space on the page the individual elements of the tensor are shown in the form $\partial \mathbf{G}^S_{row,column}$.

$$\mathbf{G}^S = \frac{1}{4\pi}\begin{bmatrix} \dfrac{(x^2-y^2)T_2}{r^3} - \dfrac{T_0 x^2}{r^2} & 2\dfrac{xyT_2}{r^3} - \dfrac{xyT_0}{r^2} & -\dfrac{xT_1}{r} \\ 2\dfrac{xyT_2}{r^3} - \dfrac{xyT_0}{r^2} & \dfrac{(y^2-x^2)T_2}{r^3} - \dfrac{T_0 y^2}{r^2} & -\dfrac{yT_1}{r} \\ \dfrac{xT_1}{r} & \dfrac{yT_1}{r} & -T_0 \end{bmatrix} \tag{B-6}$$

$$\frac{\partial \mathbf{G}^S_{1,1}}{\partial x} = \frac{1}{4\pi}\left[\frac{(5xy^2-x^3)T_2}{r^5} - \frac{(x^4+x^2y^2)\dfrac{\partial T_0}{\partial x}}{r^4} + \frac{(x^4-y^4)\dfrac{\partial T_2}{\partial x}}{r^5} - 2\frac{T_0 xy^2}{r^4}\right] \tag{B-7}$$

$$\frac{\partial \mathbf{G}^S_{1,2}}{\partial x} = \frac{1}{4\pi}\left[\left(2\frac{x^2y}{r^4} - \frac{y}{r^2}\right)T_0 + \left(2\frac{y}{r^3} - 6\frac{x^2y}{r^5}\right)T_2 + 2\frac{xy\dfrac{\partial T_2}{\partial x}}{r^3} - \frac{xy\dfrac{\partial T_0}{\partial x}}{r^2}\right] \tag{B-8}$$

$$\frac{\partial \mathbf{G}^S_{1,3}}{\partial x} = \frac{1}{4\pi}\left[\frac{(-x^3-xy^2)\dfrac{\partial T_1}{\partial x}}{r^3} - \frac{T_1 y^2}{r^3}\right] \tag{B-9}$$

$$\frac{\partial \mathbf{G}^S_{2,1}}{\partial x} = \frac{1}{4\pi}\left[\left(2\frac{x^2 y}{r^4} - \frac{y}{r^2}\right)T_0 + \left(2\frac{y}{r^3} - 6\frac{x^2 y}{r^5}\right)T_2 + 2\frac{xy\dfrac{\partial T_2}{\partial x}}{r^3} - \frac{xy\dfrac{\partial T_0}{\partial x}}{r^2}\right] \tag{B-10}$$

$$\frac{\partial \mathbf{G}^S_{2,2}}{\partial x} = \frac{1}{4\pi}\left[\frac{\left(x^3 - 5\,y^2 x\right)T_2}{r^5} - \frac{\left(x^2 y^2 + y^4\right)\dfrac{\partial T_0}{\partial x}}{r^4} + \frac{\left(y^4 - x^4\right)\dfrac{\partial T_2}{\partial x}}{r^5} + 2\frac{T_0\, y^2 x}{r^4}\right] \tag{B-11}$$

$$\frac{\partial \mathbf{G}^S_{2,3}}{\partial x} = \frac{1}{4\pi}\left[\frac{yT_1\, x}{r^3} - \frac{y\dfrac{\partial T_1}{\partial x}}{r}\right] \tag{B-12}$$

$$\frac{\partial \mathbf{G}^S_{3,1}}{\partial x} = \frac{1}{4\pi}\left[\frac{\left(x^3 + xy^2\right)\dfrac{\partial T_1}{\partial x}}{r^3} + \frac{T_1\, y^2}{r^3}\right] \tag{B-13}$$

$$\frac{\partial \mathbf{G}^S_{3,2}}{\partial x} = \frac{1}{4\pi}\left[-\frac{yT_1\, x}{r^3} + \frac{y\dfrac{\partial T_1}{\partial x}}{r}\right] \tag{B-14}$$

$$\frac{\partial \mathbf{G}^S_{3,3}}{\partial x} = \frac{1}{4\pi}\left[-\frac{\partial T_0}{\partial x}\right] \tag{B-15}$$

$$\frac{\partial \mathbf{G}^S_{1,1}}{\partial y} = \frac{1}{4\pi}\left[\frac{\left(y^3 - 5\,x^2 y\right)T_2}{r^5} - \frac{\left(y^2 x^2 + x^4\right)\dfrac{\partial T_0}{\partial y}}{r^4} + \frac{\left(x^4 - y^4\right)\dfrac{\partial T_2}{\partial y}}{r^5} + 2\frac{T_0\, x^2 y}{r^4}\right] \tag{B-16}$$

$$\frac{\partial \mathbf{G}^S_{1,2}}{\partial y} = \frac{1}{4\pi}\left[\left(2\frac{y^2 x}{r^4} - \frac{x}{r^2}\right)T_0 + \left(2\frac{x}{r^3} - 6\frac{y^2 x}{r^5}\right)T_2 + 2\frac{xy\dfrac{\partial T_2}{\partial y}}{r^3} - \frac{xy\dfrac{\partial T_0}{\partial y}}{r^2}\right] \tag{B-17}$$

$$\frac{\partial \mathbf{G}^S_{1,3}}{\partial y} = \frac{1}{4\pi}\left[\frac{xT_1\, y}{r^3} - \frac{x\dfrac{\partial T_1}{\partial y}}{r}\right] \tag{B-18}$$

$$\frac{\partial \mathbf{G}^S_{2,1}}{\partial y} = \frac{1}{4\pi}\left[\left(2\frac{y^2 x}{r^4} - \frac{x}{r^2}\right)T_0 + \left(2\frac{x}{r^3} - 6\frac{y^2 x}{r^5}\right)T_2 + 2\frac{xy\dfrac{\partial T_2}{\partial y}}{r^3} - \frac{xy\dfrac{\partial T_0}{\partial y}}{r^2}\right] \tag{B-19}$$

$$\frac{\partial \mathbf{G}^S_{2,2}}{\partial y} = \frac{1}{4\pi}\left[\frac{\left(5x^2 y - y^3\right)T_2}{r^5} - \frac{\left(x^2 y^2 + y^4\right)\dfrac{\partial T_0}{\partial y}}{r^4} + \frac{\left(y^4 - x^4\right)\dfrac{\partial T_2}{\partial y}}{r^5} - 2\frac{T_0\, yx^2}{r^4}\right] \tag{B-20}$$

$$\frac{\partial \mathbf{G}^S_{2,3}}{\partial y} = \frac{1}{4\pi}\left[\frac{\left(-x^2 y - y^3\right)\dfrac{\partial T_1}{\partial y}}{r^3} - \frac{T_1\, x^2}{r^3}\right] \tag{B-21}$$

$$\frac{\partial \mathbf{G}^S_{3,1}}{\partial y} = \frac{1}{4\pi}\left[-\frac{xT_1\, y}{r^3} + \frac{x\dfrac{\partial T_1}{\partial y}}{r}\right] \tag{B-22}$$

$$\frac{\partial \mathbf{G}^S_{3,2}}{\partial y} = \frac{1}{4\pi}\left[\frac{\left(yx^2 + y^3\right)\dfrac{\partial T_1}{\partial y}}{r^3} + \frac{T_1\, x^2}{r^3}\right] \tag{B-23}$$

$$\frac{\partial \mathbf{G}^S_{3,3}}{\partial y} = \frac{1}{4\pi}\left[-\frac{\partial T_0}{\partial y}\right] \tag{B-24}$$

The remaining partial tensor derivatives of $\mathbf{G}^S$ with respect to z, h, σ_k, and t_k have an equivalent structure defined below as $\partial \mathbf{G}^S / \partial \Omega$ where $\Omega = z, h, \sigma_k, t_k$. The elements $\partial T_0 / \partial \Omega$, $\partial T_1 / \partial \Omega$, and $\partial T_2 / \partial \Omega$ are given in Section B.3.

$$\frac{\partial \mathbf{G}^S_{1,1}}{\partial \Omega} = \frac{1}{4\pi}\left[-\frac{x^2}{r^2}\frac{\partial T_0}{\partial \Omega} + \frac{\left(x^2 - y^2\right)}{r^3}\frac{\partial T_2}{\partial \Omega}\right] \tag{B-25}$$

$$\frac{\partial \mathbf{G}^S_{1,2}}{\partial \Omega} = \frac{1}{4\pi}\left[-\frac{xy}{r^2}\frac{\partial T_0}{\partial \Omega} + 2\frac{xy}{r^3}\frac{\partial T_2}{\partial \Omega}\right] \tag{B-26}$$

$$\frac{\partial \mathbf{G}^S_{1,3}}{\partial \Omega} = \frac{1}{4\pi}\left[-\frac{x}{r}\frac{\partial T_1}{\partial \Omega}\right] \tag{B-27}$$

$$\frac{\partial \mathbf{G}_{2,1}^{S}}{\partial \Omega} = \frac{1}{4\pi}\left[-\frac{xy}{r^2}\frac{\partial T_0}{\partial \Omega} + 2\frac{xy}{r^3}\frac{\partial T_2}{\partial \Omega}\right] \tag{B-28}$$

$$\frac{\partial \mathbf{G}_{2,2}^{S}}{\partial \Omega} = \frac{1}{4\pi}\left[-\frac{y^2}{r^2}\frac{\partial T_0}{\partial \Omega} + \frac{\left(y^2 - x^2\right)}{r^3}\frac{\partial T_2}{\partial \Omega}\right] \tag{B-29}$$

$$\frac{\partial \mathbf{G}_{2,3}^{S}}{\partial \Omega} = \frac{1}{4\pi}\left[-\frac{y}{r}\frac{\partial T_1}{\partial \Omega}\right] \tag{B-30}$$

$$\frac{\partial \mathbf{G}_{3,1}^{S}}{\partial \Omega} = \frac{1}{4\pi}\left[\frac{x}{r}\frac{\partial T_1}{\partial \Omega}\right] \tag{B-31}$$

$$\frac{\partial \mathbf{G}_{3,2}^{S}}{\partial \Omega} = \frac{1}{4\pi}\left[\frac{y}{r}\frac{\partial T_1}{\partial \Omega}\right] \tag{B-32}$$

$$\frac{\partial \mathbf{G}_{3,3}^{S}}{\partial \Omega} = \frac{1}{4\pi}\left[-\frac{\partial T_0}{\partial \Omega}\right] \tag{B-33}$$

B.3 Hankel transform integral partial derivatives

The three Hankel transform integrals are shown below followed by their partial derivatives with respect to the geometric parameters $x,\ y,\ z,\ h$ and the conductivity of the kth layer σ_k, and thickness of the kth layer t_k. The required terms $\partial \mathsf{R}_0\left(\lambda\right)\big/\partial \sigma_k$ and $\partial \mathsf{R}_0\left(\lambda\right)\big/\partial t_k$ are shown in the section B.4 following.

For ease of nomenclature it is convenient to first define the general integral,

$$I_{n\nu} = \int_0^\infty \mathsf{R}_0\left(\lambda\right)\lambda^n e^{-\lambda\left(z+h\right)} J_\nu\left(\lambda r\right) d\lambda \tag{B-34}$$

T_0, and its partial derivatives.

$$T_0 \quad = -\int_0^\infty \mathsf{R}_0\left(\lambda\right)\lambda^2 e^{-\lambda\left(z+h\right)} J_0\left(\lambda r\right) d\lambda = -I_{20} \tag{B-35}$$

$$\frac{\partial T_0}{\partial x} = \int_0^\infty \frac{x}{r}\mathsf{R}_0\left(\lambda\right)\lambda^3 e^{-\lambda\left(z+h\right)} J_1\left(\lambda r\right) d\lambda = \frac{x}{r} I_{31} \tag{B-36}$$

$$\frac{\partial T_0}{\partial y} = \int_0^\infty \frac{y}{r}\mathsf{R}_0\left(\lambda\right)\lambda^3 e^{-\lambda\left(z+h\right)} J_1\left(\lambda r\right) d\lambda = \frac{y}{r} I_{31} \tag{B-37}$$

$$\frac{\partial T_0}{\partial z} = \frac{\partial T_0}{\partial h} = \int_0^\infty \mathsf{R}_0\left(\lambda\right)\lambda^3 e^{-\lambda\left(z+h\right)} J_0\left(\lambda r\right) d\lambda = I_{30} \tag{B-38}$$

$$\frac{\partial T_0}{\partial \sigma_k} = -\int_0^\infty \frac{\partial \mathsf{R}_0\left(\lambda\right)}{\partial \sigma_k}\lambda^2 e^{-\lambda\left(z+h\right)} J_0\left(\lambda r\right) d\lambda \tag{B-39}$$

$$\frac{\partial T_0}{\partial t_k} = -\int_0^\infty \frac{\partial \mathsf{R}_0\left(\lambda\right)}{\partial t_k}\lambda^2 e^{-\lambda\left(z+h\right)} J_0\left(\lambda r\right) d\lambda \tag{B-40}$$

T_1 and its partial derivatives.

$$T_1 = -\int_0^\infty \mathsf{R}_0\left(\lambda\right)\lambda^2 e^{-\lambda\left(z+h\right)} J_1\left(\lambda r\right) d\lambda = -I_{21} \tag{B-41}$$

$$\begin{aligned}\frac{\partial T_1}{\partial x} &= -\int_0^\infty \frac{x}{r^2}\mathsf{R}_0\left(\lambda\right)\lambda^2 e^{-\lambda\left(z+h\right)}\left(J_0\left(\lambda r\right)\lambda r - J_1\left(\lambda r\right)\right) d\lambda \\ &= -\frac{x}{r} I_{30} + \frac{x}{r^2} I_{21}\end{aligned} \tag{B-42}$$

$$\begin{aligned}\frac{\partial T_1}{\partial y} &= -\int_0^\infty \frac{y}{r^2}\mathsf{R}_0\left(\lambda\right)\lambda^2 e^{-\lambda\left(z+h\right)}\left(J_0\left(\lambda r\right)\lambda r - J_1\left(\lambda r\right)\right) d\lambda \\ &= -\frac{y}{r} I_{30} + \frac{y}{r^2} I_{21}\end{aligned} \tag{B-43}$$

$$\frac{\partial T_1}{\partial z} = \frac{\partial T_1}{\partial h} = \int_0^\infty \mathsf{R}_0\left(\lambda\right)\lambda^3 e^{-\lambda\left(z+h\right)} J_1\left(\lambda r\right) d\lambda = I_{31} \tag{B-44}$$

$$\frac{\partial T_1}{\partial \sigma_k} = -\int_0^\infty \frac{\partial \mathsf{R}_0\left(\lambda\right)}{\partial \sigma_k}\lambda^2 e^{-\lambda\left(z+h\right)} J_1\left(\lambda r\right) d\lambda \tag{B-45}$$

$$\frac{\partial T_1}{\partial t_k} = -\int_0^\infty \frac{\partial \mathsf{R}_0\left(\lambda\right)}{\partial t_k}\lambda^2 e^{-\lambda\left(z+h\right)} J_1\left(\lambda r\right) d\lambda \tag{B-46}$$

T_2 and its partial derivatives.

$$T_2 = -\int_0^\infty \mathsf{R}_0\left(\lambda\right)\lambda\, e^{-\lambda\left(z+h\right)} J_1\left(\lambda r\right) d\lambda = -I_{11} \tag{B-47}$$

$$\begin{aligned}\frac{\partial T_2}{\partial x} &= -\int_0^\infty \frac{x}{r^2}\mathsf{R}_0\left(\lambda\right)\lambda\, e^{-\lambda\left(z+h\right)}\left(J_0\left(\lambda r\right)\lambda r - J_1\left(\lambda r\right)\right) d\lambda \\ &= -\frac{x}{r} I_{20} + \frac{x}{r^2} I_{11}\end{aligned} \tag{B-48}$$

$$\frac{\partial T_2}{\partial y} = -\int_0^\infty \frac{y}{r^2}\mathsf{R}_0\left(\lambda\right)\lambda e^{-\lambda\left(z+h\right)}\left(J_0\left(\lambda r\right)\lambda r - J_1\left(\lambda r\right)\right)d\lambda \\ = -\frac{y}{r}I_{20} + \frac{y}{r^2}I_{11} \tag{B-49}$$

$$\frac{\partial T_2}{\partial z} = \frac{\partial T_2}{\partial h} = \int_0^\infty \mathsf{R}_0\left(\lambda\right)\lambda^2 e^{-\lambda\left(z+h\right)}J_1\left(\lambda r\right)d\lambda = I_{21} \tag{B-50}$$

$$\frac{\partial T_2}{\partial \sigma_k} = -\int_0^\infty \frac{\partial \mathsf{R}_0\left(\lambda\right)}{\partial \sigma_k}\lambda e^{-\lambda\left(z+h\right)}J_1\left(\lambda r\right)d\lambda \tag{B-51}$$

$$\frac{\partial T_2}{\partial t_k} = -\int_0^\infty \frac{\partial \mathsf{R}_0\left(\lambda\right)}{\partial t_k}\lambda e^{-\lambda\left(z+h\right)}J_1\left(\lambda r\right)d\lambda \tag{B-52}$$

B.4 Complex reflection coefficient partial derivatives

In Section 2.5.2 it was explained how the complex reflection coefficient R_0, whose partial derivatives, is equivalent to the quotient of two elements of a propagation matrix such that $\mathsf{R}_0 = P_{21}/P_{11}$. The propagation matrix,

$$\mathbf{P} = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} = \mathbf{M}_1 \prod_{k=2}^{N_L} \mathbf{M}_k\,, \tag{B-53}$$

was defined as the product of individual layer matrices $\mathbf{M}_k$ defined as,

$$\mathbf{M}_k = \frac{1}{2}\begin{bmatrix} \left(1+\frac{\mu_{k-1}u_k}{\mu_k u_{k-1}}\right) & \left(1-\frac{\mu_{k-1}u_k}{\mu_k u_{k-1}}\right) \\ \left(1-\frac{\mu_{k-1}u_k}{\mu_k u_{k-1}}\right)e^{-2u_{k-1}t_{k-1}} & \left(1+\frac{\mu_{k-1}u_k}{\mu_k u_{k-1}}\right)e^{-2u_{k-1}t_{k-1}} \end{bmatrix},\ k = 2,3,\cdots,N_L\,. \tag{B-54}$$

and for the special case $k = 1$ the layer matrix simplified to,

$$\mathbf{M}_1 = \frac{1}{2}\begin{bmatrix} \left(1+\frac{\mu_0 u_1}{\mu_1 u_0}\right) & \left(1-\frac{\mu_0 u_1}{\mu_1 u_0}\right) \\ \left(1-\frac{\mu_0 u_1}{\mu_1 u_0}\right) & \left(1+\frac{\mu_0 u_1}{\mu_1 u_0}\right) \end{bmatrix}. \tag{B-55}$$

The terms $u_k = \sqrt{(\lambda^2 + i\mu_k \sigma_k \omega)}$ and t_k have the same meaning as in Chapter 2. The partial derivative of R_0 with respect to the a property v_k of the kth layer is given by,

$$\frac{\partial \mathsf{R}_0}{\partial v_k} = \frac{\partial}{\partial v_k}\left(\frac{P_{21}}{P_{11}}\right) = \frac{1}{P_{11}}\frac{\partial P_{21}}{\partial v_k} - \frac{P_{21}}{P_{11}^2}\frac{\partial P_{11}}{\partial v_k}. \tag{B-56}$$

We therefore require,

$$\frac{\partial \mathbf{P}}{\partial v_k} = \begin{cases} \mathbf{M}_1\ \mathbf{M}_2 \ldots \left(\mathbf{M}_k \dfrac{\partial \mathbf{M}_{k+1}}{\partial v_k} + \dfrac{\partial \mathbf{M}_k}{\partial v_k}\mathbf{M}_{k+1}\right)\mathbf{M}_{k+2}\ldots\mathbf{M}_N & 1 \le k < N_L \\ \mathbf{M}_1\ \mathbf{M}_2 \ldots\ldots \mathbf{M}_{k-1}\dfrac{\partial \mathbf{M}_k}{\partial v_k}, & k = N_L \end{cases}. \tag{B-57}$$

Then for the specific derivatives that are required, that is the cases where $v_k = \sigma_k$ and $v_k = t_k$, the following four partial derivative expressions have been derived which are substituted into Equation B-57 and then in turn into Equation B-56.

$$\frac{\partial \mathbf{M}_k}{\partial \sigma_k} = \frac{1}{4}\begin{bmatrix} \dfrac{i\mu_{k-1}\,\omega}{u_k\, u_{k-1}} & \dfrac{-i\mu_{k-1}\,\omega}{u_k\, u_{k-1}} \\ \dfrac{-i\mu_{k-1}\,\omega}{u_k\, u_{k-1}\, e^{2u_{k-1}\,t_{k-1}}} & \dfrac{i\mu_{k-1}\,\omega}{u_k\, u_{k-1}\, e^{2u_{k-1}\,t_{k-1}}} \end{bmatrix}, \tag{B-58}$$

$$\frac{\partial \mathbf{M}_{k+1}}{\partial \sigma_k} = \frac{1}{4}\begin{bmatrix} -a & a \\ \left[a - b(1-c)\right]e^{-2u_k\,t_k} & \left[a - b(1+c)\right]e^{-2u_k\,t_k} \end{bmatrix}, \tag{B-59}$$

where, $a = \dfrac{i\mu_k{}^2 u_{k+1}\,\omega}{\mu_{k+1}\, u_k{}^3}$, $\quad b = i2t_k\,\mu_k\,\omega$, $\quad c = \dfrac{\mu_k\, u_{k+1}}{\mu_{k+1}\, u_k}$,

$$\frac{\partial \mathbf{M}_k}{\partial t_k} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \tag{B-60}$$

$$\frac{\partial \mathbf{M}_{k+1}}{\partial t_k} = \begin{bmatrix} 0 & 0 \\ \dfrac{\left(-\mu_{k+1}\, u_k + \mu_k\, u_{k+1}\right)}{\mu_{k+1} e^{2\, u_k\, t_k}} & \dfrac{\left(-\mu_{k+1}\, u_k - \mu_k\, u_{k+1}\right)}{\mu_{k+1} e^{2\, u_k\, t_k}} \end{bmatrix} \tag{B-61}$$

Appendix C

Uniform bicubic B-splines

Uniform bicubic B-splines are widely used in computer graphics applications, particularly for the approximation of surfaces in 3D. They have also been used in seismic tomography for the parameterization of geological interfaces (e.g., Rawlinson, 2000). A complete description of the theory relating to B-splines is provided by Bartels et al. (1987).

A uniform bicubic B-spline surface is formed by mosaicing together successive surface patches to form a composite surface in much the same way as a patchwork quilt is constructed. The surface is defined by a set of control vertices that must be arranged to form a topologically rectangular 2D mesh, however there is no restriction on the actual physical location of the vertices. The control vertex at the *i,j*th mesh intersection has position vector coordinate $\mathbf{p}_{i,j} = (x_{i,j}, y_{i,j}, z_{i,j})$. The *i,j*th surface patch is bounded by the lines joining vertices $\mathbf{p}_{i,j}, \mathbf{p}_{i+1,j}, \mathbf{p}_{i+1,j+1}$ and $\mathbf{p}_{i,j+1}$. Independent parameters u and v are defined over the *i,j*th patch in local mesh coordinates such that $u = 0$ and $v = 0$ at the *i,j*th mesh intersection and $u = 1$ and $v = 1$ and at the *i+1,j+1*th mesh intersection.

The B-spline surface on the *i,j*th surface patch is defined everywhere on the patch as,

$$\mathbf{B}_{i,j}(u,v) = \sum_{k=-1}^{2} \sum_{l=-1}^{2} b_k(u)\, b_l(v)\, \mathbf{p}_{i+k,j+l}\,, \tag{C-1}$$

where the functions $b_n(t)$ are the set of four uniform cubic B-spline basis functions,

$$\left.\begin{aligned} b_{-1}(t) &= \tfrac{1}{6}(1 - 3t + 3t^2 - t^3) \\ b_{0}(t) &= \tfrac{1}{6}(4 - 6t^2 + 3t^3) \\ b_{1}(t) &= \tfrac{1}{6}(1 + 3t + 3t^2 - 3t^3) \\ b_{2}(t) &= \tfrac{1}{6}t^3 \end{aligned}\right\}. \qquad \text{(C-2)}$$

Equation C-1 represents the general expression for defining a 3D surface from a set of arbitrarily located control vertices. Through the use of arbitrarily located vertices, complex folded and closed surfaces can be represented by uniform bicubic B-splines.

For the present application we do not need to represent such complex surfaces. We only need to represent a physical property distribution that is a single valued function $S(x, y)$ of horizontal spatial location. This allows us to simplify Equation C-1 by firstly choosing the uv control mesh axes to be aligned with the xy spatial coordinate axes. Furthermore because we are not representing a spatial surface, rather than conceptualize $z_{i,j}$ (the z component of the vertex $\mathbf{p}_{i,j}$) as a spatial z-ordinate, it is more intuitive to describe $z_{i,j}$ to simply be the coefficient $c_{i,j}$ of a mesh node.

In this simplified description, which is illustrated in Figure C.1, the mesh is rectangular in horizontal spatial coordinates (i.e., in plan view). It is defined by m columns and n rows of nodes and thus the composite surface has m-1 columns and n-1 rows of surface patches. The ith column of nodes is located at x_i and the jth row of nodes is located at y_j. The spacing between the nodes is $\triangle x$ and $\triangle y$ in the x and y axes directions respectively.

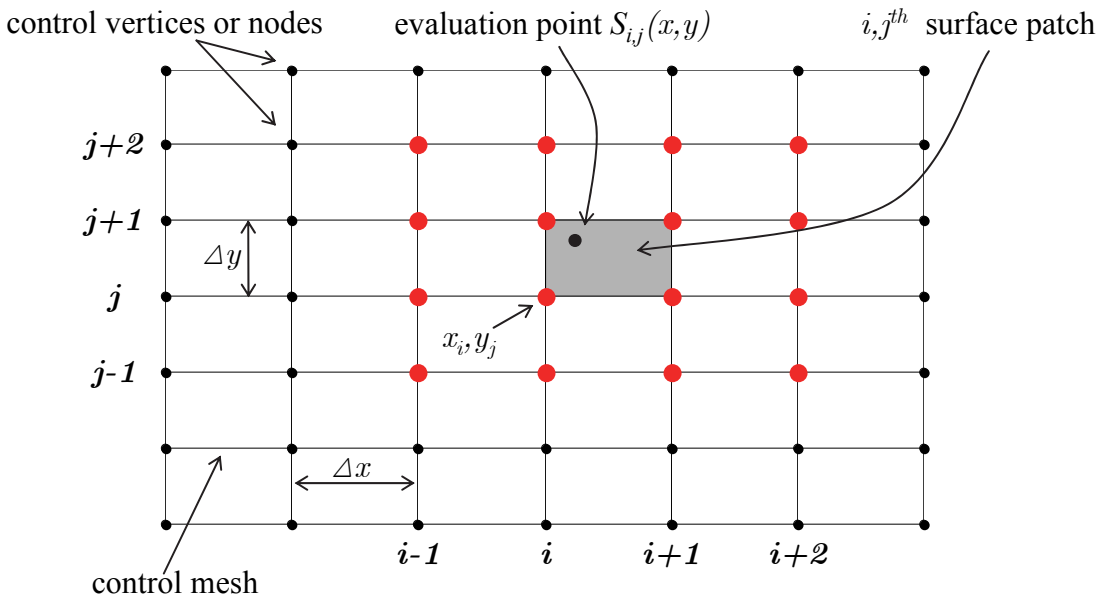

Figure C.1 *Plan view of a portion of a spline mesh showing the mesh, the nodes at the mesh intersections, a surface patch, and an evaluation point on that surface patch.*

Under these conditions the value of the physical property to be represented on the i,jth surface patch can be evaluated by the expression,

$$S_{ij}(x,y) = \sum_{p=i-1}^{i+2} \sum_{q=j-1}^{j+2} w_{ij}(x,y)\, c_{ij} \text{ , where} \tag{C-3}$$

$$w_{ij}(x,y) = b_{p-i}\left(\frac{x-x_i}{\Delta x}\right) \times b_{q-j}\left(\frac{y-y_j}{\Delta y}\right). \tag{C-4}$$

It is immediately clear that the value of the surface at any point on a patch is the weighted sum of the coefficients of the sixteen mesh nodes that are immediately surrounding the patch. All other node coefficients have zero influence on the value of the surface over the patch. For the surface patch shown on Figure C.1 it is the sixteen red nodes that contribute to the surface value. Figure C.2 illustrates the magnitude of the weights $w_{i,j}(x,y)$, calculated from Equation C-4, for each of the sixteen nodes that contribute to the surface over the i,jth surface patch.

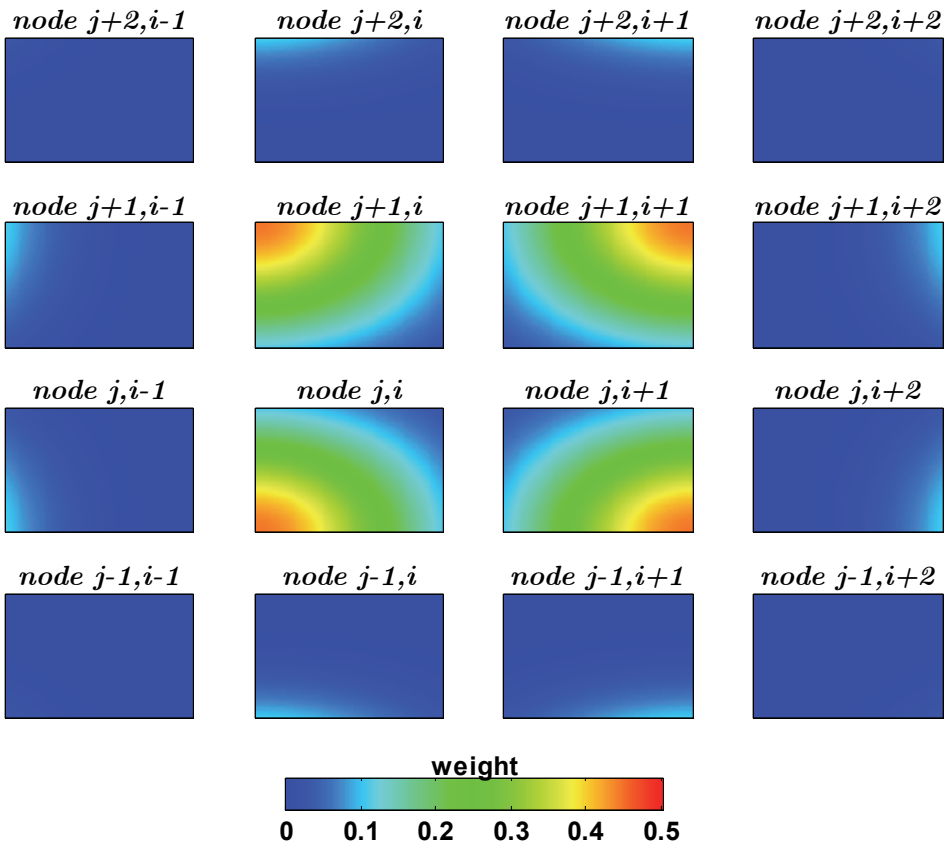

Figure C.2 *The weights that are applied to the coefficients of the sixteen spline nodes that contribute to the value of the surface over the i,jth surface patch.*

So far, the surface is not defined on the patches around the very edge of the composite surface because, to evaluate Equation C-3 we would require an extra perimeter of nodes around the extremities of the mesh. To allow evaluation of the surface on these border patches, an end condition that the second partial derivative be zero at the edge of the composite surface can be applied (Barsky, 1982; Rawlinson, 2000). This is achieved by imagining that phantom nodes exist on the outside perimeter of the defined mesh (i.e., for columns $i=0$ and $i=m+1$ and rows $j=0$ and $j=n+1$). Then, the end condition is satisfied if the phantom node coefficients are the following linear combinations of the defined node coefficients,

$$\left.\begin{aligned}
c_{0,0} &= 4c_{1,1} - 2c_{1,2} - 2c_{2,1} + c_{2,2} \\
c_{0,j} &= 2c_{1,j} - c_{2,j} && j = 1,\ldots,n \\
c_{0,n+1} &= 4c_{1,n} - 2c_{1,n-1} - 2c_{2,n} + c_{2,n-1} \\
c_{i,n+1} &= 2c_{i,n} - c_{i,n-1} && i = 1,\ldots,m \\
c_{m+1,n+1} &= 4c_{m,n} - 2c_{m,n-1} - 2c_{m-1,n} + c_{m-1,n-1} \\
c_{m+1,j} &= 2c_{m,j} - c_{m-1,j} && j = 1,\ldots,n \\
c_{m+1,0} &= 4c_{m,1} - 2c_{m,2} - 2c_{m-1,1} + c_{m-1,2} \\
c_{i,0} &= 2c_{i,1} - c_{i,2} && i = 1,\ldots,m
\end{aligned}\right\}. \qquad \text{(C-5)}$$

Substitution of Equation C-5 into Equation C-3 therefore allows evaluation of the surface over all the border patches.

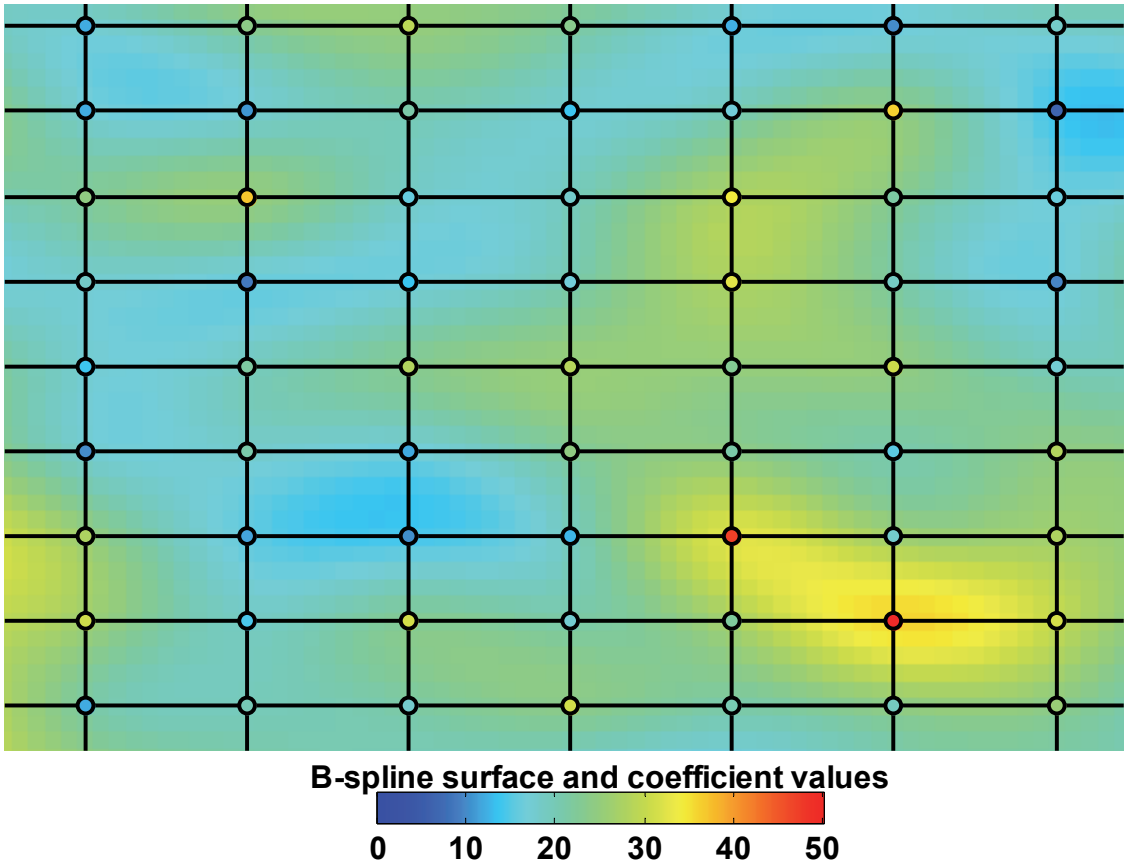

Figure C.3 *A portion of a bicubic B-spline composite surface showing the mesh lines, node coefficients values (enclosed coloured circles) and the resulting surface that was evaluated from the coefficients (background colour).*

The results of the evaluation of a uniform bicubic B-spline surface from the node coefficients are illustrated in Figure C.3. At the mesh intersections, the nodes are displayed and coloured according to their coefficient values. The surface value (Equation C-3) has been evaluated from the coefficients over all surface patches and is

plotted as an image using the same colour lookup table as the coefficients. Two properties of the spline surface become immediately obvious from the figure.

Firstly, unlike some other classes of splines, uniform bicubic B-splines are non-interpolating, that is, the value of the surface at the nodes does not equal the node coefficient. For understanding their use in the holistic inversion method, this is important to the extent that the node coefficients should not be thought of as being directly comparable to the physical property distribution that they are meant to parameterize, but rather that they are abstract quantities.

The second readily recognized property is that the surface is everywhere continuous on each individual surface patch and is also continuous between adjacent patches, i.e. the patches match along the edges where they join. A further property of these surfaces, ensured by their mathematical construction but not readily apparent from Figure C.3, is that they are everywhere C^2 continuous in every direction. This means that their first and second derivatives with respect to distance in any direction are continuous, and this is the reason why the surface can accurately be described as being implicitly smooth and continuous.