Geophysical Prospecting | EAGE European Association of Geoscientists & Engineers

Geophysical Prospecting, 2018, **66**, 1402–1414 doi: 10.1111/1365-2478.12630

A parallel computing thin-sheet inversion algorithm for airborne time-domain data utilising a variable overburden

Tue Boesen[1]*, Esben Auken[1], Anders Vest Christiansen[1], Gianluca Fiandaca[1] and Cyril Schamper[2]

[1]*Aarhus University, 8000 Aarhus, Denmark, and* [2]*Sorbonne Universités, UPMC Univ Paris 06, UMR 7619 METIS, 75005 Paris, France*

Received September 2017, revision accepted February 2018

ABSTRACT

Accurate modelling of the conductivity structure of mineralisations can often be difficult. In order to remedy this, a parametric approach is often used. We have developed a parametric thin-sheet code, with a variable overburden. The code is capable of performing inversions of time-domain airborne electromagnetic data, and it has been tested successfully on both synthetic data and field data. The code implements an integral solution containing one or more conductive sheets, buried in a half-space with a laterally varying conductive overburden. This implementation increases the area of applicability compared to, for example, codes operating in free space, but it comes with a significant increase in computational cost. To minimise the cost, the code is parallelised using OpenMP and heavily optimised, which means that inversions of field data can be performed in hours on multiprocessor desktop computers. The code models the full system transfer function of the electromagnetic system, including variable flight height. The code is demonstrated with a synthetic example imitating a mineralisation buried underneath a conductive meadow. As a field example, the Valen mineral deposit, which is a graphite mineral deposit located in a variable overburden, is successfully inverted. Our results match well with previous models of the deposit; however, our predicted sheet remains inconclusive. These examples collectively demonstrate the effectiveness of our thin-sheet code.

Key words: Electromagnetics, Inversion, Modelling.

INTRODUCTION

Electromagnetic (EM) data are commonly used to detect conductivity and applied in subsurface detection and modelling. One of the earliest uses was the detection of sulphide mineralisations, which often produces signals with orders of magnitudes larger than the non-mineralised background. While detection of such deposits is easy, accurate modelling remains a challenge. In cases where sharp boundaries and large conductivity contrasts are indicated, it can often be beneficial to employ a parametric formulation, where the anomalies are modelled using locally defined discrete objects, described with only a few key parameters. This is often beneficial from a computational point of view, since a parametric approach drastically reduces the number of variables used in a model.

Historically, there have been several parametric approaches made in EM geophysics. Price (1948) introduced the thin-sheet formulation; since then, various authors (Schmucker 1971; Annan 1974; Lajoie and West 1976; Vasseur and Weidelt 1977; Weidelt 1981; Walker and West 1991; Fainberg, Pankratov and Singer 1993; Avdeev, Kuvshinov and Pankratov 1998) have presented numerically stable integral equation formulations of the thin-sheet problem. The formulation by Weidelt (1981) has been the foundation for several other papers that, in some way, extend the

*E-mail: tue@geo.au.dk

formulation, including those by Zhou (1989) and Song, Kim and Lee (2002), as well as this paper.

There are two common areas in controlled-source EM where thin-sheet modelling has been utilised: (i) mineral exploration and (ii) seabed oil exploration. In mineral exploration, free-space sheet modelling is commonly used today. Thin-sheet modelling for oil exploration is possible using resistive sheets in conductive backgrounds (Constable and Weiss 2006; Swidinsky and Edwards 2010).

While thin-sheet modelling is a well-known parametric model used in EM geophysics, other alternative parametrisations exist as well: Dorn, Miller and Rappaport (2000) and Dorn and Lesselier (2006) used a parametric model to simulate plumes from landfills; Aghasi, Kilmer and Miller (2011) developed a general parametric inversion framework using radial basis functions; and McMillan *et al.* (2015) developed a 3D parametric hybrid inversion scheme for airborne EM data using skewed Gaussian ellipsoids to represent the target anomalies.

Focusing on thin-sheet modelling, the simplest and most common approach is to consider a thin sheet in free space (Duncan 1987; Macnae *et al.* 1998). In mineral exploration, this is often a reasonable approximation, since the host rock is usually several orders of magnitude more resistive than the mineral deposit. However, even though the free-space approximation is valid in many cases, there are many examples where the approximation breaks down (Wolfgram, Hyde and Thomson 1998; Reid, Fitzpatrick and Godber 2010).

In order to extend the domain where sheet codes are applicable and improve their accuracy, we have gone beyond the free-space approximation and have extended a layered sheet code, originally developed for far-field data in the frequency domain by Zhou (1989), to consider a two-layered earth where the top layer incorporates a variable overburden. The code, which will be described in detail in the following sections, is implemented in Fortran within an already established and robust inversion engine (Auken *et al.* 2014) and is optimised using task parallelisation OpenMP directives, which provide superior CPU balancing compared to traditional loop parallelisation. Furthermore, a sophisticated adaptive frequency sampler has been developed, in order to minimise the number of frequencies needed to transform the response to time domain. The code is capable of handling multiple sheets and has the option of locking sheets together, in what is commonly referred to as a thick sheet. The code models the full system transfer function of any transient EM system. It has been extensively validated, both against 1D algorithms, as well as against the thin-sheet code presented in Raiche (1998). Finally, we demonstrate the effectiveness of our sheet code by successfully inverting both a synthetic example as well as the Valen mineral deposit.

METHODOLOGY

Governing equations

Starting with Maxwell's equations, under the assumption that displacement currents are negligible and that the medium is non-magnetisable, the thin-sheet theory is formulated following the approach of Weidelt (1981) and Zhou (1989). In the space–frequency domain, Faraday's and Ampere's laws are given as

$$\nabla \times \mathbf{E}(\mathbf{r}) = -i\omega\mu_0\mathbf{H}(\mathbf{r}), \tag{1}$$

$$\nabla \times \mathbf{H}\,(\mathbf{r}) = \sigma(\mathbf{r})\mathbf{E}(\mathbf{r}) + \mathbf{J}_e(\mathbf{r}), \tag{2}$$

where $\mathbf{E}$ and $\mathbf{H}$ are the electric and magnetic fields, ω is the angular frequency, μ_0 is the vacuum permeability, $\boldsymbol{\sigma}$ is the conductivity, and $\mathbf{J}_e$ is the source current density. Equations (1) and (2) can be combined to give

$$\nabla \times \nabla \times \mathbf{E}(\mathbf{r}) + i\omega\mu_0\boldsymbol{\sigma}(\mathbf{r})\mathbf{E}(\mathbf{r}) = -i\omega\mu_0\mathbf{J}_e(\mathbf{r}). \tag{3}$$

This is combined with the thin-sheet approximation:

$$\mathbf{n} \times (\mathbf{H}_+ - \mathbf{H}_-) = \tau(\mathbf{r})\mathbf{E}, \tag{4}$$

where $\mathbf{n}$ is a normal vector to the sheet, $\mathbf{H}_+$ and, $\mathbf{H}_-$ are the magnetic fields above and below the thin sheet, and $\tau(\mathbf{r})$ is the conductance of the sheet.

From equations (3) and (4), an expression for the electric field anywhere in space can be found (Zhou 1989):

$$\mathbf{E}(\mathbf{r}) = \mathbf{E}_n(\mathbf{r}) - i\omega\mu_0 \sum_{j=1}^{N_S} \int_{S_j} \tau_j(\mathbf{r}')\mathbf{G}(\mathbf{r},\mathbf{r}') \cdot \mathbf{E}_{s_j}(\mathbf{r}')dS', \tag{5}$$

where $\mathbf{E}_n(\mathbf{r})$, is the primary electric field:

$$\mathbf{E}_n(\mathbf{r}) = -i\omega\mu_0 \int_V \mathbf{J}_e(\mathbf{r}') \cdot \mathbf{G}(\mathbf{r},\mathbf{r}')dV', \tag{6}$$

where N_s is the number of sheets, S_j is the surface of the jth sheet, τ_j is the conductance of the j'th sheet, $\mathbf{G}(\mathbf{r},\mathbf{r}')$ is the dyadic Green's function, $\mathbf{E}_{S_j}$ is the tangential electric field on the j'th sheet, V is the whole space volume excluding the sheets, and $\mathbf{r}'$ is the corresponding position vector.

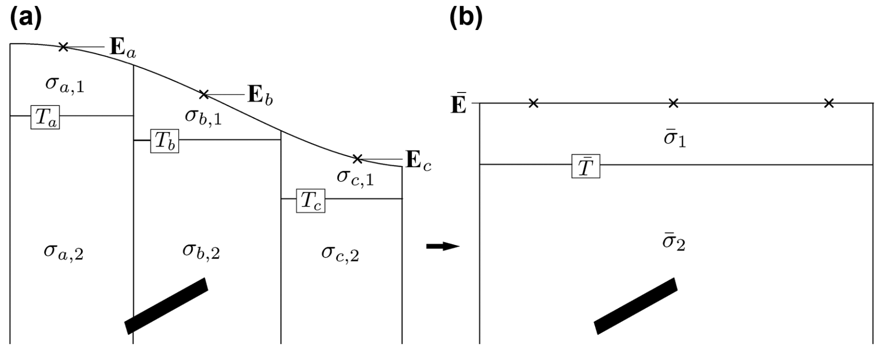

Figure 1 Mapping from individual 1D soundings to an average background model used for sheet calculation. (A) Three individual 1D models, with unique elevation, E, conductivities, σ, and thicknesses, T, for all soundings. (B) The resulting average background model used for the sheet calculation, where all the shown parameters are now weighted mean values calculated in accordance with equation (10).

Since equation (5) is valid anywhere, it can also be used to find the electric field on the sheets; thus, the electric field on the *i*th sheet is given by the following closed form:

$$\mathbf{E}_{S_i}(\mathbf{r}) = \mathbf{E}_n(\mathbf{r}) - i\omega\mu_0 \sum_{j=1}^{N_s} \int_{S_j} \tau_j(\mathbf{r}')\mathbf{G}(\mathbf{r}, \mathbf{r}') \cdot \mathbf{E}_{S_j}(\mathbf{r}')dS'. \tag{7}$$

The electric field on the thin sheet is found by discretising the sheet into rectangular cells and treating each cell as a dipole. By introducing an inductive and a channelling current, as suggested by Weidelt (1981), a coupling between the dipoles can be found. With the coupling matrix known, the response from the sheet to an external field can be found. Finally, in order to get the response in the time domain, a fast Hankel transform is utilised (Johansen and Sørensen 1979).

Sheet-modelling with a varying overburden

The computation of the sheet response with a smooth varying conductive overburden is done by computing one common layered background from all the 1D layered models in the modelling domain. Each of the 1D models represents a sounding measured by the airborne electromagnetic system. The secondary electric fields derived from the sheets are then computed using equation (7), and the fields are added to the layered response from each of the soundings. Computing one common layered background is an approximate mapping that relies on the assumption that the elevation and background resistivity do not vary abruptly, such that the average background model, in which the sheet is calculated, is representative of the actual background. An illustration of this mapping can be seen in Fig. 1 and will be explained more thoroughly in the following.

The mapping is made by giving each sounding an individual weight for each sheet, where the weight, w_{ij}, of the i'th sounding in comparison to the j'th sheet is given as

$$w_{ij} = \frac{\tau_j}{d_{ij}^2}, \tag{8}$$

with τ being the conductance of the sheet and d being the distance between the sounding and the sheet. The conductance contributes to the weight to quantify the importance of the sheet in comparison to other sheets, whereas the distance quantifies the importance of the sounding in comparison to other soundings. Thus, the weight ensures that soundings close to the sheet have a larger impact on the common background model, whereas soundings far away play a negligible role.

The normalised weights, $\hat{w}$, are then given as

$$\hat{w}_{ij} = \frac{w_{ij}}{\sum_{m=1}^{N_{\text{sound}}} \sum_{K=1}^{N_{\text{sheet}}} w_{mk}}, \tag{9}$$

and the weighted average parameters, $\bar{m}$, are calculated as

$$\bar{m} = \sum_{i=1}^{N_{\text{sound}}} \sum_{j=1}^{N_{\text{sheet}}} \widehat{w_{ij}} m_i. \tag{10}$$

Equation (10) is used to determine the average parameter values of elevation, **E**, thickness, T, and conductivity of layers 1 and 2, σ_1, σ_2—all of which are needed in order to create the common background model used for the sheet calculation, as seen in Fig. 1.

Once the common background model has been determined, the sheet response is calculated based on equations (5)–(7), and their constituent magnetic field equations. The variable overburden is included in the total magnetic field, $\mathbf{H}_T$, by adding the secondary magnetic field from the sheet, $\mathbf{H}_S$, to the layered magnetic field responses, $\mathbf{H}_L$, based on 1D models at each sounding position.

$$\mathbf{H}_T = \mathbf{H}_S + \mathbf{H}_L. \tag{11}$$

One way to look at this is that a variable overburden is included by dropping the common background response from the sheet calculation and, instead, recalculating this using the individual 1D models. Thus, the additional cost of using a variable overburden, in comparison to a two-layer homogenous background, is a computation of all the 1D-background responses for each iteration, which is almost negligible in comparison to the cost of the thin-sheet calculation.

Modelling the system response

The full system response is calculated following Auken *et al.* (2014). This implies that the actual bandwidth of the receiver coil and the receiver instrument is modelled using Butterworth filters (Effersø, Auken and Sørensen 1999) and the current waveform is modelled with piecewise linear elements (Fitterman and Stewart 1986). The shape of the transmitter loop is modelled by integrating the response from horizontal electric dipoles following the path of the transmitting wire.

Inversion

The objective function to be minimised in the inversion problem is given as

$$q = q_{\mathrm{obs}} + q_{\mathrm{prior}} + q_{\mathrm{reg}}, \tag{12}$$

with q_{obs} being the observed data misfit, q_{prior} being the misfit to the prior model information about both the layered models as well as sheets, and q_{reg} being the misfit to the regularisation of the layered models. A least squares solution (L2-norm) is used to minimise the objective function incorporating a Tikhonov regularisation scheme. The scheme is an extension of the inversion algorithm described in Auken and Christiansen (2004) and Auken *et al.* (2014).

The n'th iterative update of the model vector $\mathbf{m}$, (which is described in greater detail in Table 1) is given as

$$\mathbf{m}_{n+1} = \mathbf{m}_n + \left(\hat{\mathbf{G}}_n^T \mathbf{C}_n^{-1} \hat{\mathbf{G}}_n + \lambda_n \Gamma_n\right)^{-1} \cdot \left(\hat{\mathbf{G}}_n^T \mathbf{C}_n^{-1} \delta \hat{\mathbf{d}}_n\right), \tag{13}$$

where λ is a damping parameter, Γ is a diagonal scaling matrix, $\delta\hat{\mathbf{d}}$ is an extended perturbed data vector, $\hat{\mathbf{G}}$ is the extended Jacobian, and $\hat{\mathbf{C}}$ is an extended covariance matrix, where the extensions are the *a priori* information and regularisation:

$$\delta\hat{\mathbf{d}} = \begin{bmatrix} \mathbf{d} - \mathbf{d}_{\mathrm{obs}} \\ \mathbf{m} - \mathbf{m}_{\mathrm{prior}} \\ -\mathbf{Rm} \end{bmatrix}, \tag{14}$$

$$\hat{\mathbf{G}} = \begin{bmatrix} \mathbf{G} \\ \mathbf{P} \\ \mathbf{R} \end{bmatrix}, \tag{15}$$

Table 1 Parameters considered by the sheet inversion algorithm. Each sheet is characterised by eight parameters, whereas each sounding is characterised by four parameters. The four sounding parameters are: ρ_1- resistivity of the first layer, ρ_2- resistivity of the second layer, $T-$ thickness of the first layer, and $A-$ flight altitude. The eight sheet parameters are: $\tau-$ conductance, $(x, y, z)-$ coordinate of the sheet centre, $L_{x/y}-$ length of sheet in the *x*/*y*-direction when all angles are 0, $\theta-$ strike angle in comparison to the *x*-axis, and $\varphi-$ dip angle

MODEL PARAMETERS	
For each sounding	For each sheet
ρ_1	τ
ρ_2	x
T	y
A	z
	L_x
	L_y
	θ
	φ

$$\hat{\mathbf{C}} = \begin{bmatrix} \mathbf{C}_{\mathrm{obs}} & 0 & 0 \\ 0 & \mathbf{C}_{\mathrm{prior}} & 0 \\ 0 & 0 & \mathbf{C}_{\mathrm{reg}} \end{bmatrix}, \tag{16}$$

where $\mathbf{d}$ is the forward response data, $\mathbf{d}_{\mathrm{obs}}$ is the observed data, $\mathbf{m}_{\mathrm{prior}}$ contains any *a priori* information about the model parameters, $\mathbf{R}$ is the roughness matrix defining which model parameters are constrained to each other, $\mathbf{G}$ is the Jacobian, and $\mathbf{P}$ defines the constraints of the *a priori* information. $\mathbf{C}_{\mathrm{obs}}$, $\mathbf{C}_{\mathrm{prior}}$, and $\mathbf{C}_{\mathrm{reg}}$ are covariance matrices for the observed data, the prior information, and the roughness matrix, respectively. The diagonal scaling matrix, Γ, is scaled individually for the different types of model parameters in the linear system (which are listed in Table 1). The reason for this is that the different types of model parameters have vastly different sensitivities, and in order to prevent any type of model parameters from being overdamped, each model parameter type is damped individually.

The model parameters of the inversion are given in Table 1. Several different characterisations of the sheet were considered, but by characterising the sheets with a central point and the strike/dip angles, a covariance analysis revealed that these parameters are the least coupled and thus, overall, best suited for inversion. Furthermore, the strike/dip angles of the sheet can often be predicted to some degree based on the data streams and can thus be given a reasonable initial value for the inversion.

Optimisation steps

Dynamic cell discretisation

When doing sheet inversion, the sheet is free to move in the subsurface and change size and shape. After the shape of the sheet is changed during the inversion process, it is imperative that the cell discretisation remains sufficiently fine. Because of this, a dynamic discretisation of the sheet was implemented, which ensures sufficient cell discretisation of the sheet while also keeping the cells as square as possible. The dynamic discretisation is given an upper and a lower bound for the number of cells each sheet can contain, as well as a desired cell size. Each sheet is then initially discretised with cells as close to square as possible, with the desired cell size. If the number of cells falls outside the allowed number of cells, the cells are then lengthened or shortened to meet this criterion as well.

The reason for keeping the cells as square as possible is that it gives higher accuracy to the integration over the sheet than if using elongated cells (not shown). During a typical inversion, roughly 200–1000 cells will be used, depending on the number of sheets, the size of the sheets, and the desired accuracy.

Adaptive frequency sampler

Transient electromagnetic modelling codes are often formulated in the frequency domain, which means that a transformation to the time domain is required. This is most efficiently done using a fast Hankel transform. For such a transformation to be accurate, ~100 frequency responses need to be calculated. In our case, a high-precision adaptive fast Hankel transformation with 14 points per decade is used, for as many decades as needed. This usually results in ~200 frequency responses that need to be determined.

A Hankel transformation can be done using either the real or the quadrature part of the magnetic field. We chose to use the quadrature part, since the adaptive sampler described below performs better on this part, due to the quadrature part being smoother, and thus easier to interpolate than the real part.

In 1D codes, computing a single-frequency response is a computationally inexpensive process, but for higher-dimensional codes, including thin-sheet modelling, calculating a frequency response becomes computationally expensive. This means that it can be computationally beneficial to calculate only selected frequencies and then interpolate to get the remaining. In order to do this, an adaptive frequency sampler

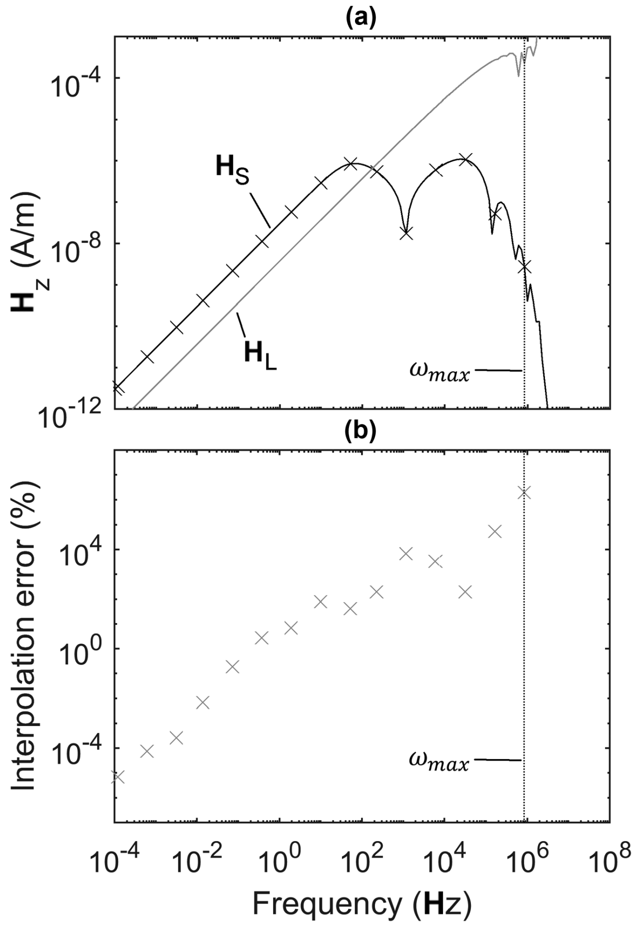

Figure 2 Adaptive frequency sampler. (A) The quadrature part of the magnetic z-component is shown as a function of frequency for both the sheet, $\mathbf{H}_S$, and the layered background, $\mathbf{H}_L$. The lines show the responses for all 14 points per decade, which are needed during the Hankel transform, whereas the crosses on the sheet response curve are the initial equidistant responses. $\omega_{\max}$ is shown as a dashed line and indicates the frequency where sheet responses become negligible and are not calculated during the adaptive sampling. (B) The interpolation error based on the initial sampling. Additional frequencies are calculated around the frequencies with the highest error.

was developed. The adaptive sampler consists of an initial stage, an adaptive stage, and a final interpolating stage. In the initial stage, the frequency sampler starts by calculating a small number of frequencies, spread equidistantly in log-space between $10^{-4} - 10^{8}$ Hz, as well as a single point at 10^{-7} Hz. The total number of initial frequencies is dynamically determined to be divisible with the number of CPUs utilised. At high frequencies, the sheet response falls rapidly towards zero and thus becomes negligible compared to the background response, as seen in Fig. 2A. From the initial calculated frequencies, a dynamically determined frequency limit, $\omega_{\max}$, is made. $\omega_{\max}$ is created such that all calculated frequencies higher than $\omega_{\max}$ fulfil the criterion that the sheet-to-background response ratio is less than 10^{-4}. For frequencies higher than $\omega_{\max}$, it is assumed that the sheet response is negligible, and thus, sheet responses are not calculated above $\omega_{\max}$ during the adaptive sampling.

The adaptive sampling can be envisioned to operate in iterations, where an iteration starts with an interpolation error sampling of the frequency spectrum within $10^{-4} - \omega_{\max}$ Hz. This is done by going through each calculated frequency except the first and the last. For each frequency, its surrounding calculated frequencies are used in an interpolation to estimate the frequency response, which is compared with the actual value. The result of such an interpolation error estimation can be seen in Fig. 2B. The frequencies with the highest interpolation error are refined by calculating two additional frequency responses around the central frequency. This is done for as many frequencies as needed for all CPUs to be working, or until all calculated frequencies have been calculated to a satisfactory level of accuracy. Once a thread has finished its calculation and there are no more available tasks to compute, the next iteration starts. This procedure is repeated until sufficient accuracy is reached or all frequencies needed for the Hankel transform, within the interval, have been calculated. For frequencies below 10^{-4} Hz, the sheet response shows an asymptotic behaviour that makes accurate interpolation possible.

The final stage consists of doing an interpolation of all the frequency responses, including the 10^{-7} Hz point, which was calculated for additional accuracy of the low-frequency interpolation.

The interpolation routine itself is adaptive and uses either a four- or a six-point natural spline interpolation. A four-point natural spline is used for the two endpoints at either end, and a six-point natural spline is used for all other points.

If the sheets are calculated with very few cells, which is sometimes done in preliminary tests, natural spline interpolation becomes too computationally demanding compared to the amount of time the sheet calculation takes. In this case, a simpler linear interpolation is used to estimate the accuracy. The reason it becomes too expensive is that the spline interpolation must be done for each receiver and because it has to be done in an OpenMP critical section, which means that no other parallel thread can access these data while it is happening. Once the adaptive sampling is completed, any frequency needed for the Hankel transform that has not explicitly been calculated during the adaptive scheme is interpolated using natural spline interpolation with all calculated frequencies included.

The benefits gained by using this adaptive frequency sampler are roughly half the number of frequency calculations performed. Thus, among the $\sim 170 - 200$ frequencies needed for the Hankel transform, $\sim 60 - 90$ frequencies with sheet response are usually calculated.

Parallelisation and scalability

Since the paradigm shift in CPUs happened around 2005, where advancements in clock frequency stagnated and multicore CPUs were introduced, parallelisation has become an increasingly important consideration for numerical modelling. Moore's law of exponential computational growth is still more or less observed over the last decade, but only if vectorisation and parallelisation are utilised. On top of this, non-uniform memory access (NUMA) systems have been introduced (Dong, Cooperman and Apostolakis 2010). NUMA systems make data placement essential, and if done incorrectly, a NUMA system can actually run slower by utilising parallelisation.

Our thin-sheet formulation operates in the frequency domain, which makes parallelisation significantly simpler than a time-domain formulation. Parallelisation across frequencies is largely independent and usually not memory bandwidth restricted, since the frequency calculations are sufficiently compute heavy in comparison to the amount of communication required between the threads. Thus, when frequency responses are calculated in an optimised loop parallelisation, the scalability is almost linear. Since the number of frequencies that needs to be calculated during the adaptive sampling is unknown, loop parallelisation would be suboptimal. Instead, OpenMP task parallelisation is employed. This allows tasks to be dynamically spawned and any available thread to perform the task. Thus, near-optimal work balancing across all threads is ensured.

Figure 3A shows the scaling of the code both when running with adaptive sampling on and without. The main reason why a linear scaling is not observed is that the total number of frequencies being calculated is, in general, not divisible with the number of threads. Thus, at the end of a sampling, some of the threads will remain idle while the last few frequencies are being calculated. When running with the adaptive sampler, this is even more pronounced since it calculates significantly fewer frequencies than when running without it. Nevertheless, the effectiveness of the sampler can be seen in Fig. 3B, where the computational time is shown.

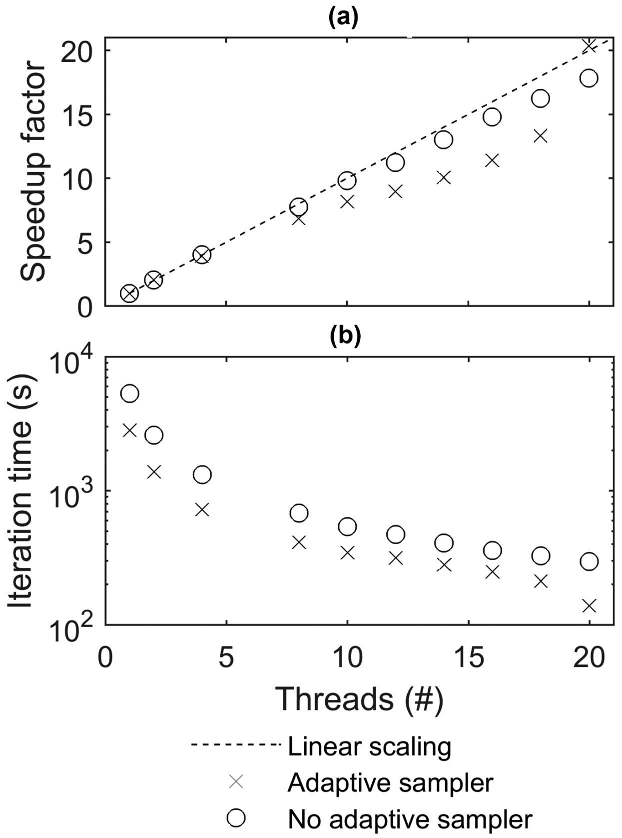

Figure 3 Parallel scaling of the thin-sheet code. Tests were performed on a modern NUMA system with two Intel Xeon E5-2650 v3 CPUs, each with 10 cores. (A) Scaling with the number of threads. (B) Iteration time. The tests shown here were performed on a system with 90 soundings and 500 sheet cells.

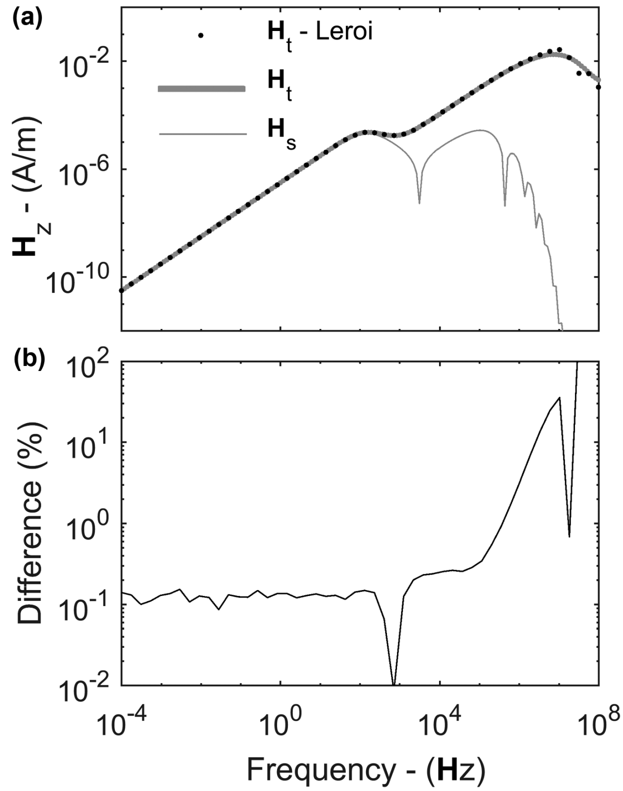

Figure 4 A comparison between Leroi Air and the sheet code. The model consists of a 100-S square sheet, placed at a 50 m depth, with a length of 100 m, in a homogenous background of 1000 Ωm. The transmitter is a 22 m × 22 m square loop, placed on the ground, 25 m off-centre in comparison to the underlying sheet, whereas the receiver is placed in the centre of the transmitter. (A) The imaginary part of the magnetic field as a function of frequency. (B) The relative difference between the total magnetic fields: $\mathbf{H}_t$ and $\mathbf{H}_t$ - Leroi.

RESULTS AND DISCUSSION

Validation

Forward modelling of the thin-sheet code has been thoroughly validated against the thin-sheet code known as Leroi Air (Raiche 1998). Differences in forward responses of around 0.1% or less were observed for all three components of the magnetic field, when the sheet response is dominant. A comparison of the quadrature part of the responses is seen in Fig. 4. Note that the high-frequency difference between the total magnetic field, $\mathbf{H}_t$, calculated by Leroi and our code happens when the sheet response, $\mathbf{H}_s$, is significantly smaller than the total magnetic field. As such, it is not a difference in the sheet response that is observed, but rather a difference in the layered background response. Our background response, however, is validated against the well-established 1D modelling code (Auken *et al.* 2014). Thus, we believe that the high-frequency difference seen when comparing with Leroi Air likely stems from uncertainties because of single-float precision in Leroi Air.

Synthetic model

A synthetic model was created to test the sheet code with a variable overburden. A 3D overview of the modelling system can be seen in Fig. 5A; it consists of a 25 m-thick overburden, with a conductive part centrally, and a highly resistive background with an almost vertical sheet placed at a depth of 120 m. The technical specifications of the SkyTEM516 system are used (see SkyTEM webpage for specifications), including both a low and a high moment, with the earliest gate being at 0.22 ms, and the latest gate being at 11 ms. The current of the low moment is 5.3 A, whereas the high moment operates

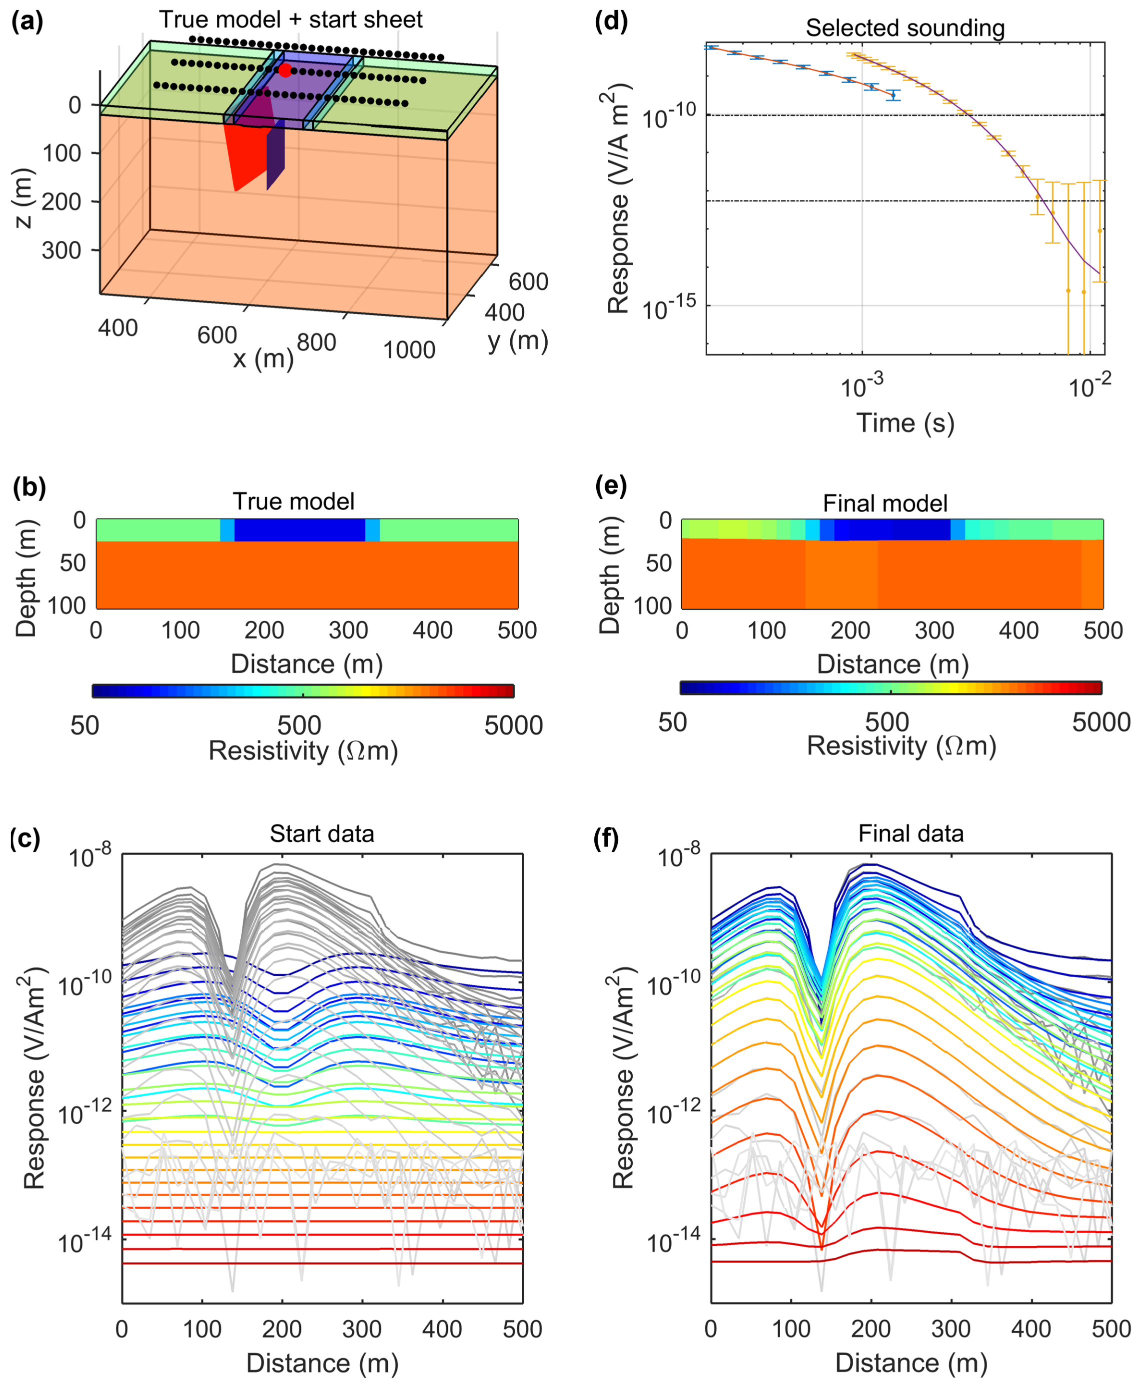

Figure 5 The synthetic model. (A) A 3D overview of the synthetic modelling system. The black dots indicate soundings, whereas the red dot is the selected sounding shown in panel (D). The red sheet is the true sheet, whereas the blue sheet shows the sheet at the beginning of the inversion (the final sheet is not shown as it coincides almost perfectly with the true sheet, see Table 2). (B, E) The true/final resistivity model along the central flight line. (C, F) Data from all gates for all soundings in the central flight line. Each line in the figure represents data from a particular gate shown across all soundings in the central flight line. The grey lines are the true data, whereas the coloured lines are the forward response from the starting/final model. (D) The true data from a single sounding station with error bars, as well as the final forward response from the inversion. The location of the selected sounding is indicated by a red dot in panel (A). Both high-moment (yellow) and low-moment (blue) data are shown. The horizontal black dashed lines indicate the noise floors for the low moment and the high moment. Note that only the z-component of the magnetic field has been used during this inversion.

Table 2 Details on the progression of the sheet parameters and the data residual during the inversion. The final sheet parameters are given for an inversion with both a variable overburden and free space. The STD factors are determined by a linearised sensitivity analysis (Auken *et al.* 2014) of the parameters at the end of the inversion. Note that the data residual is normalised with the noise

Sheet			Variable overburden		Free-space
Parameter	Starting	True	Final	STD factor	Final
τ (S)	15	30	30	1.01	24
x (m)	590	550	550	1.00	557
y (m)	510	470	469	1.00	468
z (m)	160	120	120	1.04	133
L_x (m)	100	200	200	1.01	220
L_y (m)	100	150	149	1.01	180
θ (°)	90	80	80	1.00	80
ϕ (°)	90	100	100	1.00	103
Data residual	5.56	–	0.11	–	1.6

at 113.9 A. The survey contains three lines of 30 soundings each, spaced 100 m apart, with 17 m between each sounding in a line, flown at a height of 30 m.

The uncertainties on the data contain 10% uniform noise as well as a 1-nV/m^2 noise floor, whereas the data have been perturbed by 10% of the uncertainty. The resulting noise floor for the low and high moments can be seen in Fig. 5D.

The starting model was set to 1000 Ωm for the overburden (15 m thick) and 2000 Ωm for the background. Constraints and standard deviations (STDs) are given as factors and follow Auken *et al.* (2014); this means that 1.0 is perfect resolution, whereas 1.1 corresponds to an STD of 10%.

The lateral constraint factors on the resistivity were set to 2 and 1.1 for the overburden and background, and the lateral constraint on the thickness was set to 1.1. An absolute prior of 0.6 m was put on the altitude with a lateral constraint of 1.01. The sheet was started with half the true conductivity and as a completely square and vertical sheet 40 m off from the true position in x, y, and z. Looking at Fig. 5C, we believe this is a fair starting point, since manual inspection clearly shows the approximate position of the sheet to well within these starting parameters. The sheet was restricted to 300–600 cells during the inversion, with a desired cell length of 5 m.

The inversion ran until a relative norm change of less than 0.7% was obtained, which took 28 iterations, and was completed in less than 15 minutes on a NUMA system with two Intel Xeon E5-2650 v3 CPUs, each with 10 cores. The results of the inversion can be found in Figs. 5D, 5E, and 5F and Table 2.

From Table 2, it is shown that the inversion with a variable overburden finds the true sheet. Furthermore, based on the STD factors, the inversion parameters are very well determined. As shown in Fig. 5E, the inversion finds the conductivity structure of both the overburden and background with only small variations from the true model. This is true for all three lines, which are practically identical.

For comparison, Table 2 also contains the sheet parameters from a similar inversion, but with the sheet in free space. In this case, the sheet is still reasonably well determined, but the added accuracy of the inclusion of the variable overburden is clearly demonstrated.

Field example—Valen mineral deposit

The Valen mineral deposit is located in the central region of the Musgrave province in the western part of South Australia. Geologically, the overall area consists of mafic units belonging to the Giles Complex, with variably conductive overburdens consisting mainly of sand or a thick regolith cover of paleovalley sediments (Eadie and Prikhodko 2013; Effersø and Sørensen 2013).

In 2011, Geotech flew a survey over what is now known as the Valen deposit (Eadie and Prikhodko 2013). The Valen deposit was modelled using Maxwell plate modelling, and several plate models were proposed based on various time gates (see Fig. 6). Based on these models, two drill holes were committed to the Valen deposit, where one of the drill holes encountered a massive 40 cm graphite layer at 89 m depth and a smaller graphite layer at ~100 m depth, both of which are attributed to the conductive response of the Valen deposit. Provided that these graphite layers are the only ones present and that they run parallel, they should be close enough to be accurately modelled as a single thin sheet.

Using the VTEM data, we have performed a sheet inversion of the Valen deposit. The data uncertainty was set to 5%. Upon inspection of the data, we noticed a possible primary field contamination in the data. In order to remedy this, we imposed a pragmatic low-pass filter of 12 kHz with a primary field damping of 90% and larger uncertainties for the early gates. Furthermore, larger uncertainties were added on the early gates, in order to remove any system response from the data. For the same reason, gates before 0.25 ms were discarded, and the earliest two gates used were given additional uncertainty. For the inversion, a section of flight lines, namely, 30380, 30390, and 30400, was processed and used. The flight lines have a spacing of 200 m between them. Though the sheet is essentially only visible in line 30390, the

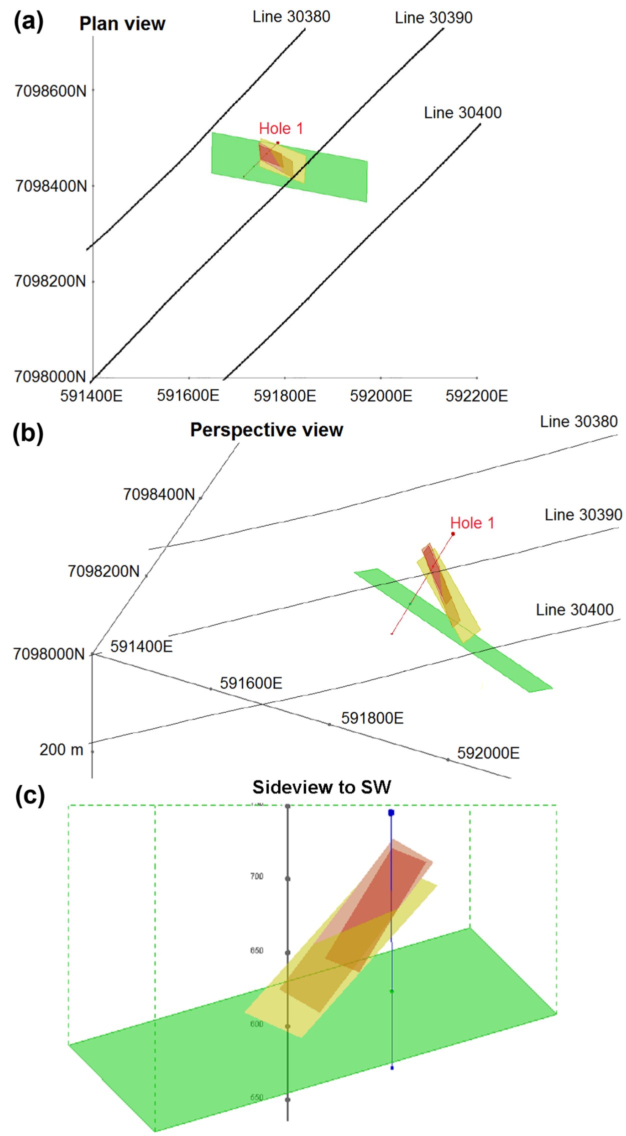

Figure 6 VTEM flight lines around the Valen mineral deposit, as well as Geotech's four proposed sheet models made in Maxwell's plate simulation program, to fit various different time-gate ranges. Based on these plate models, a drill hole dubbed Hole 1 was proposed. (A) Plan view. (B) Perspective view. The figures are modified from Blundell (2012).

other lines were kept to help constrain the sheet. The VTEM data were originally sampled at 10 Hz but were stacked to create 2-Hz soundings, for a total of 96 soundings, with an average inline sounding distance of ~17 m and with the transmitter positioned in an altitude of ~40 m.

Based on 1D modelling of the data, it was clear that the regional geology changes from moderately resistive in the northeast to highly resistive hardrock in the southwest.

Table 3 The progression of the sheet parameters and the data residual during the inversion of the data from the Valen deposit. The STD factors are determined by a sensitivity analysis of the parameters at the end of the inversion in order to calculate how well determined the parameters are. *For completeness sake, it should be noted that the STD factors of UTMX and UTMY are calculated in a local coordinate system, where the final sheet position is given as $(x, y) = (458, 545)$

	Starting	Final	STD factor
Conductance (S)	300	205	1.02
UTMX (km)	591.844	591.776	1.00*
UTMY (km)	7098.463	7098.508	1.00*
Depth (m)	150	167	1.03
Width (m)	150	172	1.02
Length (m)	150	178	1.06
Strike angle (°)	231	232	1.00
Dip angle (°)	40	68	1.00
Data residual	7.5	1.9	–

Consequently, the starting model was chosen to reflect this. The starting resistivities were chosen to be 275 Ωm for the soundings in the northeast and 1500 Ωm overburden with 2000 Ωm background for the soundings in the southwest. Between these two areas, a group of five soundings is set with an intermediate value of 800 Ωm. The thickness of the overburden was uniformly set to 50 m. *A priori* uncertainties of 2 and 1.1 were placed on the resistivities of the overburden and background, respectively. Lateral constraints were placed on the resistivities and thickness in order to enforce approximate homogeneity for the background model, with values of 2, 1.1, and 1.1 being chosen for resistivities of the overburden, background, and thickness, respectively. The altitude was given an absolute prior constraint of 2 m.

The starting values for the sheet can be seen in Table 3 and were chosen based on the suggested Maxwell models for the Valen deposit, but were otherwise given complete freedom to change. The sheet was restricted to 300–600 cells during the inversion, with a desired cell length of 5 m.

The inversion took six iterations and was done in ~13 minutes on a NUMA system with two Intel Xeon E5-2650 v3 CPUs, each with 10 cores. The results from the inversion can be seen in Fig. 7 and Table 3.

Figure 7 only shows details from the central line, since this is the only line where the sheet response is detectable. The two surrounding lines merely provide constraints for the sheet but do not notably help pinpoint the exact location of the sheet.

The high conductance closest to the sheet location in the resistive background in Fig. 7E suggests that a single sheet

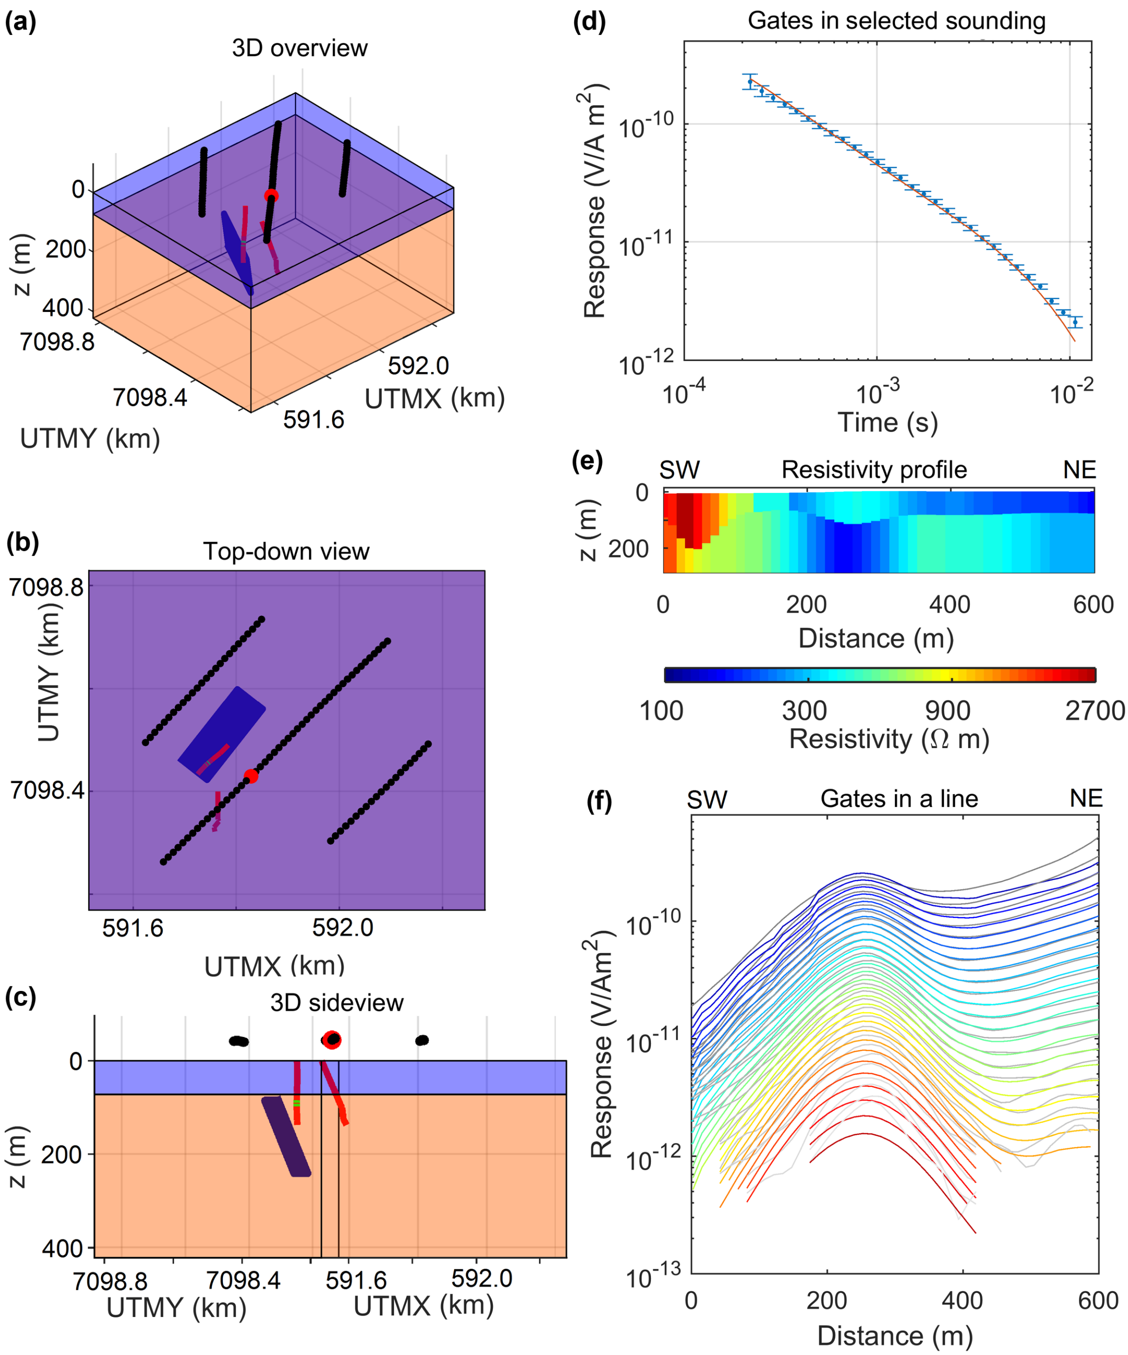

Figure 7 Inversion results from the Valen deposit. (A–C) A 3D overview of the flight lines, the final sheet, as well as the two drill holes conducted in the area. The black dots indicate soundings, whereas the red dot is the selected sounding shown in panel (D), the two red lines indicate the two drill holes, with the green stripes being the location of the graphite mineralisation encountered in the drill hole. Note that only one of the two drill holes encountered the graphite mineralisation. (D) A comparison of forward and observed data, with error bars, for all gates in the selected sounding. (E) The final 1D resistivity models below the central line. (F) A comparison of forward and observed data, for all gates in the central flight line. Each line represents a specific time gate. All coloured lines are forward model data, whereas the grey lines are observed data. Note that only the z-component of the magnetic field has been used during this inversion.

is not sufficient to account for the entire signal found. This is further seen in Fig. 7D where a deviation is observed at late times, suggesting that part of the sheet has a higher conductance.

Since the drill holes do not intersect our predicted sheet model, a comparison of the ground truth and our predicted sheet model remains inconclusive. It is possible that there is an even larger dominating mineralisation at the location predicted by our model. However, it is also possible and perhaps more likely that the graphite found during the drill hole is the only vein and that our model is off by $\sim$50 m. If that is the case, then several factors could be contributing to this. First, it could be that the mineralisation is simply not shaped in a way that can be well modelled by a thin sheet. Second, the sheet is only measured on one flight line, and since we only had access to the z-component from that line, it is reasonable to expect a relatively high uncertainty on the position of the predicted sheet.

Comparing our model with the Maxwell models, the models are, overall, in agreement about the position, shape, and angle of the sheet, though the conductance of our sheet is significantly lower than theirs, which is likely because Maxwell uses a free-space background, and hence, their sheet encompasses the background signal as well. While our model corresponds well with the Maxwell model, we believe that our method is superior due to several reasons. First, we are able to present a single sheet, which reasonably models the Valen mineralisation for all the used gates, and since our results are derived through inversions, they are, in that sense, more objective. Furthermore, our approach provides a data misfit, which we believe is a valuable objective parameter when determining the trustworthiness of a model.

CONCLUSION

We have developed an advanced thin-sheet code capable of performing a full non-linear inversion of airborne electromagnetic data. The thin sheet is calculated in a two-layered background, and a variable overburden is furthermore included. The background and overburden are based on 1D models, which are mapped from 1D local models to a regional average sheet background as a weighted mean. Significant time and resources have been spent in parallelising and optimising the performance of the sheet code, which includes dynamic cell discretisation and an OpenMP task parallelised adaptive frequency sampler. This means that the code is capable of doing inversion of field data in a matter of hours on a desktop computer. The effectiveness of the code was shown both on a synthetic example emulating a variable overburden, where the code showed significant improvement over free-space sheet modelling, as well as the Valen mineral deposit, which is located in an area with changing background geology.

ACKNOWLEDGEMENTS

This study was supported by Innovation fund Denmark, through the AirTech4Water project. Furthermore, the authors would like to thank colleague Kristoffer Andersen for the many helpful discussions and insights about the underlying theory, as well as Tim Munday and CSIRO for the field data.

REFERENCES

Aghasi A., Kilmer M. and Miller E.L. 2011. Parametric level set methods for inverse problems. *SIAM Journal on Imaging Sciences* **4**, 618–650.

Annan A.P. 1974. *The equivalent source method for electromagnetic scattering analysis and its geophysical application*. PhD thesis, Memorial University of Newfoundland, Canada.

Auken E. and Christiansen A.V. 2004. Layered and laterally constrained 2D inversion of resistivity data. *Geophysics* **69**, 752–761.

Auken E., Christiansen A.V., Kirkegaard C., Fiandaca G., Schamper C., Behroozmand A.A. *et al.* 2014. An overview of a highly versatile forward and stable inverse algorithm for airborne, ground-based and borehole electromagnetic and electric data. *Exploration Geophysics* **46**, 223–235.

Avdeev D.B., Kuvshinov A.V. and Pankratov O.V. 1998. An imaging of buried anomalies, using multi-sheet inversion. *Earth, Planets and Space* **50**, 417–422.

Blundell K. 2012. Deering Hills Project. November 2011 TEM Survey. Valen Model Summary. pp. 1. Musgrave Minerals.

Constable S.C. and Weiss C.J. 2006. Mapping thin resistors and hydrocarbons with marine EM methods: insights from 1D modeling. *Geophysics* **71**, G43–G51.

Dong X., Cooperman G. and Apostolakis J. 2010. Multithreaded Geant4: semi-automatic transformation into scalable thread-parallel software. *European Conference on Parallel Processing*, Berlin, Heidelberg, pp. 287–303.

Dorn O. and Lesselier D. 2006. Level set methods for inverse scattering. *Inverse Problems* **22**, R67–R131.

Dorn O., Miller E.L. and Rappaport C.M. 2000. A shape reconstruction method for electromagnetic tomography using adjoint fields and level sets. *Inverse Problems* **16**, 1119–1156.

Duncan A. 1987. Interpretation of down-hole transient EM data using current filaments. *Exploration Geophysics* **18**, 36–39.

Eadie T. and Prikhodko A. 2013. VTEM–SkyTEM Survey Comparison over Valen Cu-Ni Deposit. Geotech Ltd.

Effersø F., Auken E. and Sørensen K.I. 1999. Inversion of band-limited TEM responses. *Geophysical Prospecting* **47**, 551–564.

Effersø F. and Sørensen K.I. 2013. SkyTEM Case Study: The Valen Survey Comparison with VTEM. SkyTEM Surveys.

Fainberg E., Pankratov O. and Singer B.S. 1993. Thin sheet modelling of subsurface and deep inhomogeneities. *Geophysical Journal International* **113**, 144–154.

Fitterman D.V. and Stewart M.T. 1986. Transient electromagnetic sounding for groundwater. *Geophysics* **51**, 995–1005.

Johansen H.K. and Sørensen K.I. 1979. Fast Hankel transforms. *Geophysical Prospecting* **27**, 876–901.

Lajoie J.J. and West G. 1976. The electromagnetic response of a conductive inhomogeneity in a layered earth. *Geophysics* **41**, 1133–1156.

Macnae J., King A., Stolz N., Osmakoff A. and Blaha A. 1998. Fast AEM data processing and inversion. *Exploration Geophysics* **29**, 163–169.

McMillan M.S., Schwarzbach C., Haber E. and Oldenburg D.W. 2015. 3D parametric hybrid inversion of time-domain airborne electromagnetic data. *Geophysics* **80**, K25–K36.

Price A.T. 1948. The induction of electric currents in non/uniform thin sheets and shells. *The Quarterly Journal of Mechanics and Applied Mathematics* **2**(3), 283–310.

Raiche A. 1998. Modelling the time-domain response of AEM systems. *Exploration Geophysics* **29**, 103–106.

Reid J., Fitzpatrick A. and Godber K. 2010. An overview of the SkyTEM airborne EM system with Australian examples. *Preview* **2010**, 26–37.

Schmucker U. 1971. Interpretation of induction anomalies above nonuniform surface layers. *Geophysics* **36**, 156–165.

Song Y., Kim H.J. and Lee K.H. 2002. An integral equation representation of wide-band electromagnetic scattering by thin sheets. *Geophysics* **67**, 746–754.

Swidinsky A. and Edwards R.N. 2010. The transient electromagnetic response of a resistive sheet: an extension to three dimensions. *Geophysical Journal International* **182**, 663–674.

Vasseur G. and Weidelt P. 1977. Bimodal electromagnetic induction in non-uniform thin sheets with an application to the northern Pyrenean induction anomaly. *Geophysical Journal International* **51**, 669–690.

Walker P.W. and West G.F. 1991. A robust integral equation solution for electromagnetic scattering by a thin plate in a conductive media. *Geophysics* **56**, 1140–1152.

Weidelt P. 1981. Report on dipole induction by a thin plate in a conductive half-space with an overburden: Federal Institute for Earth Science and Raw Materials. Hannover, Germany.

Wolfgram P., Hyde M. and Thomson S. 1998. How to find localised conductors in GEOTEM® data. *Exploration Geophysics* **29**, 665–670.

Zhou Q. 1989. *Audio-frequency electromagnetic tomography for reservoir evaluation*. PhD thesis, University of California, USA.