Large-scale parallel 3D inversion of frequency and time domain AEM data

Martin Čuma[1,2], Leif Cox[2], and Michael S. Zhdanov[1,2], [1]University of Utah, [2]TechnoImaging*

Summary

3D inversion of airborne electromagnetic data is a challenging task due to the large amounts of data collected over relatively large areas. In this paper we detail an inversion algorithm based on a moving sensitivity domain approach using the integral equation method coupled with a multistep regularized conjugate gradient inversion. To tackle the computational demands, along with the reduction of the problem due to the moving sensitivity domain approach, we also parallelize the problem over the data using Message Passing Interface (MPI) and OpenMP. The workflow of the interpretation includes 1D inversion to obtain a background structure that serves as an input to the 3D inversion. The background is either a half-space, unique under each data point, in the case of frequency domain, or layered background in the case of time domain inversion. We demonstrate the effectiveness of the developed method and computer software by 3D inversion examples of frequency and time domain airborne EM surveys.

Introduction

The mineral and groundwater exploration depends on large regional surveys which can detect small-scale ore bodies or resources such as perched water tables. Airborne electromagnetic (AEM) surveying is one of the few methods which can economically cover large areas with the resolution required for such exploration. Historically, simple techniques, like conductivity depth transforms (Macnae et al., 1998) and 1D inversions (layered earth inversions (e.g., Viezzoli et al., 2009), were used for interpretation of the airborne data. Advances are still being made with respect to these 1D methods to make them very large scale and fast with parallelization (e.g. Kirkegaard and Auken, 2014). More advanced transforms have also been developed to extend the approximate inversion methods to 2D (e.g., Guillemoteau, 2012). An excellent comparison of these methods with each other and with 3D inversion is given in Ley-Cooper et. al. (2014).

The difficulties in performing full 3D inversion for AEM surveys stems from the necessity to solve as many large linear systems of equations as there are transmitter positions in the survey. However, it is widely known that AEM data are only sensitive to a limited sensitivity domain (footprint) (e.g., Liu and Becker, 1990; Beamish, 2003; Reid et al., 2006). An AEM system's sensitivity domain is defined as the lateral extent of the sensitivity for the AEM system, and is typically in the order of hundreds of meters to a kilometer. For a single transmitter-receiver pair, there is no need to calculate the responses or sensitivities beyond the AEM's sensitivity domain. The sensitivity matrix for the entire 3D model can then be constructed as the superposition over the entire inverse model of the Fréchet derivatives from all transmitter-receiver pairs for corresponding sensitivity subdomains. This combined sensitivity matrix can be stored as a sparse matrix with memory and computational requirements reduced by several orders of magnitude. The number of nonzero elements in each row of the sensitivity matrix is just the number of elements within each footprint (in an order of hundreds to thousands) rather than the total number of elements in the model (hundreds of thousands to millions).

The concept of a moving sensitivity domain was introduced in Cox and Zhdanov (2007), Cox et al. (2010, 2012), and Zhdanov and Cox (2013). This concept made possible a 3D inversion of frequency-domain (FD) AEM survey data that did not rely on any approximations in the modeling or inversion kernels.

In this paper we implement and evaluate parallel integral equation-based 3D inversion of frequency and time domain (TD) AEM data.

Modeling and Inversion

Time-domain AEM modeling can be accomplished either by direct time-domain solutions or by Fourier transformation of frequency-domain solutions. The latter offers three distinct advantages. First, the effects of frequency-dependent conductivity, such as induced polarization, can be modeled. Second, artificial dispersion effects that arise in direct time-domain solutions are avoided. Third, the matrix equations for multiple right-hand side source terms can be rapidly solved with iterative solutions. Our approach therefore calculates the forward modeling response in the frequency domain, and in the case of time domain data, this response is then transformed to the time domain. In the forward modeling, we use the integral equation (IE) method where the background field is a field generated by the given sources in the model with a background distribution of conductivity σ_b, and the anomalous field is produced by the anomalous conductivity distribution $\Delta\sigma(r)$, $r \in D_a \subset R^3$ (Hursán and Zhdanov, 2002; Zhdanov, 2009).

Inversion is the process where we seek to recover the 3D conductivity distribution from the AEM data. However, AEM surveys are finite in their spatial and frequency content, and are contaminated with noise. This means that AEM inversion is ill posed; i.e., solutions are non-unique and unstable. Regularization must be introduced so as to obtain a unique and stable solution, by minimization of the Tikhonov parametric functional, $P^\alpha(\sigma)$:

$$P^\alpha(\sigma) = \varphi(\sigma)^2 + \alpha\left\|\sigma - \sigma_{apr}\right\|^2 \rightarrow min \qquad (1)$$

where σ is the N_m length vector of conductivities, d is the N_d length vector of observed data, σ_{apr} is the N_m length vector of the a priori conductivities, and ‖…‖ denotes the respective Euclidean norm. The first term of equation (1) describes the misfit functional (or residual errors) between the predicted d_{pred} and observed d_{obs} AEM data.

$$\varphi(\sigma) = \left\| w_d\left(d_{pred} - d_{obs}\right)\right\|^2 \qquad (2)$$

where

$$w_d = \frac{1}{\varepsilon},\ \varepsilon = \left|d_{obs}\frac{\varepsilon_{per}}{100}\right| + \varepsilon_{abs} \qquad (3)$$

where ε is a vector of the estimated errors in each data point, ε_{per} is the estimated error in each data point in percent and ε_{abs} is the estimated absolute error in data units. This method normalizes the importance of all data channels with respect to their uncertainty.

The second term of equation (1) describes the stabilizing functional, which in this case is written as a minimum norm stabilizer. The choice of a stabilizer determines the class of the solutions from which a model is sought (Zhdanov, 2002, 2015). The regularization parameter, α, provides a balance (or bias) between the misfit and stabilizing functionals.

The parametric functional is minimized iteratively, with either the steepest descent or conjugate gradient method, using a two-level minimization approach. After each forward modeling update (higher level iteration), we perform a number of conductivity model updates using the same Fréchet derivative, until a threshold of difference between the current conductivity and the conductivity used in the previous modeling step is reached. This triggers another forward modeling update. If this threshold is reached only over a subset of inversion domain cells, forward modeling is performed only for the data points which include these cells. We call this approach adaptive forward modeling.

Each data point is sensitive only to a limited number of cells in the 3D model. In Figure 1 we show percent of total response (as calculated from integrated sensitivities) as a function of distance and half-space resistivity. The frequency domain RESOLVE system resolution is limited to a few hundred meters while the response of the time domain TEMPEST system is considerably larger.

With a moving sensitivity domain, the Fréchet matrix can be constructed as a sparse matrix with memory and computational requirements reduced by several orders of magnitude. The number of nonzero elements in each row of the sensitivity matrix is just the number of elements within each sensitivity domain (in an order of hundreds or thousands) rather than the total number of elements in the domain (hundreds of thousands to millions).

Due to potentially large variations in the conductivities over the AEM survey areas, it is advantageous to allow data points and their MSDs to have different background conductivity structure. We call this variable background (VB). In the AEM modeling and inversion setting, we use two kinds of variable backgrounds. One VB is related to each data point itself and it is unique for each data point. It is used throughout the inversion and for calculation of the receiver background fields and domain to receiver Greens' tensors. This data point VB can be either half-space, or layered.

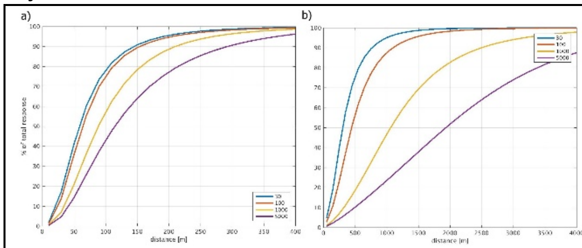

Figure 1: Percent of total response (as calculated from integrated sensitivities) as a function of footprint size for the different half-space resitivities for the (a) RESOLVE and (b) TEMPEST systems. Note the 10x larger distance scale in the TEMPEST plot

In the forward modeling, a half-space background is used, obtained as about 5-10 logarithmically spaced conductivity values from a range of the data point backgrounds. Each data point is assigned one of these backgrounds, the closest to its background value (or average of the background in the case of a layered background). Having a limited number of forward modeling backgrounds allows us to limit the amount of pre-calculation and storage of the background domain fields and Green's tensors in the MSD, but still keeps the anomalous conductivities in an appropriate range.

Parallelization

Our AEM modeling and inversion program is parallelized using Message Passing Interface (MPI) and OpenMP. In the AEM, moving sensitivity domain size is limited and as such modeling computation and storage requirements are relatively small for each data point. This lends to parallel distribution over the data (soundings), while keeping the problem scalable. The advantage of such parallelization is limited inter-process communication, as modeling of each data point is independent. Loops in each MPI task are shared-memory parallelized using OpenMP.

We distribute the data points evenly across the MPI tasks, but, since we use the iterative solver, the number of iterations to solution can vary, which can lead to load imbalance. The load imbalance can be made worse by the adaptive forward modeling, recalculating the response only if the conductivity model under each data point changes more than certain threshold. We alleviate this imbalance by round-robin distribution of the data points, but, in the future, we will consider exploring adaptive load balancing by on demand migration of data points between the MPI tasks.

In Figure 2 we show parallel scaling on a subset of the frequency domain data discussed below. We use one to eight 24 core nodes with two Intel Xeon CPUs. On a single node (4 MPI tasks) we vary number of OpenMP threads from 1 to 6, on multiple nodes we use 4 MPI tasks per node and 6 OpenMP threads per MPI task. The OpenMP scaling in the figure is relative to one thread performance, while the MPI scaling is relative to one node performance.

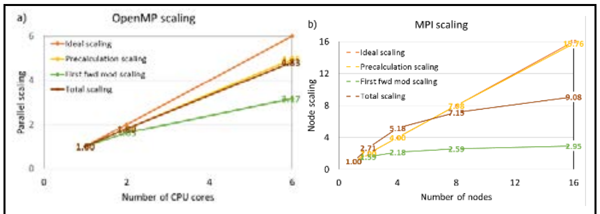

Figure 2: OpenMP (a) and MPI (b) parallel scaling.

We look at three different scaling characteristics. The pre-calculation includes once-per-inversion calculation of background fields and domain to receiver Green's tensors. Since this calculation is independent for each data point and frequency, it exhibits linear or nearly linear scaling both for OpenMP and MPI. Then we look at the first forward modeling calculation, where the scaling is less than linear. In the case of OpenMP, there are several factors. One is limited memory bandwidth with increased thread count. Another is the size of the footprint, which limits the amount of calculation available for OpenMP parallelization. In the frequency domain the footprint is about an order of magnitude smaller than in the time domain, where we observe improved OpenMP scaling by factor of 30-50 % as compared to the frequency domain. In the case of MPI, the poorer scaling is mostly due to the load discussed above. Finally, the scaling of the whole inversion is improved, as the subsequent modeling iterations converge with less variability. In the case of MPI, we even notice super-linear scaling for 2 and 4 nodes, which we attribute to the reduction in memory bandwidth contention as the problem is distributed on more than one node.

Case study – inversion of the FD AEM data

We inverted a Fugro RESOLVE data set obtained for the USGS in the Yukon Flats area near Fort Yukon, Alaska (Minsley et. al., 2012). The goal of this survey was permafrost mapping. According to Minsley et. al., 2012, the uppermost unfrozen Eolian silt and sands have an expected resistivity of 100-200 Ohm-m. At larger depth there are frozen fluvial gravels with resistivity greater than 1000 Ohm-m, below which are lacustrine silts and clays with resistivity near or below 100 Ohm-m. Within the survey there are numerous water bodies, including the Yukon River and Twelvemile Lake. Water resistivity in Twelvemile Lake was measured at a range between 2.5 and 18 Ohm-m.

The FD AEM data consisted of 1200 line km with six frequencies between 0.4 and 129 kHz, covering about a 300 km^2 area and was inverted on a 10×25 m horizontal grid and 24 vertical cells ranging from 1 m at the surface to 15 m at depth, to a total of 155 m depth, with nearly 30 million cells. We used a data point every 10 meters, which ended up being 81,185 receivers with 6 frequency readings each. The MSD was set to a 400 m diameter.

We used the following workflow for the FD inversion:

- Perform half-space inversion to obtain the best-fit half-space conductivity under each data point.
- Extrapolate and smooth this model over the 3D model cells; this constitutes the half-space background for each data point and the optional initial model.
- Find maximum and minimum 3D model conductivity and create a logarithmically spaced set of conductivities, 4 per decade, bound by this minimum and maximum; this constitutes the half-space background model for forward modeling.
- Run the 3D inversion.

The 3D inversion was run on 120 nodes with two six-core Intel Xeon X5660 2.8 GHz CPUs and took 3.5 hours to reach convergence at RMS 1.67.

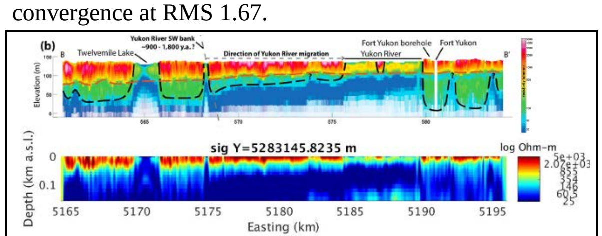

Figure 3: Vertical cross sections of Ft. Yukon survey along profile B-B' obtained by Minsley et al (2012) (upper panel) and by this study (lower panel).

In Figure 3 we compare vertical cross section of our resistivity model with that of Minsley et. al., 2012 obtained with 1D inversion. The models are very similar, which one would expect for the layered permafrost formations. At the surface we notice unfrozen area with resistivity 100 Ohm-m and less, which follows the Yukon river sediments. More conductive features include the Yukon River itself and numerous lakes, the largest of which is Twelvemile Lake in the left of the picture. The rest of the surface is highly resistive and consists of frozen silts and sands. With increasing depth, the leftmost third of the area is resistive and frozen with occasional conductors caused by unfrozen areas under water bodies. About the center third is less resistive suggesting partially frozen sediments, followed by a conductive unfrozen area under the Yukon River.

Case study – inversion of the TD AEM data

The TEMPEST time-domain survey was conducted by MMG for mineral exploration. The fixed-wing AEM system recorded 13 channels of in-line and vertical B data from 6.5 μs to 6.5 ms. The survey was inverted on a horizontal grid of 20×20 m and 24 vertical cells ranging from 5 m at the surface to 70 m at depth to total depth of 730 m. This equals about 6.7 million cells. 12 time channels of 15,788 measurement positions were used with spacing about 40 m resulting in 189,456 data points.

In time domain inversion, we used the following workflow:

- Perform 1D half space inversion to obtain best fit half space conductivity under each data point.
- Extrapolate and smooth this 1D model; this constitutes the half space background under each data point and the initial model for 1D layered inversion.
- Perform 1D layered inversion to obtain layered 1D model.
- Extrapolate and smooth the 1D layered model over 3D model cells, this is an (optional) initial model for 3D inversion and layered background conductivity under each data point.
- Find maximum and minimum 3D model conductivity and create logarithmically spaced set of conductivities, 4 per decade, bound by this minimum and maximum; this constitutes the half space background model for forward modeling.
- Run the 3D inversion.

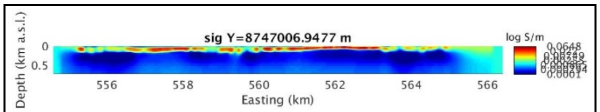

Figure 4: Vertical slice through 3D inversion result. The conductive overburden is well imaged with the time domain data.

As compared to the frequency domain, the time domain adds the extra step of 1D layered inversion, which is implemented as an option in our parallel inversion program.

The inversion was run on 43 nodes with 24 CPU cores each (Xeon E5-2680 v3) and took 18.5 hours to achieve RMS misfit of 3.6. The increased computer resources in the time domain inversion, as compared to frequency domain, are needed for larger MSD (1200×800 m diameter) and to compute a larger number of frequencies (32 in this case, in the range of 0.1 Hz to 100 kHz).

The targets for the inversion were mineralized black shale units. These are conductive and up to 100 m thick, which makes a great airborne target. Also in the area is a conductive overburden of variable thickness and uneconomic near-surface conductive lineaments. This can be easily confused with the mineralized shale if accurate interpretation is not done. The plunge, dip, and general geometry of the black shale was also of interest to the client. Figure 4 shows the conductive overburden of variable thickness in the area, which are not of economic interest but show in the data as conductors. The conductive features of economic interest are shown clearly in Figure 5, which is a horizontal slice at 325 m below the surface. The black shale units are clearly imaged.

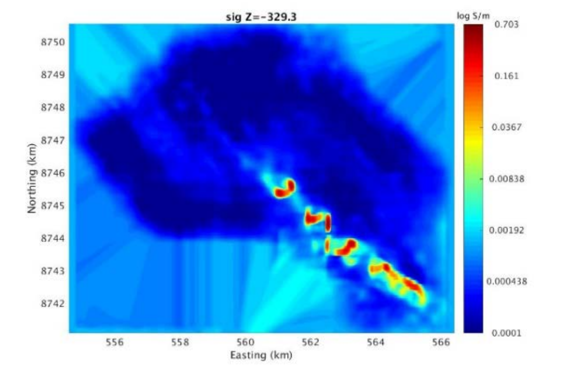

Figure 5: Horizontal slice through 3D inversion result at 329 meters below the surface. The economic mineralization is being shown here by association with the black shales

Conclusions

In this paper we have introduced a method and optimized workflows for the large-scale inversion of frequency and time domain AEM data. Utilization of a moving sensitivity domain along with multilevel parallelization allows us to invert large AEM surveys for a finely discretized model. The convergence of the inversion is improved with variable background under each data point, which is obtained from a smoothed 1D inversion result.

Our implementation is based on the 3D integral equation method for computing data and sensitivities, as well as the re-weighted regularized conjugate gradient method for minimizing the parametric functional, and has been generalized in such a way that it can be applied to any AEM system.

The 3D model for the frequency domain case study corresponds well to the published 1D inversion model, which was expected for the case of the horizontally-layered geoelectrical formations representing permafrost.. The model produced for the time domain case study provides a good example of true 3D inversion capabilities, clearly imaging the conductors of interest at a depth of several hundreds of meters.

Acknowledgments

The authors are grateful to USGS and Stefaans du Plessis of MMG for providing the field data and permission to publish. The authors acknowledge TechnoImaging and the University of Utah's Consortium for Electromagnetic Modeling and Inversion (CEMI) for support of this research and permission to publish the results. We also acknowledge the University of Utah's Center for High Performance Computing for support and computer time allocation.

EDITED REFERENCES

Note: This reference list is a copyedited version of the reference list submitted by the author. Reference lists for the 2016 SEG Technical Program Expanded Abstracts have been copyedited so that references provided with the online metadata for each paper will achieve a high degree of linking to cited sources that appear on the Web.

REFERENCES

Auken, E., A. V. Chistiansen, B. H. Jacobsen, N. Foged, and K. L. Sorensen, 2005, Piece-wise 1D laterally constrained inversion of resistivity data: Geophysical Prospecting, **53**, 497–506. http://dx.doi.org/10.1111/j.1365-2478.2005.00486.x.

Beamish, D., 2003, Airborne EM footprints: Geophysical Prospecting, **51**, 49–60. http://dx.doi.org/10.1046/j.1365-2478.2003.00353.x.

Cox, L. H., and M. S. Zhdanov, 2007, Large-scale 3D inversion of HEM data using a moving footprint: 7th Annual International Meeting, SEG, Expanded Abstracts, 467–470, http://dx.doi.org/10.1190/1.2792464.

Cox, L. H., G. A. Wilson, and M. S. Zhdanov, 2010, 3D inversion of airborne electromagnetic data using a moving footprint: Exploration Geophysics, **41**, 250–259, http://dx.doi.org/10.1071/EG10003.

Cox, L. H., G. A. Wilson, and M. S. Zhdanov, 2012, 3D inversion of airborne electromagnetic data: Geophysics, **77**, no. 4, WB59–WB69, http://dx.doi.org/10.1190/geo2011-0370.1.

Cox, L. H., M. Endo, B. Siemon, and M. S. Zhdanov, 2015, Large-scale 3D inversion of airborne modeling based on the hybrid IE-FE method and the moving sensitivity domain approach: 14th Biennial South African Geophysical Association Conference Proceedings.

Hursán, G., and M. S. Zhdanov, 2002, Contraction integral equation method in three-dimensional electromagnetic modelling: Radio Science, **37**, 1-1–1-13. http://dx.doi.org/10.1029/2001RS002513.

Liu, G., and A. Becker, 1990, Two-dimensional mapping of sea-ice keels with airborne electromagnetic: Geophysics, **55**, 239–248. http://dx.doi.org/10.1190/1.1442832.

Macnae, J., A. King, N. Stolz, A. Osmakoff, and A. Blaha, 1998, Fast AEM data processing and inversion: Exploration Geophysics, **29**, 163–169. http://dx.doi.org/10.1071/EG998163.

Reid, J. E., A. Pfaffling, and J. Vrbancich, 2006, Airborne electromagnetic footprints in 1D earths: Geophysics, **71**, no. 2, G63–G72. http://dx.doi.org/10.1190/1.2187756.

Viezzoli, A., E. Auken, and T. Munday, 2009, Spatially constrained inversion for quasi 3D modelling of airborne electromagnetic data — An application for environmental assessment in the Lower Murray Region of South Australia: Exploration Geophysics, **40**, 173–183. http://dx.doi.org/10.1071/EG08027.

Zhdanov, M. S., and L. Cox, 2013, Method of subsurface imaging using superposition of sensor sensitivities from geophysical data acquisition systems: U.S. Patent US 2013/0173163.

Zhdanov, M. S., 2002, Geophysical inverse theory and regularization problems: Elsevier, 609 p.

Zhdanov, M. S., 2009, Geophysical electromagnetic theory and methods: Elsevier, 848.

Zhdanov, M. S., 2015, Inverse theory and applications in geophysics: Elsevier, 704.