Paper 48

Paradigm Change in Interpretation of AEM Data by Using a Large-Scale Parallel 3D Inversion and Moving Sensitivity Domain Approach

Ĉuma, M. [1,2], Cox, L. [1], Zhdanov, M.S. [1,2]

1. TechnoImaging, LLC
2. Consortium for Electromagnetic Modeling and Inversion (CEMI), University of Utah

ABSTRACT

Three-dimensional 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 present a 3D inversion algorithm based on a moving sensitivity domain approach using the integral equation method coupled with a multistep regularized conjugate gradient inversion. The developed method can be used for 3D inversion of both frequency domain and time-domain electromagnetic data. The time-domain data are inverted following transformation of the frequency domain fields to the time-domain. 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, and layered background in the case of time-domain inversion. We demonstrate the effectiveness of the developed method and computer software by a frequency domain example of permafrost mapping near Ft. Yukon, Alaska, regional airborne time-domain survey in Kamiskotia, Ontario, and a time-domain mineral exploration survey.

INTRODUCTION

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, 2015). More advanced transforms have also been developed to extend the approximate inversion methods to 2D (e.g., Guillemoteau, 2012). Despite the advances, these methods will always be approximate and the development needs to be directed toward full 3D solutions. 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 metres to a kilometre. This is significantly smaller than the area of even a small AEM survey. 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), Zhdanov et al. (2016), and Zhdanov and Cox (2017). This concept made possible a 3D inversion of frequency-domain AEM that did not rely on any approximations in the modelling or inversion kernels. Since then, others have also attempted to utilize this approach to introduce full 3D inversion codes based on finite difference (Yang et al., 2014), finite element (Haber and Schwarzbach, 2014) and hybrid FE-IE (Cox et al., 2015) solutions. In this paper, we implement and evaluate parallel integral equation-based 3D inversion of frequency and time-domain data.

INVERSION METHODOLOGY

Modelling and Inversion

Time-domain AEM modelling 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.

In "Proceedings of Exploration 17: Sixth Decennial International Conference on Mineral Exploration" edited by V. Tschirhart and M.D. Thomas, 2017, p. 679–690

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 modelling 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 modelling, we use the integral equation (IE) method, with the EM field represented as a sum of the background, $\{\boldsymbol{E}^b, \boldsymbol{H}^b\}$, and anomalous, $\{\boldsymbol{E}^a, \boldsymbol{H}^a\}$ fields:

$$\boldsymbol{E} = \boldsymbol{E}^b + \boldsymbol{E}^a, \boldsymbol{H} = \boldsymbol{H}^b + \boldsymbol{H}^a, \quad (1)$$

where the background field is 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, $\Delta\sigma$.

Then, the electric and magnetic fields can be obtained by the following integral equations:

$$\begin{aligned}\boldsymbol{E}(\boldsymbol{r}') - \boldsymbol{E}^b(\boldsymbol{r}') &= \iiint_{D_a} \boldsymbol{G}_E(\boldsymbol{r}', \boldsymbol{r})\Delta\boldsymbol{\sigma}\boldsymbol{E}(\boldsymbol{r})dV \\ &= \boldsymbol{G}_E^{D_a}(\Delta\boldsymbol{\sigma}\boldsymbol{E}), \end{aligned} \quad (2)$$

$$\begin{aligned}\boldsymbol{H}(\boldsymbol{r}') - \boldsymbol{H}^b(\boldsymbol{r}') &= \iiint_{D_a} \boldsymbol{G}_H(\boldsymbol{r}', \boldsymbol{r})\Delta\boldsymbol{\sigma}\boldsymbol{E}(\boldsymbol{r})dV \\ &= \boldsymbol{G}_H^{D_a}(\Delta\boldsymbol{\sigma}\boldsymbol{E}), \end{aligned} \quad (3)$$

where $\boldsymbol{G}_E$ and $\boldsymbol{G}_H$ are the electric and magnetic Green's tensors for a layered model with conductivity, σ_b.

In equations (2) and (3), the symbols $\boldsymbol{G}_E^{D_a}$ and $\boldsymbol{G}_H^{D_a}$ denote the electric and magnetic Green's operators with a volume integration of D_a. The process of solving the forward electromagnetic problem according to equations (2) and (3) consists of two parts. First, it is necessary to find the electric fields inside domain D_a (where $\Delta\sigma \neq 0$), which requires the solution of an integral equation *(domain equation)* (2). Second, using integral *data equation*, (3), we calculate the magnetic field in the receiver's domain (Hursán and Zhdanov, 2002; Zhdanov, 2009).

Inversion is the process whereby 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 nonunique and unstable. Regularization must be introduced to obtain a unique and stable solution, by minimization of the Tikhonov parametric functional, $P^a(\boldsymbol{\sigma})$.

$$\begin{aligned}P^a(\boldsymbol{\sigma}) = \|\boldsymbol{W}_d(\boldsymbol{A}(\boldsymbol{\sigma}) - \boldsymbol{d}_{obs})\|^2 + \\ \alpha\left\|\boldsymbol{W}_{\boldsymbol{m}}(\boldsymbol{\sigma} - \boldsymbol{\sigma}_{apr})\right\|^2 \rightarrow \min, \end{aligned} \quad (4)$$

where $\boldsymbol{A}$ is the nonlinear forward modelling operator, $\boldsymbol{\sigma}$ is the vector of conductivities, $\boldsymbol{d}_{\boldsymbol{obs}}$ is the vector of observed data, $\boldsymbol{\sigma}_{\boldsymbol{apr}}$ is the vector of the a priori conductivities, and $\|...\|$ denotes the respective Euclidean norm. The data and model weights can be introduced to equation (4) through data and model weighting matrices, $\boldsymbol{W}_{\boldsymbol{d}}$ and $\boldsymbol{W}_{\boldsymbol{m}}$, respectively. The first term of equation (4) describes the misfit functional between the predicted and observed AEM data. The data weights are calculated as follows:

$$\boldsymbol{W}_d = \frac{1}{\epsilon}, \quad (5)$$

where $\boldsymbol{\epsilon}$ is a vector of the estimated errors in each data point, as follows

$$\epsilon = \left|\boldsymbol{d}_{\boldsymbol{obs}}\frac{\epsilon_{per}}{100}\right| + \epsilon_{abs}, \quad (6)$$

where 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. Hence, data points of large magnitude are not made overly important, yet small magnitude noisy data are not fitted to tight tolerances.

The second term of equation (4) 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, and need not be restricted to the so-called "smooth" stabilizers only. In other words, one can use focusing stabilizers as well (Zhdanov, 2002, 2015). The regularization parameter, α, provides a balance 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 modelling 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 modelling step is reached. This triggers another forward modelling update.

Furthermore, if this threshold is reached only over a subset of inversion domain cells, new forward modelling is performed only for the data points which include these cells. We refer to this approach as *adaptive forward modelling.*

Each data point is sensitive to a very limited number of cells in the 3D model only. In Figure 1 we show percent of total response as a function of distance and half-space resistivity. The frequency domain RESOLVE system resolution is limited to a few hundred metres even in a resistive background. The size of the sensitivity domain of the time-domain TEMPEST system is considerable larger.

These sensitivity plots can be useful for estimating the maximum line spacing; i.e., for 3D inversion, there should be overlap in the sensitivities from different lines. These sensitivity sizes also demonstrate how small the sensitivity domain of a single sounding is.

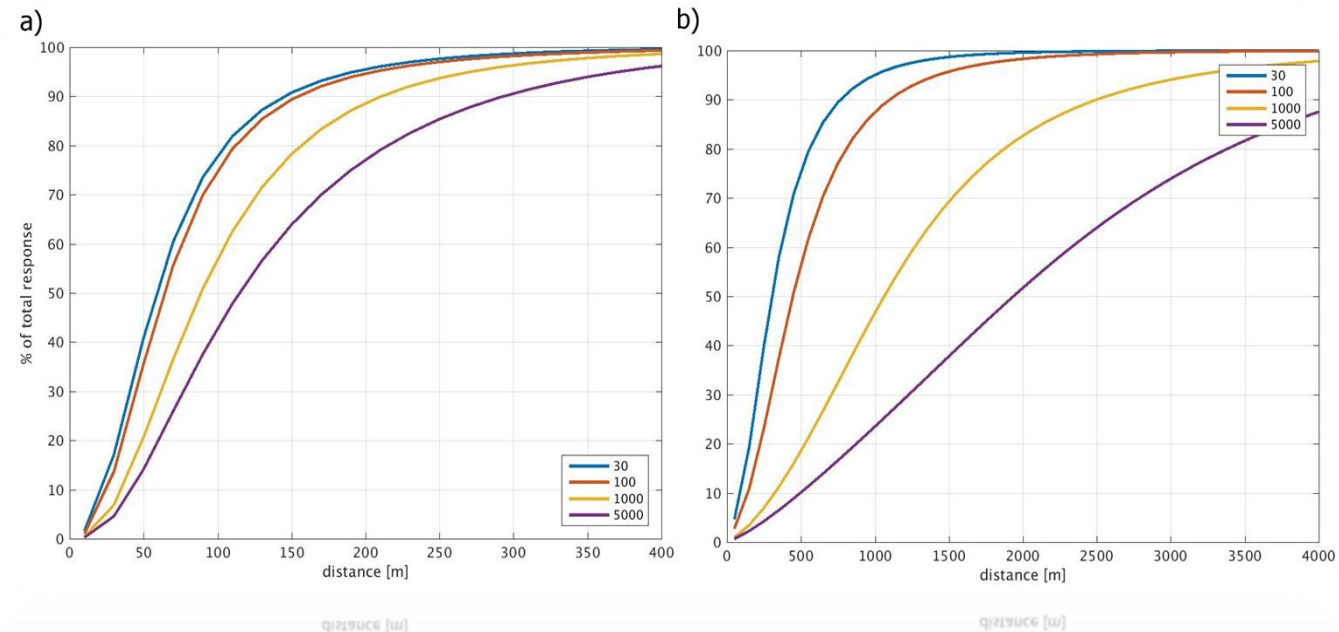

Figure 1: Percent of total response (as calculated from integrated sensitivities) as a function of the size of the sensitivity domain for the different half-space resistivities for (a) RESOLVE and (b) TEMPEST systems. Note the 10x larger distance scale in the TEMPEST plot.

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 modelling and inversion setting, we use two different 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 Green's tensors. This data point VB can be either half-space, or layered (the latter being especially useful in time-domain inversion).

In the forward modelling, 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 modelling backgrounds allows us to limit the amount of precalculation and storage of the background domain fields and Green's tensors in the MSD, but still keeps the anomalous conductivities within a reasonable range.

Parallelization

Our AEM modelling and inversion software is parallelized using Message Passing Interface (MPI) and OpenMP. Due to the relatively high frequencies used in AEM, the size of the moving sensitivity domain is limited and as such, modelling computation and storage requirements are relatively small for each data point. Furthermore, these requirements remain constant with increasing inversion domain size and number of observation data. This allows us to primarily distribute the problem to MPI tasks over the data, while keeping the problem scalable. The advantage of observed data parallelization is limited to inter-process communication, as modelling of each data point is independent. Tasks that run over the modelling or inversion domain and are the same for the data points, e.g., calculation of the domain-to-domain Green's tensors for each frequency, are also task parallelized over the domain. Loops in each MPI task are shared-memory parallelized using OpenMP. This allows us to run on commodity computer clusters using hundreds of nodes using the mixed MPI/OpenMP parallelization, generally mapping two or four MPI processes per node. The MPI process vs. OpenMP thread count depends on the type of calculation performed. For example, in 1D inversion, the modelling is implemented as one cell MSD, which limits the loops' trip counts in modelling and time-domain transform; therefore, using fewer OpenMP threads is more efficient. In the 3D modelling, the MSD contains thousands of cells, which lends to efficient OpenMP parallelization over more threads.

Nodes	1	1	1	1	2	4	8	16
Procs	4	4	4	4	4	4	4	4
Threads/proc	1	2	6	6	6	6	6	6
Precalculation time (s)	1524	840	307	307	154	76	38	19
Precalculation scaling	1	1.81	4.95	1.00	2.00	4.00	7.98	15.76
First forward modelling time (s)	902	555	284	284	179	130	109	96
First forward modelling scaling	1	1.63	3.17	1	1.59	2.18	2.59	2.95
Total time (s)	21416	11886	4438	4438	1634	857	621	489
Total scaling	1.00	1.80	4.83	1.00	2.71	5.18	7.15	9.08

Table 1: Parallel scaling of the AEM inversion code in frequency domain.

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 an option in the program which recalculates the forward response only if the conductivity model under each data point changes more than a 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 Table 1 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 E5-2680 v3 (Haswell) 2.5 GHz CPUs. For single node (4 MPI tasks) runs we vary the number of OpenMP threads from 1 to 6, with multiple MPI tasks we use 6 OpenMP threads per MPI task, thus running 4 tasks per node. In the first three columns of the table we run on one node changing the number of threads from 1 to 6. That gives us an idea of the thread-based OpenMP parallel scaling--albeit running on 4 MPI tasks, as the problem is still fairly large for a single task single thread.

The remaining columns increase the node count, thus evaluating distributed memory MPI scaling. The parallel scaling is shown related to one thread in the first three columns, and relative to one node in the remaining columns, to denote separately the OpenMP and MPI scaling. We look at three different scaling characteristics. The precalculation includes once-per-inversion run 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 modelling 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. Table 1 shows frequency domain where the footprint includes several thousand cells. In time-domain, the footprint is larger and we observe improved OpenMP scaling by a factor of 30-50%. In case of MPI, the poorer scaling is mostly due to the imbalance created by different convergence times for the iterative solver for each data point. Finally, looking at the performance of the whole inversion, we notice improvement in scaling, as the subsequent modelling iterations show less imbalance. 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 STUDIES

Frequency Domain AEM surveys

Commercial frequency domain systems have a transmitter and receiver offset in the horizontal in-line direction by a few metres. The transmitter and receiver loops are either vertical coaxial or horizontal coplanar and are housed in a small torpedo-like "bird". The systems are flown at approximately 30 m of terrain clearance slung under a helicopter. The low altitude and slow speed, coupled with the fact that the highest frequency is typically close to 100 kHz, give an accurate and high-resolution image of the subsurface formations. However, the frequency domain operations and the small transmitter-receiver offset limit the depth of penetration of the frequency domain EM field, especially in areas of conductive overburden. This creates a relatively small sensitivity domain and depth of exploration. Since only 5 - 6 frequencies are measured, the inversion of the frequency domain data is computationally considerably faster than the time-domain inversion.

For frequency domain inversion, we use the following workflow:

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

In the demonstration of the frequency domain AEM inversion, we used the Fugro RESOLVE data set obtained for the USGS in the Yukon Flats area near Fort Yukon, Alaska (Ball et. al., 2011; Minsley et. al., 2012). The data were collected over ~1200 line-km with six frequencies between 0.4 and 129 kHz, covering about a 300 km^2 area; line spacing approximately 350 m; and data spacing approximately 7 m. Figure 2 shows the location and geology in the area, with the AEM survey denoted by a small grey block.

In many frequency domain AEM surveys, including this one, flight-line spacing can be larger than the sensitivity extent, which leads to striping in the resulting 3D conductivity model. This is demonstrated in a 15-m deep horizontal section of a subset of the data in Figure 3(a). We generally tend to perform a certain degree of cross-line smoothing in the model, typically via gradient regularization, and boxcar smoothing across half the cross-line distance. This results in the more continuous model shown in Figure 3(b), although the smoothing tends to increase the misfit to some extent.

The choice of variable background leads to faster and better convergence and elimination of conductivity model artifacts caused by the large difference between the background and actual model conductivity. This is shown for the same model subset in Figure 3(c), where we used a homogeneous half-space background of 30 Ohm-m. The final RMS misfit of this inversion was 8.2 as compared to 2.7 for the variable background 3(b). Notice also several model artifacts in the constant background model.

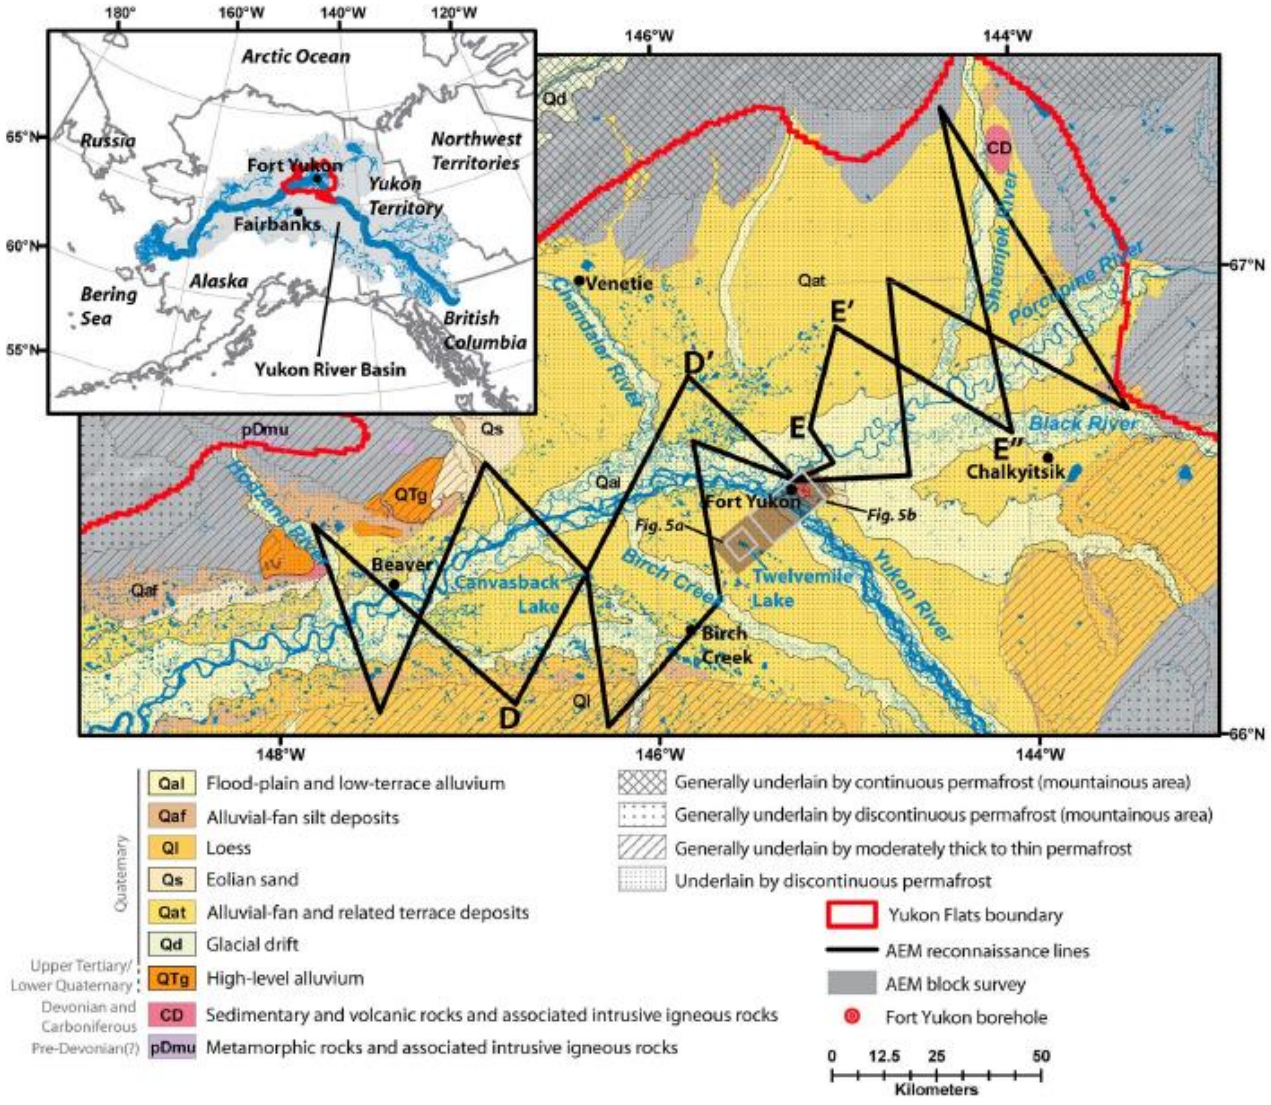

Figure 2: Location and geology of the Fort Yukon survey area. From Minsley et. al. (2012).

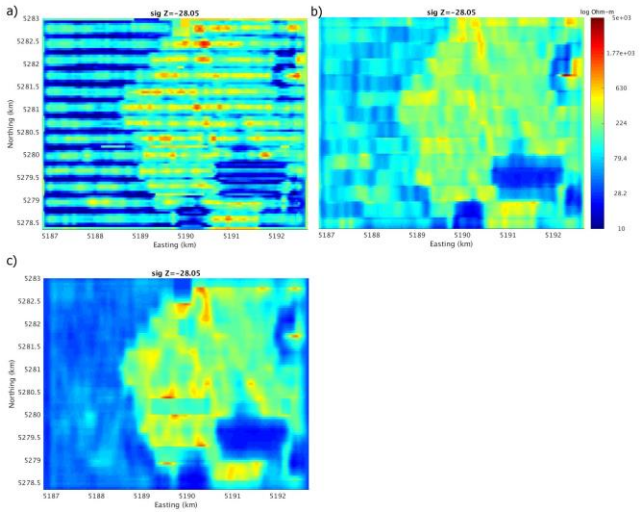

Figure 3: Horizontal cross sections of Ft. Yukon survey subset at 15 m depth obtained with: (a) no line smoothing and variable background, (b) line smoothing and variable background, (c) line smoothing and constant 30 Ohm-m half-space background.

The full data set was inverted on a 10 m x 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 m, which ended up being 81,185 receivers with 6 frequency readings each. The MSD was set to a 400 m diameter. The 3D inversion was run on 120 nodes with two six-core Intel Xeon X5660 2.8 GHz CPUs and 24 GB of RAM, and took 3.5 hours to complete with 25 forward modelling updates needed for convergence.

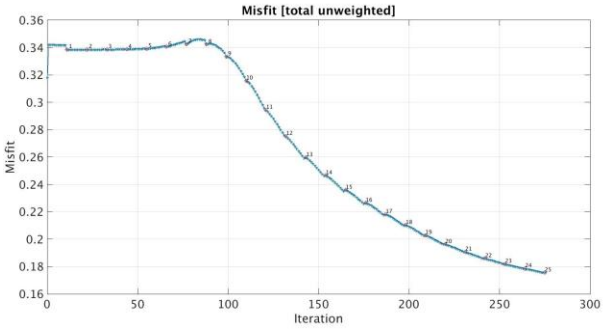

Figure 4: Misfit curve of the Yukon River FD AEM inversion.

Figure 4 shows the misfit decrease during the inversion. The total number of inversion iterations was about 275 with 25 forward modelling updates marked by the numbers in the graph. An example of the data fit for one line (line C-C' in Figure 9) is shown in Figure 5. The fit is very good for the middle frequency range and deteriorates for the lowest and the highest frequencies. As an illustration of a good fit throughout the whole survey, in Figures 6 and 7 we show a map view of the observed and predicted data for the real and imaginary parts of the data at a coplanar frequency of 8200 Hz.

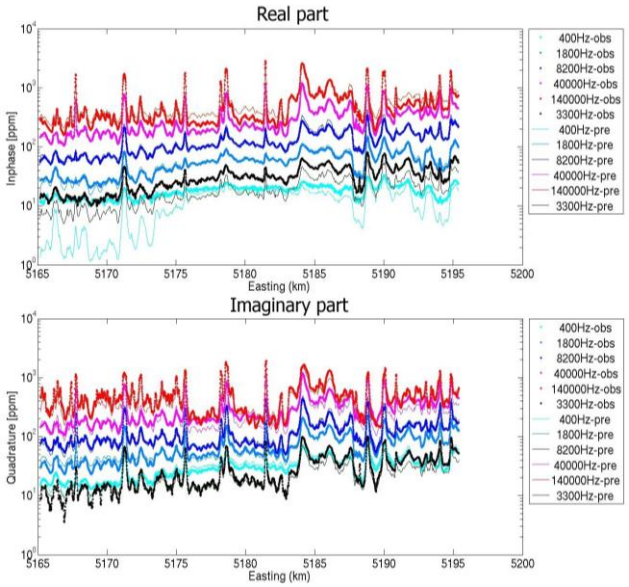

Figure 5: Observed (obs) and predicted (pre) data for Line 5 of the FD AEM Ft. Yukon survey.

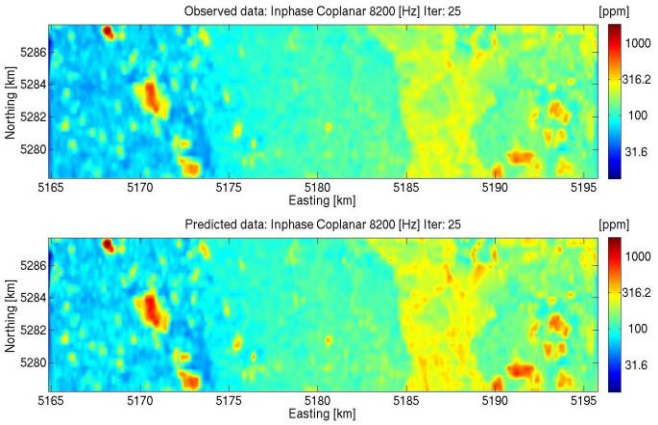

Figure 6: Map view of observed and predicted data of the real part of the coplanar field component at 8200 Hz.

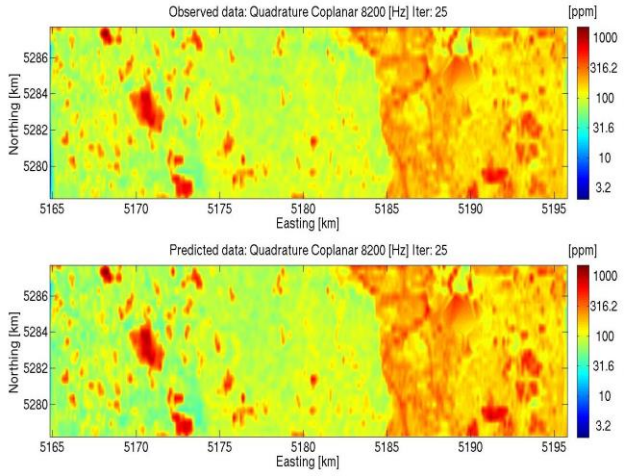

Figure 7: Map view of observed and predicted data of the imaginary part of the coplanar field component at 8200 Hz.

The goal of this survey was permafrost mapping. The area is near the boundary of continuous permafrost to the north, and discontinuous permafrost to the south, making it a good location for studying permafrost dynamics. According to Minsley et. al. (2012), the uppermost unfrozen Eolian silt and sands have an expected resistivity of 100–200 Ohm-m. At greater depths, there are frozen fluvial gravels with a resistivity greater than 1000 Ohm-m, below which are lacustrine silts and clays with a resistivity near or below 100 Ohm-m. Within the survey area there are numerous water bodies, including the Yukon River and Twelvemile Lake. Water resistivity in Twelvemile Lake was measured at 18 Ohm-m, and the lowest water resistivity in the area was recorded at 2.5 Ohm-m. Other frequencies show similar behavior, with the lowest and highest ones, as mentioned earlier, showing minor discrepancies.

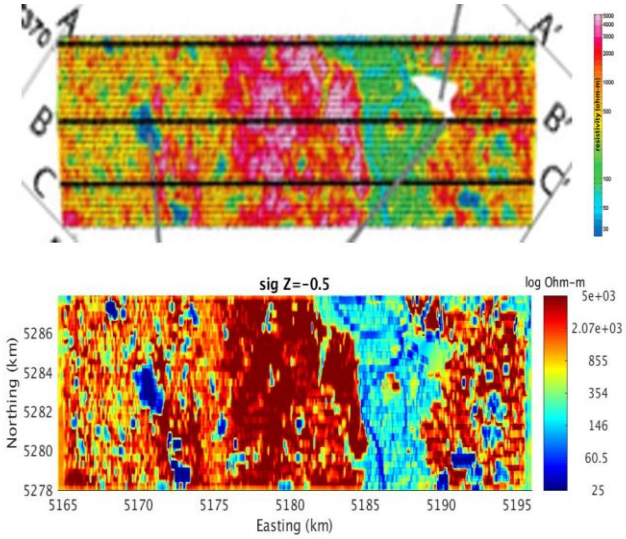

Figure 8: Horizontal cross sections of Ft. Yukon survey at 0 m depth obtained by Minsley et al. (2012) (upper panel) and by this study (lower panel).

In Figures 8–11 we compare horizontal and vertical cross sections of our resistivity model with that of Minsley et. al. (2012) obtained with 1D inversion. The models are very similar. At the surface, we notice an area having a resistivity of 100 Ohm-m and less that suggests an unfrozen area, 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 centre of the figures. The rest of the surface is highly resistive and consists of frozen silts and sands. In a depth slice at 45 m (Figure 9) the leftmost third of the area represents a frozen resistive background with occasional conductors caused by unfrozen areas under water bodies. Roughly the central third of the area is less resistive, suggesting partially frozen sediments, and this is flanked to the east by a conductive unfrozen area under the Yukon River.

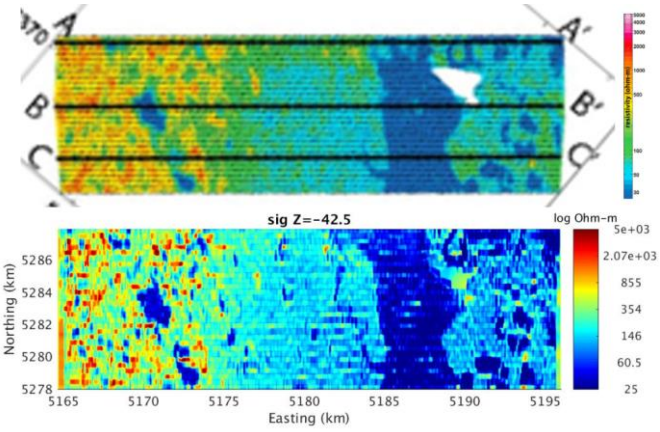

Figure 9: Horizontal cross sections of Ft. Yukon survey at 45 m depth obtained by Minsley et al. (2012) (upper panel) and by this study (lower panel).

The vertical slices, shown along the profiles in the horizontal slice figures, offer a complementary view. The Yukon River channel and the water bodies are clearly shown as conductive through the resolution limit of the AEM data. More continuous permafrost to the northwest is presented as an area of deeper resistivity, with less depth to the east, which is interpreted as a result of migration of the Yukon River in the past few thousand years.

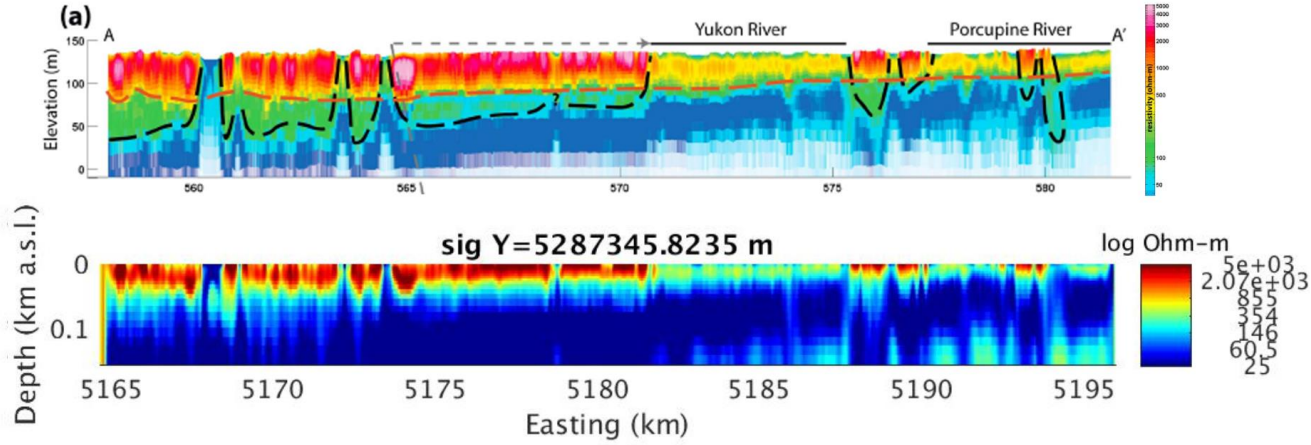

Figure 10: Vertical cross sections of Ft. Yukon survey along profile A-A' obtained by Minsley et al. (2012) (upper panel, interpreted lithologic and permafrost boundaries plotted as dashed lines) and by this study (lower panel).

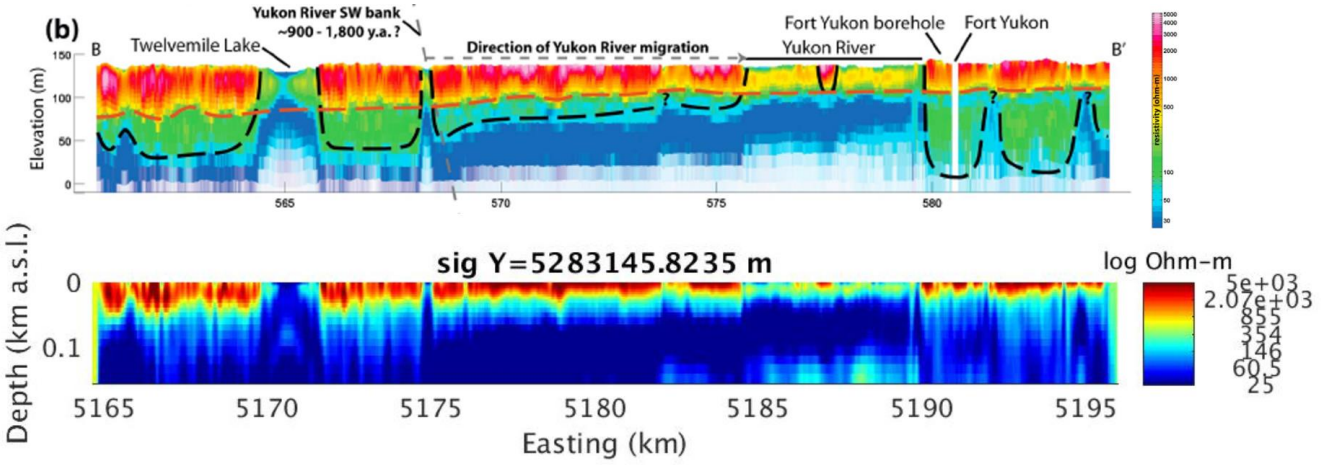

Figure 11: Vertical cross sections of Ft. Yukon survey along profile B-B' obtained by Minsley et al. (2012) (upper panel, interpreted lithologic and permafrost boundaries plotted as dashed lines) and by this study (lower panel).

Time-domain AEM surveys

Time-domain systems have a much wider range of configurations than frequency domain systems. However, regarding inversion, all time-domain systems transmit at much lower base frequencies than the frequency domain systems, typically 20–50 Hz instead of 100s of Hz. This, combined with the measurements being made in the time-domain, potentially produces a much larger sensitivity domain. The depth of investigation and number of layers that can be resolved by the time-domain systems, also increase, leading to a slightly more complex workflow as compared to the frequency domain systems.

In the time-domain inversion, we use the following workflow:

- Perform 1D half-space inversion to obtain the best fit half-space conductivity under each data point; this constitutes the half-space background under each data point.
- Extrapolate and smooth this 1D model; this constitutes 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, four per decade, bound by this minimum and maximum; this constitutes the half-space background model for forward modelling.
- Run the 3D inversion.

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.

We have two case histories to demonstrate the time-domain inversion capabilities. The first is from Kamiskotia area, Ontario, Canada from a MegaTEM survey. The second one uses the TEMPEST system.

MegaTEM II TDEM

The MegaTEM II survey was flown in 2003 by Fugro in the Kamiskotia area of Ontario, Canada. They acquired 3700 line-kilometres of electromagnetic data with a 90 Hz base frequency with a half-sine waveform. Fifteen off-time channels were used. The receiver was 128 m behind and 50 m below the transmitter. We inverted this data set using the parallel 3D inversion software with moving sensitivity domain, described above. This inversion provides an excellent example of the power of the moving sensitivity domain approach. The inversion domain was discretized into 50 m x 50 m cells horizontally and logarithmically spaced with depth. Approximately three million data points were inverted to an inversion domain with 15 million active cells. This inversion used a half-space model. The variable background model was used, but the conductivity varied only horizontally and not vertically. Thus, a best-fit homogenous half-space for each transmitter-receiver pair was found and included in the background and a priori model.

Figure 12 shows a horizontal slice extracted from the 3D inversion. The white areas with "missing" conductivity are areas with significant powerline contamination that were removed from the model. The insert in the figure shows a closeup of the inversion result. This feature shows multiple faults, which corresponds well to the known geology, and excellent fidelity.

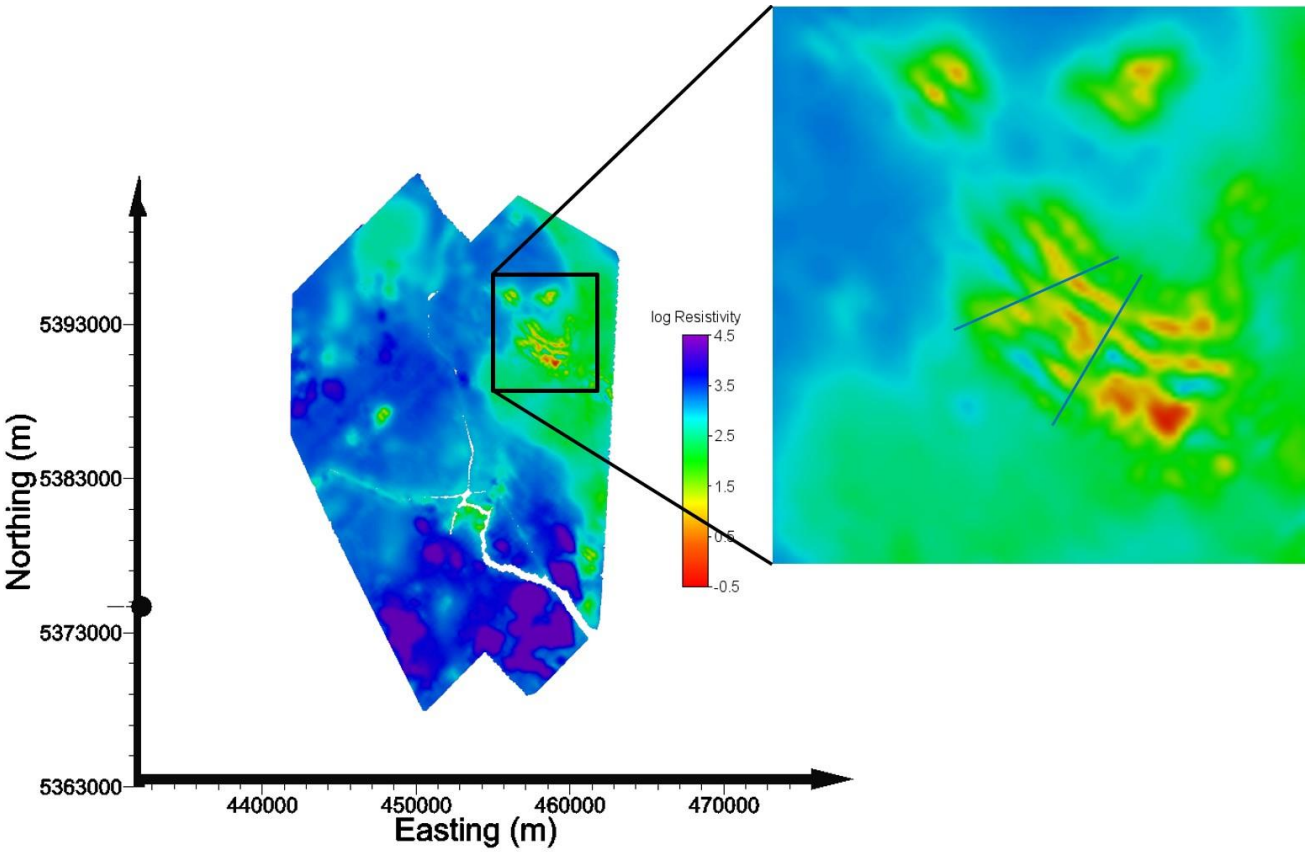

Figure 12: A horizontal slice extracted from the 3D inversion at a depth of 350 m below the surface. The blue lines in the insert are interpreted faults that are also mapped on the ground.

TEMPEST TDEM

The TEMPEST time-domain fixed-wing system was configured as a 75 Hz base frequency with a 100% duty-cycle square-wave transmitter waveform. The system recorded 13 channels of in-line and vertical B data 84 m behind and 54 m below the transmitter. B-field channels from 6.5 µs to 6.5 ms were recorded.

The survey was inverted on a horizontal grid of 20 m x 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. Twelve time channels of 15,788 measurement positions were used with a spacing of about 40 m, resulting in 189,456 data points. The inversion was run on 43 nodes with 24 CPU cores each (Xeon E5-2680 v3) and took 18.5 hours to achieve a 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 m x 800 m) and to compute a larger number of frequencies (32 in this case, in the range of 0.1 Hz to 100 kHz). Misfit decrease is shown in Figure 13. The model converges quite quickly thanks to a 1D fitted layered background.

Figure 14 shows the data fit for line 48 of the survey located approximately at northing 8746.5 km. The short arrival times fit quite well, with the fit deteriorating as the time increases.

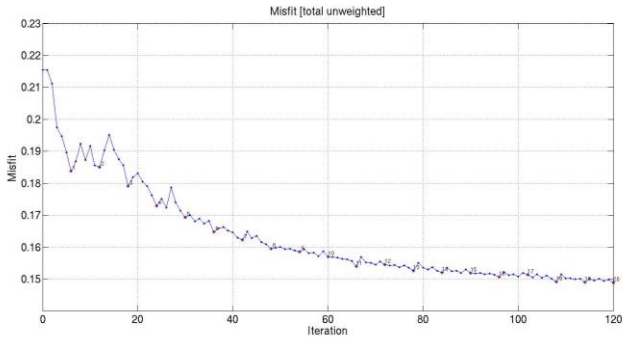

Figure 13: Misfit curve of the time-domain inversion.

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 are a conductive overburden of variable thickness and uneconomic near-surface conductive lineaments. These 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 15 shows the conductive overburden of variable thickness in the area.

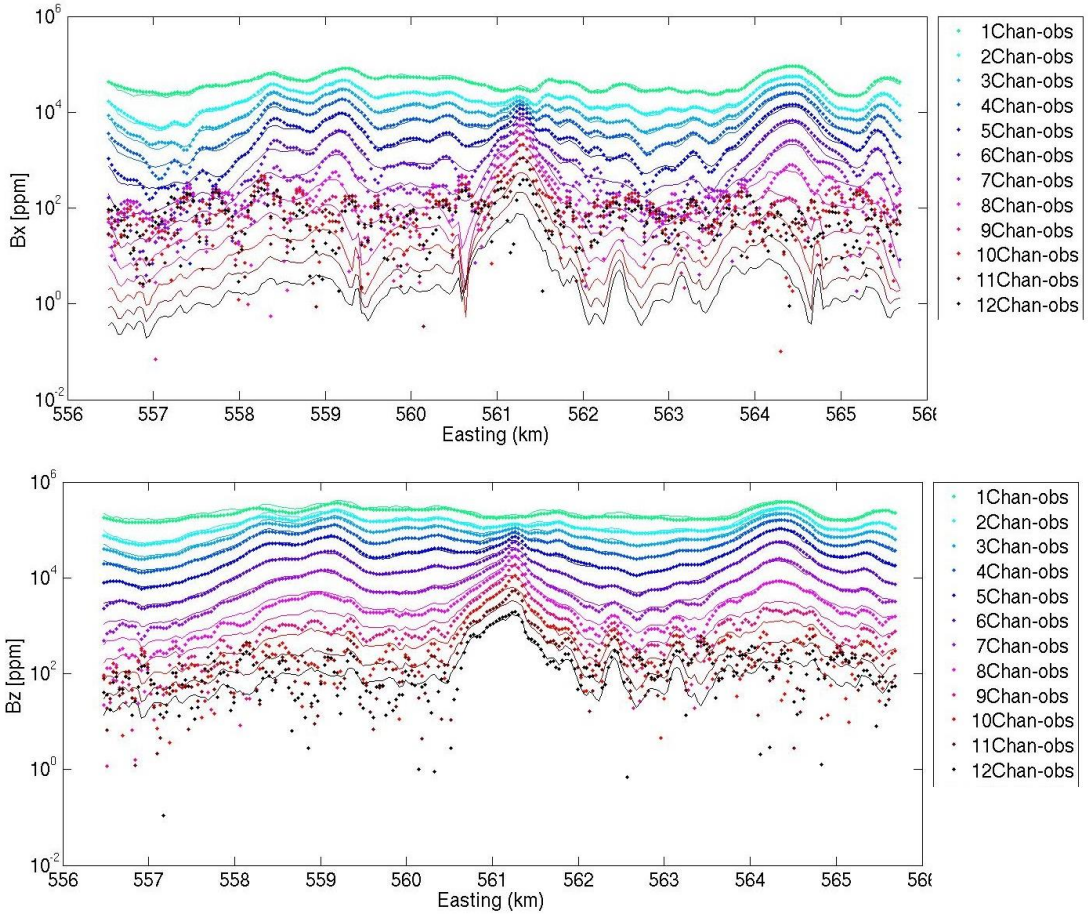

Figure 14: Observed and predicted data plot for line 48 of the time-domain survey.

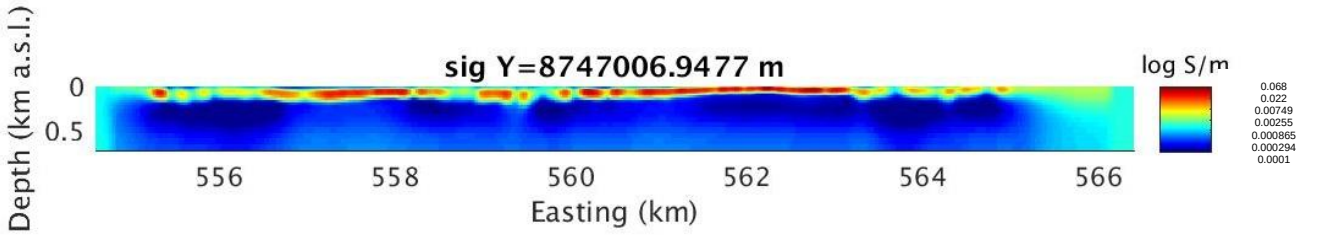

Figure 15: Vertical slice through 3D inversion result. The conductive overburden is well imaged with the time-domain data.

Figure 16 shows a very near-surface conductive feature at 2 m depth. There is a small drainage channel which flows through the property. Figure 17 shows the near-surface (~50 m below the surface) lineaments that are not of economic interest, but are shown in the data as conductors. The conductive features of economic interest are shown clearly in Figure 18, which is a horizontal slice at 325 m below the surface. The black shale units are clearly imaged in this figure.

CONCLUSIONS

In this paper, we have introduced a method and optimized workflows for the inversion of frequency and time-domain AEM data. Utilization of a moving sensitivity domain approach 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 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 metres.

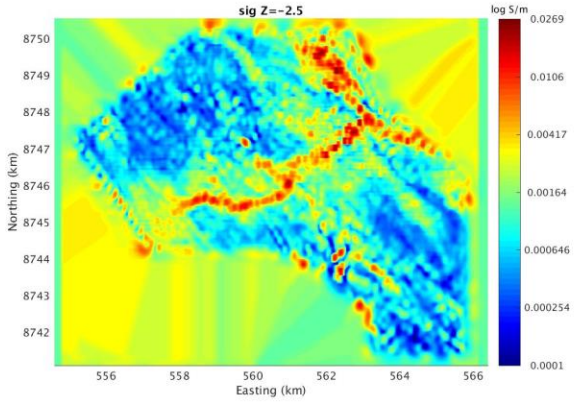

Figure 16: Horizontal slice through 3D inversion result at 2 metres below the surface. A superficial water feature is shown as a conductor.

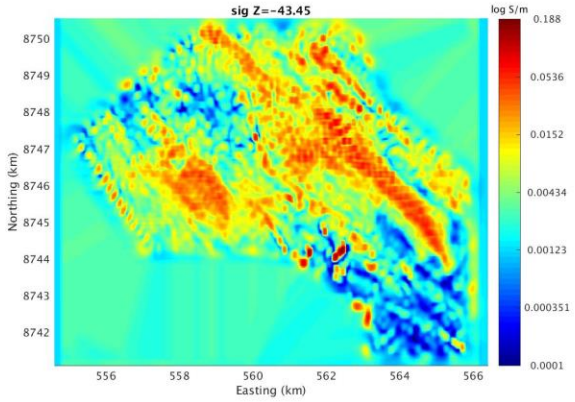

Figure 17: Horizontal slice through 3D inversion result at 50 metres below the surface. Geological near-surface conductive lineaments are visible. These are not of economic interest.

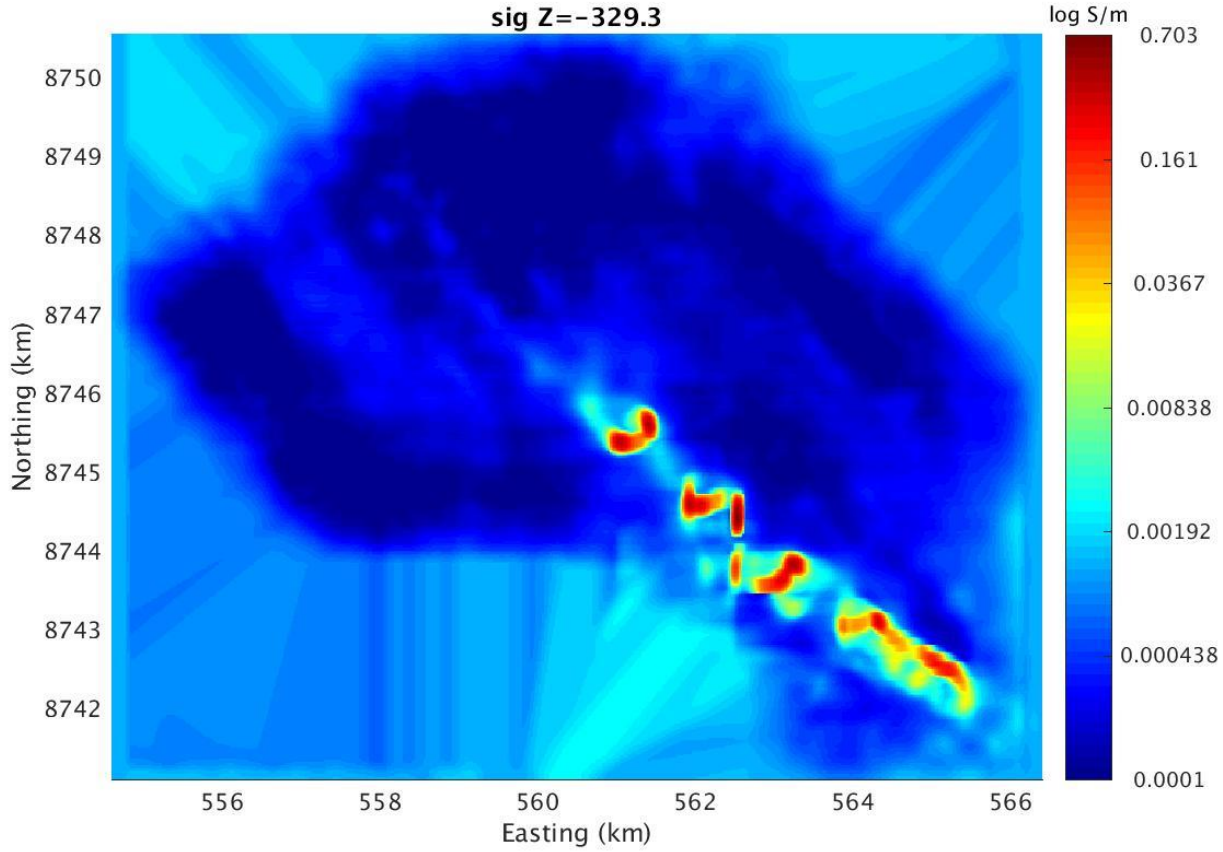

Figure 18: Horizontal slice through the 3D inversion results at 325 m below the surface. The economic mineralization is being shown here by association with the black shales.

Finally, we suggest that the developed novel approach to rigorous 3D inversion of the AEM data can be used not only for the interpretation of future ADM surveys, but for reprocessing historical AEM data, collected in the past. This will make it possible to extract valuable information from the existing AEM data, which was not possible to achieve using the traditional layered-earth inversions.

ACKNOWLEDGMENTS

The authors acknowledge TechnoImaging for support of this research and for permission to publish.

We also acknowledge the support of The University of Utah's Consortium for Electromagnetic Modeling and Inversion (CEMI) for support of this research and Center for High Performance Computing for computer time allocation.

The authors are grateful to FUGRO and MMG for providing the field data for the case studies presented in this paper.

REFERENCES

Ball, L. B., B.D. Smith, B. J. Minsley, J. D. Abraham, C. I. Voss, M. Deszcz-Pan, and J. C. Cannia, 2011, Airborne electromagnetic and magnetic survey data of the Yukon Flats and Ft. Wainwright areas, central Alaska, June 2010: U.S. Geological Survey Open File Rep., 2011-1304.

Beamish, D., 2003, Airborne EM footprints: Geophysical Prospecting, 51, 49-60.

Cox, L. H. and M. S. Zhdanov, 2007, Large-scale 3D inversion of HEM data using a moving footprint: 77th Annual International Meeting, SEG, Expanded Abstracts, 467-470.

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.

Cox, L. H., G. A. Wilson, and M. S. Zhdanov, 2012, 3D inversion of airborne electromagnetic data: Geophysics, 77 (3), 316-321.

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.

Guillemoteau, J., P. Sailhac, and M Behaegel, 2012, Fast approximate 2D inversion of airborne TEM data: Born approximation and empirical approach: Geophysics, 77, WB89-WB97.

Haber, E. and C. Schwarzbach, 2014, Parallel inversion of large-scale airborne time-domain electromagnetic data with multiple OcTree meshes: Inverse Problems, 30.

Hursán, G. and M.S. Zhdanov, 2002, Contraction integral equation method in three-dimensional electromagnetic modelling: Radio Science, 37.

Kirkegaard, C. and E. Auken, 2015, A parallel, scalable and memory efficient inversion code for very large-scale airborne electromagnetics surveys: Geophysical Prospecting, 63, 495-507.

Ley-Cooper, A.Y., A. Viezzoli, J. Guillemoteau, G. Vignoli, J. Macnae, L. Cox, and T. Munday, 2014, Airborne electromagnetic modelling options and their consequences in target definition: Exploration Geophysics, 46, 74-84.

Liu, G. and A. Becker, 1990, Two-dimensional mapping of sea-ice keels with airborne electromagnetics: Geophysics, 55, 239-248.

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

Minsley, B. J., J. D. Abraham, B. D. Smith, J. C. Cannia, C. I. Voss, M. T. Jorgenson, M. A. Walvoord, B. K. Wylie, L. Anderson, L. B. Ball, M. Deszcz-Pan, T. P. Wellman, and T. A. Ager, 2012, Airborne electromagnetic imaging of discontinuous permafrost: Geophysical Research Letters, 39, L02593.

Reid, J. E., A. Pfaffling, and J. Vrbancich, 2006, Airborne electromagnetic footprints in 1D earths: Geophysics, 71, G63-G72.

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.

Yang, D., D. W. Oldenburg, and E. Haber, 2014, 3-D inversion of airborne electromagnetic data parallelized and accelerated by local mesh and adaptive soundings: Geophysical Journal International, 196, 1492-1507.

Zhdanov, M. S., 2002, Geophysical Inverse Theory and Regularization Problems: Elsevier, Amsterdam, 609 p.

Zhdanov, M. S., 2009, Geophysical Electromagnetic Theory and Methods: Elsevier, 848 p.

Zhdanov, M. S., 2015, Inverse Theory and Applications in Geophysics: Elsevier.

Zhdanov, M. S., G. Wilson, and L. Cox, 2016, Method of real time subsurface imaging using electromagnetic data acquired from moving platforms: U. S. Patent No. 9322910.

Zhdanov, M. S., and L. Cox, 2017, Method of subsurface imaging using superposition of sensor sensitivities from geophysical data acquisition systems: U. S. Patent No. 9542359.