GEOPHYSICS, VOL. 77, NO. 4 (JULY-AUGUST 2012); P. WB37–WB45, 16 FIGS.
10.1190/GEO2011-0367.1

Large-scale inversion of ZTEM data

Elliot Holtham[1] and Douglas W. Oldenburg[1]

ABSTRACT

A Z-Axis Tipper Electromagnetic Technique (ZTEM) survey is an airborne natural source electromagnetic survey that relates the vertical magnetic field to the horizontal magnetic fields measured at a reference station on the ground. For large airborne surveys, the high number of cells required to discretize the entire area at a reasonable resolution can make the computational cost of inverting the data set all at once prohibitively expensive. We present an iterative methodology that can be used to invert large natural source surveys by using a combination of coarse and fine meshes as well as a domain decomposition that allows the full model area to be split into smaller subproblems, which can be run in parallel. For this procedure, the entire data set is first inverted on a coarse mesh. The recovered coarse model and computed fields are used as starting models and source terms in the subsequent tiled inversions. After each round of tiled inversions, the tiles are merged together to form an update model, which is then forward modeled to determine if the model achieves the target misfit. Following this procedure, we first invert the data computed from a large synthetic model of the Noranda mining camp. The inverted models from this example are consistent among our different tiling choices. The recovered models show excellent large-scale agreement with the true model and they also recover several of the mineralized zones that were not apparent from the initial coarse inversion. Finally, we invert a 30×30 km block of the 2010 ZTEM survey collected over the porphyry Pebble Deposit in Alaska. The inverted ZTEM results are consistent with the results obtained using other electromagnetic methods.

INTRODUCTION

Because most shallow and outcropping deposits have already been discovered, many of the earth's remaining natural resources are buried deep and under cover, making them difficult to find. To satisfy the strong global resource demand, techniques to explore for these big targets must be developed. Electromagnetic methods can be used to map electrical conductivity that can then be linked to geologic features of interest. When exploring for large buried targets, natural source electromagnetic methods can be advantageous over controlled source methods because of the deeper penetration of plane-wave sources.

Although the traditional natural source method, the magnetotelluric (MT) method, has been effectively applied to mining and hydrocarbon exploration, a practical limitation of the MT technique is that surveys are costly and time-consuming because many expensive stations must be installed to measure the needed electromagnetic field components on the earth's surface. It was the desire to collect airborne natural source data that prompted the development of Audio Frequency Magnetics (AFMAG) (Ward, 1959). Unfortunately, because the direction and strength of the inducing source fields varies with time, AFMAG results were not always repeatable (limitations of AFMAG are outlined in Ward et al. [1966]). More recently, Geotech Ltd. has modified the AFMAG concept and developed the Z-Axis Tipper Electromagnetic Technique (ZTEM) (Lo and Zang, 2008) technique. In the ZTEM technique, the vertical component of the magnetic field is recorded above the entire survey area, while the horizontal magnetic fields are recorded at a ground-based reference station. MT-processing techniques yield frequency domain transfer functions typically between 30 and 720 Hz that relate the vertical fields over the survey area to the horizontal fields at the reference station. Although the ZTEM method has no sensitivity to purely 1D conductivity structures (purely 1D structures produce zero vertical magnetic field component), the method is highly

Manuscript received by the Editor 29 September 2011; revised manuscript received 9 January 2012; published online 9 July 2012.
[1]University of British Columbia, Department of Earth and Ocean Sciences, UBC-Geophysical Inversion Facility, Vancouver, British Columbia, Canada. E-mail: eholtham@eos.ubc.ca; doug@eos.ubc.ca.

sensitive to lateral conductivity contrasts. If ZTEM inversions are started with good estimates for the background conductivity, then the inversions can recover reasonable conductivity estimates.

Large ZTEM data sets can be effective exploration tools, particularly when exploring on the district and regional scale where subsurface geology is hidden under cover. However, to justify the financial expense of collecting the data, viable interpretation methods must exist. This requires that inversions be fairly fast, especially for field data where multiple inversions are often performed using different parameters. Unfortunately, for surveys that cover large areas, and where even moderate inverted resolution is desired, the discrete 3D Maxwell systems for solving the electromagnetic (EM) problem quickly become very large. Although computer technology continues to improve, clock speed and instruction level parallelism have started to flatten since 2003 (Shalf, 2007). This means that most future improvements in the size of inverse problem that can be tackled must come from improved methodologies and increased parallelism.

One viable method to solve large inverse problems is to reduce the modeling domain by using footprint methods, where computations are simplified by only considering model cells that influence the data above some threshold criterion. These methods have been used to invert large airborne controlled source electromagnetic data sets (Cox et al., 2010) and MT data sets (Gribenko et al., 2010). Another solution is to use domain decomposition methods (DDM) and split the computational domain into smaller manageable subproblems that can be solved quickly in parallel before being merged together to form the final solution. In this paper, we present a highly parallel and effective tiling procedure to invert large natural source surveys.

ZTEM DATA

The ZTEM data relate the vertical magnetic fields computed above the earth to the horizontal magnetic field at some fixed reference station. This relation is given by

$$H_z(r) = T_{zx}(r, r_0)H_x(r_0) + T_{zy}(r, r_0)H_y(r_0), \quad (1)$$

where r is the location for the vertical field, and r_0 is the location of the ground-based reference station. Our source functions for the natural fields are random, and as with MT, we need two polarizations. The transfer functions for each polarization are given by

$$\begin{pmatrix} H_z^{(1)}(r) \\ H_z^{(2)}(r) \end{pmatrix} = \begin{pmatrix} H_x^{(1)}(r_0) & H_y^{(1)}(r_0) \\ H_x^{(2)}(r_0) & H_y^{(2)}(r_0) \end{pmatrix} \begin{pmatrix} T_{zx} \\ T_{zy} \end{pmatrix}, \quad (2)$$

where the superscripts (1) and (2) refer to the source field polarization in the x and y directions, respectively.

INVERSION METHODOLOGY

Although this paper focuses on inverting ZTEM data, the approach and methodologies are equally applicable to MT as well as a combination of MT and ZTEM data (Holtham and Oldenburg, 2010a). Our original natural source inversion algorithm is that of Farquharson et al. (2002), and our ZTEM inversion algorithm, Holtham and Oldenburg (2010b), was created by modifying this code. We wish to solve the inverse problem by minimizing the objective function

$$\Phi = \|\mathbf{W_d}(\mathbf{F}(\mathbf{m}) - \mathbf{d}^{\text{obs}})\|_2^2 + \beta(\alpha_{s,x,y,z}\|\mathbf{W}_{s,x,y,z}(\mathbf{m} - \mathbf{m}^{\text{ref}})\|_2^2), \quad (3)$$

where $\mathbf{m}$ is the full fine cell model, $\mathbf{m}^{\text{ref}}$ is a reference model, $\mathbf{F}$ is the forward modeling operator, $\mathbf{d}^{\text{obs}}$ is the observation vector, $\mathbf{W}_d$ is a diagonal matrix whose elements are the reciprocals of the standard deviations of the data errors, $\mathbf{W}_s$ is a diagonal matrix, and $\mathbf{W}_x$, $\mathbf{W}_y$, and $\mathbf{W}_z$ are the first-order finite-difference matrices in the x-, y-, and z-directions. The αs are adjustable parameters: α_s controls the closeness of the recovered model to the reference, and $\alpha_{x,y,z}$ determine the smoothing in the x-, y-, and z-directions. Here, β is the regularization parameter that is reduced throughout the inversion process until the desired data misfit ϕ_d^* has been achieved. Although the goal is to solve this problem, it is difficult to tackle directly because of the size of the problem and the cost of evaluating $\mathbf{F}(\mathbf{m})$. To avoid inverting the full domain directly, we can first invert the full data set on a coarse mesh to determine an initial coarse inversion result $\mathbf{m}_c$. This coarse mesh result can be interpolated onto the desired fine mesh to form an initial model

$$\mathbf{m} = \mathbf{L}_c^f \mathbf{m}_c, \quad (4)$$

to the full inverse problem. Here, $\mathbf{L}_c^f$ is an interpolation operator that goes from the coarse mesh to the fine mesh. Next, we can decompose the full fine-scale model domain $\mathbf{m}$, and the data set into T smaller tiles,

$$\tilde{\mathbf{m}}_t = \mathbf{P}_t^m \mathbf{m}, t = 1, .., T, \quad (5)$$

and

$$\tilde{\mathbf{d}}_t = \mathbf{P}_t^d \mathbf{d}^{\text{obs}}, t = 1, .., T, \quad (6)$$

where $\mathbf{P}_t^m$ and $\mathbf{P}_t^d$ are projection operators that form the model and data subdomains, respectively. Here, $\mathbf{P}_t^m$ and $\mathbf{P}_t^d$ are constructed such that the $\tilde{\mathbf{m}}_t$s and $\tilde{\mathbf{d}}_t$s are overlapping. We also define the projection operator,

$$\tilde{\mathbf{u}}_t = \mathbf{P}_t^u \mathbf{u}, t = 1, .., T, \quad (7)$$

which maps the fields $\mathbf{u}$ from the full conductivity structure to each subdomain. We would like to invert these T overlapping subproblems separately and in parallel; however, in doing so directly we would neglect interactions between subdomains. To mitigate this problem, we can incorporate the domain interactions through source terms obtained on a full domain model. To accomplish this, the coarse model $\mathbf{m}_c$ is first interpolated onto the full fine mesh to form an initial solution $\mathbf{m}^1$. This model is then forward modeled to compute initial fields $\mathbf{u}^1$ at each mesh location. Both $\mathbf{m}^1$ and $\mathbf{u}^1$ are then projected onto the T domains to form $\tilde{\mathbf{m}}_t^1$ and $\tilde{\mathbf{u}}_t^1$. Although these new fields and models are only defined on an individual subdomain, they were still obtained by solving the inverse problem on the full model domain and hence contain information about the conductivity outside their respective domains.

To understand how the fields computed from a full conductivity structure can be exploited in the subdomain inversions, we examine our forward modeling procedure of Farquharson et al. (2002) (the solution of Maxwell's equations is that of Haber et al.[2000]) where the total fields are computed by summing computed primary and secondary fields. The electric field is decomposed into vector

and scalar potentials $\mathbf{E} = \mathbf{A} + \nabla\phi$, and the Coulomb gauge condition $\nabla \cdot \mathbf{A} = 0$ is imposed for uniqueness. The secondary fields $\mathbf{A}_s$ and ϕ_s can then be computed as

$$\begin{pmatrix} \mathbf{L} + i\omega\mu_0\mathbf{S} & i\omega\mu_0\mathbf{SG} \\ \mathbf{DS} & \mathbf{DSG} \end{pmatrix} \begin{pmatrix} \mathbf{A}_s \\ \phi_s \end{pmatrix} = \begin{pmatrix} -iw\mu_0(\mathbf{S} - \mathbf{S}_p)\mathbf{E}_p \\ -\mathbf{D}(\mathbf{S} - \mathbf{S}_p)\mathbf{E}_p \end{pmatrix}. \tag{8}$$

Here, $\mathbf{L}$ represents the discretization of the Laplacian operator, $\mathbf{S}$ represents the harmonically averaged cell conductivities, $\mathbf{G}$ and $\mathbf{D}$ are the discretizations of the gradient and divergence operators, $\mathbf{S}_p$ contains the averaged conductivities of the primary field model, and $\mathbf{E}_p$ is the electric field computed on the primary conductivity structure. We can write the block matrix system in equation 8 more compactly as,

$$\mathbf{B}\mathbf{u}_s = \mathbf{q}(\mathbf{S} - \mathbf{S}_p)\mathbf{Q}\mathbf{u}_p, \tag{9}$$

where $\mathbf{Q}$ is some operator that maps the $A - \phi$ solution $\mathbf{u}_p$ to $\mathbf{E}_p$. In this approach, the primary conductivity structure and fields are obtained from a full conductivity structure and used as source terms on the right hand side of equation 8. Using full model domain source terms ensures that the large-scale physics of the problem will be modeled more accurately. For each subdomain mesh, we minimize the objective functions,

$$\tilde{\Phi}_t = \phi_d(\tilde{\mathbf{m}}_t) + \beta\phi_m(\tilde{\mathbf{m}}_t) \tag{10}$$

to solve the tth inverse problem. Once the $\tilde{\mathbf{m}}_t$s have been determined, the kth update to the full model $\mathbf{m}^k$ is given by merging each subdomain model together,

$$\mathbf{m}^k = M(\tilde{\mathbf{m}}^k_{t=1:T}), \tag{11}$$

where M is an operator that merges the subdomain conductivity structures. Additional details on how the tiles are merged are discussed later in the subsection, model updates, and merging tiled inversions. We can determine if the current model $\mathbf{m}^k$ is a satisfactory solution to our full inverse problem by forward modeling this new conductivity structure and computing the data misfit ϕ_d. If the data misfit falls within that allowed by the initial inverse problem in equation 3, then a solution to the desired inverse problem has been obtained. If the data have not been sufficiently fit, then we can perform another iteration of this procedure where we now use the fields and models from $\mathbf{m}^k$ as the source terms in the $k^{th} + 1$ subdomain inversions. This procedure, outlined in algorithm 1, can be repeated until the data has been sufficiently fit.

SYNTHETIC EXAMPLE

In this section, a synthetic model is used to develop a workflow type procedure by examining in detail a few elements of the algorithm described in the previous section. The synthetic model is from the Noranda District in Canada, home to 20 economic volcanogenic massive sulfide deposits (VMS), 19 orogenic gold deposits, and several intrusion-hosted copper-molybdenum deposits (Gibson and Galley, 2007). The original model, provided courtesy of the Xstrata mining group, contained 12.7 million cells covering an area of almost 20×20 km. The 38 geologic units in the model were converted into expected conductivities, and the entire 12.7 million cell model was forward modeled at 30, 45, 90, 180, 360, and 720 Hz. The data were then corrupted with noise to form the observed synthetic data. Two slices at −275 and −475 m of the synthetic model can be seen in Figures 1 and 2.

Algorithm 1 Large-scale inversion

$\mathbf{m}_c \leftarrow \text{argmin}_{\mathbf{m}_c}(\phi_d(\mathbf{m}_c) + \beta\phi_m(\mathbf{m}_c))$
Initialize: $\tilde{\beta}$
$\tilde{\mathbf{d}}_{t=1:T} = \mathbf{P}^d_t\mathbf{d}^{\text{obs}}$
$\mathbf{m}^1 = \mathbf{L}^f_c\mathbf{m}_c$
$\mathbf{u}^1 \leftarrow \mathbf{F}(\mathbf{m}^1)$
for $k = 1, 2, \ldots$ **do**
  for $t = 1 \ldots T$ (in parallel) **do**
    Set primary model:
    $\tilde{\mathbf{m}}^k_{tp} = \mathbf{P}^m_t\mathbf{m}^k$
    $\tilde{\mathbf{u}}^k_{tp} = \mathbf{P}^u_t\mathbf{m}^k$
    Invert each tile: Initialize:
    $\tilde{\mathbf{m}}^k_t = \mathbf{P}^m_t\mathbf{m}^k$
    Solve: $\tilde{\mathbf{m}}^k_t \leftarrow \text{argmin}_{\tilde{\mathbf{m}}^k_t}(\phi_d(\tilde{\mathbf{m}}^k_t) + \tilde{\beta}\phi_m(\tilde{\mathbf{m}}^k_t))$
    where: $\tilde{\mathbf{B}}(\tilde{\mathbf{m}}^k_t)\tilde{\mathbf{u}}^k_{st} = \mathbf{q}(\tilde{\mathbf{S}}^k_t - \tilde{\mathbf{S}}^k_{tp})\mathbf{Q}\tilde{\mathbf{u}}^k_{tp}$
    from equation 9
  end for
  Merge models:
  $\mathbf{m}^{k+1} = M(\tilde{\mathbf{m}}^k_{t=1:T})$
  Compute fields:
  $\mathbf{u}^{k+1} \leftarrow \mathbf{F}(\mathbf{m}^{k+1})$
  Check misfit:
  if $\phi_d(\mathbf{m}^{k+1}) <= \phi^*_d$ **then**
    break
  end if
end for

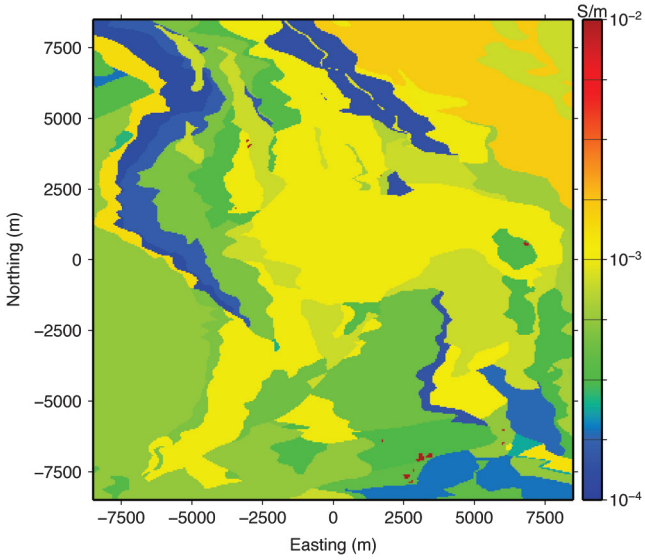

Figure 1. Depth slice at −275 m above mean sea level (AMSL) of the true model. The topography for the synthetic model ranges from −25 to −250 m AMSL. The conductivity of the ore was chosen to be 0.2 S/m; however, the color scale on the model has been clipped at 0.01 S/m to improve the visualization of the other geologic units.

Preliminary coarse inversion

The first step to find an approximate solution to the full inverse problem is to perform a coarse inversion to determine the large-scale conductivity structure that will be used to compute the primary fields and starting models for subsequent finer cell size inversions. The earth is discretized into relatively large cells such that the total number of cells in the mesh is small and the inverse problem can be solved quickly. In this example, the coarse mesh contained $42 \times 42 \times 73$ cells. The cell lengths in the x- and y-directions was 500 m. Working initially on a coarse mesh, which for this synthetic example can be inverted in approximately 30 minutes, allows multiple inversions to be simultaneously run with different parameters and starting models. The coarse inversion result can be seen in Figure 3. Because of the large cell dimensions, some geologic structures such as accurate body boundaries and fine scale features may not be recovered by the discretization. Therefore, it is important not to overfit the data and risk adding discretization artifacts into the inversion result. This is particularly true for higher frequency data that will have smaller skin depths and contain more information about fine scale features. In fact, because the initial goal of the coarse scale inversion is to quickly determine the large-scale conductivity features, some data, particularly the higher frequency data, may be omitted to reduce the computational cost and prevent the large cell sizes of the coarse mesh from violating the shorter skin depths at the higher frequencies. In this example, the 360 and 720 Hz data were omitted for the coarse mesh inversion.

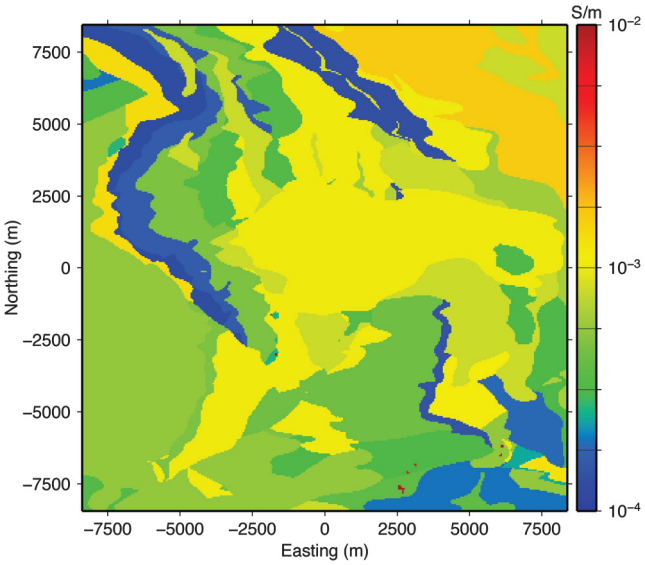

Figure 2. Depth slice at −475 m (AMSL) of the true model. The conductivity of the ore was chosen to be 0.2 S/m; however, the color scale on the model has been clipped at 0.01 S/m to improve the visualization of the other geologic units.

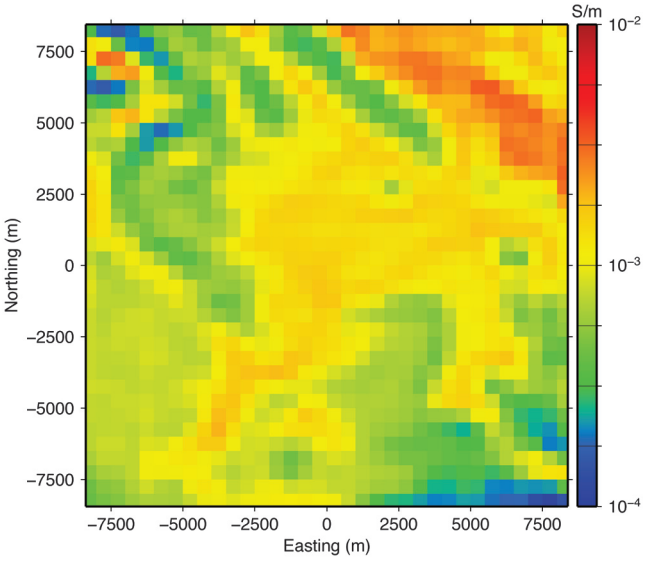

Figure 3. Coarse inversion result at −475 m (AMSL) shown on the same color scale as the true model. Even with coarse 500×500 m cells in the x- and y- directions, the main geologic features of the model have been recovered.

Interpolating models and computing primary fields

To compute the secondary fields on each subdomain tile, the primary fields are needed according to equation 8. These computed primary fields contain information about the large-scale structure and are first obtained by linearly interpolating the coarse inversion model onto the fine model domain. Linear averaging when interpolating from a coarse mesh to a fine mesh will create a smoother fine model that is consistent with our smooth model regularization in equation 3. The new interpolated fine model is then forward modeled to compute the primary electric fields that are used (equation 8) to calculate the total fields on each subdomain mesh. Whereas forward modeling operations on the full fine mesh can be time-consuming and resource intensive, this operation only needs to be performed once for each frequency and outer tile iteration. Our experience shows that working with primary fields computed using the interpolated fine model is superior to the fields directly from the coarse mesh and the additional cost of forward modeling on the fine mesh is worthwhile.

Setting up the subdomains

The coarse inversion result can be used to generate an initial starting model and to compute the primary fields for subsequent finer discretized inversions. Once the initial models and primary fields have been computed, the subdomain meshes must be designed. When decomposing the computational domain and designing the subdomain meshes, there are two critical elements that must be considered. First, the decomposition should satisfy the physics between domain interactions. Second, in the final inverted model there should be continuity of conductivity structures across the various domain boundaries. Incorporating domain interactions is done by using primary fields in the tiled inversions that have been computed from full domain conductivity structures. Continuity of the conductivity structure across domain boundaries is accomplished by using data and mesh overlap. Generally, the core regions of the meshes are chosen to overlap by slightly less than a skin depth (additional padding cells are added outside this overlap region) at the lower frequencies. Here, a tradeoff must be achieved between increased overlap and increased computational requirements. Another tradeoff that must be considered is the number of tiles used to decompose the domain. As the number of tiles increases, although the

individual inversion time for each subproblem might decrease, the percentage of overlap cells versus total number of cells in each mesh increases and the decomposition may become less efficient. Here, it is up to the user to choose a reasonable number of tiles such that each tile runs efficiently on their available computer hardware.

For this synthetic example, two different sets of meshes, one with three tiles and the other with nine tiles, are used to demonstrate the scalability and robustness of the methodology. For the three tiles example, each subdomain mesh contained approximately 3.5 million cells with $353 \times 140 \times 70$ cells in the x-, y-, and z-directions. Each mesh in the nine tile subdomain example contained approximately 1.4 million cells with $139 \times 140 \times 70$ cells in the x-, y-, and z-directions. The three and nine tile examples contained the same sized cells (50×50 m cells in the x- and y-directions) in the core and overlapping regions. In the vertical direction, the cells' dimensions started at 25 m and then expanded with depth. Both sets of meshes used 20 overlapping cells on each mesh boundary. The mesh layouts and overlaps for the three and nine tile examples can be seen in Figures 4 and 5.

Perform subdomain inversions

Inverting large field data sets without some sort of decomposition method is not feasible. Selecting the optimal tile size is a complex process that depends on many factors such as system CPU, RAM, and inversion parameters. As previously stated, as the tile size becomes smaller, although the individual inversion time for each subproblem might decrease, the percentage of overlap cells versus total number of cells in each mesh increases and the decomposition may become less efficient. The user must decide between working with smaller tiles with quicker inversion times and increased parallelism (because all of the tiles can be run simultaneously in parallel) and larger tiles that may take longer to run but have fewer overlap cells versus total number of cells and the full problem has less overlapping tile boundaries. Generally, our methodology is to select the largest tile size that can be inverted in an acceptable time. For the synthetic example, the three tiles and nine tiles examples were inverted on Intel Xeon X5660 processors, each with 64 GB of RAM. For the nine tiles example, there was sufficient memory on each computer to run the six frequencies in parallel. Each beta iteration took around 36 hours. For the three tiles example, only three frequencies could be run in parallel on each computer with the 64 GB of available memory. Here, each beta iteration took around 145 hours. During the subdomain inversions, the initial beta is chosen to be the final beta from the initial coarse inversion result, and the same cooling scheme is used on the tiled inversions as in the initial coarse inversion.

Reference station considerations

An additional complication when tiling ZTEM inversions is the reference station that is used to normalize the data because of the unknown source field amplitude. Although the location of the reference station may be included in full data set inversions, when the full modeling domain is subdivided, the true reference station location cannot practically be included in each domain. If a separate location for the reference station was used for each tile, then the local conductivity structure from each tile would influence the horizontal fields and the computed data would not be consistent between the tiles. To circumvent this problem, the horizonal field values are fixed to those obtained from the initial coarse inversion, which should be a good approximation to the true fields.

Model updates and merging tiled inversions

Once the subdomain inversions have been completed, a full model update can be obtained by merging the inverted tiles. For updating model cells that correspond to only one inversion tile, it is clear which conductivity value should be assigned to the update model. The situation becomes more complicated for update cells that correspond to multiple subdomain cells (overlapping tile regions). A simple merging technique that will mitigate discontinuities across tile boundaries is to use averaging schemes that assign the conductivity of the update cell to be the weighted average of the overlapping conductivity tiles. The weights for each cell can be based on the distance from that cell to the respective tile boundary. From our practical experience, this method works quite well; however, we can improve on this method by trying to incorporate some estimate of each cell's sensitivity to the individual tiles. First, it is assumed that the sensitivity J_{ij} of the jth cell to the ith datum has a geometric decay that drops off as $1/r^2$ according to Biot-Savart's Law. Secondly, it is assumed that there is an attenuation component that drops off as $\exp(-r/\delta)$ where δ is some approximation to the skin depth. Although clearly the skin depth will vary with frequency and model conductivities, to first order we can gain some insight into the

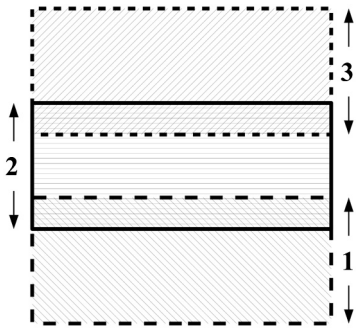

Figure 4. Mesh layouts for the example using three overlapping domains. Each subdomain mesh contained approximately 3.5 million cells with $353 \times 140 \times 70$ cells in the x-, y-, and z-directions. The central core region of each domain overlapped by 20 cells, each with dimensions of 50 m, in the x- and y-directions.

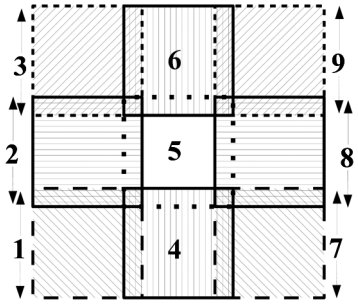

Figure 5. Mesh layouts for the example using nine overlapping domains. Each mesh contained approximately 1.4 million cells with $139 \times 140 \times 70$ cells in the x-, y-, and z-directions. The central core region of each domain overlapped by 20 cells, each with dimensions of 50 m, in the x- and y-directions.

relative sensitivity of a particular cell to each data by using a reasonable fixed skin depth. That is, we may approximate the sensitivity to be proportional to,

$$J_{ij} = \frac{K}{r_{ij}^2} \exp\left(\frac{-r_{ij}}{\delta}\right), \tag{12}$$

where K is the constant of proportionality. For cells that fall in an overlapping region, we can determine weighting schemes based on the relative sensitivity of a particular cell to each tile of interest. For some cell j, the normalized sensitivities for each tile (using a simple example of two overlapping tiles) can be given as a sum over the N_1 and N_2 data points in each tile (data points that overlap data points in other tiles are omitted in the calculation)

$$W_j^1 = \frac{1}{N_1} \sum_{i=1}^{N_1} J_{ij}^{(1)} \tag{13}$$

and

$$W_j^2 = \frac{1}{N_2} \sum_{i=1}^{N_2} J_{ij}^{(2)}. \tag{14}$$

Here, we have assumed that each data should be weighted equally. The conductivity for the jth cell $\bar{\sigma}_j$ can then be given as a weighted average of the overlapping inverted tile conductivities σ^1 and σ^2,

$$\bar{\sigma}_j = \frac{1}{W^1 + W^2}(W^1\sigma_j^1 + W^2\sigma_j^2). \tag{15}$$

This method is easily extended to more complex tile overlaps. The model updates for the synthetic and field data examples were computed using this update scheme. The final inversion models from the three and nine tiles examples in Figures 6 and 7 recover the large-scale geologic features as well as several of the larger known mineralized regions that were not evident from the coarse inversion. The inversion results of the three and nine tiles examples are very similar, illustrating that the inversion results are not highly dependant on the choice of subdomain meshes. Figure 8 compares the recovered high conductivity regions (mineralized zones) with

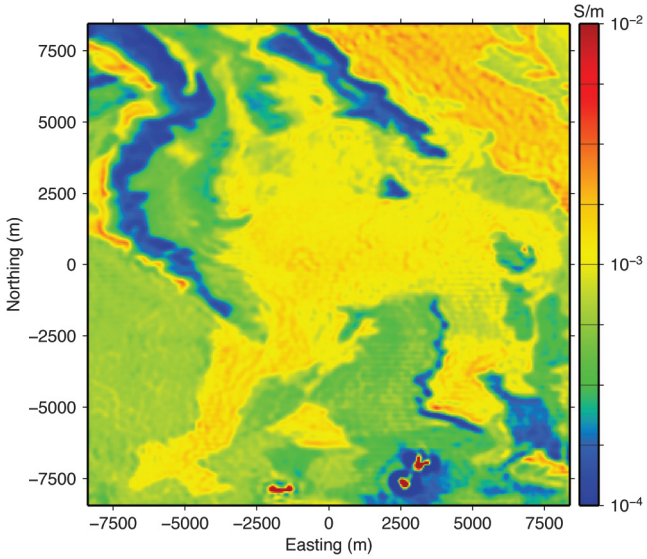

Figure 6. Depth slice at −275 m (AMSL) of the inversion result after merging the three tiles together shown on the same color scale as the true model. The resolution in this model is greatly improved compared to the coarse result. Several of the larger known deposits have now been recovered; they were not visible from the initial coarse inversion.

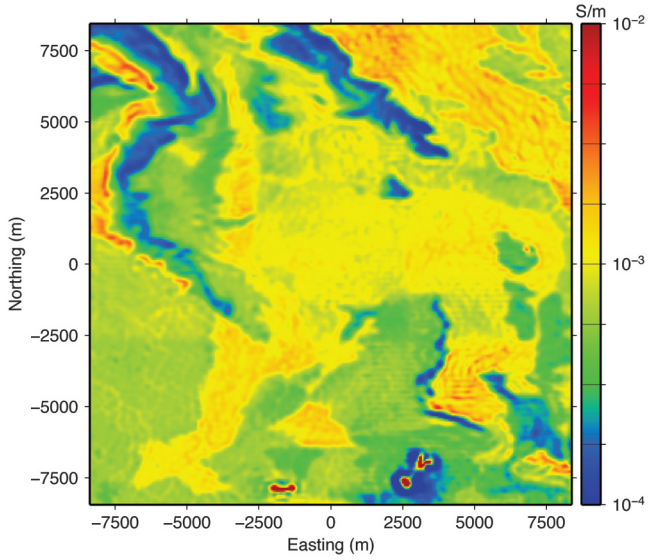

Figure 7. Depth slice at −275 m (AMSL) of the inversion result after merging the nine tiles together shown on the same color scale as the true model.

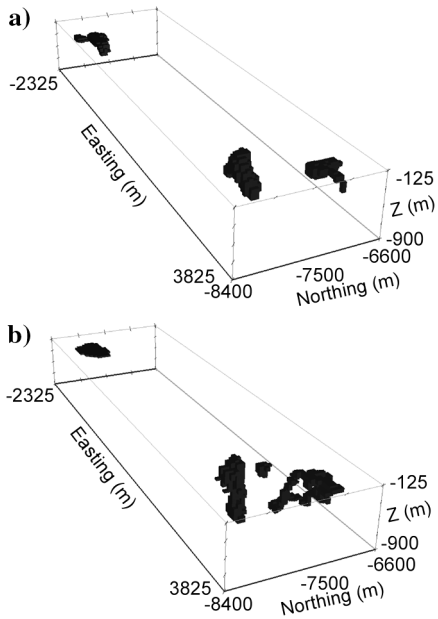

Figure 8. Comparison of the larger southern mineralized regions. (a) Recovered bodies more conductive than 0.01 S/m that corresponds to the mineralized bodies in the southern section of the inverted model. (b) Actual mineralized bodies.

the true mineralized zones. Although one cannot expect to recover the exact complex geometry of these relatively small mineralized zones, the inverted model does a good job of recovering the general geometry of the ore bodies.

PEBBLE FIELD DATA EXAMPLE

The Pebble deposit is a world-class calc-alkalic copper-gold molybdenum porphyry deposit located in the Bristol Bay region of southwest Alaska (geographic location shown in Figure 9). The deposit is managed by the Pebble Partnership, a joint venture between Northern Dynasty Mines and Anglo-American plc. Because of the size of the deposit and the quality and diversity of geophysical data collected over the area, this deposit is an ideal test site for geophysical techniques and inversion methodologies. Over the years, many groups have looked at the geologic and geophysical data from the area, but in particular Paré and Legault (2010) provide a fairly comprehensive overview of the general geology and the available geophysical data sets. In summary, from Paré and Legault (2010), the Pebble deposit consists of two contiguous zones, the Pebbel west and east zones. The Pebble west zone occurs around small granodioritic stocks that intrude the country rock, and the Pebble east zone occurs within a granodioritic stock and in sills that cut the country rocks. Because the Pebble west zone extends to the surface, the deposit was discovered in 1986, far before its deeper and higher grade neighbor, which was not discovered until 2005. Pebble east is completely hidden by up to 600 m of volcano-sedimentary cover, which explains why it took almost two decades to discover and highlight the importance to develop deep probing geophysical techniques. More recently, the far east zone has been discovered at a depth of greater than 1.5 km. The deposit hosts K-silicate alteration associated with quartz-sulfide veins, overprinted by phyllosilicate alteration. The higher grade mineralization at Pebble east is associated with advanced argillic alteration. From Rebagliati et al. (2009), the measured and indicated resource of both zones is given at a 0.3% equivalent copper cut-off grade of 5.096 billon tons at 0.43% Cu, 0.35 g/t Au, and 256 ppm Mo. A plan view and cross section of the Pebble west and east zones, can be seen in Figures 10 and 11.

Over the years, a wealth of geophysical data has been collected over the Pebble deposits. Along with potential field data, multiple DC-IP surveys have been collected along with two small MT surveys. In the summer of 2009, a 3840 line-km Spectrem survey (Leggatt et al., 2000) was performed. In the same summer, a small 250 line-km ZTEM survey was flown over the Pebble deposit. In the summer of 2010, a much larger 6504 line-km ZTEM survey was flown over the deposit and surrounding areas. The 2010 survey used 250 m line spacing and collected 30, 45, 90, 180, 360, and 720 Hz data. In this section of the paper, we apply our large-scale inversion methodology outlined in

Figure 9. Location of the Pebble deposit in the Bristol Bay region of Alaska.

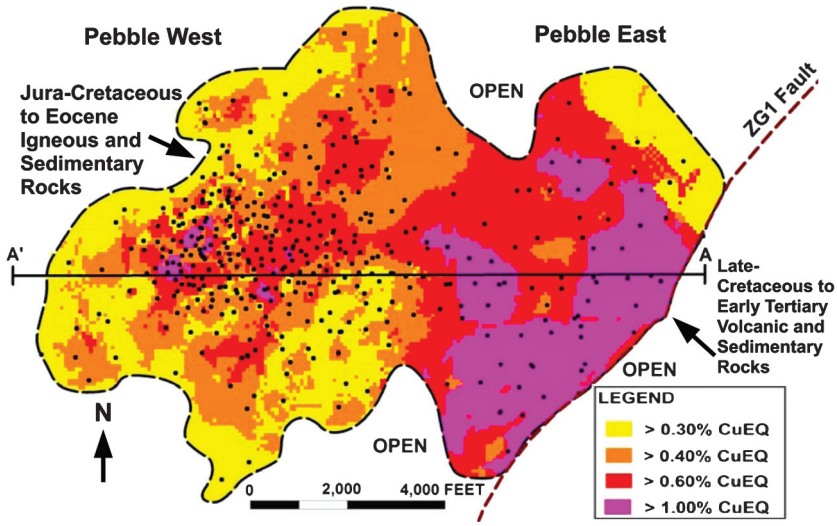

Figure 10. Plan view of the Pebble deposit showing the deposit outline and grade distribution across the deposit (modified after Rebagliati et al. [2009]).

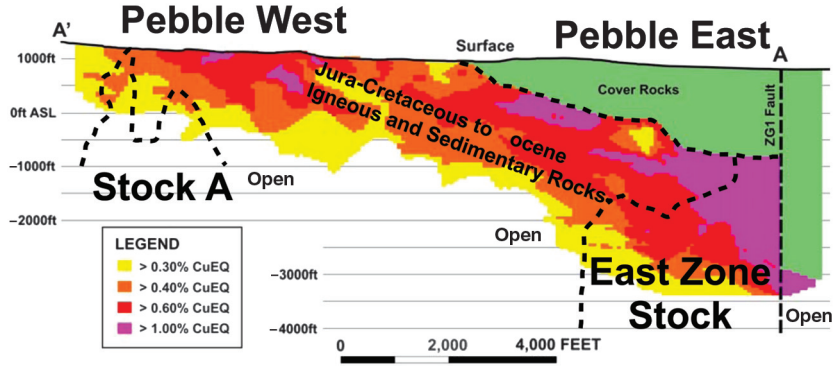

Figure 11. Cross section of the Pebble deposit showing the grade distributions and the intrusive stocks (modified after Rebagliati et al. [2009]).

the previous sections to invert a large 30×30 km section of the newer 6504 line-km ZTEM survey. The results are compared with deposit scale 2D DC resistivity inversions and the regional scale time constant analysis from the Spectrem AEM data.

A central (30×30 km) block of the 2010 ZTEM survey was chosen to be inverted in 3D. The ZTEM survey flight path geometry and the inverted subset of the full survey is shown in Figure 12. Initially the lower frequency (30, 45, 90, and 180 Hz) data were inverted on a coarse $83 \times 83 \times 52$ mesh. The cell dimensions in the central core region were 400×400 m in the x- and y-directions and started at 25 m in the vertical direction before expanding with depth. Although the 400 m cell sizes in the horizontal dimensions may not be adequate to recover the fine deposit scale features, the coarse discretization allowed the full 30×30 km block to be inverted in around two hours on a single dual Xeon X5660 processor. The coarse inversion result at an elevation of 100 m AMSL can be seen in Figure 13. Next, the modeling domain was split into three tiles, which were then run in parallel. Each tile contained 314, 112, and 52 cells in the x-, y-, and z-directions. Each cell in the coarse inversion was subdivided into 16 cells by using 100×100 m horizontal cells in the tiled inversion. The tiles took around 36 hours to run on dual Xeon X5660 processors. The tiles were merged using the approximate sensitivity procedure outlined in the synthetic inversion section of this paper. Figure 14 shows the fine scale model inverted by merging the three tiles together. The regional scale inversion result shows good agreement with the Spectrem AEM z-component time constant tau plot shown in Figure 15. On the deposit scale, Figure 16 shows a comparison between the ZTEM

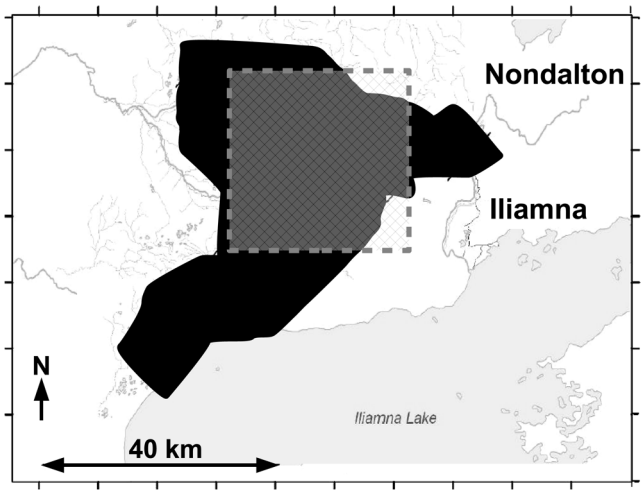

Figure 12. ZTEM survey block of 6504 line-km flown in the summer of 2010. The 30×30 km block inverted in this section is outlined by the dashed square.

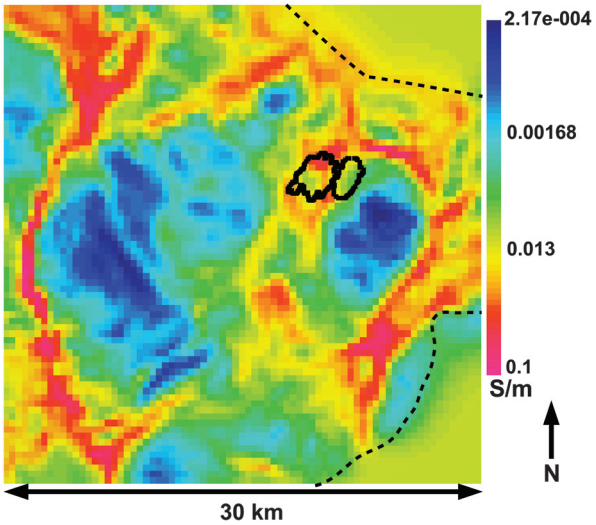

Figure 13. Coarse ZTEM inversion result at a depth of 100 m (AMSL). The edge of the data coverage is shown by the dashed lines and the deposit is outlined by the solid lines.

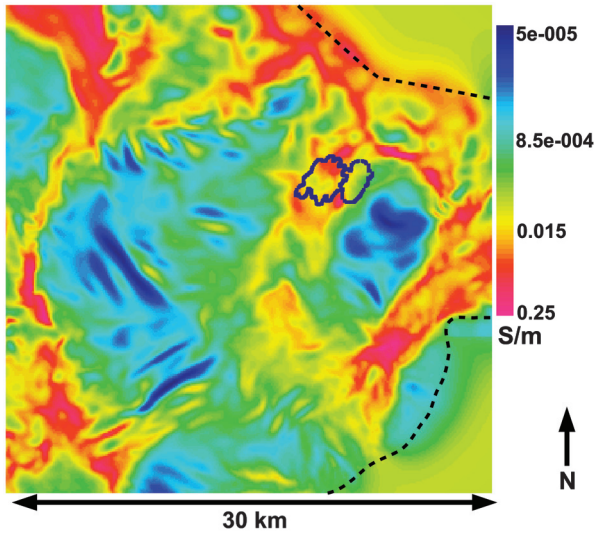

Figure 14. Fine ZTEM inversion result at an elevation of 100 m (AMSL) formed by merging the three subproblem tiles together. The edge of the data coverage is shown by the dashed lines and the deposit is outlined by the solid lines.

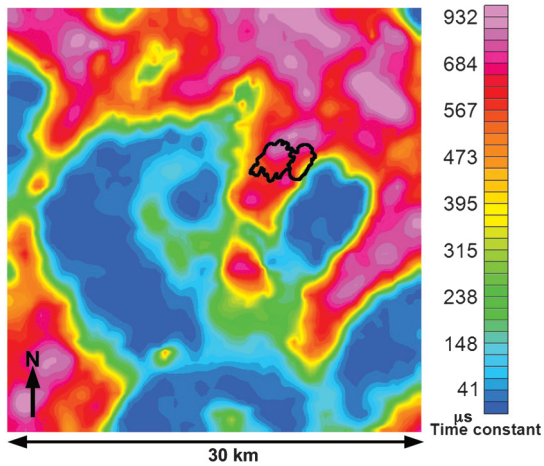

Figure 15. Spectrem z-component time constant plot over the same region as the inverted ZTEM data from Klinkert et al. (2009). The outline of the deposit is outlined in black.

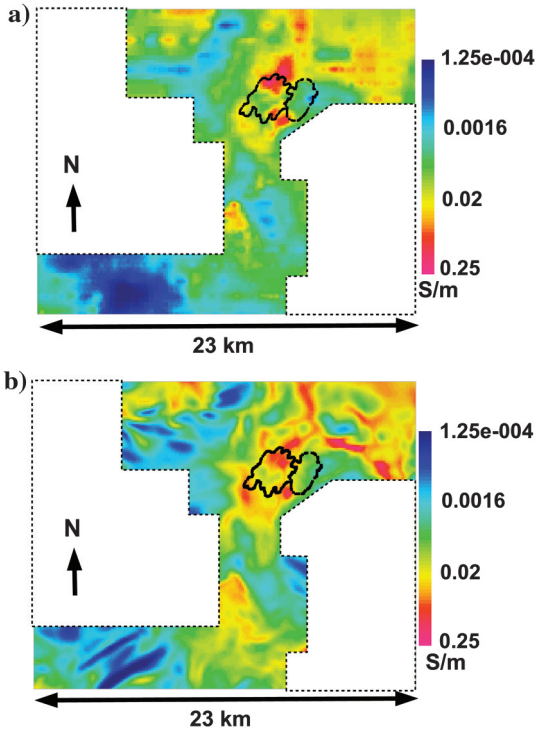

Figure 16. Inversion results both shown on the same color scale and elevation of 100 m (AMSL). The deposit is outlined in black. (a) 3D DC Inversion (b) ZTEM inversion shown over the same area as the DC resistivity data coverage. To first order, the results seem consistent.

inversion result and a 3D DC inversion performed using a UBC OcTree DC/IP inversion code. To first order, there seems to be good agreement on the regional scale between the Spectrem and ZTEM and on the deposit scale between the ZTEM data and the DC data.

CONCLUSIONS

In this paper, we show how data from a large-scale ZTEM survey can be inverted by using a strategy involving coarse and fine meshes as well as a domain decomposition that splits the computational domain into smaller manageable subproblems, which can be solved in parallel. We have presented an algorithm and practical workflow procedure to carry this out. The methodology shows promising results on a synthetic example from the Noranda mining camp. The procedure was also used to invert a 30×30 km block of the 6504 line-km ZTEM survey collected over the Pebble deposit in the Bristol Bay region of southwest Alaska. The results over this world class porphyry deposit are encouraging and seem consistent with the other electromagnetic data sets.

ACKNOWLEDGMENTS

The authors would like to thank Mike Allard from Xstrata Zinc, Louis Martin from Xstrata Copper, and Gervais Peron from Mira Geosciences for providing the geologic shape model for the Noranda mining camp that has formed the basis of our synthetic example. The authors would also like to thank Pascal Paré of Anglo American and the Pebble Partnership for providing the data and support on the project.

REFERENCES

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, doi: 10.1071/EG10003.

Farquharson, C. G., D. W. Oldenburg, E. Haber, and R. Shekhtman, 2002, An algorithm for the three-dimensional inversion of magnetotelluric data: 72nd Annual International Meeting, SEG, Expanded Abstracts, 649–652.

Gibson, H., and A. G. Galley, 2007, Mineral deposits of Canada: A synthesis of major deposit-types, district metallogeny, the evolution of geological provinces: Mineral Deposits Division, **5**, 533–552.

Gribenko, A., A. M. Green, M. Cuma, and M. S. Zhdanov, 2010, Efficient 3D inversion of MT data using integral equations method and the receiver footprint approach: Application to the large-scale inversion of the Earthscope MT data: 80th Annual International Meeting, SEG, Expanded Abstracts.

Haber, E., U. M. Ascher, D. A. Aruliah, and D. W. Oldenburg, 2000, Fast simulation of 3D electromagnetic problems using potentials: Journal of Computational Physics, **163**, 150–171, doi: 10.1006/jcph.2000.6545.

Holtham, E., and D. Oldenburg, 2010a, Three-dimensional inversion of MT and ZTEM data: 80th Annual International Meeting, SEG, Expanded Abstracts, 655–659.

Holtham, E., and D. W. Oldenburg, 2010b, Three-dimensional inversion of ZTEM data: Geophysical Journal International, **182**, 168–182, doi: 10.1111/j.1365-246X.2010.04634.x.

Klinkert, P., S. Letts, J. Smit, and L. Polome, 2009, Spectrem survey of the pebble area (Alaska): Pebble Limited Partnership, internal report by Spectrem Air Ltd, 75.

Leggatt, P. B., P. S. Klinkert, and T. B. Hage, 2000, The spectrem airborne electromagnetic system — Further developments: Geophysics, **65**, 1976–1982, doi: 10.1190/1.1444881.

Lo, B., and M. Zang, 2008, Numerical modeling of ZTEM (airborne AFMAG) responses to guide exploration strategies: 78th Annual International Meeting, SEG, Expanded Abstracts, **27**, 1098–1102.

Paré, P., and J. M. Legault, 2010, Ground IP-resistivity, and airborne spectrem and helicopter ZTEM survey results over pebble copper-moly-gold porphyry deposit, Alaska: 80th Annual International Meeting, SEG, Expanded Abstracts, 1734–1738.

Rebagliati, C. M., J. Lang, E. Titley, D. Rennie, L. Melis, D. Barratt, and S. Hodgson, 2009, Technical report on the 2008 program and update on mineral resources and metallurgy, pebble copper-goldmolybdenum project, Iliamna Lake area, southwestern Alaska, U.S.A.: Technical report, For: Northern Dynasty Minerals Ltd.

Shalf, J., 2007, The new landscape of parallel computer architecture: Journal of Physics: Conference Series, **78**, doi: 10.1088/1742-6596/78/1/012066.

Ward, S. H., 1959, AFMAG — Airborne and ground: Geophysics, **24**, 761–787, doi: 10.1190/1.1438657.

Ward, S. H., J. O'Donnell, R. Rivera, G. H. Ware, and D. C. Fraser, 1966, AFMAG — Applications and limitations: Geophysics, **31**, 576–605, doi: 10.1190/1.1439795.