Massive parallelization of 3D electromagnetic inversion using local meshes

Dikun Yang[1], Douglas Oldenburg[1], Eldad Haber[1]
[1]Geophysical Inversion Facility, University of British Columbia

SUMMARY

Despite recent advances of algorithms and computing techniques, speed remains one of the major concerns for practical 3D EM modeling and inversion. To reduce run times, we develop a generic parallelization scheme for modeling EM fields in the diffusive regime using local meshes that have fine-scale discretization near the locations of transmitter and receiver. If the sizes of transmitter and receiver are relatively small, a local mesh that exclusively serves one transmitter-receiver pair can be much smaller and easier to solve than a universal (global) mesh that is good for every transmitter and receiver. Local meshes thus allow a large EM problem to be decomposed to many small highly independent problems, and massive parallelization can be effectively applied. Two control-source EM surveys are considered in this paper: an airborne EM survey and a large ground loop EM survey, for which different types of local mesh are designed according to the configurations of the systems. We show that the local mesh method is capable of providing adequate accuracy in forward modeling and in the sensitivity computation. A synthetic example of inversion shows nearly linear speedup as the number of processors increases. We illustrate the technique with an airborne EM field data example.

Keywords: electromagnetic, forward modeling, inversion, parallelization, local mesh

INTRODUCTION

Rigorous 3D EM modeling (forward modeling and inversion) usually involves a large number of cells in a 3D discretized domain and it is very time-consuming for most realistic applications. The difficulties are exacerbated when many sources are used in the survey. Numerous efforts have been made to accelerate the numerics by using improved solvers and sophisticated spatial/temporal discretizations. With the rapid development of computer technology, especially the emerging GPU computing, massive parallelization offers another way to boost the efficiency of 3D EM modeling. Alumbaugh et al. (1996), Newman & Alumbaugh (1997) and Commer & Newman (2004) parallelized their finite difference method by assigning each processor a sub-domain of the whole problem. Their method showed great acceleration when a relatively small number of processors were used, but it suffered from excessive inter-processor communication due to the cell-level parallel design when more processors were added. Some other domain decomposition methods, depending on the number of sub-domains, may not be able to take full advantage of modern parallel computing devices with hundreds of thousands of processors. In order to tackle the multi-source problem, Oldenburg et al. (2013) used a direct solver, MUMPS, to factor the Maxwell matrix. The work for factorization and storage of the factored matrices was distributed over a computing cluster. The use of a direct solver was shown to have great benefit but for very large problems the amount of memory required, and the communication time between processors, limited the size of problem that could be tackled efficiently.

Modern high performance computing systems usually have a large number of processors, each with reasonable memory, and each individual processor/node is powerful enough to carry out the modeling for a small or mid-sized problem. It is our intention to design the parallelization so that as many processors as possible can be utilized and have the processors work independently or with minimal communication between parallel workers. This requires a parallel design at higher level. The important ingredient is that for EM fields in the quasi-static approximation, the spatial discretization needs to be fine only near the location of the transmitter and receiver. Due to the diffusive nature of the fields, the discretization far away from the transmitter and receiver can be coarse. This motivates the design of a local mesh that is specifically discretized to serve one particular source-receiver pair (or only receiver in natural source EM), in contrast to a global mesh traditionally used in EM modeling that has fine cells everywhere for all the distributed sources and receivers. Such a local mesh is much smaller and faster than a global mesh, so although this may lead to a number of local meshes, the two essential ingredients of inversion, forward responses and sensitivity, can be more efficiently computed concurrently on many local meshes that are independent to each other. Once the forward solutions and the sensitivities are available on the local meshes, the inversion is carried out on the fine global mesh with regular techniques of model regularization. Using more processors simply reduces the number of local mesh problems each processor needs to solve and adds very little overhead to the run time compared to the parallelization at cell level.

Two scenarios of commonly-used transmitter-receiver configuration are investigated here. The first is airborne time-domain EM, in which the transmitter and receiver are magnetic dipoles and almost coincident. This results in a straightforward local mesh with small cells near the sounding location and larger cells expanding outward. Then the similar idea is applied to a large ground loop time-domain EM survey. To make small local meshes for large loop up to a few square kilometers, we break down the transmitter wire to small segments and then pair the transmitter segments with receivers. The forward modeled responses and the sensitivity can be obtained by superposition. The efficiency of the local mesh method is demonstrated by both synthetic and field data examples.

LOCAL MESH IN AIRBORNE EM

An airborne EM survey consists of many distributed soundings, one for each transmitter-receiver pair, over a large area. An airborne sounding can be modeled with a magnetic dipole, or small current loop, source and a receiver that measures the EM fields at the same or nearby location. Because of the large numbers of model parameters and sources for the entire survey, airborne EM is one of the most difficult EM modeling problems in exploration geophysics. Here we retain the finite volume technique (Oldenburg et al., 2013) for solving Maxwell's equations. The parameters required to generate the local mesh (cell sizes, expansion factor, truncation distance, etc.) need to be determined but a semi-automated method is used to ensure that fields are sufficiently accurate and that the number of cells is kept small. Figure 1 shows an example of a local mesh for the sounding at the center of the survey area. The conductivities of coarse cells in Figure 1b are calculated using volume-weighted averaging.

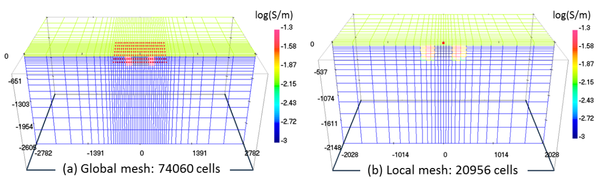

Figure 1. Design of a local mesh for airborne EM. (a) Conductivity model on the global mesh; the grid of red dots indicates sounding locations. (b) Local mesh for the sounding at the mesh center indicated by the red dot.

The vertical component of dB/dt data is forward modeled on both global and local meshes (Figure 2). While they provide essentially the same numerical values, the modeling on the local mesh (17s) is much more efficient than that on the global mesh (179s for one source). This difference will increase if the size of the global mesh increases.

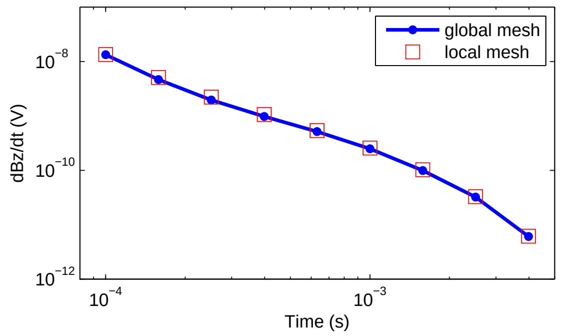

Figure 2. Forward modeled data using the global and local mesh.

An inversion also requires the computation of sensitivity. The sensitivity is mostly smooth in space and has fine-scale variations near the transmitter and receiver, so the same local mesh in forward modeling can also be used for the sensitivity. We use an interpolation to convert the sensitivity on a local mesh to the global mesh,

$$\mathbf{J}_g = \mathbf{J}_l \cdot \mathbf{V}_l^{-1} \cdot \mathbf{Q} \cdot \mathbf{V}_g, \tag{1}$$

where $\mathbf{J}_g$ is the global sensitivity, $\mathbf{J}_l$ is the local sensitivity, $\mathbf{V}_l$ is a diagonal matrix of the local cell sizes, $\mathbf{Q}$ is a sparse matrix mapping the sensitivity function from the local to the global mesh and $\mathbf{V}_g$ is a diagonal matrix of the global cell sizes. $\mathbf{J}_l$ is dense but small, so once the sensitivity is *compressively* computed and stored on a local mesh, the most computationally expensive operation, $\mathbf{J}_g$ times a vector, can be rapidly carried out by multiplying the right-hand-side matrices in equation (1) for different soundings in parallel. Parallel workers compute and store the forward responses and sensitivity locally, so they do not need to exchange any information. Communication between the host and parallel workers occurs only for passing vectors that resulted from intensive computations. Figure 3 compares the global sensitivity computed on the global mesh and interpolated from a local mesh. The difference is minimal and the interpolated sensitivity is adequate for the inversion.

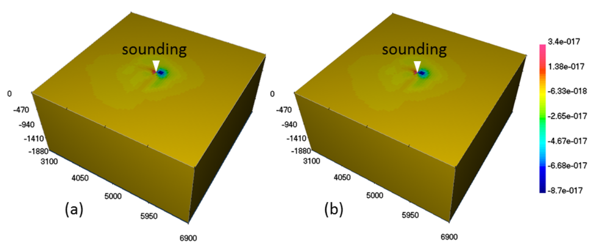

Figure 3. Sensitivity of one time channel for a random conductivity model. (a) Directly computed on the global mesh. (a) Computed on a local mesh and then interpolated to the global mesh.

LOCAL MESH IN GROUND LOOP EM

The same ideas for local meshes in airborne EM are also valid for ground loop EM surveys. However, the large transmitter loop employed in a ground EM survey may result in a local mesh that still has too many fine cells. Therefore, we break the transmitter loop into many line segments and compute the forward response at one receiver location by summing the responses due to the individual segments. Figure 4a shows a local mesh finely discretized near the receiver and one of the segments for the 1km-by-1km transmitter. The entire loop is decomposed into 84 50m-long segments. We note the mesh in Figure 4a is much smaller than the global mesh with fine cells everywhere under the large loop and around all the receivers. In this example, supposing 100 receivers in the survey, there are 8400 source-receiver pairs (local meshes), which are good for massive parallelization. Figure 4b provides a comparison of forward modeled responses between directly using the global mesh and the sum of local mesh results.

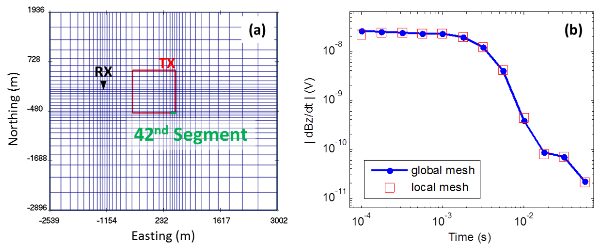

Figure 4. Local mesh method for ground EM. (a) A local mesh for the 42^{nd} transmitter segment (green) and a receiver; (b) Forward modeled responses on the global mesh and on the local meshes.

The sensitivities of a receiver due to all the segments are computed on the same local meshes as forward modeling. The global sensitivity, while not explicitly computed, is available for matrix-vector multiplication through

$$\mathbf{J}_g = (\sum_{i=1}^{N} \mathbf{J}_{li} \cdot \mathbf{V}_{li}^{-1} \cdot \mathbf{Q}_i) \cdot \mathbf{V}_g, \qquad (2)$$

where N is the number of transmitter segments, $\mathbf{J}_{li}$ is the local sensitivity for the i^{th} segment, $\mathbf{V}_{li}$ is a diagonal matrix of the cell sizes in the i^{th} local mesh and $\mathbf{Q}_i$ is a 3D interpolation matrix for the i^{th} local mesh. The global sensitivity of one dBz/dt channel for a random conductivity model computed on local meshes is compared to the sensitivity directly from the global mesh in Figure 5.

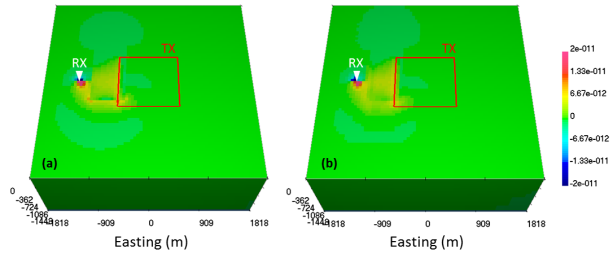

Figure 5. Sensitivity of one dBz/dt time channel to a random model in a ground EM survey. (a) The sensitivity directly computed on the global mesh; (b) The sensitivity interpolated from local meshes and computed using equation (2).

SYNTHETIC INVERSION USING LOCAL MESH

We now use a synthetic model to illustrate an inversion using local meshes and also test the scalability of our parallel design. The synthetic model consists of two conductive blocks buried in a uniform half-space as shown in Figure 6a. An airborne TEM survey covering an area measured 1.2 by 3.6 km with soundings spaced 100m apart was simulated on a global mesh using the forward modeling code in Oldenburg et al. (2013). The data are 7 time channels of the vertical component of dB/dt at 481 soundings. The global mesh used for generating synthetic data and as base mesh for local mesh inversion has 155820 cells. One forward modeling on this global mesh takes about 30 minutes on 24 processors, so a complete inversion would need a few hundred hours. In the first run of synthetic inversion, we use 12 processors with 481 local meshes; then add 12 more processors in the second run. The recovered model is shown in Figure 6b.

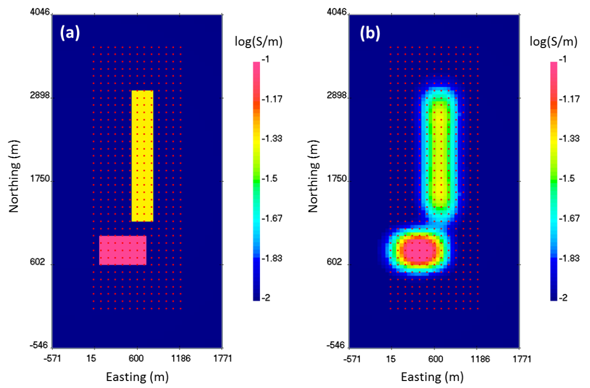

Figure 6. True (a) and recovered (b) model at depth of 150m for the synthetic inversion; red dots indicate the locations of airborne soundings.

In Figure 7 we compare the run time at every iteration in the first and second run . The second inversion, with 24 processors, gains almost 2 times speed-up and finishes in 7.5 hours as a result of doubled number of processors.

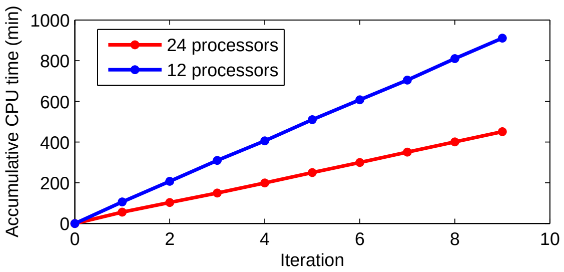

Figure 7. Run time of local mesh inversions on 12 processors and 24 processors.

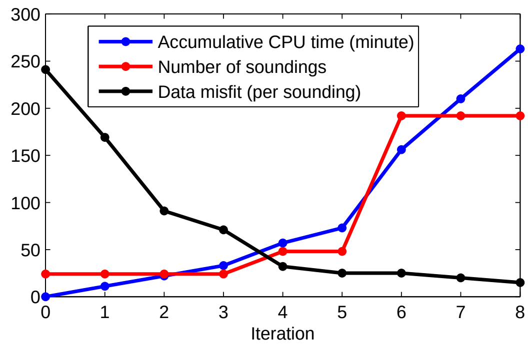

Figure 8. Field data inversion using local mesh and adaptive soundings.

FIELD DATA INVERSION USING LOCAL MESHES

We now invert an airborne field data set that has 14362 soundings covering a total area of 7.3 km^2. Resolution analysis showed the finest scale was about 50m. The finely discretized global mesh consisted of 433520 cells. Inversion of the entire survey on the global mesh using the code in Oldenburg et al. (2013) would be extremely time-consuming and we therefore implement the local mesh method. However, 14362 local meshes taxes our capabilities because we are limited to 24 processors and that means each processor must store, and work with, about 600 local meshes. To overcome this bottleneck, we introduce a further alteration. The magnetic fields observed on the surface or in the air are mostly smooth functions in space and most EM surveys are redundant due to oversampling. We have developed an adaptive random downsampling scheme that matches the number of data to the degree of regularization in inversion (Yang et al., 2012). This scheme ensures, based on physics, that only the necessary amount of data are modeled at different stages of inversion, regardless of the total number of data produced by the instruments. At each iteration only a restricted number of randomly chosen soundings are used to generate an update. Figure 8 summarizes three key parameters of the inversion on 24 processors: the accumulative run time, the number of soundings randomly selected at each iteration and the data misfit. For early iterations only 24 soundings were needed to build up the large-scale conductivity distribution; more small-scale features were built up by adding more soundings; the number of soundings used in the final iteration is only 192. During the inversion, 744 local meshes were constructed and solved for the forward responses and sensitivity and most importantly, the whole process only took 4.3 hours. We anticipate more speedup if more processors were available.

CONCLUSIONS

We propose a method to decompose an EM survey to many independent source-receiver pairs, each of which is solved on a mesh that is finely discretized only near the dipole source and receiver. This type of mesh, referred as a *local mesh*, exclusively serves one particular source-receiver pair and is much smaller than a global mesh used to simulate data for every transmitter and receiver in a survey. Our decomposition at the sounding-level is ideal for massive parallelization because there are usually hundreds or thousands of source-receiver pairs in an EM survey and local meshes are highly independent to each other. The transmitter and receiver configurations of airborne and ground loop surveys are investigated. We show, for both scenarios, that local mesh method can compute the forward responses and the sensitivity efficiently with adequate precision. The linear speedup with the number of processors is demonstrated with a synthetic example. The local mesh method is also applied to a field data set inversion that was traditionally considered too time-consuming to be practical.

REFERENCES

Alumbaugh, D. L., Newman, G. A., Prevost, L., & Shadid, J. N. (1996). Three-dimensional widehand electromagnetic modeling on massively parallel computers. *Radio Science*, *31*, 1-23.

Commer, M., & Newman, G. A. (2004). A parallel finite-difference approach for 3D transient electromagnetic modeling with galvanic sources. *Geophysics*, *69*, 1192-1202.

Newman, G. A., & Alumbaugh, D. L. (1997). Three-dimensional massively parallel electromagnetic inversion - I. Theory. *Geophysical Journal International*, *128*, 345-354.

Oldenburg, D. W., Haber, E., & Shekhtman, R. (2013). Three dimensional inversion of multi-source time domain electromagnetic data. *Geophysics*, in press.

Yang, D., Oldenburg, D. W., & Haber, E. (2012). Fast 3D inversion of airborne time-domain electromagnetic data using localized meshes and adaptive soundings. *Abstract GP31C-05 presented at 2012 Fall Meeting, AGU, San Francisco, Calif., 3-7 Dec.*