

Massively Parallel Modeling and Inversion of Electrical Resistivity Tomography data using PFLOTRAN

Piyooosh Jaysaval, Glenn E. Hammond, and Timothy C. Johnson

Pacific Northwest National Laboratory, 902 Battelle Blvd., Richland, WA 99352, USA

Correspondence: Piyooosh Jaysaval (piyooosh.jaysaval@pnnl.gov)

Abstract.

Electrical resistivity tomography (ERT) is a broadly accepted geophysical method for subsurface investigations. Interpretation of field ERT data usually requires the application of computationally intensive forward modeling and inversion algorithms. For large-scale ERT data, the efficiency of these algorithms depends on the robustness, accuracy, and scalability on high performance computing resources. In this regard, we present a robust and highly scalable implementation of forward modeling and inversion algorithms for ERT data. The implementation is publicly available and developed within the framework of PFLOTRAN, an open-source, state-of-the-art massively parallel subsurface flow and transport simulation code. The forward modeling is based on a finite volume discretization of the governing differential equations, and the inversion uses a Gauss–Newton optimization scheme. To evaluate the accuracy of the forward modeling, two examples are first presented by considering layered (1D) and 3D earth conductivity models. The computed numerical results show good agreement with the analytical solutions for the layered earth model and results from a well-established code for the 3D model. Inversion of ERT data, simulated for a 3D model, is then performed to demonstrate the inversion capability by recovering the conductivity of the model. To demonstrate the parallel performance of PFLOTRAN’s ERT process model and inversion capabilities, large-scale scalability tests are performed by using up to 131,072 processes on a leadership class supercomputer. These tests are performed for the two most computationally intensive steps of the ERT inversion: forward modeling and Jacobian computation. For the forward modeling, we consider models with up to 122 million degrees of freedom (DOFs) in the resulting system of linear equations, and demonstrate that the code exhibits almost linear scalability on up to 10,000 DOFs per process. On the other hand, the code shows superlinear scalability for the Jacobian computation, mainly because all computations are fairly evenly distributed over each process with no parallel communication.

20 1 Introduction

Direct current electrical resistivity tomography (ERT) is one of the key geophysical methods for shallow-subsurface investigations, having applications in areas such as groundwater (Dahlin, 2001; Johnson et al., 2012; Meyerhoff et al., 2014; Park et al., 2016; Greggio et al., 2018; Alshehri and Abdelrahman, 2021), mineral exploration (Badmus and Olatinsu, 2009; Bery et al., 2012; Uhlemann et al., 2018; Martínez et al., 2019), environmental monitoring and remediation (Rosales et al., 2012; Rucker et al., 2013; Gabarrón et al., 2020; Rockhold et al., 2020; Kessouri et al., 2022), and engineering problems (Dahlin,

1996; Rizzo et al., 2004; Lysdahl et al., 2017). It can also be used for large-scale deep-subsurface investigations, e.g., for geothermal systems, active volcano imaging, and tectonic studies (Storz et al., 2000; Caputo et al., 2003; Johnson et al., 2010; Richards et al., 2010). Additional applications can also be found in detailed reviews by Slater (2007), Revil et al. (2012), Loke et al. (2013), Singha et al. (2022), and the references therein. In recent years, availability of multielectrode and multichannel
30 instrumentations permitted acquisition of massive amounts of ERT data consisting of tens of thousands or even hundreds of thousands of observations. Efficient inversion and interpretation of such massive datasets requires fast, accurate, and highly scalable forward modeling and inversion algorithms to simulate and invert ERT data in arbitrary 3D conductivity structures. This paper presents an open-source implementation of such ERT modeling and inversion algorithms.

ERT data simulation is performed by solving the electrostatic Poisson equation. Numerical methods are needed to solve
35 the governing equation in arbitrary 3D conductivity structures. The most used numerical methods for 3D ERT modeling are the finite-difference (FD) method (Dey and Morrison, 1979; Spitzer, 1995; Penz et al., 2013), finite-element (FE) method (Coggon, 1971; Li and Spitzer, 2002; Rucker et al., 2006; Blome et al., 2009; Johnson et al., 2010; Ren et al., 2018), and integral equation method (Lee, 1975; Schulz, 1985; Méndez-Delgado et al., 1999). Among these methods, the FD and FE methods are the attractive choices for highly heterogeneous distributions of electrical conductivity in the subsurface. However,
40 application of the FD method is mostly limited to simple model geometries due to its requirement of Cartesian grids (Rucker et al., 2006). On the contrary, the FE method has proven to be effective in accounting for complex geometries, specifically by using unstructured meshes (Rucker et al., 2006; Blome et al., 2009; Johnson et al., 2010; Ren et al., 2018).

The finite-volume (FV) method can also be used effectively to account for complex geometries (Jahandari and Farquharson, 2014). The FV method is usually seen relatively close to the FD method as it inherits the simplicity similar to the FD method
45 in its implementation. Unlike the FD method which discretizes the differential form of the governing equation, the FV method, however, directly discretizes its integral form. Nevertheless, the FV method has rarely been considered for ERT simulation problems. In our literature search, we came across only Cockett et al. (2015) which implements the method for the ERT modeling.

On the other hand, ERT data inversion is a non-linear optimization problem to minimize a cost function that represents a
50 measure of the difference between observed and simulated ERT data. For large-scale 3D ERT data, the inversion is often performed iteratively by linearizing the optimization problem at each iteration, calculating a model update, updating a conductivity model, and subsequently simulating ERT data for the updated model to examine the new cost function. The process continues until the cost function reduces to a predefined tolerance level. The model update is typically calculated using a gradient-based local optimization method, e.g., the steepest-descent, conjugate gradient, or Gauss–Newton method (Park and Van, 1991; Ellis and Oldenburg, 1994; Zhdanov and Keller, 1994; Zhang et al., 1995; Günther et al., 2006; Johnson et al., 2010). The Gauss–
55 Newton method has usually been preferred due to its high convergence rate. Indeed, the fast convergence occurs at the cost of intensive computations because the method requires building the Jacobian (or sensitivity) and/or Hessian matrices at each inversion iteration (Nocedal and Wright, 2006).

For large-scale 3D ERT problems with tens of millions of degrees of freedom (DOFs), forward modeling and inversion
60 are computationally demanding jobs and may need hours, if not days, of computing time if the underlying algorithms are not

properly implemented to scale well on high performance computing (HPC) resources or supercomputers. For such datasets, it is computationally infeasible to use desktop based ERT data processing software without exploiting HPC resources. Recently, a few open-source 3D ERT modeling and inversion codes have been developed to handle moderately sized ERT problems composed of up to a few million DOFs, e.g., ResIPy (Blanchy et al., 2020), SimPEG (Cockett et al., 2015), and pyGIMLI (Rücker et al., 2017). However, these codes still lack full-scale parallel implementations to efficiently distribute computations over a supercomputer. On the other hand, another open-source code, E4D, is highly scalable on HPC resources (Johnson et al., 2010; Johnson and Wellman, 2015). It can be run on hundreds or even thousands of processes (or cores) on a supercomputer (Johnson et al., 2010). But despite this, the maximum number of processes that it can exploit is limited to the number of electrodes in a survey.

In this paper, we present an implementation of massively parallel 3D ERT modeling and inversion algorithms. The forward modeling is based on the FV method and inversion employs the Gauss–Newton method. The algorithms are implemented within PFLOTRAN (www.pflotran.org), an open source, massively parallel code that leverages parallel data structures and solvers within PETSc (Portable, Extensible Toolkit for Scientific Computation) (Balay et al., 2021) to simulate subsurface flow and transport processes on supercomputers (Hammond et al., 2012, 2014). Although we are presenting only ERT modeling and inversion capabilities here, the main motivation of implementing these capabilities in PFLOTRAN is to perform coupled modeling and inversion of ERT, flow, and/or transport problems in the future.

The paper is structured as follows: we first outline a detailed background theory on the ERT forward modeling including the governing Poisson equation and FV method to solve it numerically. We then describe the ERT inversion by defining it as a nonlinear optimization problem, and the Gauss–Newton method to find a minimizer of the problem along with details on computing the Jacobian matrix. Thereafter, we discuss the parallel implementations of the modeling and inversion algorithms. We subsequently benchmark numerical results computed using PFLOTRAN against analytical solutions for a layered (1D) earth model and numerical solutions for a 3D earth model. Next, a 3D inversion result is presented to illustrate the inversion capability. Scalability tests are then performed to demonstrate the parallel performance of the PFLOTRAN ERT process model, before discussing the results and drawing our final concluding remarks.

2 Forward ERT modeling

Forward modeling is a way of simulating ERT data for any given 1D, 2D, or 3D electrical conductivity models of the earth. For 1D conductivity models, the electrical potentials generated by an induced point current source can be obtained by using analytical or semi-analytical methods (Das, 1995; Pervago et al., 2006). However, for multi-dimensional heterogeneous conductivity models with complex geometries, one must use numerical methods to compute the electrical potential. This section describes the numerical computations of the electrical potentials in a 3D medium and their superposition to produce simulated ERT data.

2.1 Poisson's equation

The ERT forward modeling is governed by the following electrostatic Poisson equation

$$\nabla \cdot \sigma(\mathbf{r})\nabla\phi(\mathbf{r}) = -I\delta(\mathbf{r} - \mathbf{r}_s), \quad (1)$$

95 where ϕ is the electrical potential at position \mathbf{r} for a given conductivity $\sigma(\mathbf{r})$ due to a current I injected through a point located at \mathbf{r}_s ; and δ is the Dirac delta function. For brevity, hereinafter, the dependencies on \mathbf{r} will be omitted except where necessary to show.

If side and bottom boundaries $\partial\Gamma$ of the 3D computational domain Γ are located at sufficiently far from the current injection location \mathbf{r}_s , the potential and the normal component of the current density $\sigma\frac{\partial\phi}{\partial n}$ asymptotically approach zero. On the top or
100 surface boundary $\partial\Gamma_S$, the normal current density $\sigma\frac{\partial\phi}{\partial n}$ is zero as no current flows through the earth surface along the outward normal vector. Consequently, we can impose zero Dirichlet or Neumann boundary conditions at the side boundaries and zero Neumann boundary at the surface boundary

$$\phi|_{\partial\Gamma} = 0 \quad \text{or} \quad \frac{\partial\phi}{\partial n}|_{\partial\Gamma} = 0, \quad \text{and} \quad \frac{\partial\phi}{\partial n}|_{\partial\Gamma_S} = 0. \quad (2)$$

To simulate ϕ in arbitrary 3D conductivity structures, Eq. (1) needs to be solved using numerical methods subject to the
105 boundary conditions in Eq. (2).

2.2 Finite volume method

The FV method is implemented to compute ϕ by solving Eq. (1). The 3D computational domain is first discretized into a set of control volumes known as cells (Fig. 1). These cells can be of arbitrary shape and size, but they must be bounded by planar surfaces for our implementation. The bounding discrete surfaces are known as cell faces.

110 Let us consider the governing Poisson equation (Eq. 1) and integrate it over the i^{th} cell of the domain Γ . This gives

$$\int_{V_i} \nabla \cdot \sigma \nabla \phi dV = - \int_{V_i} I \delta(\mathbf{r} - \mathbf{r}_s) dV, \quad (3)$$

where V_i is the volume of the i^{th} cell.

By applying the Gauss divergence theorem in Eq. (3) and using the translation property of the Dirac delta function, we get

$$\int_{S_i} \sigma \nabla \phi \cdot \hat{\mathbf{n}} dS = -I|_{\mathbf{r}=\mathbf{r}_s}, \quad (4)$$

115 where S_i is the surface area of the bounding faces of i^{th} cell, and dS is a differential area on the bounding surface with a unit surface normal $\hat{\mathbf{n}}$ pointing outward.

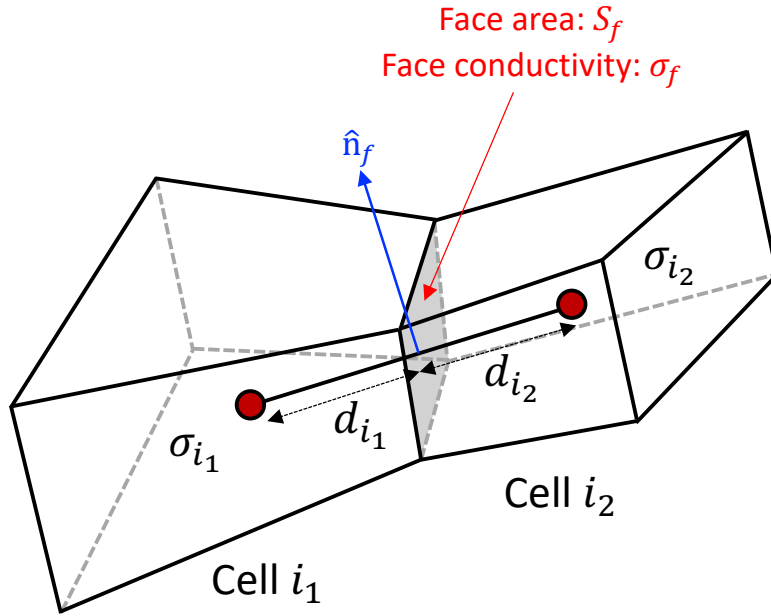


Figure 1. Two adjacent FV cells i_1 and i_2 having conductivity σ_{i_1} and σ_{i_2} , respectively. The distances of the common face f from the center of cells i_1 and i_2 , respectively, are d_{i_1} and d_{i_2} .

If the i^{th} cell is bounded by N_f faces, Eq. (4) can be replaced by an equivalent discrete form

$$\sum_{f=1}^{N_{f,i}} \int_{S_f} \sigma \nabla \phi \cdot \hat{\mathbf{n}} dS = \sum_{f=1}^{N_{f,i}} \sigma_f (\nabla \phi)_f \cdot \hat{\mathbf{n}}_f S_f = -I|_{\mathbf{r}=\mathbf{r}_s}, \quad (5)$$

where f represents faces and S_f is the area of the bounding face f .

120 Using a two-point flux approximation for each face term in Eq. (5) by considering two adjacent cells i_1 and i_2 shown in Fig. 1, we get

$$\sum_{f=1}^{N_{f,i}} \sigma_f (\nabla \phi)_f S_{nf} = -I|_{\mathbf{r}=\mathbf{r}_s}, \quad (6)$$

where S_{nf} is the area of the common face f projected onto the plane normal to the vector connecting centers of cells i_1 and i_2 .

125 The conductivity at the face, σ_f , can be obtained by averaging the conductivity of cells i_1 and i_2 . We use the following harmonic distance weighted scheme to calculate σ_f

$$\sigma_f = \frac{\sigma_{i_1}\sigma_{i_2}(d_{i_1} + d_{i_2})}{\sigma_{i_1}d_{i_2} + \sigma_{i_2}d_{i_1}}, \quad (7)$$

where σ_{i_1} and σ_{i_2} are the conductivities of cells i_1 and i_2 , and d_{i_1} and d_{i_2} are the distances of the projected face from the center of cells i_1 and i_2 , respectively (Fig. 1).

130 The value of the flux at the face f , i.e., $(\nabla\phi)_f$, is obtained by using the approximation

$$(\nabla\phi)_f = \frac{\phi_{i_2} - \phi_{i_1}}{d_{i_1} + d_{i_2}}. \quad (8)$$

Substituting Eqs. (7) and (8) in Eq. (6), and assuming that the computational domain Γ is subdivided into N_m FV cells, i.e., $\Gamma = \cup_{i=1}^{N_m} \Gamma^i$ where $\Gamma^i \subset \mathbb{R}^3$ is the i^{th} cell domain, we have

$$\sum_{i=1}^{N_m} A_{i,k} \phi_k = s_i, \quad (9)$$

135 where $A_{i,k}$ is the local coefficient matrix and s_i is the source vector for the i^{th} cell.

The assembly of the local coefficient matrix into a global system, upon carrying out the summation in Eq. (9), results into the system of linear equations

$$\mathbf{A}\Phi = \mathbf{s}, \quad (10)$$

where $\mathbf{A} \in \mathbb{R}^{N_m \times N_m}$ is the symmetric system matrix resulting from the FV discretization of Eq. (1), $\Phi \in \mathbb{R}^{N_m}$ is a vector
 140 containing the unknown electric potential for all cells, and $\mathbf{s} \in \mathbb{R}^{N_m}$ is the source vector resulting from the right-hand side of Eq. (1). To solve this linear system, we utilize various options of iterative solvers available from PETSc. Each electrode in an ERT survey results in a different vector \mathbf{s} . Thus, for N_e electrodes, we need to solve Eq. (10) for N_e times with a different right-hand side. The solutions provide electrical potential at each cell center of the discrete computational domain and for each current injection location. Finally, the potential distribution for any combination of source/sink electrodes can be obtained by
 145 subtracting the potential associated with the sink electrodes from the source electrodes potential (Johnson et al., 2010).

3 Inversion

ERT Inversion is a procedure of estimating electrical conductivity of the earth from a set of observed ERT data and/or some imposed constraints. The ERT inverse problem is a non-linear optimization problem. For large-scale 3D ERT data, the inversion is often performed iteratively by linearizing the problem at each iteration. This section details our approach of performing
 150 inversion of ERT data.

3.1 Optimization problem

We pose the inverse problem as a non-linear minimization problem of finding a conductivity model

$$\mathbf{m}_{\text{inv}} = \arg \min_{\mathbf{m} \in \mathbf{M}} \psi(\mathbf{m}), \quad (11)$$

where $\mathbf{m} = (m_1, m_2, \dots, m_{N_m})^T$ is an unknown model parameters vector in the model set \mathbf{M} . Since the conductivity within
 155 subsurface can vary over several order of magnitudes, the inversion operations are performed for logarithmic conductivity, i.e.,
 we assume $\mathbf{m} = \ln \sigma$. This also enforces positivity of conductivity values. The quantity ψ is the cost function, defined as

$$\psi(\mathbf{m}) = \psi_d(\mathbf{m}) + \beta \psi_m(\mathbf{m}), \quad (12)$$

where ψ_d is the data cost function that represents a measure of the misfit between the observed d^{obs} and synthetic d^{syn}
 ERT data; ψ_m is the regularization cost function representing a measure of the variation of the model parameters, and β is a
 160 regularization parameter.

A least-squares method is used to minimize ψ (Nocedal and Wright, 2006). Thus, ψ_d and ψ_m can be expressed as

$$\psi_d = [\mathbf{d}^{\text{obs}} - \mathbf{d}^{\text{syn}}(\mathbf{m})]^T \mathbf{W}_d^T \mathbf{W}_d [\mathbf{d}^{\text{obs}} - \mathbf{d}^{\text{syn}}(\mathbf{m})], \quad (13)$$

$$\psi_m = (\mathbf{m} - \mathbf{m}_{\text{ref}})^T \mathbf{W}_m^T \mathbf{W}_m (\mathbf{m} - \mathbf{m}_{\text{ref}}), \quad (14)$$

where vectors \mathbf{d}^{obs} and \mathbf{d}^{syn} are N_d dimensional vectors containing the observed and synthetic ERT data. $\mathbf{W}_d \in \mathbb{R}^{N_d \times N_d}$
 165 is the data weighting matrix, usually a diagonal matrix, whose elements are estimated based on the standard deviation of
 the noise; $\mathbf{W}_m \in \mathbb{R}^{N_m \times N_m}$ is the regularization matrix, and \mathbf{m}_{ref} is a N_m dimensional vector containing the logarithmic of
 reference conductivity model parameters.

3.2 Gauss-Newton method

To find a minimizer \mathbf{m}_{inv} of the cost function in Eq. (12), we use a Gauss-Newton minimization approach. This is an iterative
 170 approach where after linearizing $\psi(\mathbf{m}_k)$ in the vicinity of a model \mathbf{m}_k for a small model perturbation $\delta \mathbf{m}_k$ at the k^{th} iteration,
 we obtain a set of linear equations or the normal equation

$$\mathcal{H}_k \delta \mathbf{m}_k = -\mathbf{g}_k, \quad (15)$$

where the Gauss-Newton Hessian $\mathcal{H}_k \in \mathbb{R}^{N_m \times N_m}$ matrix is given by

$$\mathbf{H}_k = \Re\{\mathcal{J}_k^T \mathbf{W}_d^T \mathbf{W}_d \mathcal{J}_k\} + \beta \mathbf{W}_m^T \mathbf{W}_m, \quad (16)$$

175 and the gradient vector by

$$\mathbf{g}_k = -\Re[\mathcal{J}_k^T \mathbf{W}_d^T \mathbf{W}_d \{\mathbf{d}^{\text{obs}} - \mathbf{d}^{\text{syn}}\}] + \beta \mathbf{W}_m^T \mathbf{W}_m \{\mathbf{m}_k - \mathbf{m}_{\text{ref}}\}. \quad (17)$$

The matrix \mathcal{J}_k is the Jacobian matrix also known as the sensitivity or the Fréchet derivative matrix. Following Johnson et al. (2010), we implement a parallel conjugate-gradient least-square (CGLS) solver (Hestenes and Stiefel, 1952) to solve the normal Eq. (15). The normal equation is first reformulated into an equivalent least-squares problem and then solved for
 180 the model update $\delta \mathbf{m}_k$ using only the Jacobian matrix \mathcal{J}_k without explicitly forming the Hessian matrix \mathbf{H}_k . As $N_m \gg N_d$ for typical ERT surveys, this approach avoids massive computational cost associated with the explicit formation of \mathbf{H}_k . We therefore only need to compute \mathcal{J}_k at each iteration, which is also a computationally intensive step. The effective computation of \mathcal{J}_k is explained in the next subsection 3.3.

The solution of Eq. (15) gives a model update vector $\delta \mathbf{m}_k$ at the k^{th} iteration such that a new model

$$185 \quad \mathbf{m}_{k+1} := \mathbf{m}_k + \alpha \delta \mathbf{m}_k \quad (18)$$

decreases the cost function ψ , i.e., $\psi(\mathbf{m}_{k+1}) < \psi(\mathbf{m}_k)$. There is a common practice to find step-length α using a line search method (Nocedal and Wright, 2006); however, from our experience, a full step-length, i.e., $\alpha = 1$, works well for the ERT data inversion.

In our implementation, the Gauss-Newton iteration usually starts with an average apparent conductivity model as an initial
 190 model and continues until the data misfit drops below a predefined tolerance level (defined by a chi-squared data misfit $\chi^2 = \phi_d/N_d \leq 1$) or it reaches a predefined maximum number of iterations. Furthermore, a predefined input conductivity model can also be used as an initial model in case it can be estimated from the existing geological and/or geophysical information. After the inversion converges, it provides the minimizer \mathbf{m}_{inv} which represents the inverted conductivity model $\sigma_{\text{inv}} = \exp(\mathbf{m}_{\text{inv}})$.

3.3 Jacobian computation

195 Calculating \mathcal{J} at each iteration is one of the most computationally intensive steps in Gauss-Newton inversion. Therefore, an efficient \mathcal{J} drives the robustness of the developed inversion scheme and maximizes its scalability. The Jacobian matrix characterizes the change in the synthetic ERT data \mathbf{d}^{syn} relative to a change in the model parameters \mathbf{m} , and is defined as the partial derivatives of \mathbf{d}^{syn} with respect to $\mathbf{m} = \ln \sigma$, i.e.,

$$\mathcal{J}_{j,k} = \frac{\partial d_j^{\text{syn}}}{\partial m_k} = \sigma_k \frac{\partial d_j^{\text{syn}}}{\partial \sigma_k}; \quad j = 1, 2, \dots, N_d \quad \text{and} \quad k = 1, 2, \dots, N_m. \quad (19)$$

200 An adjoint state method is used to effectively compute the Jacobian matrix. Let A and B be the source and sink electrodes located at \mathbf{r}_A and \mathbf{r}_B , and M and N be the two measuring or receiving electrodes located at \mathbf{r}_M and \mathbf{r}_N . The potential difference between M and N can be obtained using the potential computed due to the source A at M and N, i.e., ϕ_{AM} and ϕ_{AN} , and the sink B at M and N, i.e., ϕ_{BM} and ϕ_{BN} . A simplified four-electrodes Wenner configuration is shown in Fig. 2. Note that, contrary to this configuration, electrodes can be placed in any other configurations, e.g., Schlumberger, pole-pole, pole-dipole, 205 dipole-dipole, and gradient arrays, among others (Dahlin and Zhou, 2004).




Figure 2. Wenner electrode configuration: four electrodes A, B, M, and N are deployed in-line with an equal spacing a between the two neighboring electrodes. Electrodes A and B act as the source and sink electrodes where current I is injected, while electrodes M and N act as the receiving electrodes measuring a potential difference $\Delta\phi$.

Using the expression derived in Appendix A (Eq. A9), an element of the Jacobian matrix is obtained by

$$\mathcal{J}_{j,k} = \frac{\partial d_j^{\text{syn}}}{\partial m_k} = (\Phi_A - \Phi_B)^T \frac{\partial \mathbf{A}}{\partial m_k} (\Phi_M - \Phi_N), \quad (20)$$

where Φ_A and Φ_B are potential distributions due to the source A and sink B, and Φ_M and Φ_N are potential distributions due to pseudo source and sink, respectively, at the receiving electrodes M and N.

210 Assuming $\Phi_S = \Phi_A - \Phi_B$ and $\Phi_R = \Phi_M - \Phi_N$, where Φ_S is the net potential due to the source and sink electrodes A and B; and Φ_R is the net potential due to source and sink assumed to be located, respectively, at the receiving electrodes M and N, this results in

$$\frac{\partial d_j^{\text{syn}}}{\partial m_k} = \Phi_S^T \frac{\partial \mathbf{A}}{\partial m_k} \Phi_R. \quad (21)$$

Equation (21) is valid for any numbers of the source/sink electrodes and the receiving electrodes as well as any of their 215 combinations. One only needs to compute the net potentials due to the source/sink electrodes Φ_S and receiving electrodes Φ_R using the superposition principle.

4 Parallel Implementation within PFLOTRAN

The parallel implementation strategy for the ERT forward modeling and inversion capabilities follows the existing PFLOTRAN's strategy for flow and transport process models. Forward ERT modeling is implemented by extending PFLOTRAN's hierarchy of Fortran classes and data structures to create a new ERT process model that constructs and solves the linear systems of equations described in Section 2. The implementation of inversion, as described in Section 3, is completely new to PFLOTRAN and requires all the necessary steps (or calculations) for ERT inversion, such as the cost-function computation, regularization, Jacobian computation, solution to the dense Gauss–Newton normal equation, and model update.

To execute simulations in parallel, domain decomposition is employed to distribute the discrete computational domain across (computer) processes using a division along the three principal axes for structured grids (calculated by PETSc or explicitly defined by the user), or in the case of unstructured grids, the ParMETIS parallel graph partitioner (Karypis and Schloegel, 2013) is employed to divide the domain. All interprocess communication is carried out with MPI (Message Passing Interface), and PETSc's parallel data structures facilitate much of the MPI communication.

Efficient partitioning and the overlapping of communication with computation (e.g., calculating floating point operations while data is being transferred between processes) helps improve parallel performance. The latter helps mask communication costs. The goal is to maximize the time spent in computation and minimize interprocess communication and enforce *load balance*, the even distribution of computation across processes employed. Efficient partitioning can be accomplished by dividing the domain evenly and minimizing the number and size of (MPI) messages being passed between processes. However, an even distribution does not necessarily guarantee load balance as certain regions of the domain with time-consuming operations (e.g., localized physics, intense input/output or I/O) may require more computation. Thus, efficient partitioning is not trivial.

PFLOTRAN's parallel file I/O is through HDF5 (Hierarchical Data Format version 5), a platform-independent binary file format that leverages MPI-IO operations to read/write a single file using multiple (potentially thousands of) processes (The HDF Group, 2022). PFLOTRAN writes ASCII-formatted observation files (files storing ERT data at specific locations in space) locally on the processes owning the observation points. Thus, the writing of observation files requires no parallel communication. For more comprehensive details on the parallel implementation strategy employed within PFLOTRAN, interested readers are referred to Hammond et al. (2012, 2014).

5 Modeling benchmarking results

To examine the accuracy of the PFLOTRAN ERT process model, simulation results are compared against solutions obtained using well established analytical and numerical methods. Two cases are considered for the comparison: three-layer earth and 3-D earth resistivity models. All benchmarking simulations were performed on the Deception supercomputer housed at the Pacific Northwest National Laboratory. It is composed of 96 compute nodes where each node has 64-core AMD EPYC 7502 processors running at 2.5 GHz boost to 3.35 GHz with 256 GB of memory.

5.1 Layered earth model

Let us consider the layered earth or 1D model shown in Fig. 3. The model is composed of three layers, whose resistivity (ρ) and thickness (t) are [100 Ωm , 30 m], [300 Ωm , 30 m], and [10 Ωm , half-space], respectively, from the top to bottom. The dimension of the model is $400 \times 400 \times 100 \text{ m}^3$. The model boundaries are extended to $\pm 3000 \text{ m}$ in the x - and y -directions, and 3000 m in the z -direction to accommodate zero Dirichlet boundary condition at the side and bottom boundaries. These boundary extensions are not shown in Fig. 3. Vertical electrical sounding (VES) data are simulated at the center of the model considering a Wenner configuration. This configuration consists of four electrodes A, B, M, and N, which are deployed in-line with an equal spacing a between the two neighboring electrodes (Fig. 2). Electrodes A and B again act as the current source and sink electrodes, while electrodes M and N as the potential receiving electrodes. Starting with a smaller value of a for a shallow depth investigation, the value is progressively increased to examine deeper depths. In our case, we start with a minimum value $a_{\min} = 8 \text{ m}$ and progressively increase it to a maximum $a_{\max} = 132 \text{ m}$ following $a = a_{\min} + i\delta a$, where $\delta a = 4 \text{ m}$ and $i = 0, 1, 2, \dots, 31$. This results in $N_d = 32$ VES sounding recordings using $N_e = 128$ electrodes.

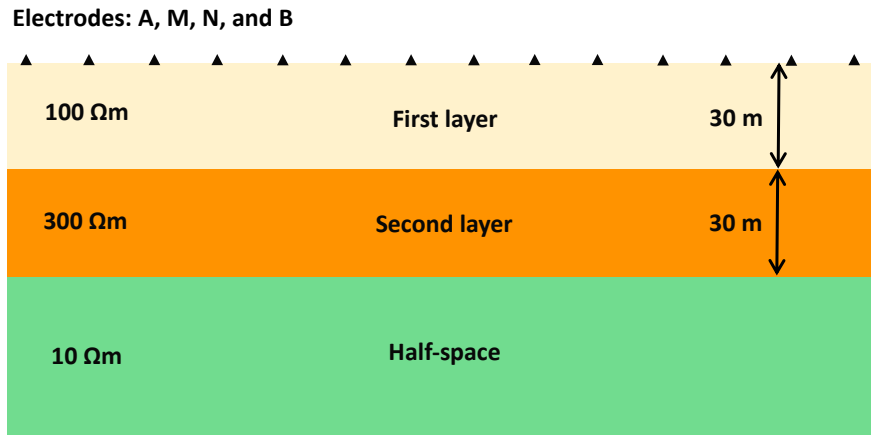


Figure 3. Vertical cross-section of a three-layer earth resistivity model used for benchmarking PFLOTRAN results against the analytical solutions.

To simulate the VES data, the model is discretized using a mix of uniform and non-uniform cell sizes in the x -, y -, and z -directions. The main computational domain is discretized with a uniform grid spacings of 2 m in the x - and y -directions and 1 m in the z -direction. The boundary regions are discretized with severely stretched non-uniform grid spacings following Jaysaval et al. (2014). The discretization results in a grid with $240 \times 240 \times 120$ cells. Therefore, using the FV method to simulate electric potentials leads to a system of linear equations with 6,912,000 DOFs. PFLOTRAN simulated the VES data using 512 processes on Deception. Total simulation time was 45 s.

The simulated VES results are compared against analytical solutions obtained from SimPEG (Cockett et al., 2015). Figure 4 shows the comparison of the apparent resistivity ρ^a responses. Analytical and PFLOTRAN results are shown by the solid

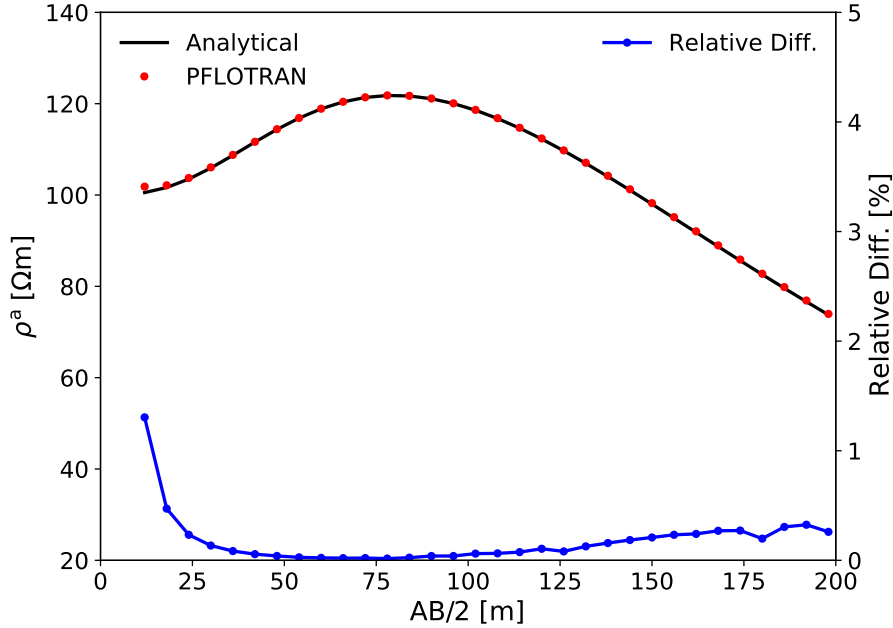


Figure 4. Apparent resistivity ρ^a responses computed for the layered earth resistivity model in Fig. 3 (upper plot). They are calculated using the analytical method (solid black lines) and PFLOTRAN (filled red circles). The lower plot (blue line) shows the relative difference (in percentage) between the analytical and PFLOTRAN results.

black line and filled red circles, respectively. We observe that both responses agree very well. To make the comparison more quantitative, we calculate the relative difference between the two responses and show the difference in Fig. 4 (blue line). The relative difference varies from 0.02% to 1.3%. We also calculate the average relative difference using

$$\epsilon = \frac{1}{N_d} \sum_i^{N_d} \frac{\rho_{1i}^a - \rho_{2i}^a}{\rho_{1i}^a}, \quad (22)$$

where ρ^a with subscripts 1 and 2, respectively, represents the apparent resistivity computed using PFLOTRAN and the analytical method from SimPEG. The average relative difference ϵ is 0.18%. These relatively small differences demonstrate the high degree of accuracy of the PFLOTRAN ERT process model for the 1D model. Note that the apparent resistivity was computed using

$$\rho^a = \frac{\Delta\phi}{IG}, \quad (23)$$

where $\Delta\phi = \phi_{AM} - \phi_{AN} - \phi_{BM} + \phi_{BN}$ and G is a geometric factor given as

$$G = \frac{1}{2\pi} \left(\frac{1}{r_{AM}} - \frac{1}{r_{AN}} - \frac{1}{r_{BM}} + \frac{1}{r_{BN}} \right), \quad (24)$$

with distances between the source electrode A and the receiving electrodes M and N are, respectively, r_{AM} and r_{AN} , while between the sink electrode B and the receiving electrodes M and N are, respectively, r_{BM} and r_{BN} . For the Wenner configuration, $r_{AM} = a$, $r_{AN} = 2a$, $r_{BM} = 2a$, and $r_{BN} = a$. Therefore, geometric factor $G = 1/2\pi a$ was used for calculating ρ^a using Eq. (23).

5.2 3D earth model

In the previous example, PFLOTRAN accuracy for simulating ERT responses was validated for a simplistic 1D earth model. In real world, however, the earth models are 3-D with varying degrees of heterogeneity. Therefore, it is important to validate the accuracy of numerical results computed using PFLOTRAN against results computed using well established methods for 3D models.




Figure 5. 3D earth resistivity model composed of two anomalous blocks embedded in a homogeneous background. ERT profiling data are simulated along a survey profile at $y = z = 0$ m from $x = -16$ m to $x = 16$ m using the Wenner configuration with increasing electrode spacing a from 1 to 6 m.

We consider the 3D model shown in Fig. 5. The model dimension is $32 \times 16 \times 8$ m³. It has a homogeneous background of conductivity 0.002 Sm⁻¹ and includes two anomalous blocks: a conductive block (0.2 Sm⁻¹) and a resistive block (0.0002 Sm⁻¹). Both anomalous blocks have a dimension of $2 \times 2 \times 2$ m³ and are separated by 4 m in the x -direction. As in the previous case, the model boundaries are extended to ± 1000 m in the x - and y -directions, and 1000 m in the z -direction to accommodate zero Dirichlet boundary condition at the side and bottom boundaries. ERT profiling data are simulated along a survey profile (yellow line) at $y = z = 0$ m from $x = -16$ m to $x = 16$ m using the Wenner configuration with increasing electrode spacing a from 1 to 6 m. Unlike the VES performed for the 1D model where data were simulated at a single observation point at the center of the model, the ERT profiling, performed for the 3D model, records data at multiple observation points along the survey profile.

The profiling, therefore, helps in investigating lateral conductivity variations in the model along the x -direction. In this case, it results in $N_d = 129$ ERT data recording by using $N_e = 32$ electrodes placed uniformly at every 1 m from $x = -15.5$ m to $x = 15.5$ m.

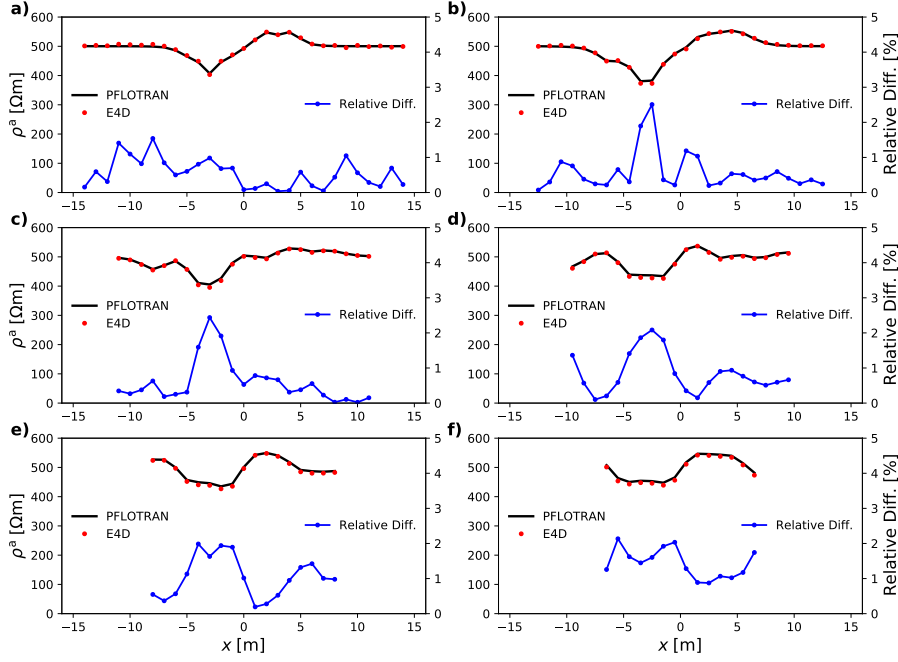


Figure 6. Apparent resistivity ρ^a responses computed for the 3D earth resistivity model in Fig. 5 using the Wenner configuration with spacings: (a) $a = 1$ m, (b) $a = 2$ m, (c) $a = 3$ m, (d) $a = 4$ m, (e) $a = 5$ m, and (f) $a = 6$ m (upper plots). They are calculated using PFLOTRAN (solid black lines) and E4D (filled red circles) for an ERT profiling along a survey profile at $y = z = 0$ from $x = -16$ m to $x = 16$ m. The lower plots (blue lines) show the relative difference (in percentage) between PFLOTRAN and E4D results.

The main computational domain is discretized with a uniform grid spacing of 0.2 m in all the three directions, while the boundary domain is again discretized with severely stretched non-uniform grid spacings. The model discretization leads to a grid with $200 \times 120 \times 60$ FV cells and, hence, 1,440,000 DOFs in the corresponding system of linear equations. PFLOTRAN simulated the ERT profiling data using 128 processes on Deception. Total simulation time was 9 s.

Figures 6(a) – (f) (upper curves) show the apparent resistivity ρ^a responses computed for the 3D model in Fig. 5 using the ERT profiling with the Wenner configuration, respectively, with electrode spacings of 1, 2, 3, 4, 5, and, 6 m. The solid black lines show the responses computed using PFLOTRAN. All responses exhibit the characteristics of the conductive and resistive anomalous blocks embedded in a homogeneous background of conductivity 0.002 Sm^{-1} or resistivity $500 \text{ }\Omega\text{m}$; see, e.g., the low apparent resistivity on the left and high apparent resistivity on the right sides of the plots. These responses are compared against responses computed using E4D (Johnson et al., 2010; Johnson and Wellman, 2015). In E4D, an unstructured-mesh FE method is implemented to perform the forward modeling. The filled red circles show the responses computed using E4D. The

310 relative differences between the two codes are shown by the lower blue lines, which vary from 0.02% to 2.5%. The average
relative difference ϵ computed using Eq. (22) is 0.77%. These relatively small differences imply a good agreement between
responses computed using both codes.

6 Inversion results

To examine the efficiency of PFLOTRAN to invert ERT data, we consider a model modified from Fig. 5 by cropping the model
315 for $|x| > 8$ m. The modified model, therefore, has the dimension of $16 \times 16 \times 8$ m³. The model is discretized using the same
grid spacings as in the previous 3D modeling example, which resulted in $120 \times 120 \times 60 = 864,000$ FV cells. We computed
synthetic ERT data for three survey lines located at the ground surface at $y = -5$ m, $y = 0$ m, and $y = 5$ m from $x = -8$ m to
 $x = 8$ m. Electrodes are placed uniformly with an electrode spacing of 1 m along the three survey lines from $x = -7.5$ m to
 $x = 7.5$ m, i.e., 16 electrodes along each survey line. Using these $N_e = 48$ electrodes, ERT data were simulated considering
320 different combinations of a set of 4 electrodes, where 2 electrodes acted as the current and the remaining 2 as the potential
receiving electrodes. These electrode combinations result in $N_d = 2070$ simulated ERT data. PFLOTRAN simulated these
ERT data using 128 processes on Deception. Total simulation time was 6 s. In this case, ERT data were simulated in the form
of the resistance obtained using $\frac{\Delta\phi}{I}$, which basically equals to the potential difference $\Delta\phi$ for unit current ($I = 1$ A).

The simulated ERT data were then inverted using PFLOTRAN starting with a model having a half-space conductivity of
325 0.0022 Sm^{-1} . The starting model conductivity was obtained by averaging the apparent conductivity computed from the given
ERT data. The previously described modeling grid also acted as the inversion grid, i.e., each cell in the model was used as an
inversion parameter. Therefore, we needed to invert for $N_m = 864,000$ model parameters. PFLOTRAN inverted the ERT data
using 512 processes on Deception. Total inversion time was 74 s.

Figure 7 shows the inverted conductivity model. The inverted model agrees well with the true model: the inversion recovered
330 the conductive and resistive anomalies in the model. The inverted resistivity, however, varies smoothly in the model because
a smooth regularization constraint is applied by minimizing the difference of the cell conductivity and the conductivity of its
neighboring cells. Further discussion on the implementation of our regularization schemes is beyond the scope of the current
paper.

The initial chi-squared data misfit χ^2 was 14.8. As the inversion progressed, χ^2 decreased as shown in Fig. 8. The inversion
335 took 11 iterations to satisfy the convergence criteria defined by $\chi^2 \leq 0.9$. To further examine the efficiency of PFLOTRAN
to fit the individual ERT data, a histogram of data-fit error is also shown in Fig. 9. The data-fit error is defined as ratio of the
residual resistance (difference of the observed and predicted resistances) and observed resistance expressed as a percentage.
The histogram implies that the ERT data were fitted comparatively well.




Figure 7. Inversion result of the simulated ERT data for a model modified from Fig. 5 by cropping the model for $|x| > 8$ m.




Figure 8. A measure of data misfit as the inversion progresses: χ^2 versus Number of Iteration.

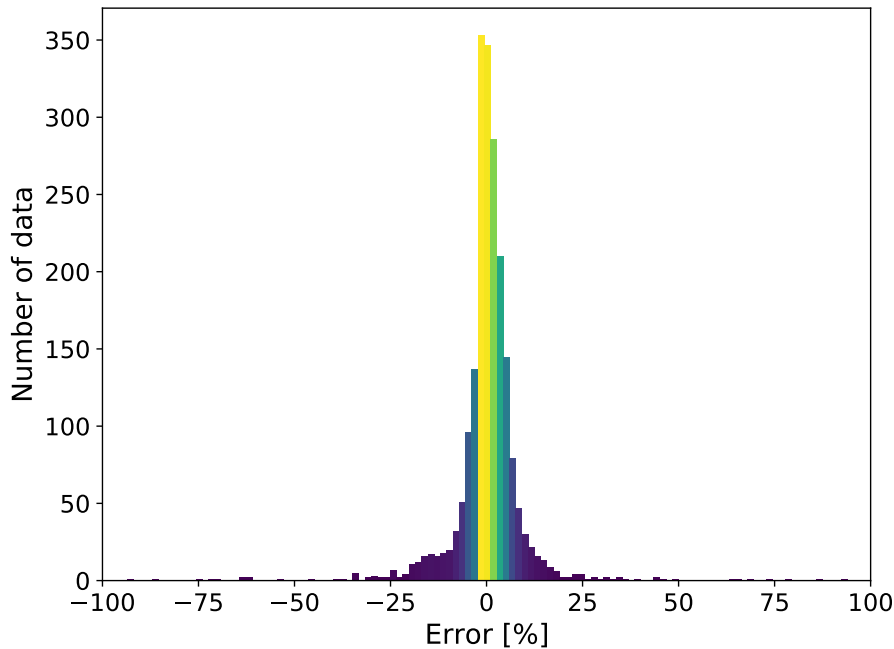


Figure 9. Histogram of data-fit error, which is defined as ratio of the residual resistance (difference of the observed and predicted resistances) and observed resistance expressed as a percentage.

7 Scalability tests

340 Scalability tests demonstrate parallel performance or how efficiently a simulator utilizes additional processes to either speed up a simulation of fixed problem size (strong scaling) or maintain speed as the problem size and number of processes increase proportionally (weak scaling). For ideal strong scaling, the speedup is proportional to the increase in the number of processes (e.g., doubling the number of processes cuts the run time in half), while ideal weak scaling results in an unaltered or constant run time. In either (ideal) case, the additional computing capacity is efficiently utilized. Ideal strong scaling is often termed
 345 linear speedup, whereas superlinear speedup is better than ideal (e.g., doubling the number of processes reduces the run time to less than half). Superlinear performance is not uncommon, as access to increased cache (efficient, high performance memory on the processor core) can generate superlinear speedup.

As discussed earlier, there are two computationally intensive steps to invert ERT data: forward modeling and Jacobian computation. To demonstrate the parallel performance of PFLOTRAN for these two steps, large-scale strong scaling tests
 350 are reported in this section by using up to 1,024 processes on the Deception supercomputer or up to 131,072 processes on the Theta supercomputer. Configuration details for Deception were previously given in the Modeling Benchmarking Results section (5). On the other hand, Theta is a leadership-class supercomputer, ranked within the top 100 fastest computers in the world and housed at the Argonne Leadership Computing Facility (ALCF) at Argonne National Laboratory (www.alcf.anl.gov/

Table 1. Test matrices obtained by rediscrctizing the three-layer earth model in Fig. 3 using different cell sizes dx , dy , and dz in the x -, y -, and z -directions. N_x , N_y , and N_z are respectively the number of FV cells in the x -, y -, and z -directions. Note that the cell sizes are given only for the main computational domain; the boundary domain is discretized using severely stretched non-uniform cell sizes.

Matrix Name	Discretization				DOFs
	$dx = dy$ (m)	dz (m)	$N_x = N_y$	N_z	
$M2$	4.0	1.00	140	120	2,352,000
$M7$	2.0	1.00	240	120	6,912,000
$M19$	1.6	0.50	290	220	18,502,000
$M35$	1.6	0.25	290	420	35,322,000
$M64$	0.8	0.50	540	220	64,152,000
$M122$	0.8	0.25	540	420	122,472,000

alcf-resources/theta). With a peak performance of 11.7 PFLOPs, it is composed of 4,392 compute nodes connected through
355 a Cray Aries network. Each node has 64-core Intel Xeon Phi 7230 processors running at 1.3 GHz with 192 GB of memory. Therefore, there are 281,088 processes available on Theta. The PFLOTRAN compilations on Deception and Theta were linked to Intel’s Math Kernel Library (MKL). The code leverages solely the distributed-memory parallel capability on the machines and not the shared-memory multi-threading.

7.1 Forward modeling scalability

360 The three-layer earth model in Fig. 3 is first rediscrctized to generate several test matrices for the analysis by considering different cell sizes in the x -, y -, and z -directions. These test matrices result in six different modeling scenarios from coarser to finer grids. Table 1 completely specifies details of the test matrices $M2$, $M7$, $M19$, $M35$, $M64$, and $M122$ along with the number of DOFs in the governing matrix equations that ranges between 2 and 122 million. The naming of these matrices is based on the number of DOFs they represent, e.g., $M19$ represents a matrix equation with about 19 million DOFs. To the
365 best of our knowledge, the modeling with 122 million DOFs stands as the largest ERT modeling reported in the geophysical literature. The next largest 3D ERT modeling problem reported in the literature was composed of about 10 million DOFs (Ren et al., 2018), while most others were often composed of less than a million DOFs (Günther et al., 2006; Rucker et al., 2006; Johnson et al., 2010; Johnson and Wellman, 2015). Having capability to solve for up to a hundred million DOFs efficiently in parallel is a crucial requirement for large-scale ERT modeling and inversion problems in the future if model dimensions are
370 significantly increased and/or finer grids are employed.

Considering the six test matrices above, we performed the forward modeling scalability test on Theta for simulating ERT data for $N_e = 128$ electrodes, the same number of electrodes as in the three-layer benchmarking example. Figure 10 illustrates the PFLOTRAN ERT process model strong scaling for each test case (wall-clock time versus the number of processes employed for a fixed problem size). The wall-clock time is shown on the vertical axis, while the horizontal axis shows the number of




Figure 10. PFLOTRAN scalability for simulating ERT data for the model in Fig. 3 for $N_e = 128$ electrodes. The test is performed on the Theta supercomputer considering six test matrices $M2$, $M7$, $M19$, $M35$, $M64$, and $M122$ with the respective number of DOFs of about 2, 7, 19, 35, 64, and 122 million (see Table 1 for details).

375 processes; both axes are shown in logarithmic base 2 scale. The line labeled "M19 Ideal" is the plot of ideal performance based on the 64 process run time for the M19 simulation:

$$T_n = \frac{T_{64} \times 64}{n}, \quad (25)$$

where T_n is the time on n processes and T_{64} is the time on 64 processes.

For all problem sizes, the wall-clock times flatten at higher numbers of processes. This phenomenon is expected as the amount of work per process (DOFs per process) decreases at higher process counts and the cost of communication between processes begins to dominate. The M19 performance begins to deviate from ideal at ~ 2048 processes, which equates to ~ 9000 DOFs per process. Similar behavior is observed for many of the other problem sizes. Hammond et al. (2014) demonstrates that for groundwater flow and reactive transport simulation PFLOTRAN runs most efficiently when at least 10,000 DOFs are assigned to each process. These results show that this 10,000 DOF threshold also applies for PFLOTRAN ERT simulations. That is, PFLOTRAN's scalability on the supercomputer should be more ideal when the number of DOFs per process is greater than 10,000.

Another factor limiting parallel performance is the limited scalability of linear Krylov solvers at high process counts. Mills et al. (2009) demonstrates that the vector inner products performed within Krylov solvers become increasingly expensive as tens of thousands of processes are employed to solve a problem. These inner products (or MPI global reduction operations) require repeated global synchronization across all processes involved. As more processes are employed to solve the problems (requiring communication across a growing network of compute nodes), the transfer of information for each vector inner

product is more expensive time-wise (i.e., increased number of hops within the high-speed network, longer distance traveled). Figure 10 reveals similar phenomenon.

7.2 Jacobian computation scalability

395 To test the scalability of PFLOTRAN for the Jacobian computation, we consider a Jacobian matrix computed for inverting the ERT data in the Inversion Result Section 6. Recalling that $N_d = 2070$ ERT data and $N_m = 864,000$ model parameters for the inversion example, the dimension of the Jacobian matrix \mathcal{J} is therefore $2070 \times 864,000$. The Jacobian matrix is computed at each inversion iteration after the forward modeling run and used in the normal equation (Eq. 15) to get a model update $\delta\mathbf{m}$.




Figure 11. PFLOTRAN strong scaling for computing the Jacobian for the ERT inversion example in Fig. 7. The dimension of the Jacobian matrix \mathcal{J} is $2070 \times 864,000$ and is distributed uniformly over each computing process. PFLOTRAN exhibits superlinear performance for larger numbers of processes.

The Jacobian scalability test is performed on the Deception supercomputer. Figure 11 illustrates the strong scalability of PFLOTRAN wall-clock time for computing \mathcal{J} . It shows that PFLOTRAN exhibits linear (and sometimes superlinear) scalability for computing \mathcal{J} with increasing numbers of processes. Consequently, the total wall-clock time for computing \mathcal{J} using 1,024 processes reduces to 0.5 s compared to 887 s when using a single process on Deception. This ideal performance in computing \mathcal{J} is most likely due to even load balance across processes and the lack of MPI communication in its construction. Moreover, as noted before, the superlinear performance is most likely due to access to increased cache with an increasing number of processes.

For large-scale ERT inversion problems, the linear scaling for computing \mathcal{J} can be exploited to reduce the overall inversion wall-clock time even if PFLOTRAN exhibits deteriorating scalability for the forward modeling on a very large number of processes. For example, for an ERT inversion problem where the \mathcal{J} calculation time is much higher than the forward modeling time, increasing the number of processes to reduce the overall inversion wall-clock time is recommended, even when forward
410 modeling exhibits significant departure from the linear scalability. One has to find a balance between the two, and the most time consuming operation should govern.

8 Discussions

In the past 20 years, a few open-source 3D ERT modeling and inversion codes were developed to handle moderate-scale ERT problems (e.g., ResIPy, SimPEG, pyGIMLI, and E4D). Although some of these codes can practically tackle most present
415 ERT modeling and inversion problems having less than 10 million DOFs, they still lack full-scale parallel implementations to efficiently distribute numerical computations over massive supercomputers. The limitation can also lead to computational challenges for large-scale ERT problems with 100 million or more DOFs, which may not be solved efficiently with the existing codes. The 3D ERT modeling and inversion capability developed within PFLOTRAN removes this limitation. To examine the accuracy of the forward modeling capability, numerical results from PFLOTRAN were compared against analytical solutions
420 for a layered earth model and numerical results obtained from E4D for a 3D earth model. The average relative differences stayed within a percent, which demonstrated the high degree of accuracy of the PFLOTRAN modeling capability. Furthermore, large-scale strong scalability tests, using up to 131,072 processes, were performed on a leadership class supercomputer to demonstrate efficient distribution of computational workloads over massive supercomputers. The scalability tests demonstrate that computational times for the Jacobian computation scale linearly with increasing number of processes. For the forward modeling, PFLOTRAN showed a linear to near-linear scaling behaviour up to 10,000 DOFs per process, beyond which
425 performance tends to degrade

Due to the efficient parallel performance of PFLOTRAN observed above, a large-scale ERT simulation problem with 122 million DOFs and 122 electrodes could be solved in about 2 minutes by using 32,768 processes as compared to taking about an hour with 512 processes. Combining this forward modeling scalability along with the ideal scaling of Jacobian
430 computation, we can aim for inverting moderate-scale ERT problems within minutes and large-scale problems within hours with the access to massive supercomputers. The moderate-scale inversion example of Section 6 demonstrated that inversion of $N_d = 2070$ data and $N_m = 864,000$ model parameters could be performed in about a minute by using 512 processes.

Finally, there were two key motivations for implementing ERT capabilities in PFLOTRAN. The first motivation was to utilize the massive parallel infrastructure implemented within PFLOTRAN for the flow and transport simulations as illustrated
435 by Hammond et al. (2012, 2014) for simulating subsurface problems composed of billions of DOFs utilizing hundreds of thousands of computing processors. This infrastructure made PFLOTRAN a natural fit to achieve our goal of massive parallel implementation for the ERT forward modeling and inversion capabilities. Researchers familiar with PFLOTRAN can now have immediate access to the added geophysical ERT capabilities with minimal learning efforts.

The second motivation was to have a native implementation of ERT modeling and inversion capabilities along with flow and transport simulations. The new capabilities will allow us to perform fully coupled hydro-geophysical simulations and inversion. The idea behind the joint hydro-geophysical simulation is to use known flow and transport parameters, such as porosity and intrinsic permeability, and simulate for hydrological variables, e.g., liquid saturation and solute concentration. Using petrophysical relationships (e.g., Archie’s law or its modifications; Slater, 2007), these parameters and variables can be mapped into electrical conductivity, which subsequently can be used for simulating ERT responses. In our previous work, a joint hydro-geophysical simulation capability was developed by combining the capabilities of PFLOTRAN and E4D (Johnson et al., 2017). However, it was not straightforward to use PFLOTRAN-E4D for joint hydro-geophysical simulations. This is because the implementation was based on obtaining flow and transport parameters and variables from PFLOTRAN using a structured mesh, which then needed to be interpolated on an unstructured tetrahedral mesh for ERT simulations by E4D. The process was cumbersome and had the potential for errors. PFLOTRAN’s native implementation of ERT avoids employing any kind of interpolation because both flow and transport and ERT simulations use the same modeling grid.

Furthermore, there are growing interests in joint hydro-geophysical inversion, i.e., estimating the distribution of subsurface flow and transport parameters, from time-lapse ERT data, see e.g., Mboh et al. (2012); Camporese et al. (2015); Rucker et al. (2017); González-Quirós and Comte (2021); Pleasants et al. (2022). The inclusion of the ERT simulation capability in PFLOTRAN fulfills the requirement of the forward modeling necessary for native implementation of hydro-geophysical inversion. Fully coupled hydro-geophysical simulation and inversion capabilities are currently under development within PFLOTRAN and will be a focus of our forthcoming papers.

9 Conclusions

Two new capabilities — forward modeling and inversion of geophysical ERT data — are developed within PFLOTRAN, a massively parallel open-source, state-of-the-art, multi-phase, multi-component subsurface flow and transport simulation code. The capabilities are accurate, robust, and highly scalable on HPC platforms. Accuracy of the forward modeling capability was demonstrated for layered earth and 3D earth conductivity models by comparing our numerical results against well-established analytical and numerical results. The average relative differences stayed within a percent, demonstrating the high degree of accuracy of the PFLOTRAN ERT process model. The inversion capability was demonstrated for a 3D synthetic case by recovering the conductivity of the model. Large-scale scalability tests illustrated that the forward modeling for models with a hundred million DOFs can be performed in a few minutes, e.g., ERT modeling for a model with 122 million DOFs and 122 electrodes was performed in 128s by using 32,768 processes on a supercomputer. The Jacobian computation, a key ingredient for the inversion, exhibited ideal or linear scaling, and as a result a Jacobian of $2070 \times 864,000$ dimension was computed in 0.5s using 1,024 processes compared to 887s using a single process. Moreover, integration of geophysical ERT capabilities within PFLOTRAN opens door to perform coupled hydro-geophysical modeling and inversion, which is a focus of our future research.

Code and data availability. The source code, input file, and results presented in this manuscript are based on PFLOTRAN version v4.0. It can be downloaded from its Git repository at <https://bitbucket.org/pflotran/pflotran> and compiled after checking out the v4.0 tag (git checkout maint/v4.0) following the instructions provided at <https://pflotran.org/documentation>. A snapshot of PFLOTRAN v4.0 is also uploaded to Zenodo repository at <https://doi.org/10.5281/zenodo.6191926> (Jaysaval et al., 2022). The corresponding version of PETSc is v3.16.4 and was configured with Intel C/C++ 2020u4 using the following configuration script: `./configure -CFLAGS='-O3' -CXXFLAGS='-O3' -FFLAGS='-O3' -with-debugging=no -download-mpich=yes -download-hdf5=yes -download-fblaslapack=yes -download-metis=yes -download-parmetis=yes -download-hdf5-fortran-bindings=yes`. The modeling and inversion results presented in the paper are also available at the above Zenodo repository in paper-examples folder. Each subfolder within it contains a README file with brief instructions to reproduce the results.

480 PFLOTRAN is available under the GNU Lesser General Public License (LGPL) version 3.0. The dependencies listed here are available under compatible open source licenses except for ParMetis which is free for research and education purposes by non-profit institutions and US government agencies.

Appendix A: Jacobian computation

Let's assume that \mathbf{p}_M and \mathbf{p}_N are two interpolation vectors and Φ_A and Φ_B are potential distributions due to the source A and sink B (Fig. 2) such that

$$\phi_{AM} = \mathbf{p}_M^T \Phi_A, \quad \phi_{AN} = \mathbf{p}_N^T \Phi_A, \quad \phi_{BM} = \mathbf{p}_M^T \Phi_B, \quad \text{and} \quad \phi_{BN} = \mathbf{p}_N^T \Phi_B. \quad (\text{A1})$$

Therefore, the j^{th} simulated potential difference between the receiving electrodes M and N can be obtained as

$$d_j^{\text{syn}} = \phi_{AM} - \phi_{AN} - \phi_{BM} + \phi_{BN}. \quad (\text{A2})$$

Using Eq. (A1) in Eq. (A2), we get

$$490 \quad d_j^{\text{syn}} = \mathbf{p}_M^T (\Phi_A - \Phi_B) - \mathbf{p}_N^T (\Phi_A - \Phi_B). \quad (\text{A3})$$

Taking derivative of Eq. (A3) with respect to conductivity m_k of k^{th} cell yields

$$\frac{\partial d_j^{\text{syn}}}{\partial m_k} = \mathbf{p}_M^T \frac{\partial}{\partial m_k} (\Phi_A - \Phi_B) - \mathbf{p}_N^T \frac{\partial}{\partial m_k} (\Phi_A - \Phi_B). \quad (\text{A4})$$

Furthermore, taking derivative of Eq. (10) with respect to m_k yields

$$\frac{\partial \mathbf{A}}{\partial m_k} \Phi + \mathbf{A} \frac{\partial \Phi}{\partial m_k} = \mathbf{0}, \quad (\text{A5})$$

495 where we used $\partial s / \partial m_k = 0$ as the source does not depend on the conductivity.

If \mathbf{A} is invertible, rearranging Eq. (A5) gives

$$\frac{\partial \Phi}{\partial m_k} = -\mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial m_k} \Phi. \quad (\text{A6})$$

Using Eq. (A6) in Eq. (A4), we get

$$\frac{\partial d_j^{\text{syn}}}{\partial m_k} = -\mathbf{p}_M^T \mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial m_k} (\Phi_A - \Phi_B) + \mathbf{p}_N^T \mathbf{A}^{-1} \frac{\partial \mathbf{A}}{\partial m_k} (\Phi_A - \Phi_B). \quad (\text{A7})$$

500 As $\partial d_j^{\text{syn}} / \partial m_k$ is a scalar quantity, using the property $\mathbf{x}^T \mathbf{y} = \mathbf{y}^T \mathbf{x}$ for two vectors \mathbf{x} and \mathbf{y} , Eq. (A7) yields

$$\frac{\partial d_j^{\text{syn}}}{\partial m_k} = -(\Phi_A - \Phi_B)^T \frac{\partial \mathbf{A}}{\partial m_k} (\mathbf{A}^{-1} \mathbf{p}_M - \mathbf{A}^{-1} \mathbf{p}_N), \quad (\text{A8})$$

where $\mathbf{A}^T = \mathbf{A}$ is used due to its symmetry.

If the same interpolation scheme is used for placing the current electrodes within the discretized cells and recording the potential at electrodes from the cell potentials, we will have $\mathbf{A}^{-1} \mathbf{p}_M = -\Phi_M$ and $\mathbf{A}^{-1} \mathbf{p}_N = -\Phi_N$. Therefore,

$$505 \frac{\partial d_j^{\text{syn}}}{\partial m_k} = (\Phi_A - \Phi_B)^T \frac{\partial \mathbf{A}}{\partial m_k} (\Phi_M - \Phi_N). \quad (\text{A9})$$

Author contributions. PJ implemented the ERT modeling and inversion capabilities with contributions from GEH. PJ performed the tests, analyzed the results, and prepared the figures. TCJ supervised the study. PJ wrote the paper with contributions from GEH, and all co-authors read and helped with editing.

Competing interests. The authors have no relevant financial or non-financial interests to disclose.

510 *Acknowledgements.* The research described in this paper is part of the INSITE (INduced Spectral Imaging for The Environment) Agile investment at Pacific Northwest National Laboratory (PNNL). It was conducted under the Laboratory Directed Research and Development Program at PNNL, a multiprogram national laboratory operated by Battelle for the U.S. Department of Energy (DOE). We gratefully acknowledge computing resources available through Research Computing at PNNL and Argonne Leadership Computing Facility (ALCF). ALCF is a DOE Office of Science User Facility supported under Contract DE-AC02-06CH11357. Finally, we would like express our grati-
515 tude to chief executive editor David Ham, topical editor Adrian Sandu, and two reviewers Mark Everett and Michael Tso, for their careful reading, suggestions, and constructive comments that significantly helped to improve the paper.

References

- Alshehri, F. and Abdelrahman, K.: Groundwater resources exploration of Harrat Khaybar area, northwest Saudi Arabia, using electrical resistivity tomography, *Journal of King Saud University-Science*, 33, 101 468, 2021.
- 520 Badmus, B. and Olatinsu, O.: Geoelectric mapping and characterization of limestone deposits of Ewekoro formation, southwestern Nigeria, *Journal of Geology and Mining Research*, 1, 008–018, 2009.
- Balay, S., Abhyankar, S., Adams, M. F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Dener, A., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., May, D. A., McInnes, L. C., Mills, R. T., Munson, T., Rupp, K., Sanan, P., Smith, B. F., Zampini, S., Zhang, H., and Zhang, H.: PETSc Users Manual, Tech. Rep. ANL-95/11 - Revision 3.15, Argonne National Laboratory, 2021.
- 525 Bery, A. A., Saad, R., Mohamad, E. T., Jinmin, M., Azwin, I., Tan, N. A., and Nordiana, M.: Electrical resistivity and induced polarization data correlation with conductivity for iron ore exploration, *The Electronic Journal of Geotechnical Engineering*, 17, 3223–3233, 2012.
- Blanchy, G., Saneiyani, S., Boyd, J., McLachlan, P., and Binley, A.: ResIPy, an intuitive open source software for complex geoelectrical inversion/modeling, *Computers & Geosciences*, 137, 104 423, 2020.
- Blome, M., Maurer, H., and Schmidt, K.: Advances in three-dimensional geoelectric forward solver techniques, *Geophysical Journal International*, 176, 740–752, 2009.
- 530 Camporese, M., Cassiani, G., Deiana, R., Salandin, P., and Binley, A.: Coupled and uncoupled hydrogeophysical inversions using ensemble Kalman filter assimilation of ERT-monitored tracer test data, *Water Resources Research*, 51, 3277–3291, 2015.
- Caputo, R., Piscitelli, S., Oliveto, A., Rizzo, E., and Lapenna, V.: The use of electrical resistivity tomographies in active tectonics: examples from the Tyrnavos Basin, Greece, *Journal of Geodynamics*, 36, 19–35, 2003.
- 535 Cockett, R., Kang, S., Heagy, L. J., Pidlisecky, A., and Oldenburg, D. W.: SimPEG: An open source framework for simulation and gradient based parameter estimation in geophysical applications, *Computers & Geosciences*, 85, 142–154, <https://doi.org/10.1016/j.cageo.2015.09.015>, 2015.
- Coggon, J.: Electromagnetic and electrical modeling by the finite element method, *Geophysics*, 36, 132–155, 1971.
- Dahlin, T.: 2D resistivity surveying for environmental and engineering applications, *First break*, 14, 1996.
- 540 Dahlin, T.: The development of DC resistivity imaging techniques, *Computers & Geosciences*, 27, 1019–1029, 2001.
- Dahlin, T. and Zhou, B.: A numerical comparison of 2D resistivity imaging with 10 electrode arrays, *Geophysical prospecting*, 52, 379–398, 2004.
- Das, U. C.: Direct current electric potential computation in an inhomogeneous and arbitrarily anisotropic layered earth, *Geophysical Prospecting*, 43, 417–432, <https://doi.org/10.1111/j.1365-2478.1995.tb00260.x>, 1995.
- 545 Dey, A. and Morrison, H. F.: Resistivity modeling for arbitrarily shaped three-dimensional structures, *Geophysics*, 44, 753–780, 1979.
- Ellis, R. and Oldenburg, D.: The pole-pole 3-D DC-resistivity inverse problem: a conjugate gradient approach, *Geophysical journal international*, 119, 187–194, 1994.
- Gabarrón, M., Martínez-Pagán, P., Martínez-Segura, M. A., Bueso, M. C., Martínez-Martínez, S., Faz, Á., and Acosta, J. A.: Electrical resistivity tomography as a support tool for physicochemical properties assessment of near-surface waste materials in a mining tailing
- 550 pond (El Gorguel, SE Spain), *Minerals*, 10, 559, 2020.
- González-Quirós, A. and Comte, J.-C.: Hydrogeophysical model calibration and uncertainty analysis via full integration of PEST/PEST++ and COMSOL, *Environmental Modelling & Software*, 145, 105 183, 2021.

- Greggio, N., Giambastiani, B., Balugani, E., Amaini, C., and Antonellini, M.: High-resolution electrical resistivity tomography (ERT) to characterize the spatial extension of freshwater lenses in a salinized coastal aquifer, *Water*, 10, 1067, 2018.
- 555 Günther, T., Rücker, C., and Spitzer, K.: Three-dimensional modelling and inversion of DC resistivity data incorporating topography—II. Inversion, *Geophysical Journal International*, 166, 506–517, 2006.
- Hammond, G. E., Lichtner, P. C., Lu, C., and Mills, R. T.: PFLOTTRAN: Reactive flow & transport code for use on laptops to leadership-class supercomputers, in: *Groundwater Reactive Transport Models*, edited by Zhang, F., Yeh, G.-T. G., and Parker, J. C., pp. 141–159, BENTHAM SCIENCE PUBLISHERS, <https://doi.org/10.2174/978160805306311201010141>, 2012.
- 560 Hammond, G. E., Lichtner, P. C., and Mills, R. T.: Evaluating the performance of parallel subsurface simulators: An illustrative example with PFLOTTRAN, *Water Resources Research*, 50, 208–228, <https://doi.org/10.1002/2012WR013483>, 2014.
- Hestenes, M. and Stiefel, E.: Methods of conjugate gradients for solving linear systems, *Journal of Research of the National Bureau of Standards*, 49, 409, <https://doi.org/10.6028/jres.049.044>, 1952.
- Jahandari, H. and Farquharson, C. G.: A finite-volume solution to the geophysical electromagnetic forward problem using unstructured grids, 565 *Geophysics*, 79, E287–E302, 2014.
- Jaysaval, P., Shantsev, D., and de la Kethulle de Ryhove, S.: Fast multimodel finite-difference controlled-source electromagnetic simulations based on a Schur complement approach, *GEOPHYSICS*, 79, E315–E327, <https://doi.org/10.1190/geo2014-0043.1>, 2014.
- Jaysaval, P., Hammond, G. E., and Johnson, T. C.: ERT modeling and inversion using PFLOTTRAN v4.0, <https://doi.org/10.5281/zenodo.6191926>, 2022.
- 570 Johnson, T. C. and Wellman, D.: Accurate modelling and inversion of electrical resistivity data in the presence of metallic infrastructure with known location and dimension, *Geophysical Journal International*, 202, 1096–1108, 2015.
- Johnson, T. C., Versteeg, R. J., Ward, A., Day-Lewis, F. D., and Revil, A.: Improved hydrogeophysical characterization and monitoring through parallel modeling and inversion of time-domain resistivity and induced-polarization data, *GEOPHYSICS*, 75, WA27–WA41, <https://doi.org/10.1190/1.3475513>, 2010.
- 575 Johnson, T. C., Slater, L. D., Ntarlagiannis, D., Day-Lewis, F. D., and Elwaseif, M.: Monitoring groundwater-surface water interaction using time-series and time-frequency analysis of transient three-dimensional electrical resistivity changes, *Water Resources Research*, 48, 2012.
- Johnson, T. C., Hammond, G. E., and Chen, X.: PFLOTTRAN-E4D: A parallel open source PFLOTTRAN module for simulating time-lapse electrical resistivity data, *Computers & Geosciences*, 99, 72–80, <https://doi.org/10.1016/j.cageo.2016.09.006>, 2017.
- Karypis, G. and Schloegel, K.: METIS/ParMETIS Parallel Graph Partitioner, Tech. rep., <https://www-users.cs.umn.edu/~karypis/metis/>, 580 accessed: 2021-11-11, 2013.
- Kessouri, P., Johnson, T., Day-Lewis, F. D., Wang, C., Ntarlagiannis, D., and Slater, L. D.: Post-remediation geophysical assessment: Investigating long-term electrical geophysical signatures resulting from bioremediation at a chlorinated solvent contaminated site, *Journal of environmental management*, 302, 113 944, 2022.
- Lee, T.: An integral equation and its solution for some two-and three-dimensional problems in resistivity and induced polarization, *Geophysical Journal International*, 42, 81–95, 1975. 585
- Li, Y. and Spitzer, K.: Three-dimensional DC resistivity forward modelling using finite elements in comparison with finite-difference solutions, *Geophysical Journal International*, 151, 924–934, 2002.
- Loke, M., Chambers, J., Rucker, D., Kuras, O., and Wilkinson, P.: Recent developments in the direct-current geoelectrical imaging method, *Journal of applied geophysics*, 95, 135–156, 2013.

- 590 Lysdahl, A. K., Bazin, S., Christensen, C., Ahrens, S., Günther, T., and Pfaffhuber, A. A.: Comparison between 2D and 3D ERT inversion for engineering site investigations—a case study from Oslo Harbour, *Near Surface Geophysics*, 15, 201–209, 2017.
- Martínez, J., Rey, J., Sandoval, S., Hidalgo, M., and Mendoza, R.: Geophysical prospecting using ERT and IP techniques to locate Galena veins, *Remote Sensing*, 11, 2923, 2019.
- Mboh, C., Huisman, J., Van Gaelen, N., Rings, J., and Vereecken, H.: Coupled hydrogeophysical inversion of electrical resistances and inflow
595 measurements for topsoil hydraulic properties under constant head infiltration, *Near Surface Geophysics*, 10, 413–426, 2012.
- Méndez-Delgado, S., Gómez-Treviño, E., and Pérez-Flores, M.: Forward modelling of direct current and low-frequency electromagnetic fields using integral equations, *Geophysical Journal International*, 137, 336–352, 1999.
- Meyerhoff, S. B., Maxwell, R. M., Revil, A., Martin, J. B., Karaoulis, M., and Graham, W. D.: Characterization of groundwater and surface
600 water mixing in a semiconfined karst aquifer using time-lapse electrical resistivity tomography, *Water Resources Research*, 50, 2566–2585, 2014.
- Mills, R. T., Hammond, G. E., Lichtner, P. C., Sripathi, V., Mahinthakumar, G., and Smith, B. F.: Modeling Subsurface Reactive Flows Using Leadership-Class Computing, in: 5th Annual Conference of Scientific Discovery through Advanced Computing (SciDAC 2009), vol. 180 of *Journal of Physics Conference Series*, <https://doi.org/10.1088/1742-6596/180/1/012062>, 2009.
- Nocedal, J. and Wright, S.: Numerical optimization, Springer Science & Business Media, 2006.
- 605 Park, S., Yi, M.-J., Kim, J.-H., and Shin, S. W.: Electrical resistivity imaging (ERI) monitoring for groundwater contamination in an uncontrolled landfill, South Korea, *Journal of Applied Geophysics*, 135, 1–7, 2016.
- Park, S. K. and Van, G. P.: Inversion of pole-pole data for 3-D resistivity structure beneath arrays of electrodes, *GEOPHYSICS*, 56, 951–960, <https://doi.org/10.1190/1.1443128>, 1991.
- Penz, S., Chauris, H., Donno, D., and Mehl, C.: Resistivity modelling with topography, *Geophysical Journal International*, 194, 1486–1497,
610 2013.
- Pervago, E., Mousatov, A., and Shevnin, V.: Analytical solution for the electric potential in arbitrary anisotropic layered media applying the set of Hankel transforms of integer order, *Geophysical Prospecting*, 54, 651–661, <https://doi.org/10.1111/j.1365-2478.2006.00555.x>, 2006.
- Pleasants, M., Neves, F. d. A., Parsekian, A., Befus, K., and Kelleners, T.: Hydrogeophysical Inversion of Time-Lapse ERT Data to Determine
615 Hillslope Subsurface Hydraulic Properties, *Water Resources Research*, 58, e2021WR031 073, 2022.
- Ren, Z., Qiu, L., Tang, J., Wu, X., Xiao, X., and Zhou, Z.: 3-D direct current resistivity anisotropic modelling by goal-oriented adaptive finite element methods, *Geophysical Journal International*, 212, 76–87, 2018.
- Revil, A., Karaoulis, M., Johnson, T., and Kemna, A.: Some low-frequency electrical methods for subsurface characterization and monitoring in hydrogeology, *Hydrogeology Journal*, 20, 617–658, 2012.
- 620 Richards, K., Revil, A., Jardani, A., Henderson, F., Batzle, M., and Haas, A.: Pattern of shallow ground water flow at Mount Princeton Hot Springs, Colorado, using geoelectrical methods, *Journal of Volcanology and Geothermal Research*, 198, 217–232, 2010.
- Rizzo, E., Colella, A., Lapenna, V., and Piscitelli, S.: High-resolution images of the fault-controlled High Agri Valley basin (Southern Italy) with deep and shallow electrical resistivity tomographies, *Physics and Chemistry of the Earth, Parts A/B/C*, 29, 321–327, 2004.
- Rockhold, M. L., Robinson, J. L., Parajuli, K., Song, X., Zhang, Z., and Johnson, T. C.: Groundwater characterization and monitoring at a
625 complex industrial waste site using electrical resistivity imaging, *Hydrogeology Journal*, 28, 2115–2127, 2020.
- Rosales, R. M., Martínez-Pagan, P., Faz, A., and Moreno-Cornejo, J.: Environmental monitoring using electrical resistivity tomography (ERT) in the subsoil of three former petrol stations in SE of Spain, *Water, Air, & Soil Pollution*, 223, 3757–3773, 2012.

- Rücker, C., Günther, T., and Spitzer, K.: Three-dimensional modelling and inversion of dc resistivity data incorporating topography—I. Modelling, *Geophysical Journal International*, 166, 495–505, 2006.
- 630 Rücker, C., Günther, T., and Wagner, F. M.: pyGIMLi: An open-source library for modelling and inversion in geophysics, *Computers & Geosciences*, 109, 106–123, 2017.
- Rucker, D. F., Myers, D. A., Cabbage, B., Levitt, M. T., Noonan, G. E., McNeill, M., Henderson, C., and Lober, R. W.: Surface geophysical exploration: developing noninvasive tools to monitor past leaks around Hanford’s tank farms, *Environmental monitoring and assessment*, 185, 995–1010, 2013.
- 635 Schulz, R.: The method of integral equation in the direct current resistivity method and its accuracy, *Journal of geophysics*, 56, 192–200, 1985.
- Singha, K., Johnson, T. C., Day-Lewis, F. D., and Slater, L. D.: *Electrical Imaging for Hydrogeology*, The Groundwater Project, Guelph, Ontario, Canada, 2022.
- Slater, L.: Near surface electrical characterization of hydraulic conductivity: From petrophysical properties to aquifer geometries—A review, *Surveys in Geophysics*, 28, 169–197, 2007.
- 640 Spitzer, K.: A 3-D finite-difference algorithm for DC resistivity modelling using conjugate gradient methods, *Geophysical Journal International*, 123, 903–914, 1995.
- Storz, H., Storz, W., and Jacobs, F.: Electrical resistivity tomography to investigate geological structures of the earth’s upper crust, *Geophysical prospecting*, 48, 455–472, 2000.
- 645 The HDF Group: Hierarchical Data Format Version 5, <http://www.hdfgroup.org/HDF5>, 2022.
- Uhlemann, S., Chambers, J., Falck, W. E., Tirado Alonso, A., Fernández González, J. L., and Espín de Gea, A.: Applying electrical resistivity tomography in ornamental stone mining: challenges and solutions, *Minerals*, 8, 491, 2018.
- Zhang, J., Mackie, R. L., and Madden, T. R.: 3-D resistivity forward modeling and inversion using conjugate gradients, *Geophysics*, 60, 1313–1325, 1995.
- 650 Zhdanov, M. S. and Keller, G. V.: The geoelectrical methods in geophysical exploration, *Methods in geochemistry and geophysics*, 31, I–IX, 1994.