Next-Generation Human Brain Neuroimaging and the Role of High-Performance Computing

Adnan Salman,* Allen Malony,* Sergei Turovets,* Vasily Volkov,† David Ozog,* Don Tucker‡

*University of Oregon {*adnan, malony, sergei, ozog*}*@cs.uoregon.edu*
†Belarusian State University *volkov@bsu.by*
‡Electrical Geodesics, Incorporated *dtucker@egi.com*

***Abstract*—Advances in human brain neuroimaging to achieve high-temporal and high-spatial resolution will depend on computational approaches to localize EEG signals to their sources in the cortex. The source localization inverse problem is inherently ill-posed and depends critically on the modeling of human head electromagnetics. In this paper we present a systematic methodology to analyze the main factors and parameters that affect the accuracy of the EEG source-mapping solutions. We argue that these factors are not independent and their effect must be evaluated in a unified way. To do so requires significant computational capabilities to explore the landscape of the problem, to quantify uncertainty effects, and to evaluate alternative algorithms. We demonstrate that bringing HPC to this domain will enable such investigation and will allow new avenues for neuroinformatics research. Two algorithms to the electromagnetics forward problem (the heart of the source localization inverse), incorporating tissue inhomogeneity and impedance anisotropy, are presented and their parallel implementations described. The head model forward solvers are evaluated and their performance analyzed.**

I. INTRODUCTION

Advances in human brain science have been closely linked with new developments in neuroimaging technology. Indeed, the integration of psychological behavior with neural evidence in cognitive neuroscience research has led to fundamental insights of how the brain functions and manifests our physical and mental reality. However, in any empirical science, it is the resolution and precision of measurement instruments that inexorably define the leading edge of scientific discovery. Human neuroscience is no exception. Brain activity takes place at millisecond temporal and millimeter spatial scales through the reentrant, bidirectional interactions of functional neural networks distributed throughout the cortex and interconnected by a complex network of white matter fibers. Unfortunately, current non-invasive neuroimaging instruments are unable to observe dynamic brain operation at these milliscales. Hemodynamic measures (functional magnetic resonance imaging (fMRI), positron emission tomography (PET)) have good 3D spatial resolution (1mm^3), but poor temporal resolution (≥ 0.5 seconds). Electromagnetic measures (electroencephalography (EEG), magnetoencephalography (MEG)) provide high temporal resolution in the order of neural activities (≤ 1 msec), but their spatial resolution lacks localization accuracy. The reason behind the limited spatial accuracy is the ambiguous nature of the electrostatic inverse problem. A given EEG signal can be explained by many sources of cortex source activity. In principle, it is impossible to solve this problem relying only on theoretical formulations alone and empirical methods are inherently underdetermined [1]. It is only by incorporating a priori knowledge and assumptions about the sources in the form of constraints that the problem can begin to be addressed. However, different assumptions and constraints give rise to different source analysis algorithms. All these algorithms are based on a head electromagnetics forward calculation that maps the cortex dipole sources in the brain to the scalp potential. Certainly, there are several factors related to modeling the human head as a volume conductor that introduce uncertainties in the forward solution [2]–[4]. Further, several factors related to EEG spatial sampling and noise level introduce other sources of error [5]. All of these factors will affect the inverse algorithms and they are likely to be highly correlated. Quantifying and ranking their effect on the source localization in a systematic way will provide insight and directions to where the research and effort should be focused.

The challenges of model space exploration, sensitivity analysis, and uncertainty quantification are found across science and engineering domains. The ability to formulate the algorithms and methodology for high-performance computing (HPC) involves the obvious tension between model complexity and computational resource availability that so often defines what is possible to achieve in practice. This paper describes our work to apply high-performance parallel computing to the domain of human brain neuroimaging. Section II describes the problem of source localization, the general methods involved, and the factors that affect source inverse solutions. Section III discusses general electromagnetic modeling of the human head. Section IV presents two algorithms we have developed, along with their verification, parallel performance, and reliability. Results and conclusions are given in Sections V and VI, respectively.

II. SOURCE LOCALIZATION

Modern dense-array EEG (dEEG) technology, such as the Geodesic Sensor Net [6] from Electrical Geodesics, Inc. (EGI) shown in Figure 1, can measure micro-volt potentials on the human scalp at up to 256 sensors every 1 msec or less.

EEG signals are the consequence of postsynaptic activities of neuronal cells. As seen in Figure 1(right), cortical neurons are arranged parallel to each other and point perpendicular to the cortical surface. It is this structural arrangement that allows currents from groups of thousands of neurons to accumulate and generate an *equivalent current dipole*. Therefore,

978-1-4799-0838-7/13/$31.00 ©2013 IEEE

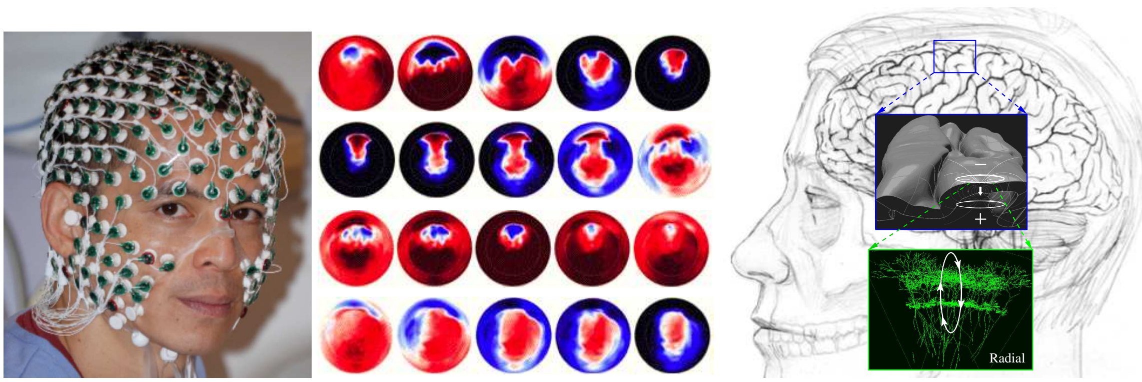

Fig. 1. EGI 256-channel Geodesic Sensor Net for dEEG recording (left). Topographical potential maps showing epileptic spike wave progression between 110-310 msec with 10 msec samples (center). Cortical neuron arrangement (right).

scalp potentials measured by dEEG can be modeled by the combined electrical potentials (called *lead fields*) produced by up to 10,000 or more cortex patches. Unfortunately, the scalp potentials are a linear superposition of all the *distributed source* lead fields and the individual EEG contributors (i.e., the distribute source dipoles), and must be disentangled to determine the dynamics of each brain region by solving source localization inverse problem. Two general approaches are used to solve the problem: the *parametric* approach and the *imaging* approach [7]–[10].

A. The Parametric Approach

The parametric approach is based on the assumption that the scalp EEG signal $\mathbf{\Phi_{EEG}}$ is generated by one or few current dipoles (less than 10) whose locations $\mathbf{r}_{qi}$ and moments $\mathbf{d}_{qi}$ (six parameters for each dipole) are unknown. These parameters are estimated by minimizing the residual energy $E(\mathbf{r}_{qi}, \mathbf{d}_{qi})$,

$$E(\mathbf{r}_{qi}, \mathbf{d}_{qi}) = ||\mathbf{\Phi_{EEG}}(\mathbf{r}) - \mathbf{\Phi_{model}}(\mathbf{r}, \mathbf{r}_{qi}, \mathbf{d}_{qi})||^2,$$

using a nonlinear search algorithm [11]–e.g., simplex or simulated annealing. Here, $\mathbf{\Phi_{model}}(\mathbf{r}, \mathbf{r}_{qi}, \mathbf{d}_{qi})$ is the lead field at sensor location $\mathbf{r}$ corresponding to a current dipole $\mathbf{d}_{qi}$ located at $\mathbf{r}_{qi}$. The search starts with a specified set of parameters and proceeds iteratively. This involves solving the forward problem at each step. Various strategies can be applied based on the number of dipoles, which parameters are fixed, and whether to consider the time-series of the EEG data [12]–[16].

The major concern about this approach is the required specification of the number of dipoles. Underestimating them causes biased results by the missing dipoles. Overestimating them causes the dipoles to fit any data and incur performance penalties due to increasing the dimensionality. Except in few cases (e.g. epileptic event), the accuracy of predicting the number of dipoles is questionable.

B. The Imaging Approach

To address the issue of identifying the optimal number of dipoles in the parametric approach, *distributed source models* are developed. In this approach, the primary current sources are assumed to be current dipoles distributed inside the brain. Since the position of each dipole is a potential location of a current source associated with a brain activity, the number of dipoles must be large enough to cover the cortex with an optimal resolution. The relationship between the current dipoles $\mathbf{J}$ and the potentials $\mathbf{\Phi}$ is defined by the linear forward equation,

$$\mathbf{\Phi} = \mathbf{K}\mathbf{J} + \epsilon, \tag{1}$$

where $\mathbf{\Phi} \in R^{(N\times1)}$ is a column vector gathering the potentials at N scalp electrodes, $\mathbf{J} \in R^{M\times3}$ is a M-vector of the magnitudes of the cortical dipoles, ϵ is a perturbation noise vector, and $\mathbf{K} \in R^{N\times M}$ is the lead field matrix (LFM). Every row in $\mathbf{K}$ is a lead field corresponding to a current dipole obtained by solving the forward problem. Given N-vector scalp EEG measurements $\mathbf{\Phi_{EEG}}$ at N electrodes and the LFM $\mathbf{K}$, the goal of the inverse problem is to invert Equation 1 to find a linear inverse operator W such that:

$$\hat{\mathbf{J}} = \mathbf{W}\mathbf{\Phi_{EEG}}, \tag{2}$$

where $\hat{\mathbf{J}}$ is an estimate of the current densities and $\mathbf{W}$ is the inverse linear operator. Since $\mathbf{J}$, and $\mathbf{\Phi}$ are linearly related, the inverse problem is reduced to finding a solution of a linear inverse problem for unknown magnitudes (vector $\mathbf{J}$). This is a well-known formulation for numerous image reconstruction problems. The problem is 1) under-determined which results in the existence of infinitely many solutions, and 2) ill-conditioned which results in unstable solutions in the presence of noise. To overcome the first issue, methods impose *a priori* constraints on the solution to select the most likely one. To overcome the second issue, methods take regularization schemes into account. Mathematically, the distributed method obtains the inverse solution by minimizing the data fitting term with an added regularization term in least-square sense,

$$F_\alpha(\mathbf{J}) = ||\mathbf{\Phi_{EEG}} - \mathbf{K}\mathbf{J}||^2 + \alpha||\mathbf{\Gamma}\mathbf{J}||^2$$

where $\alpha||\mathbf{\Gamma}\mathbf{J}||^2$ is the constraints and regularization term, and $||\mathbf{\Phi_{EEG}} - \mathbf{K}\mathbf{J}||^2$ is the data fitting term. The difference between different distributed inverse methods is in the choice and application of the constraints and the regularization scheme. From this formulation solving the inverse problem is achieved in four steps:

1) Defining the solution space by deciding *a priori* the number and locations of the distributed dipoles inside the brain.

2) Specifying the number and locations of electrodes where EEG signals are sampled.
3) Computing the lead field of the distributed dipoles at the electrode locations; mapping the solution space to the scalp space.
4) Applying an inverse algorithm to find estimates of the dipole moments that best describe a given EEG signal.

Based on these steps, we can identify and classify the factors that influence the source localization solution and require further research and exploration. The aim is to find optimal choices and direct the research to improve the most influential factors for each inverse algorithm. In the following three sections we discuss these choices in more detail, then outline our approach in studying their influence.

III. Electromagnetics Forward Model

Modeling the human head as a volume conductor defines the relationship between current source generators in the brain and the measured electrical potentials on the scalp. Given a volume conductor $\mathbf{\Omega}$ with boundary $\Gamma_{\mathbf{\Omega}}$, current sources within the volume induce electric and magnetic fields which can be calculated on the surface. If the conductivities σ and the current sources S are known, the electric and magnetic fields inside the volume are fully described by the quasi-static approximation of Maxwell's equations–Poisson equation,

$$\nabla \cdot \sigma(x, y, z) \nabla \phi(x, y, z) = S, \tag{3}$$

in Ω with no-flux Neumann boundary conditions on the scalp:

$$\sigma(\nabla \phi) \cdot n = 0. \tag{4}$$

Here, n is normal to Γ_Ω and $\sigma = \sigma_{ij}(x, y, z)$ is the conductivity tensor. The solution of Eq. 3 depends on the volume conduction properties, geometry and conductivity. A complete realistic volume conductor model that captures the fine details is not expected. However, specification of its configuration is a key factor in improving the source localization accuracy. The aim is to identify and rank the most influential factors on source localization.

A. The geometry factor

The simplest head model consists of a single homogeneous sphere [17]. It is far from reality given the significant difference between bone and fluid tissues. As a first improvement, three-shell models representing brain, skull and scalp are introduced [18]. They show good qualitative agreement with general EEG observations [19]. For further improvement, models that include Cerebrospinal fluid (CSF) and gray matter in four and five-shell models [20], [21] and analytic solutions that handle radial-to-tangential anisotropy are available [22], [23]. These models capture the major tissue layers, and their simple geometry allows for analytic solutions [24]. However, they have obvious limitations. The head tissues do not have uniform thickness and conductivities [25] [26], and the skull contains characteristics which are difficult to represent, such as sutures.

Structural imaging such as magnetic resonance imaging (MRI) and computational tomography (CT) provide images of anatomical details with resolution better than 1mm^3. These images can be segmented to a number of tissues where each is assumed to have uniform electrical properties. The quality and accuracy of the geometric model is directly related to the imaging modality and the quality of segmentation. MRI is sensitive to soft tissue, while CT is sensitive to bones. Forward models obtained from these images have better accuracy compared to spherical models [27], [28]. However, their computational complexity is significantly higher.

Several questions require answers and more exploration. How many tissues should be considered and which tissues? How do different geometric models affect source localization? How important is it to model geometric variations such as fissures and different types of bones? What is the required level of detail? The answers to these questions should be in the context of source localization. The influence of these factors impacts various algorithms differently. The importance of one factor can not be understood in isolation of the others. Therefore, answering these and similar questions must be done via simulations that leverage HPC.

B. The conductivity factor

Once the tissues are identified from the imaging modalities, their conductivity model and values must be assigned. Unfortunately, the conductivities of the head tissues are poorly known, especially for the skull. In general, the conductivity of a biological tissue is related to its concentration of fluid [29]. Tissues with higher fluid concentration are more conductive. Cell-free fluids such as CSF have a uniform and high conductivity [19], [30], [31], while compact bones have the lowest.

The brain consists of gray matter and white matter. Gray matter contains neurons' cell bodies and is accepted to be homogeneous and isotropic. White matter contains nerves that connect different parts of the cortex. The conductivity along the nerve is 9 times greater than in the perpendicular direction (anisotropic).

The skull conductivity has been poorly known, and the published data are not consistent. Skull bones can be classified according to the fluid concentration in their material into: compact bones having low fluid concentration and spongy bones having higher fluid concentration. Sutures are composed of materials that are highly rich with fluids. Therefore, and experiment confirmed, sutures are highly conductive, and spongy bones are more conductive than compact bones [32]–[35]. The cortical part has a layered structure consisting of a spongy middle layer sandwiched between two compact bones. Measurements show that the lower layer is more conductive than the upper layer and the middle layer is much more conductive than the outer layers [33]. In addition to variations in the bone types, structural variations within the skull such as openings and thin regions have a large impact on the effective conductivity of the skull. These holes and openings are filled with nerves and body fluids, which provide electrical current paths to pass through the skull and consequently increase the effective conductivity. The structural variations effect becomes significantly important in infants and young children, where the skull is not completely developed [36], [37].

The electrical properties of different head tissues are inhomogeneous and anisotropic. The skull as well as the scalp have a multilayer structure with each layer having different electrical properties. This structure can be either modeled as a multilayer structure with isotropic layers or it can be modeled as a single anisotropic tissue. How important is it to include these fine details in the model, and which features are most important? This question requires further analysis and investigations, so an efficient anisotropic solver is needed to enable thorough analysis. Section IV describes a new algorithm which accounts for anisotropic tissues.

C. Numerical methods

Realistic head models become standard for their improvement over spherical models. Boundary Element Method (BEM) solves the surface integral equations instead of the volume partial equation. This approach reduces the dimensionality to 2D which improves performance. However, it is restricted to handle only homogeneous, isotropic and closed tissue compartments. In contrast, Finite Difference Method (FDM) and Finite Element Method (FEM) are based on digitizing the whole volume into small volumetric elements. Consequently, various modeling properties such as inhomogeneity and anisotropy can be handled. The price for this flexibility is performance. The stiffness matrix becomes large and only iterative methods can be used [38], [39]. Iterative methods are relatively inefficient since they require repeated application for each source configuration; however, using the reciprocity theorem overcomes this [40], [41].

FEM is computationally more effecient than FDM due to the freedom in the choice of the computational points compared to the FDM's regularly fixed points. However, constructing a FEM mesh from MRI image is difficult and can be inaccurate due to the complex head geometry. With FDM, the regular cubed images map directly to the computational grid without any effort. The FDM is also accurate, reliable, able to handle non-uniform conduction characteristics and computationally efficient.

IV. Methods and Materials

In this section we describe our general approach to solving the source localization problem, the two primary computational hurdles that necessitate large-scale HPC solutions, and the two FDM solvers that enable new scientific study in head modeling and simulation.

Our approach towards source localization is based on the idea of providing generic LFMs (gLFM) that serve as generators of LFMs. A gLFM maps the amplitudes of generic dipolar sources to generic electrodes potentials. The generic distributed dipoles are placed at every voxel in the gray matter and the generic electrodes are placed at 1mm^3 inter-spacings on the scalp. Different gLFMs are constructed for different volume conduction characteristics. Once gLFMs are computed, many different LFMs can be sampled based on different constraints or resolution imposed on the sources (e.g. constraint dipoles on the cortex), the number and locations of the electrodes, and the volume conduction characteristics. This can be achieved efficiently by sampling from the rows and columns of the gLFM appropriately. The idea is to factor out the common and computationally intensive part of the analysis from the application of different inverse algorithms. Since different gLFMs captures different volume conduction parameters, the influence of these factors, number and distribution of electrodes, and different constraints on dipoles can be analyzed in a unified way using, for example, sensitivity analysis procedures.

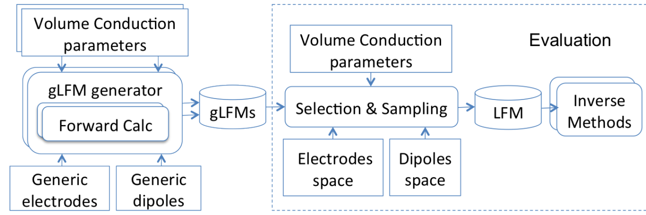

Fig. 2. Generic LFM generation using HPC, each gLFM maps a generic dipoles location to a generic electrodes locations for a given volume conduction characteristics. Once computed, different distributed dipoles algorithms or equivalent dipole parametric search can be applied and reapplied under different conditions and constraints in the evaluation.

In the application of the inverse algorithm, a gLFM is selected based on volume conduction characteristics shown in Figure 2. Then LFMs are sampled from the gLFM by choosing the appropriate rows corresponding to the electrodes map. In the case of distributed dipole models, the appropriate columns corresponding to imposing constraints on the sources are sampled as well. Then different distributed dipoles algorithms can be applied. In case of the parametric approach, the non-linear search proceeds on the columns of the LFM corresponding to the location of the dipoles and different orientations are considered by the linear superposition of lead fields corresponding to the three orthogonal directions weighted by the dipole orientation.

A. Generic lead field matrix generation.

Computing a gLFM at the highest resolution is computationally intensive. Assuming 5K electrodes and 500K dipole locations with 3 orthogonal orientations for each dipole. Each gLFM requires 1.5 million forward calculations which is not practical. However, using the reciprocity principle, the number of forward calculations is reduced significantly to the number of scalp electrodes which is 5K per gLFM. To perform global sensitivity analysis, at least 1000 gLFMs are required. This means 5 million forward calculations are required. Fortunately, all these calculations are independent and can be computed concurrently. In principle, assuming the availability of infinite resources, the time required to compute all these gLFMs is equal to the time required for a single forward solution. In practice the available resources are limited and consequently the performance of the forward solver becomes the limiting factor. Figure 2 shows the gLFM computation factored out from the evaluation analysis.

B. Conductivity inverse model.

The other crucial problem in the individualized head modeling is the determination of a subjects' unique internal head-tissue conductivities. One approach to find these values is the bounded Electrical Impedance Tomography (bEIT) method. In bEIT, low-frequency alternating currents are injected into the head through pairs of electrodes attached to the scalp. Then the response is measured on the other electrodes. Once an appropriate objective function describing the difference between the measured scalp potentials, V, and the predicted potentials ϕ^p, is defined (e.g., least square norm), a search for the global minimum is undertaken using nonlinear optimization algorithms (e.g., simulated annealing [42], [43] or simplex search). Using either optimization method, the search for the optimal conductivities requires a large number of forward calculations, in the order of 3K for a single current injection pair. Typically, we consider 60 pairs, which require 200K forward calculations. Since Poisson's equation is non-linear regarding the conductivities, the reciprocity principle cannot be applied in this case. Three levels of parallelism are applied in these calculations. At the highest level, current injection pairs are processed via a cluster job scheduler. Each injection pair performs a simulated annealing optimization in parallel using an MPI application. Finally, the independent forward solvers run on GPUs using CUDA or shared memory using OpenMP.

In both gLFM generation and conductivity optimization, a large number of forward solutions is required. Therefore, any forward solver must efficient, robust and accurate. In the following we present two FDM algorithms to solve Poisson's equation (Eq. 3). The first is limited to isotropic conductivity of the tissues and based on the alternating direction implicit (ADI) method. The second can handle anisotropic properties of tissues. A parallel implementation of both algorithms in shared memory and GPU architecture is described.

C. ADI algorithm

The ADI method finds the solution of Eq. 3 as the steady state of the appropriate evolution problem,

$$\frac{\partial \phi}{\partial t} + \nabla \cdot \sigma(x,y,z)\nabla\phi(x,y,z) = S. \tag{5}$$

At steady state, the time derivative is zero and the solution corresponds to the original problem. At every iteration step the spatial operator is split into the sum of three 1D operators, which are evaluated alternatively at each substep. For example, the difference equations in x direction is given as [44],

$$\frac{\phi_i^{n+1} - \frac{1}{3}(\phi_i^n + \phi_j^n + \phi_k^n)}{\tau} + \delta_x\phi_i^{n+1} + \delta_y\phi_j^n + \delta_z\phi_k^n = S, \tag{6}$$

where τ is a time step and $\delta_{x,y,z}$ is the appropriate 1D second order spatial difference operator. The finite-difference scheme is used over the solution domain by using a rectangular grid with spatial spacings of h_x, h_y, h_z in the x, y, z directions, and τ in time. Using the notation $x_i = ih_x$, $y_j = jh_y$, $z_k = kh_z$ and $t_n = n\tau$ for integer values of i, j, k, and n, the electrical potential at a grid point, (i,j,k), at time, t_n, is written as $\phi_{ijk}^n = \phi(x_i, y_j, z_k; t_n)$. The notation ϕ_q^n means the solution along the direction q in the time step n. Such a scheme is accurate to $O(\tau^2 + \Delta x^2 + \Delta y^2 + \Delta z^2)$. In contrast with the classic ADI method, the multi-component ADI does not require the operators to be commutative. In addition, it uses the regularization (averaging) for evaluation of the variable at the previous instant of time.

The ADI algorithm consists of a time iteration loop in which each time step is split into three substeps. In each sub-step, a tridiagonal system of equations is solved along x, y and z directions. For instance, in the first sub-step the spatial operator acts only on the x direction. So, all N_yN_z equations along the $x-$direction are independent and can be solved concurrently. Similarly, in the second sub-step, all N_xN_y equations along the $y-$direction and, in the third sub-step, all N_xN_y along the $z-$directions are independent and can be solved concurrently. At the end of each sub-step all equations must be solved before proceeding to the next sub-step. The parallel algorithm pseudo-code is shown below:

Algorithm 1: ADI parallel algorithm

while *not terminate* **do**
    Solve N_yN_z tridiagonal systems concurrently;
    Barrier
    Solve N_xN_z tridiagonal systems concurrently;
    Barrier
    Solve N_xN_y tridiagonal systems concurrently;
    Barrier
    Update(terminate)

Implementation of the parallel algorithm in shared memory architecture is straight forward, where the time loop runs sequentially and then in each sub-step all OpenMP threads cooperate in solving the independent tridiagonal system of equations concurrently. Similarly, within a GPU architecture, the time loop runs sequentially on the host, and a grid of blocks of threads is executed to solve the tridiagonal systems of equations on the device in each sub-step. Each thread solves a tridiagonal system of equations. The performance of the GPU code is mainly limited by the global memory access. Threads are coalesced in accessing global memory when solving in the y and z-directions. However, they are not coalesced when solving in x-direction. We used shared memory and intermediate computations which improved performance when computing in the x direction.

D. VAI algorithm

In the 3D anisotropic case, we use the Vector Additive Implicit (VAI) algorithm as introduced in [45]. In this algorithm, a 13-point stencil is used to approximate the differential operator and order the variables. It includes two diagonally-adjusted cells with one common symmetry point as shown in Figure 3.

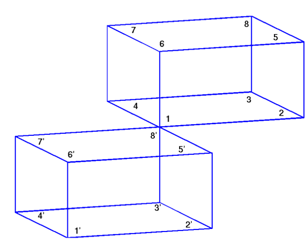

Fig. 3. The 13-point VAI stencil

For ordering variables, the calculation domain is split into a set of rectangular cells ordered akin to a 3D checkerboard. A subset of uniformly colored cells is considered each time. Each cell has eight corners (grid points). Each corner belongs to two adjusted cells. Eight components of the approximate solution correspond to the eight points of each cell. Two components of an approximate solution are related to each grid point. The stencil and ordering of variables are adapted for a two-component vector-additive method for solving the linear system, $Ay = f$, with $A = A_1 + A_2$, with the form,

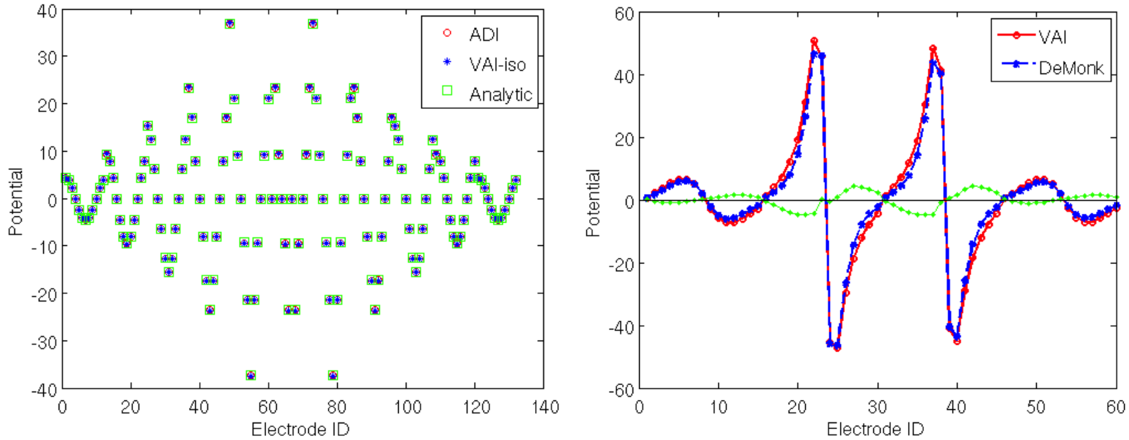

Fig. 4. ADI and VAI-isotropic are compared to the analytical solution for a 4-shell isotropic spherical model (left). VAI in anisotropic setting with tangential to radial conductivity of 10 compared with anisotropic sphere [46].

$$\frac{y_1^{n+1} - \tilde{y}_1}{\tau} + A_1 y_1^{n+1} + P_{21} A_2 y_2^n = f,$$
$$\frac{y_2^{n+1} - \tilde{y}_2}{\tau} + P_{12} A_1 y_1^n + A_2 \tilde{y}_2^{\,n+1} = f,$$
$$\tilde{y}_1 = (y_1^n + P_{12} y_2^n)/2,$$
$$\tilde{y}_2 = (P_{12} y_2^n + y_2^n)/2,$$

where τ is iterative parameter, P_{12} and P_{21} are permutation matrices. The matrices A_1 and A_2 are in block-diagonal form with 8×8 diagonal blocks. These matrices are composed from coefficients of finite-difference scheme and they are complemented parts of the finite-difference operator cells of the stencil.

The structure of the VAI method is similar to the implicit block Jacobi method with a preconditioner in the form of a block-diagonal matrix with 8×8 diagonal blocks. Because each block can be processed independently, this approach is highly parallelizable. In a shared memory architecture, the iterative loop runs sequentially and then at each iteration step the OpenMP threads cooperate in computing the 8×8 blocks until the termination condition is satisfied. Similarly, in our GPU implementation, the iterative loop runs on the host and at each iteration step, a grid of blocks of threads is executed on the GPU, where each thread performs the computation of one 8×8 block. At the end of each iteration step, the host checks the convergence criteria. Since all blocks are homogenous and have the same size of a multiple of four, accessing the global memory is efficient when using float4 and int4 CUDA data types.

E. *Reciprocity*

The idea of the reciprocity theorem is that the electric field throughout a volume conductor caused by injecting a unit current between two electrodes on the scalp can be used to determine the potential difference between these two electrodes caused by current dipole sources in the volume. This theorem reduces the calculation of the potential difference between two electrodes on the scalp caused by any dipole at any location and with any orientation to one forward calculation. This will reduce the required number of forward calculations to equal the number of scalp sensors. Mathematically, the potential difference between a recording electrode A and the reference electrode R on the scalp due to a dipole source at location $\mathbf{r}$ can be written as,

$$\Phi_A - \Phi_R = \frac{\mathbf{E}_{AR}(\mathbf{r}) \cdot \mathbf{d}}{I_{AR}},$$

where $\mathbf{d}$ is the dipole moment, $\mathbf{E}_{AR}(\mathbf{r}) = -\nabla\Phi(\mathbf{r})$ is the electric field at location $\mathbf{r}$ caused by I_{AR}, and I_{AR} is the current flowing between the source and sink electrodes.

V. Results

Evaluating the influence of the main factors that affect the accuracy of source localization and the extraction of conductivities of the head tissues through the bEIT technique requires a forward solver that meets three main requirements: 1) accuracy, in that it solves the Poisson equation accurately with complex geometries, 2) efficiency, in that it allows conducting these studies in a practical amount of time, and 3) reliability, in the sense that it can handle several volume conduction characteristics, such as anisotropy and fine details of anatomical structure. In this section we evaluate these aspects of the ADI and VAI forward solvers and demonstrate that they meet these requirements.

A. *Verification*

The ADI and VAI method implementations should first be verified with respect to a known solution. The source localization field has long used concentric k-shell spherical models ($k = 3, 4$) as a theoretical standard of reference (each shell represents a head tissue), because analytical solutions are known for the isotropic and anisotropic case [22], [47]. Using a 4-sphere testcase with 200 x 200 x 200 voxels, Figure 4 shows the perfect correspondence between the theoretical isotropic, ADI and VAI-isotropic results for a set of shell conductivities. The plot compares potentials for each of the 132 electrodes. We obtained the same correspondence by experimenting with tissue conductivities, different radii and different current source-sink locations.

Analytical solutions for anisotropic spherical models [47] are also available for VAI verification. These results are shown

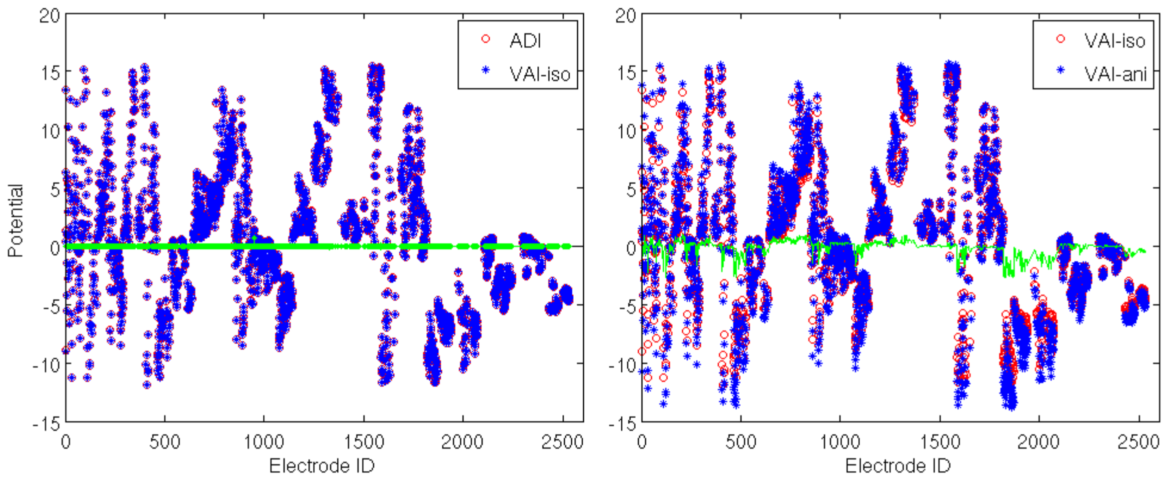

Fig. 5. Cross verification of ADI and VAI-isotropic results of generic LFM computation. The figure shows a random column from ADI LFM compared to the corresponding column from VAI-isotropic LFM. Reciprocity principle is used in the computation of both matrices(left). The potentials for one column of the LFMs generated using ADI (isotropic) and VAI (anisotropic) are shown(right). Differences between values are plotted in green.

in Figure 4 (right). The accuracy with respect to the spherical model in both cases is very good, lending strong confirmation that the algorithm is working properly.

Of course, nobody's head is shaped like a 4-shell sphere. However, verifying the algorithms with real human heads is more challenging. The colin27 MRI dataset was segmented at 2mm^3 and 1mm^3 resolutions into five tissue: scalp, skull, CSF, gray matter, and white matter. We used cross verification of ADI and VAI in isotropic setting to compute a LFM for each resolution case for known conductivities and current source. Although the numerical methods are different, we expect agreement for the isotropic case, as is verified in Figure 5.

Using the computed LFMs, we compared isotropic versus anisotropic methods by taking one column of each LFM and plotting the two projections of the same activated dipole. Figure 5 (right) shows the potential differences at the sensor locations. Our future work will answer the question of how the potential differences affect source localization accuracy when considering the number of scalp electrodes, the solution space configuration, and the conductivities of the tissues.

B. Computational Performance

The previous section shows that the ADI and VAI solvers produce accurate results, but computational performance is also important to enable our large-scale computational requirements. A Matlab version of both the ADI and VAI forward solvers takes several hours to compute a single solution for a 1mm^3 head model, which prohibits the computation of even a single LFM. Figure 6 gives performance results for our OpenMP and CUDA versions of the ADI and VAI for colin27 at 1mm^3 in terms of average iteration time.

Both ADI and VAI are iterative solvers and will run at a minimum 400 iterations before convergence with a 1mm^3 head model. Averages of 500 iterations for ADI and 1000 for VAI are common in our experience. Increasing the cores for OpenMP up to eight continues to deliver performance improvement on the compute nodes we tested. It is clear that the GPUs deliver the best performance returns. While this is true, many nodes don't have GPUs. Thus, both OpenMP and CUDA implementations are important. The memory footprint for each solver is about 800MB for the colin27 1mm^3, making both of them appropriate for most available GPUs.

Forward solvers are the core computational components for the conductivity inverse and LFM calculations. The conductivity inverse problem will need to process the bEIT measurements for up to 64 current injection pairs in the general case. Depending on the number of conductivity unknowns, each conductivity search for a single pair will require many thousands of forward solutions to be generated. Simulated annealing is currently used as the optimization strategy [43] and our parallel implementation will support up to twelve simultaneous forward solves. Clearly, conductivity results for each pair can also be done in parallel. The results from all pairs are then analyzed to determine final tissue conductivity estimates. The total computational requirements are prodigious, requiring over 180,000 forward solutions.

Computing gLFMs for all current dipoles is computationally intensive. Because a gLFM requires capturing scalp potentials corresponding to dipoles at any position in the gray matter and in any orientation, it is necessary to calculate the potentials corresponding to the three orthogonal x-, y-, and z-orientations for each dipole location. Then the potential corresponding to any orientation can be constructed by superposition of the potentials corresponding to the three basis vectors. However, by using the reciprocity principle, we only need N_e forward solutions to construct such a gLFM, where N_e is the number of electrodes. Each forward solution provides the potentials at an electrode corresponding to all dipoles in the gray matter. We created an isotropic gLFM and an anisotropic gLFM for colin27 based on 1,925 generic recording electrodes. This required 1,925 forward solutions to be computed for each gLFM, by placing a current source at each electrode and the sink at a common reference electrode and calculating the potentials at that electrode corresponding to every dipole location. Thus, each LFM is 2.1e6 x 1,925 in size.

From a computational viewpoint, the LFM generation is fully parallelized since the computation of every dipole forward solution (or when using reciprocity, the computation

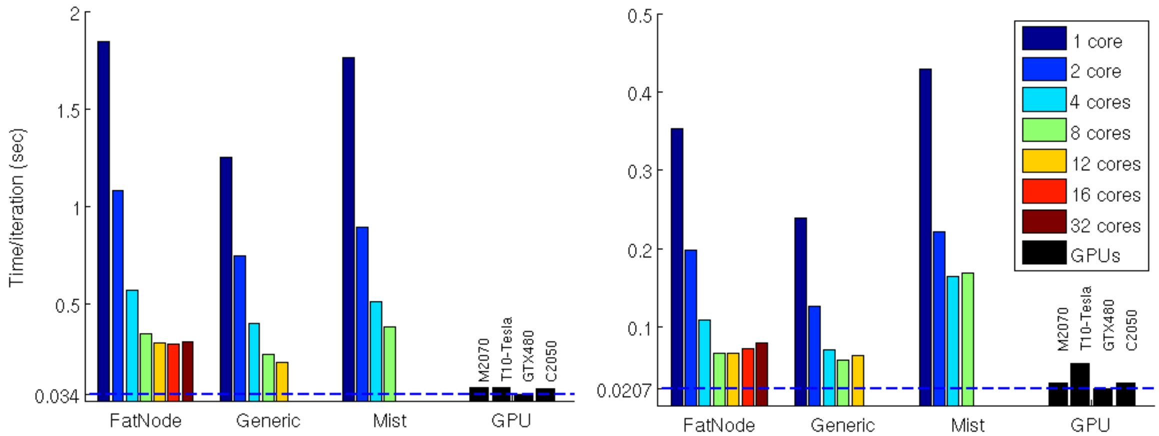

Fig. 6. Single-node performance for 1 iteration of the ADI (left) and VAI (right) solvers on ACISS and Mist clusters(see Fig. 7). VAI completes a single iteration in ∼20% the time of an ADI iteration for OpenMP and ∼60% for CUDA. Multi-node MPI performance is not shown, but forward solutions are independent of each other, so scaling is good for both gLFM and conductivity inverse problems when using either OpenMP or GPUs.

of the potentials corresponding to all dipoles at every electrodes) is independent, so the application is quite scalable. For instance, a run on the ACISS machine at the University of Oregon which utilized 98 GPUs in an ADI LFM calculation was calculated in approximately 11 minutes.

C. Reliability

Both ADI and VAI solvers are reliable in the sense that anatomical structure of the geometric model of the human head can easily be captured without any pre-processing or mesh generation of the structural MRI and/or CT images. Both solvers handle accurately any fine details of geometric features such as skull holes at the available image resolution (currently at 1mm^3). Once higher resolution images become available, no pre-processing or modification is required on these solvers. This flexibility is important, as it allows for studying the influence of structural details on the source localization solution. Further, in both solvers the conductivity values can be assigned at the voxel level which allows differentiating the electrical properties at 1mm^3 scale. This is important if we wanted to consider the influence of fine-detail characteristics, such as sutures. Also, placing dipoles anywhere inside the brain is a matter of placing a current source and sink separated by a voxel. In addition to this flexibility, the VAI algorithm allows for assignment of the anisotropic conductivity tensor at the voxel level. Of course, the price of all this flexibility is the computational performance. However, with access to sufficient computational resources this scientific workflow could potentially scale to a level that would enable source localization for a large number of individualized head models.

VI. Conclusion and Future Directions

This paper presents four main contributions to the neuroscience domain: 1) We identified and classified the main factors influencing accuracy of source localization solutions and which of these factors require further research and investigations, 2) We provided an approach to study the effect of these factors in a unified manner using different inverse approaches or algorithms; our approach is based on factoring out the common and computationally intensive part from the analysis,

Mist (University of Oregon - UO): 24 Dell 1950 (2x 2.33 GHz quadcore Intel Xeon w/ 16GB), 192 total cores; 2 NVIDIA Tesla S1070 (4x Tesla GPU), 8 total GPUs
Aciss-fatnodes (UO): 16 compute nodes (4x 2.27GHz 8-core Intel X7560 CPUs w/ 384GB DDR3 RAM), 512 total cores
Aciss-generic (UO): 128 compute nodes (2x 2.67GHz 6-core Intel X5650 w/ 72GB DDR3 memory), 1536 total cores
Aciss-gpunodes (UO): 52 compute nodes (2x 2.67GHz 6-core Intel X5650 w/ 72GB DDR3 memory), 624 total cores; 3 NVIDIA Telsa M2070 GPU, 156 GPUs total

Fig. 7. Parallel computer platforms used for experiments.

which allows for the application of different inverse algorithms and different constraints, 3) We demonstrated that HPC enables the creation of such constructs (gLFM) at high resolution and detail, and 4) We provided two accurate, efficient, and reliable FDM-based forward solvers parallelized using OpenMP in shared memory and CUDA on GPUs.

References

[1] D. H. Fender, “Source localization of brain electrical activity,” pp. vol. 1, pp. 355403, 1987.

[2] S. van den Broek, F. Reinders, M. Donderwinkel, and M. Peters, “Volume conduction effects in eeg and meg,” *Electroencephalography and Clinical Neurophysiology*, vol. 106, no. 6, pp. 522 – 534, 1998.

[3] W. K., S. G., G. L., C. JF, and F. A., “Effects of dipole position, orientation and noise on the accuracy of eeg source localization,” *Biomedical Engineering Online*, 2003.

[4] C. H. Wolters, A. Anwander, X. Tricoche, D. Weinstein, M. A. Koch, and R. S. MacLeod, “Influence of tissue conductivity anisotropy on EEG/MEG field and return current computation in a realistic head model: a simulation and visualization study using high-resolution finite element modeling,” *Neuroimage*, vol. 30, no. 3, pp. 813–826, Apr 2006.

[5] C.-G. Benar and J. Gotman, “Non-uniform spatial sampling in eeg source analysis,” in *Engineering in Medicine and Biology Society, 2001. Proceedings of the 23rd Annual International Conference of the IEEE*, vol. 1, pp. 903–905 vol.1.

[6] D. Tucker, “Spatial sampling of head electrical fields: the geodesic sensor net,” *Electroencephalography and Clinical Neurophysiology*, vol. 87, no. 3, pp. 154–163, 1993.

[7] F. Darvas, D. Pantazis, E. KucukaltunYildirim, and R. Leahy, “Mapping human brain function with meg and eeg: methods and validation,” *NeuroImage*, vol. 23, pp. S289–S299, 2004, darvas2004.pdf.

[8] S. Baillet, J. Mosher, and L. R.M., "Electromagnetic brain mapping," *IEEE Signal Processing Magazine*, vol. 18, no. 6, pp. 14–30, 2001.

[9] R. Grech, T. Cassar, J. Muscat, K. P. Camilleri, S. G. Fabri1, M. Zervakis, P. Xanthopoulos, V. Sakkalis, and B. Vanrumste, "Review on solving the inverse problem in eeg source analysis," *NeuroEngineering and Rehabilitation*, vol. 5, no. 25, pp. 1–33, November 2008, http://www.jneuroengrehab.com/content/5/1/25.

[10] R. Pascual-Marqui, "Review of methods for solving the eeg inverse problem," *International Journal of Bioelectromagnetism*, vol. 1, pp. 75–6, 1999.

[11] K. Uutela, M. Hamalainen, and R. Salmelin, "Global optimization in the localization of neuromagnetic sources." *IEEE Trans Biomed Eng*, vol. 45, pp. 716–23, 1998.

[12] M. Scherg and D. Von Cramon, "Evoked dipole source potentials of the human auditory cortex." *Electroencephalogr Clin Neurophysiol*, vol. 65, pp. 344–60, 1986.

[13] J. Foxe, M. McCourt, and D. Javitt, "Right hemisphere control of visuospatial attention: line-bisection judgments evaluated with high-density electrical mapping and source analysis," *Neuroimage*, vol. 19, pp. 710–26, 2003.

[14] M. Scherg, T. Bast, and P. Berg, "Multiple source analysis of interictal spikes: goals, requirements, and clinical value (review)," *J Clin Neurophysiol*, vol. 16, pp. 214–24, 1999.

[15] A. Rodriguez-Rivera, B. Van Veen, and R. Wakai, "Statistical performance analysis of signal variance–based dipole models for meg/eeg source localization and detection," *IEEE Transactions on Biomedical Engineering*, vol. 50, no. 2, pp. 137–149, 2003.

[16] J. De Munck, "The estimation of time varying dipoles on the basis of evoked potentials," *Electroencephalography and Clinical Neurophysiology*, vol. 77, p. 156, 1990.

[17] E. Frank, "Electric potential produced by two point current sources in a homogeneous conduction sphere," *Applied Physics*, vol. 23, no. 11, pp. 1225–1228, 1952.

[18] S. Rush and D. Driscoll, "current distribution in the brain from surface electrodes," *Anesth. Analg.*, vol. 47, no. 6, pp. 717–723, 1968.

[19] P. Nunez, *Electric Fields of the Brain: The neurophysics of the EEG*. Oxford University Press, 2006.

[20] J. de Munck, B. van Dijk, and H. Spekreijse, "Mathematical dipoles are adequate to describe realistic generators of human brain activity," *IEEE Trans. Biomed. Eng.*, vol. 35, no. 11, pp. 960–966, November 1988.

[21] J. de Munck and M. Peters, "A fast method to compute the potential in the multisphere model," *IEEE Transactions on Biomedical Engineering*, vol. 40, no. 11, pp. 1166–1174, 1993.

[22] T. Ferree, K. Eriksen, and D. Tucker, "Regional head tissue conductivity estimation for improved eeg analysis," *IEEE Trans. Biomed. Eng.*, vol. 47, no. 12, pp. 1584–92, 2000.

[23] R. Uitert, D. Weinstein, and C. Johnson, "Can a spherical model substitute for a realistic head model in forward and inverse meg simulations?" in *13th Int. Conf. on Biomagnetism,*, 2002, pp. 798–800.

[24] J. Ary, S. Klein, and D. Fender, "Location of sources of evoked scalp potentials: Corrections for skull and scalp thicknesses," *IEEE Transactions on Biomedical Engineering*, vol. 28, no. 6, pp. 447–452, 1981.

[25] R. Uitert, D. Weinstein, and C. Johnson, "Volume currents in forward and inverse magnetoencephalographic simulations using realistic head models," *Ann. Biomed. Eng.*, vol. 31, pp. 21–31, 2003.

[26] C. Wolters, A. Anwander, X. Tricoche, D. Weinstein, M. Koch, and R. MacLeod, "Influence of tissue conductivity anisotropy on eeg/meg field and return current computation in a realistic head model: A simulation and visualization study using high-resolution finite element modeling," *NeuroImage*, vol. 3, pp. 813–826, 2006.

[27] G. Huiskamp, M. Vroeijenstijn, R. van Dijk, G. Wieneke, and A. Huffelen, "The need for correct realistic geometry in the inverse eeg problem," *IEEE Trans. Biomed. Eng.*, vol. 46, pp. 1281–1287, Nov 1999.

[28] B. Yvert, O. Bertrand, M. Thvenet, J. F. Echallier, and J. Pernier, "A systematic evaluation of the spherical model accuracy in eeg dipole localization," *Electroencephalogr. Clin. Neurophysiol.*, vol. 102, pp. 452–459, 1997.

[29] J. Kosterich, K. Foster, and S. Pollack, "Dielectric properties of fluid saturated bone: The effect of variation in conductivity of immersion fluid," *IEEE Trans. Biomed. Engr.*, vol. 31, pp. 369–375, 1984.

[30] P. Nicholson, "Specific impedance of cerebral white matter," *Experimental Neurology*, vol. 13, pp. 386–401, 1965.

[31] D. Gullmar, J. Haueisen, M. Eiselt, F. Giessler, L. Flemming, A. Anwander, T. Knosche, C. Wolters, M. Dumpelmann, D. Tuch, and J. Reichenbach, "Influence of anisotropic conductivity on eeg source reconstruction: Investigations in a rabbit model," *IEEE Trans. Biomed. Eng.*, vol. 53, pp. 1841–1850, 2006.

[32] S. Law, "Thickness and resistivity variations over the upper surface of the human skull," *Brain Topography*, vol. 6, pp. 99–109, 1993.

[33] M. Akhtari, H. Bryant, A. Mamelak, E. Flynn, L. Heller, J. Shih, M. Mandelkern, A. Matlachov, D. Ranken, E. Best, M. DiMauro, R. Lee, and W. Sutherling, "Conductivities of three-layer live human skull," *Brain Topography*, vol. 14, no. 3, pp. 151–167, 2002.

[34] M. Akhtari, H. Bryant, D. Emin, W. Merrifield, A. Mamelak, E. Flynn, J. Shih, M. Mandelkern, A. Matlachov, D. Ranken, E. Best, M. DiMauro, R. Lee, and W. Sutherling, "A model for frequency dependence of conductivities of the live human skull," *Brain Topography*, vol. 16, no. 1, 2003.

[35] C. Tang, F. You, G. Cheng, D. Gao, F. Fu, and G. Yang, "Correlation between structure and resistivity variations of the live human skull," *IEEE Transactions on Biomedical Engineering*, vol. 55, no. 9, pp. 2286–2292, Sept 2008.

[36] L. Flemming, Y. Wang, A. Caprihan, M. Eiselt, J. Haueisen, and Y. Okada, "Evaluation of the distortion of eeg signals caused by a hole in the skull mimicking the fontanel in the skull of human neonates," *Clin. Neurophysiol*, vol. 116, no. 5, pp. 1141–1152, 1984.

[37] L. Jing, "Effects of holes on eeg forward solutions using a realistic geometry head model," *Journal of neural engineering*, vol. 4, no. 3, pp. 197–204, 2007.

[38] M. Mohr and B. Vanrumste, "Comparing iterative solvers for linear systems associated with the finite difference discretisation of the forward problem in electro-encephalographic source analysis," *Med. Biol. Eng. Comput.*, vol. 41, pp. 75–84, 2003.

[39] R. Hoekema, K. Venner, J. Struijk, and J. Holsheimer, "Multigrid solution of the potential field in modeling electrical nerve stimulation," *Computers and Biomedical Research*, vol. 31, pp. 348–362, 1998.

[40] S. Rush and D. Driscoll, "Eeg electrode sensitivity – an application of reciprocity," *IEEE Trans. Biom. Eng.*, vol. 16, no. 1, pp. 15–22, 1969.

[41] H. Hallez, B. Vanrumste, P. Van Hese, Y. D'Asseler, I. Lemahieu, and R. Van de Walle, "A finite difference method with reciprocity used to incorporate anisotropy in electroencephalogram dipole source localization," *Physics in Medicine and Biology*, vol. 50, pp. 3787–3806, 2005.

[42] S. Kirkpatrick, C. Gelatt, and M. Vecchi, "Optimization by simulated annealing," *Science, Number 4598, 13 May 1983*, vol. 220, 4598, pp. 671–680, 1983.

[43] A. Salman, A. Malony, S. Turovets, and D. Tucker, "Use of Parallel Simulated Annealing for Computational Modeling of Human Head Conductivity," in *International Conference on Computational Science (ICCS 2007)*, Y. S. et al., Ed., vol. LNCS 4487. Springer-Verlag, 2007, pp. 86–93, (Best paper award).

[44] J. Douglas, Jim, "Alternating direction methods for three space variables," *Numerische Mathematik*, vol. 4, pp. 41–63, 1962. [Online]. Available: http://dx.doi.org/10.1007/BF01386295

[45] V. Volkov, A. Zherdetsky, S. Turovets, and A. Malony, "A 3d vector-additive iterative solver for the anisotropic inhomogeneous poisson equation in the forward eeg problem," in *International Conference on Computational Science (1)*, ser. Lecture Notes in Computer Science, vol. 5544. Springer, 2009, pp. 511–520.

[46] J. De Munck, "The potential distribution in a layered anisotropic spheroidal volume conductor," *Appl. Phys.*, vol. 64, no. 2, pp. 464–470, 1988.

[47] J. de Munck, T. Faes, A. Hermans, and R. M. Heethaar, "A parametric method to resolve the ill-posed nature of the eit reconstruction problem: a simulation study," *Annals of the New York Academy of Sciences*, vol. 873, pp. 440–453, 1999.