

Use of Parallel Simulated Annealing for Computational Modeling of Human Head Conductivity

Adnan Salman¹, Allen Malony¹, Sergei Turovets¹, and Don Tucker²

¹ NeuroInformatics Center, 5219 University of Oregon, Eugene, OR 97403, USA

² Electrical Geodesic, Inc., 1600 Millrace Dr, Eugene, OR 97403, USA

Abstract. We present a parallel computational environment used to determine conductivity properties of human head tissues when the effects of skull inhomogeneities are modeled. The environment employs a parallel simulated annealing algorithm to overcome poor convergence rates of the simplex method for larger numbers of head tissues required for accurate modeling of electromagnetic dynamics of brain function. To properly account for skull inhomogeneities, parcellation of skull parts is necessary. The multi-level parallel simulated annealing algorithm is described and performance results presented. Significant improvements in both convergence rate and speedup are achieved. The simulated annealing algorithm was successful in extracting conductivity values for up to thirteen head tissues without showing computational deficiency.

1 Introduction

Accurate knowledge of the geometry of human head tissues and their conductivities plays a fundamental role in high-resolution neuroimaging. This knowledge is necessary and essential to create computational models of the electromagnetic characteristics of the human head, allowing precise monitoring of brain dynamics in both space and time. Dense-array electroencephalography (EEG) up to 256 sensor channels can be projected to cortex locations using computational head models (this is known as the *inverse* or *source mapping* problem), but tradeoffs concerning simulation complexity and solution accuracy and speed will determine model precision. While tissue geometry is observable through structural MRI or CT measures, the regional conductivities of the human head tissues are largely unknown. Because the skull is the most resistive tissue, the lack of accurate skull conductivity estimates is particularly problematic given also the developmental variations in the human skull from infancy through adolescence. Without an accurate forward model specifying the volume conduction from cortex to scalp, even advanced inverse efforts cannot achieve precision with EEG data as the error of source localization due to the conductivity uncertainty may reach a few centimeters [1].

Our group has developed a high-performance modeling framework for simulation of human head electromagnetics based on finite difference methods (FDM) with realistic head shape. In our 2005 ICCS paper [2], we combined FDM modeling with a parameterized electrical impedance tomography (EIT) measurement procedure and a downhill simplex algorithm to extract three- and four-tissue conductivities in simulations with reasonable accuracy. We parallelized the simplex search using a multi-start technique and good speedup was achieved. However, two factors argue for a new approach. First,




Fig. 1. The Geodesic Sensor Net, current injection between selected electrode pairs and simultaneous acquisition of the return potentials from the sensor array

experimental studies report that skull is anisotropic and highly heterogeneous, and can not be modeled as a uniform tissue[6]. Second, if we increased the number of segmented tissues to more accurately model skull inhomogeneities, the viability of the simplex method diminishes quickly. Six tissues are enough to reduce the probability of convergence to less than 10% and of finding optimal solutions to much less.

This paper reports on our work replacing the simplex method with a simulated annealing algorithm that can discover conductivities up to thirteen segmented tissues. We have parallelized the conductivity search problem and evaluated its convergence and scalability attributes on the SDSC DataStar machine. We have also prototyped an exponential parallelization strategy and tested its scalability potential. The sections that follow provide more background on the problem domain, describe the computational design, and present our performance results. The paper concludes with future work.

2 Methods

To non-invasively determine head tissue conductivities, a tomographic-based search procedure must be used to optimally match EEG data measured from a subject's head to simulated solutions of the electromagnetic response. Small currents are injected into the head and electrical response measured at dense-array sensors on the scalp (see Figure 1). Given these measurements, we can search for conductivity solutions using a FDM model of the subject's head. Once a set of conductivities is chosen, "forward" calculations simulate the electromagnetic effects when current is injected at the chosen locations. The simulated electrical potentials are computed and compared to those measured. Based on the error, the "inverse" calculation attempts to improve the next choice of conductivity parameters. Together, the forward and inverse calculations define the (indirect) conductivity modeling problem. A complete formal description of the forward and inverse problems can be found in our early ICCS paper [2]. Here we provide only a brief review.

2.1 Forward and Inverse Problems

The electrical forward problem can be stated as follows: given the positions and magnitudes of current sources, as well as geometry and electrical conductivity of the head

volume Ω calculate the distribution of the electrical potential on the surface of the head (scalp) Γ_Ω . Mathematically, it means solving the linear Poisson equation [3]:

$$\nabla \cdot \sigma(x, y, z) \nabla \phi(x, y, z) = S, \quad (1)$$

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

$$\sigma(\nabla \phi) \cdot n = 0 \quad , \quad (2)$$

on Γ_Ω . Here n is the normal to Γ_Ω , $\sigma = \sigma_{ij}(x, y, z)$ is an inhomogeneous tensor of the head tissues conductivity and S is the source current. We have built a FD forward problem solver for Eq. (1) and (2) based on the multi-component alternating directions implicit (ADI) algorithm [4]. It is a generalization of the classic ADI algorithm [5].

The inverse problem for the electrical imaging modality has the general tomographic structure. From the assumed distribution of the head tissue conductivities, σ_{ij} , and the given injection current configuration, S , it is possible to predict the set of potential measurement values, ϕ^p , given a forward model F (Eq. (1), (2)), as the nonlinear functional [5]:

$$\phi^p = F(\sigma_{ij}(x, y, z)). \quad (3)$$

Then an appropriate objective function is defined, which describes the difference between the measured, V , and predicted data, ϕ^p , and a search for the global minimum is undertaken using advanced nonlinear optimization algorithms. In this paper, we used the simple least square error norm:

$$E = \left(\frac{1}{N} \sum_{i=1}^N (\phi_i^p - V_i)^2 \right)^{1/2}, \quad (4)$$

where N is a total number of the measuring electrodes. With the constraints imposed by the segmented MRI data, one needs to know only the average regional conductivities of a few tissues, for example, scalp, skull, cerebrospinal fluid (CSF), and brain, which significantly reduces the dimensionality of the parameter space in the inverse search.

In our earlier work, to solve the nonlinear optimization problem in Eq. (4), we employed the downhill simplex method of Nelder and Mead. Our observation was that simplex search performs well when the number of parameters is few (three or four parameters), while it fails completely when the number of parameters is larger than six. To avoid the local minima in the simplex search, we used a statistical approach. The inverse procedure was repeated for hundreds sets of conductivity guesses from appropriate physiological intervals, and then the solutions closest to the global minimum solutions were selected using a simple error threshold criteria $E < E_{threshold}$.

2.2 Skull Inhomogeneities

The human skull includes eight cranial bones and fourteen facial bones in addition to the sutures. Figure 2 shows how the head bones are parcellated based on anatomical features. It is expected that anatomically different parts of the skull have different conductivities values and experiments bear this out [6]. For instance, conductivity of trilayer




Fig. 2. Anatomically parcellated-skull into 11 major bones

bones increases linearly with thickness, the parts with absence of cancellous bones are less conductive, and sutures are highly conductive. Thus, it is important to characterize the skull inhomogeneities as much as possible in the conductivity modeling.

However, increasing the number of modeled tissues also increases search and computational complexity. For this reason, it has been impossible to date to determine the tradeoffs of tissue dimensionality, simulation time-to-solution, and conductivity model accuracy. We know that the simplex method is not viable beyond six tissues. When skull parts are separately modeled, a more powerful technique is required. The results here are based on the simulated annealing algorithm [7]. For evaluation purposes, we computed scalp potentials for thirteen preset tissue conductivities and tested how well the model performed as the number of tissues varies. To address the increased computational demands, we parallelized the conductivity search based on simulated annealing. The computational design is described in the next section.

3 Computational Design

With higher resolution and more detailed parcellation of the skull, the computational requirements of the forward and inverse calculation increase significantly. Beyond thirteen search parameters and $1mm^3$ MRI/CT resolution, the computation quickly becomes impractical. Certainly, to pursue higher dimensionality, it is clear we must replace the simplex method used in our earlier work. We chose the advanced simulated annealing algorithm [7] which has been shown to be more robust for optimization across complex multi-variate search spaces. This algorithm has allowed us to extend the conductivity modeling to study the impact of skull inhomogeneities on the conductivity modeling convergence and performance.

The simulated annealing method consists of three nested loops:

Temperature cooling: The outer loop controls the temperature cooling, the temperature is reduced by a factor of r after executing N_t search-radius loops.

Search radius: The intermediate loop, the search radius loop, executes the inner loop N_s times before the maximum step length adjusted. The maximum step lengths are adjusted such that approximately half of the moves are accepted.

Control point: The inner loop considers new moves in all directions by perturbing the control point in all direction; the perturbation is constraint between 0 and the maximum step length in each direction.

```

input : Initial Temperature  $T_0$  and initial Point  $X_0$ 
output:  $X_{optimal}$ 
 $T = T_0, X = X_0, F = Cost(X_0)$ 
while  $T > 0$  do
  for  $i \leftarrow 0$  to  $N_t$  do
    for  $j \leftarrow 0$  to  $N_s$  do
      for  $k \leftarrow 0$  to  $N$  do
         $X_k \leftarrow \text{purturb}(X, k)$ 
         $F_k \leftarrow Cost(X_k)$ 
         $\Delta F \leftarrow F_k - F$ 
        if  $\Delta F < 0$  then accept  $X_k$  and update  $X_{optimal}$ 
        else accept  $X_k$  with probability  $\propto \exp(\Delta F/T)$ 
      end
    end
    adjust-maximum-step-length
  end
  check-termination-condition
  reduce temperature  $T = rT$ 
end

```

(a)

```

input : Initial Temperature  $T_0$  and initial Point  $X_0$ 
output:  $X_{optimal}$ 
 $T = T_0, X = X_0, F = Cost(X_0)$ 
 $N_{task} = N_s / \text{Number of tasks}$ 
while  $T > 0$  do
  for  $i \leftarrow 0$  to  $N_t$  do
     $X_{best} = X, F_{best} = F$ 
    for  $j \leftarrow 0$  to  $N_{task}$  do
      for  $k \leftarrow 0$  to  $N$  do
         $X_k \leftarrow \text{purturb}(X, k)$ 
         $F_k \leftarrow Cost(X_k)$ 
         $\Delta F \leftarrow F_k - F$ 
        if  $\Delta F < 0$  then accept  $X_k$  and update  $X_{best}$ 
        else accept  $X_k$  with probability  $\propto \exp(\Delta F/T)$ 
      end
    end
    MPI communication
    adjust-maximum-step-length
  end
  Master task gets each task best point
  Update  $X_{optimal}$ 
  check-termination-condition
  Master task Broadcast  $X_{optimal}$ 
   $X = X_{optimal}, F = F_{optimal}$ 
  reduce temperature  $T = rT$ 
end

```

(b)

Fig. 3. Pseudo code for simulated annealing algorithm: a) Serial search about the control point with parallel forward calculations, and b) Parallelization along the search radius with parallel forward calculations

Transitions that lower the cost function are always accepted, while transitions that raise the cost function are accepted with probability based on the temperature and the size (cost) of the move (Metropolis criteria). The acceptance of the uphill moves allows the system to overcome local minima and make the algorithm insensitive to the initial starting values. The simulated annealing converges when the minimum become stable that does not change for more than epsilon after several temperature reduction loop iterations. To reduce the number of forward calculation, we added a second termination condition when the value of the cost function becomes less than some tolerance value. The complete algorithm is given in Figure 3(a).

The inner loop spends most of the time in solving the forward problem, since the forward solver is the highest cost component of conductivity modeling. In our earlier work, we parallelized the forward solver using OpenMP and achieved speedups of 5.8 on eight processors and 8.2 on sixteen processors on a 16-processor IBM p690 system (see [2]). Here we also parallelized the simulated annealing algorithm along the search-radius, modifying the parallel MPI-based methods described in [7] for our problem domain. Figure 4(a) shows a high-level view of the parallel simulated annealing approach, with the intermediate loop distributed across several nodes (1 task per node), each running parallelized forward calculations. All tasks start with the same initial values, but with different random generator seeds. Each task perform a random search around the control point by perturbing the control point in all direction. The perturbation is constrained to be within the *maximum step length*. At the end of the search radius loop, the master task gathers every task best solution and updates the optimal solution and all nodes communicate to adjust the maximum step length. At the end of the temperature reduction loop all tasks updates the control point with optimal solution. The complete algorithm is given in Figure 3(b).




Fig. 4. Alternative parallel simulated annealing algorithms: a) Two-level parallelism – parallel execution occurs on the level of the intermediate loop, where each task runs the parallel forward calculation, and b) Three-level parallelism – in addition to the parallel forward calculation, each task computes the inner loop in parallel

Most of the performance numbers reported in the next section are for the parallel simulated annealing algorithm in Figure 4(a). Choosing twelve tasks and 16-way forward solves allows the amount of parallel execution to reach 192 processors. There does not appear to be much benefit in increasing the number of tasks beyond twelve and the OpenMP performance flattens beyond sixteen processors. To create greater potential parallelism, we decided to enumerate all possible perturbation paths around the control point, increasing the number of candidate points from N to 2^N , and breaking the serialization introduced in the former algorithm when the next perturbation point depends on the acceptance or rejection of the prior point. This new algorithm (shown graphically in Figure 4(b)) effectively multiplies the degree of parallelism by $2^N/N$ with a potential performance gain of N . The numbers in Figure 4 for each case indicate the different types of parallelism. Interestingly, the nature by which this algorithm was created guarantees it to produce at least as optimal results as the former, but it requires significantly more resources to achieve greater performance benefits.

4 Computational Results

We conducted a series of experiments to test both the convergence properties and the performance of the conductivity modeling based on parallel simulated annealing. All experiments were performed on the San Diego Supercomputing Center DataStar system, a cluster of shared-memory nodes consisting of 16-processor IBM p690 and 8-processor IBM p655 machines. All results presented below were performed using 8-way OpenMP tasks running forward calculations each on a separate p655 node.

As mentioned earlier, we preset thirteen tissue conductivity values (eleven skull parts, scalp, brain) and ran experiments to test the conductivity model accuracy on fewer tissue numbers. We started with eleven tissue to verify convergence to acceptable values. The simulated annealing search starts with initial random conductivities selected from the biomedical ranges and stops when one of three criteria is met as described in the computational design. Our results verify the ability of simulated annealing to extract eleven tissues with good accuracy and precision. Figure 5(right) shows the dynamics of the 11-tissue inverse search convergence, giving the temperature cooling, the cost function, and one tissue's conductivity. This calculation was done on a single 8-processor p655 node in our lab and took approximately 31 hours to complete.

Having verified convergence for a large number of tissues, we decided to limit the number of tissues to five (three skull parts, scalp, brain) to test the performance




Fig. 5. The Speedup, the mean and the standard deviation of the execution time of the two-level parallel simulated annealing solver. Each point corresponds to three experimental runs using 5 tissues (left). Simulated annealing dynamics, the temperature cooling schedule, cost function, and the retrieved conductivity of frontal bone (right).

properties of the parallel simulated annealing algorithm. This will allow us also to contrast performance with the earlier simplex outcome. The execution time and performance speedup for a $2mm^3$ size problem from one to twelve tasks (each task run on a 8-processor p655 node) is shown in Figure 5(left). The speedup is almost linear with the number of nodes. Three experiments were run on each data point to show the performance variation due to the random number generator sequence.

Even though we have excellent speedup, the degree of parallelism is limited. The second version significantly increases parallelism by generating N random numbers, one for perturbing the control point in each direction, and generated all possible points by enumerating the perturbation in all directions. Then the cost function at these points is evaluated in parallel. After computing all possible paths, the simulated annealing criteria is applied. In theory this parallelism can speedup the computation by a factor of N . In addition, we can get further speedup by selecting the best point from all points that was computed and not only from those points on a simulated annealing path. This speedup is due to speeding the convergence.

For verification purposes, we compared this new algorithm with the former for a problem where the conductivities of three tissues are found. We used a single task and two processors for the forward calculation. Thus, in the former algorithm, the parallelism degree is two. In the new algorithm, the parallelism is of degree sixteen (2^3 perturbed points by two processors in the forward calculation), allowing two 8-processor p655 nodes to be used, with a speedup potential of three (the number of tissues). Our experimental results show an overall performance improvement of 1.98. The best point selection gives a convergence speedup of 8%, while the inner loop parallelism produces a speedup of 1.77. We believe better speedup can be achieved by eliminating communication overheads. For larger numbers of tissues, the number of processors needed to realize the potential parallelism in the new algorithm increases by a power of two. When this resource scaling is unavailable, the method can be throttled to use a smaller degree of perturbation fanout.

5 Conclusion

We provided an efficient computational environment that made the reconstruction of the human head tissues conductivities possible. The simulated annealing algorithm proved to be stable in extracting up to 13 tissues without any sign of failure. Remarkably, using the parallel simulated annealing algorithm, we were able to extract five tissues in 40 minutes, compared to over three days using our previous simplex methods. More importantly, the new algorithms are very robust. In all the testing for this paper, the simulated annealing runs never failed to converge.

There are still improvements to come. The potential parallelism in our new algorithm begs for large-scale testing. For instance, with eleven tissues, over 2048-way parallelism is possible. We intend to perform experiments on IBM BG/L and Cray XTE machines soon. This level of performance is important because our research group at the Neuroinformatics Center will begin creating a head model database for several demographic populations in the next year.

References

1. Huiskamp, G., Vroeyenstijn, M., van Dijk, R., Wieneke, G., Hufferlen, A.C.: The need for correct realistic geometry in the inverse EEG problem. *IEEE Transactions on Biomedical Engineering* **46** (1999) 121–1287
2. Salman, A., Turovets, S., Malony, A., Eriksen, J., Tucker, D.: Computational Modeling of Human Head Conductivity. *International Conference on Computational Science* **3514**, (2005) 631–638
3. Gulrajani, R.M.: *Bioelectricity and Biomagnetism*. John Wiley & Sons, New York (1998)
4. Abrashin, V.N., Egorov, A.A., Zhadaeva, N.G.: On the Convergence Rate of Additive Iterative Methods. *Differential Equations* **37** (2001) 867–879
5. Hielscher, A.H., Klose, A.D., Hanson, K.M.: Gradient Based Iterative Image Reconstruction Scheme for Time-Resolved Optical Tomography. *IEEE Transactions on Medical Imaging* **18**(1999) 262–271
6. Law, S.: Thickness and resistivity variations over the upper surface of the human skull. *Brain Topography* **6** (1993) 99–109
7. Higginson, J.S., Neptune, R.R., Anderson, F.C.: Simulated parallel annealing within a neighborhood for optimization of biomechanical systems. *Journal Biomechanics* **38** (2005) 1938–1942