Multilevel Methods for Inverse Bioelelectric Field Problems

C.R. Johnson[1], M. Mohr[2], U. Rüde[2], A. Samsonov[1], and K. Zyp[1]

[1] Scientific Computing and Imaging Institute, University of Utah
[2] Lehrstuhl für Informatik 10 (Systemsimulation),
Friedrich–Alexander–Universität Erlangen–Nürnberg

Abstract. The reconstruction of bioelectric fields from non-invasive measurements can be used as a powerful new diagnostic tool in cardiology and neurology. Mathematically, the reconstruction of a bioelectric field can be modeled as an inverse problem for a potential equation. This problem is ill-posed and requires special treatment, in particular either regularization or an otherwise suitable restriction of the solution space.

The differential equation itself can be discretized by finite differences or finite elements and thus gives rise to a large sparse linear systems for which multigrid is one of the most efficient solvers, but regularization, adaptive mesh refinement, and efficient solution techniques must be combined to solve the inverse bioelectric field problem efficiently. While multigrid algorithms can reduce the compute times substantially, new local regularization techniques can be used to improve the quality of the reconstruction. Local mesh refinement can be used to increase the resolution in domains of increased activity, but must be used with care because refined meshes worsen the ill-conditioning of the inverse problem.

1 Introduction

Hundreds of times each second, the brain sends electrical impulses racing through the body's web of nerve cells to the motor neurons, where they initiate the chemical reactions that cause muscles to contract. Several decades ago, scientists recognized that these excitation currents produce an electrical field that can be detected as small voltages in the skin or scalp. By measuring changes in the pattern of the body's electrical activity, they could detect some forms of heart diseases and neurological disorders. Electrocardiograms (ECGs), for the heart, or electroencephalograms (EEGs), for the brain, measure these voltages. However, because they are relatively coarse and imprecise techniques, they provide physicians with mere snapshots of heart or brain activity. These glimpses help doctors spot disorders but are often insufficient for accurate diagnosis. For the latter, doctors turn to other techniques — even, in rare cases, to exploratory surgery. For example, to determine whether a patient who is not responding to medication has an operable form of the disorder known as temporal lobe epilepsy, neurosurgeons open the cranium and attach electrodes directly to the brain to identify whether the abnormal electrical activity is highly localized (thus operable) or diffused over the entire brain. Similarly, in clinical cardiology there are forms of abnormal heart rhythm (arrhythmias) which do not respond to drug treatment or to artificial stimulation and must be treated surgically. Determining just where to operate, however, is seldom possible with standard ECGs. Instead, such a determination requires either further investigation via a catheter or an open chest examination with a roving, hand-held electrode or a flexible grid of surface electrodes. The localization process is often lengthy and may be unsuccessful [10].

Electrocardiographic imaging, uses an inverse solution applied to electric potentials recorded on the body surface. The goal of the computation is then to reconstruct the complete distribution of electric potentials throughout the torso and the cardiac current source characteristics that produce those distributions. This gives a more complete and informative description of the source of abnormal cardiac activity. To produce such images it is necessary to have accurate individual geometric patient models, to compute the solution of field equations for the potential distributions. Regularization is needed to deal with the extremely ill-posed nature of the problem.

The detection of focal electrical activity in the brain can be characterized as *an inverse source localization* problem. Given a subset of electrostatic potentials measured on the scalp and the geometry and electrical conductivity properties within the head, the inverse problem is to calculate the current vectors and potential fields within the brain. Mathematically this can be stated as *inversely* solving Poisson's equation for the internal current sources. Clinical applications of inverse source localization and imaging from the EEG are already part of practice in advanced neurological and radiological departments, but there remain significant limitations to the technique. Present estimates of the best-case accuracy of a localized dipole in the brain lie in the range of 7–10 mm, a relatively coarse error given the size scale of the brain.

In this paper we will describe the models used in bioelectric field computations and will report on research applying multilevel methods to solve large-scale ill-posed bioelectric field problems using realistic geometries. In particular, multigrid methods can be used as very efficient solvers for the associated forward problem and we will present our experience with algebraic multigrid methods (AMG) which are well suited to deal with complex geometries, as are characteristic of this application. In Sect. 5 we will extend standard regularization to become a local regularization technique and in Sect. 4 we will discuss adaptive mesh refinement methods and their use in inverse problems. In Sect. 7 we will present our numerical results.

2 Mathematical Models

In this article we discuss two types of inverse problems. The first inverse problem we consider is *electrocardiographic imaging*, which can be formulated

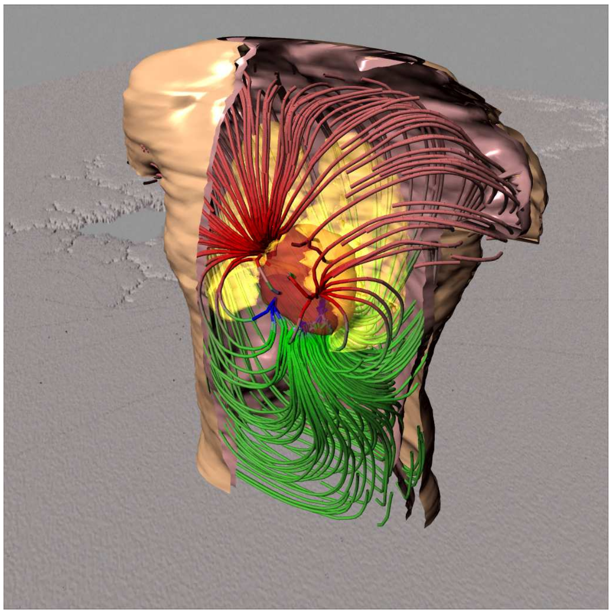

Fig. 1. A visualization of the geometry and electrical current flow in a model of the human thorax for a single time step in the cardiac cycle.

as the problem of finding the potentials (voltages) on the epicardial surface as a function of those on the surface of the thorax, together with knowledge of the geometric and resistive features of the intervening volume conductor.

The second inverse problem we consider is *source localization* and arises in neurology and neurosurgery planning and intervention. Given a subset of electrostatic potentials measured on the scalp and the geometry and electrical conductivity properties within the head, the inverse problem is to calculate internal current sources that are mathematically modeled as dipoles with unknown magnitude, position, and orientation.

In both cases the physical and mathematical model is based on a general formulation in terms of the primary current sources within the torso or brain described by a Poisson-like equation for electrical conduction:

$$\nabla \cdot \sigma \nabla \Phi = -I_v \quad \text{in } \Omega \tag{1}$$

with the boundary condition:

$$\sigma \nabla \Phi \cdot \mathbf{n} = 0 \quad \text{on } \Gamma_T = \partial \Omega \tag{2}$$

where Φ is the electrostatic potential, σ is the conductivity tensor, I_v are the current sources per unit volume, and Γ_T and Ω represent the surface and the volume of the thorax or head, respectively.

Typically, we can assume Cauchy boundary conditions:

$$\Phi = \Phi_0 \quad \text{on } \Sigma \subseteq \Gamma_T \quad \text{and} \quad \sigma \nabla \Phi \cdot \mathbf{n} = 0 \quad \text{on } \Gamma_T \tag{3}$$

where the Neumann condition is describing that no current flows out of the body, and the Dirichlet condition is a mathematical abstraction of potential measurements on Σ, which are in reality obtained from only a finite number of electrodes.

The simplest source localization problem uses the assumption that $I_v = Q(m, x, y, z, \phi, \theta)$ is a single dipole, characterized by six parameters, corresponding to magnitude m, position (x, y, z), and orientation (ϕ, θ). These parameters are now adjusted such that the resulting electrostatic potential best matches the given measured data. More general models allow for several dipoles and consequently increased number of unknown parameters, or — even more generally — the reconstruction of I_v without assuming anything specific about its shape. Unfortunately this latter problem is generally not uniquely solvable. Solutions can only be found with suitable regularization or a restriction of the solution space by assuming I_v of a special form, where the dipole model is the most common one.

Electrocardiographic imaging describes cardiac sources in terms of the epicardial surface potentials such that instead of Poisson's equation we solve a generalized Laplace's equation:

$$\nabla \cdot \sigma \nabla \Phi = 0 \tag{4}$$

again with boundary conditions (3), but now Γ_T is only a part of $\partial\Omega$ and recovering the unknown (Dirichlet) boundary values on the epicardial surface Γ_E is the primary task.

The solution to the inverse ECG problem in the above form is unique. However, all these problems share a characteristic of all electrocardiographic inverse problems in that they are ill-posed in the Hadamard sense; the solution does not depend continuously on the data. Thus small errors in the measurement of the torso potentials or thorax geometry can yield unbounded errors in the solution. The origins of this ill-posedness are biophysical and lie in the attenuation of potentials as we move away from the source and the fact that the potential at any point on the torso surface is a weighted superposition of all the individual sources within the heart. Hence, the ECG or EEG represents an integration of many sources, the influence of each of which decreases sharply with distance. To find the inverse solution, we must perform the complementary operations of amplification and differentiation on the data, but also on the inevitable noise that accompanies it. The result is exquisite sensitivity to any fluctuations in the input signals or geometric models [13].

3 Discretization and Geometric Models

In order to solve the boundary value problem in (1) or (4) in terms of the bioelectric potentials we need to pose the problem in a computationally tractable way. With current technology it has become feasible to employ the real geometry of an individual patient's head or torso in the setup of Ω. The necessary information is obtained from magnetic resonance imaging (MRI scans), which in some cases is supplemented by computer tomography (CT). The data are used to form a 3D model (e.g. a mesh of tetrahedra, or a rectangular box of voxels).

MRI scans return a collection of 2D slices containing pixel information. Thus a simple and obvious approach to get to a volume representation is to extend pixels from each 2D slice to 3D voxels. In the segmentation process each such voxel is then assigned a tissue information. Usually one restricts oneself to a limited number of tissue types and assumes isotropic conductivities in each compartment, see Fig. 2. For our purposes at least air, scalp, skull and brain have to be distinguished, where the conductivity of air is assumed to be zero. For a more sophisticated segmentation, see e.g. [24].

The next step is to find an approximate solution Φ_h of (1). Using a derivation similar to finite volume theory, we assume that Φ_h is a constant Φ_l within each voxel V_l for $l = 1, \ldots, N$. Using the technique of boxed integration we approximate (1) for every voxel V_l with a difference equation

$$\sum_{V_k \in \mathcal{N}(V_l)} c_{l,k}\Phi_k + c_l \Phi_l = F_l \quad \text{for } l = 1, \ldots, N \tag{5}$$

Here $\mathcal{N}(V_l)$ denotes the *orthogonal* neighbors of V_l, i.e voxels which share a common face with V_l.

The linear system (5) has the form $A\xi = b$ where the system matrix A is sparse, symmetric, and positive semidefinite. Since we are solving a pure Neumann problem, the matrix is singular with a one-dimensional null space. However, the problem can still be solved by standard iterative methods. Due to the sparsity each relaxation step is cheap, but the convergence rate of relaxation methods deteriorates quickly with the problem size.

For the work on the inverse ECG problem a finite element discretization has been used. The torso was modeled using a tetrahedral mesh the details of which can be found in Sect. 7. A 2D variant is shown in Fig. 2. The finite element approximation of the original boundary value problem (1) can be stated as:

$$\text{Find } \Phi_h \in V_h \text{ such that } (\sigma\nabla\Phi_h, \nabla\overline{\Phi}) = -(I_v, \overline{\Phi}) \ . \tag{6}$$

In our case $(.,.)$ denotes the L_2 inner product and V_h is a finite dimensional space of functions living on the tetrahedral mesh. The problem can equivalently be expressed as a system of N equations with N unknowns, $\xi_1, \ldots, \xi_N$ (the electrostatic potentials at the nodes in the domain). In matrix form, the above system is again expressed as $A\xi = b$, where here $A = (a_{ij})$ is the global stiffness matrix and has elements $(a_{ij} = (\sigma_{ij} \nabla \Psi_i, \nabla \Psi_j))$, while $b_i = -(I_v, \Psi_i)$ is the vector of source contributions. The major advantage of finite elements compared to finite difference is a greater flexibility in handling complicated geometries and an extensive theory for error estimation.

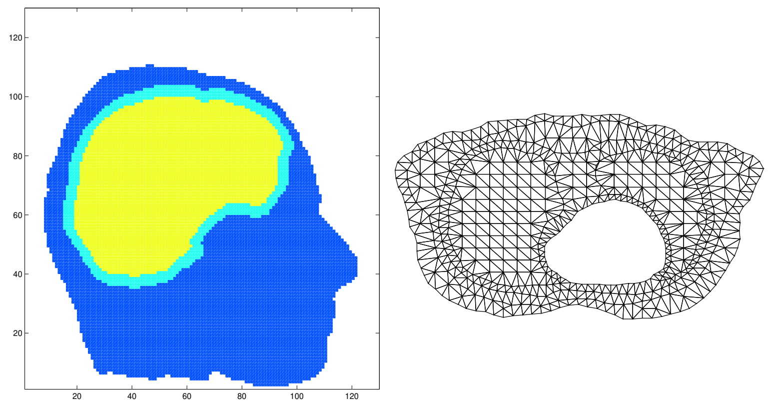

Fig. 2. 3-compartment voxel model of a human head (left), triangular mesh of a torso cross-section (right).

4 Adaptive Mesh Refinement

The most straightforward technique to define a computational grid is to set it up once and then use it throughout the computation. Typically all voxels or elements are then roughly of the same size and they are held constant over time, independent of any changes in the bioelectric potentials.

More advanced techniques provide for an *adaption* of the meshes. Finer meshes are to be placed in regions of increased activity, where otherwise discretization errors could sharply increase. This is especially true in source localization problems where naturally much of the potential activity is focused around the location of the dipoles. But also in cardiac imaging there is more activity close to the heart and therefore seemingly a good reason to use finer meshes there.

Generally, automatic mesh refinement techniques are based on *a posteriori* error *indicators* or error *estimators* [18,22]. A wide variety of such error estimators has been suggested in the literature. In contrast to the standard a-priori estimates for discretization methods, such techniques take the actual shape of the solution into account and can be used to identify regions with large error, where consequently a finer mesh will lead to an improved accuracy in the overall problem.

Additionally, the remeshing has to be done carefully to avoid a too high computational cost on the one side and also avoid a degeneration of the meshes. For the benefit of multilevel solvers the refinement process should generate a sequence of nested meshes. This implies that the finite element spaces are also nested, and this in turn can be used to define an multilevel iterative solver. However, finite element theory also requires that the interior angles of a mesh not become to large. This puts additional constraints on any automatic mesh refinement algorithm [20].

It has always been assumed that improving the accuracy of the forward solution also improves the subsequent inverse solution. This seems to be true up to a point since further refinement increases the number of degrees of freedom and correspondingly increases the condition number of the transfer matrix. Finding the "optimal" point at which the added degrees of freedom stop increasing the accuracy of the inverse solution is an open research problem.

5 Inverse Problems and Regularization

Dealing with the ill-posed nature of the inverse problem has become the primary goal of a great deal of recent research as it remains the single largest obstacle to medical implementation. The most common approach is to apply constraints to the inverse solution in order to reduce its dependence on boundary conditions. The points of difficulty lie in the choice of constraints and the corresponding weight each receives in determining the best — in whatever sense we wish to define "best" — solution (see [12] for a recent review).

Traditional schemes for solving the inverse problem, (4), have involved reformulating the linear equation, $A\xi = b$, into $\xi_T = K\xi_E$, where ξ_T and ξ_E are the voltages on the body surface (torso) and heart's surface (epicardium) respectively, and K is the $T \times E$ transfer (lead field) matrix of coefficients relating the measured torso voltages to the voltages measured on the heart. From this statement of the forward problem, the inverse problem can then be expressed as $\xi_E = K^{-1}\xi_T$.

Unfortunately, K is ill-conditioned, and regularization techniques are necessary to restore continuity of the solution back onto the data. Most regularization schemes treat all parts of K uniformly and apply Tikhonov, singular value decomposition (SVD), or Twomey algorithms with a single regularization parameter. Briefly, one wishes to find an approximate vector, $\xi_{E\varepsilon}$, that represents a balance between the residual error $\|\xi_T - K\xi_{E\varepsilon}\|$, which is corrupted by error and the ill-posed nature of the problem, and whatever side constraints we place on the solution based on *a priori* knowledge. Rather than applying the same level of regularization to the entire matrix K, we have found that it is possible, and advantageous, to first decompose K into submatrices and then apply regularization differently to each component of the resulting expression.

The rationale for such a local approach is that the discontinuities in the inverse solution appear irregularly distributed throughout the solution domain. Tikhonov and other regularization schemes are, in effect, filters, which restore continuity by attenuating the high (spatial) frequency components of the solution. Since regularization is usually applied globally, the results can leave some regions overly damped or smoothed while others remain poorly constrained.

We have two ideas regarding multilevel methods applied to ill-posed problems. The first idea is to apply regularization only to regions (represented by the sub-matrices of K or some other partition) that require it and to apply different amounts of regularization to different sub-matrices/partition, in effect a hierarchical, more local regularization scheme. Our second idea is to use the natural filtering of high frequencies in multigrid and construct a multigrid method that can effectively solve ill-conditioned systems.

Since the goal of the bioelectric inverse problems is to accurately and noninvasively (i.e., non-surgically) estimate accurate solutions, we need a method to choose an *a priori* optimal regularization parameter. Traditional schemes have been based on discrepancy principles [16] and quasi-optimality criterion [16]. Recently, Hansen [6] has extended the observations of Lawson and Hanson [11] and proposed a new method for choosing the regularization parameter based on an algorithm that locates the "corner" of a plot of the norm (or semi-norm) of the side constraint, $\|Lx_\alpha\|$, versus the norm of the corresponding residual vector, $\|Ax_\alpha - b\|$. In this way, one can evaluate the compromise between the residual error and the effect of the side constraint on the solution. We have utilized this *L-curve* algorithm along with our local regularization scheme to improve the accuracy of solutions to the inverse problem. We have also utilized an AMG approach (described in more detail in Section 7) to help in regularizing the ill-conditioned inverse problems.

6 Multilevel Methods

Many approaches to solve inverse bioelectric field problems involve the repeated numerical solution of related forward problems, usually for different boundary conditions and/or right hand sides. Prominent examples are the setup of the lead field matrix [21,23] or the preconditioning of Krylov subspace methods for the inverse problem. In both our applications, the forward problem is of elliptic type. Multigrid algorithms belong to the fastest solution algorithms for such problems and thus can be used to speed up the solution process. Nevertheless, multigrid algorithms have not yet found widespread use for this application. While theoretically well suited for the task, the implementation is complicated by the problem geometry and the jumping and/or anisotropic conductivities. These problems have been individually solved with multigrid in other application fields [2,3], but to get full multigrid efficiency, a substantial development effort in each individual case is required.

A possibility to simplify the development of multigrid is the use of so-called algebraic multigrid methods. These algorithms solely work on the system matrix and do not need further geometric information. AMG was originally developed in the early 1980's [19], but has only recently become a mainstream topic of multigrid research. AMG deals with many of the difficulties that arise in employing standard multigrid in an automatic way and is thus much easier to apply in practical situations. In particular, AMG automatically finds coarser grids: there is no need for the user to provide a mesh hierarchy as is necessary for conventional multigrid algorithms. AMG is also robust against jumps in the coefficients, the type of boundary conditions, or geometric features of the solution domain. Therefore this paper will focus on exploring the suitability of AMG methods for the bioelectric field problems. Compared to standard multigrid, the AMG storage requirements are usually significantly higher than for a comparable geometric multigrid. General sparse matrix structures are also less storage efficient than grid-oriented data structures, which could otherwise be used in our finite difference models. Furthermore, the operation count for AMG is higher and in most cases the convergence rates are not quite as good as for geometric multigrid. The further development of AMG is currently a very active field of research.

The discretization of elliptic PDEs leads in most cases to linear systems that can be handled with Conjugate Gradient methods [5]. The convergence of this method crucially depends on the spectrum (ratio and distribution) of eigenvalues of the problem matrix. Thus preconditioning is employed to improve convergence. Geometric and algebraic multigrid methods can be used as preconditioners, and arguably, this combination of CG and multilevel techniques is currently the most advanced technique to solve sparse linear systems. For more details, see Section 7. Techniques for applying multilevel methods directly to the inverse problem are described in [15].

7 Results

Finite Difference Head Model. Here we will report on results for different iterative solvers for the EEG forward problem discretized by finite differences, as described in Sect. 3. The original task is to locate a dipole inside the brain from EEG data. The algorithm proposed for this task in [21] consists of two stages. In a preparatory stage the potential distribution for $(N-1)$ forward problems has to be computed, where N is the number of electrodes / measurement points. In the second stage the dipole localization can be performed without further forward solves using the reciprocity theorem of Helmholtz.

We used a three-compartment (scalp, skull, brain) geometric head model with constant, isotropic conductivities in each compartment, cf. Sect. 3. Our experiments use two different data sets, a small one (A) with a resolution of 65^3 cells and a larger one (B) with 129^3 cells. Note that in both cases the head occupies only about 25% of the surrounding cube of voxels. We compare the following linear solvers:

– SOR: Successive over-relaxation on the pure Neumann problem with a relaxation parameter $\omega = 1.94$.
– CG: The conjugate gradient method applied to the pure Neumann problem (1), or to the case, where the value at one cell is fixed at zero (2).
– AMG: Algebraic multigrid method applied to the definite problem.
– PCG(AMG): AMG preconditioned CG with one AMG-V-cycle per CG step with one pre- and post-smoothing Gauss-Seidel step.

The SOR method is implemented in a highly optimized code described in [25] and specifically developed for the given problem, while the AMG/PCG solver is a general purpose code and is taken from the Pebbles package [1]. The performance comparison is summarized in Tab. 1. In all cases, the iteration is stopped when the L_2-norm of the residual drops below 10^{-8}. The iteration counts are given for a single forward solve, while the user times are taken for the whole first part of the algorithm. This includes setup of system matrix/stencil coefficients, the solution of 26 forward problems corresponding to 27 electrodes, and, in case of AMG and PCG(AMG), additionally the construction of the matrix hierarchy. Creation of the latter takes about 6s for data set (A) and 79s for (B). The figures in Tab. 1 show that at least for smaller problems a simple, but well implemented SOR algorithm is still reasonably competitive, at least in terms of user time. It outperforms the CG method without preconditioning even in the number of iterations. Clearly, however, the number of iterations for SOR is much higher than for the multilevel algorithms. On the other hand, a single SOR iteration is much cheaper than an AMG cycle. Additionally the Pebbles AMG implementation has larger memory requirements than the SOR code. Overall, at least for the larger data set, AMG in combination with CG acceleration is more efficient

User time in seconds				
Method	Data set			
	A		B	
	Setup	Total	Setup	Total
SOR	0.2	139	1.2	5857
CG(1)	0.2	400	1.2	5358
AMG	6.2	218	79.0	8156
PCG(AMG)	6.2	136	79.0	1924

Number of Iterations		
Method	Data set	
	A	B
SOR	314	1644
CG(1)	408	741
CG(2)	533	987
AMG	18	98
PCG(AMG)	11	20

Table 1. User times (left) for an AMD Athlon K7, 700MHz, 768MB, running linux, iteration counts (right) for the two data sets

than either SOR or plain CG. Note, however, that the Pebbles code does not yet perform with asymptotic optimality, as multigrid theory predicts for an optimally designed multigrid method. The number of iterations is not independent of the mesh size. For plain AMG, almost 100 iterations are required for the large data set. CG acceleration cuts that down to 20, but this is still twice as much as for the small data set and thus clearly not yet fully scalable.

Thus these experiments show that multilevel algorithms have a good potential for accelerating the forward solution process. However, we believe that the full potential of multigrid for this application is not yet exploited. We therefore plan to implement a geometric multigrid method for the forward problem, using data structures similar to those used in the SOR implementation. This should lead to faster run times for each iteration and reduce the memory requirements. Using techniques described e.g. in [2] for treating discontinuous coefficients will also lead to improved convergence rates.

Finite Element Heart Model. To study forward and inverse problems in electrocardiography, we developed a series of two- and three-dimensional boundary element and finite element models based upon magnetic resonance images from a human subject. Each of 116 MRI scans were segmented into contours defining torso, fat, muscle, lung, and heart regions. Additional node points were added to each layer and pairs of layers tessellated into tetrahedra using a Delaunay triangulation algorithm [8]. The resulting model of the human thorax contained approximately 675,000 volume elements, each with a corresponding conductivity tensor [4]. A visualization of a forward simulation from epicardial potentials is shown in Fig. 1.

We created a two-dimensional model from the full 3D model shown in Fig. 1, located approximately 4 cm above the apex of the heart, Fig. 2. This provided a more tractable geometry for the initial studies on the effects of adaptive control of errors and local regularization.

Local Regularization. To test the local regularization technique, we applied epicardial potentials recorded during open chest surgery from a cardiac arrhythmia patient as the Dirichlet boundary conditions of the two-dimensional model described above. The tissue conductivities were assigned as follows: fat = 0.045 S/m, epicardial fat-pad = 0.045 S/m, lungs = 0.096 S/m, skeletal muscle (in the fiber direction) = 0.3 S/m, skeletal muscle (across the fiber direction) = 0.1 S/m, and an average thorax value = .24 S/m. Forward solutions calculated using the adapted mesh served as torso boundary conditions, where $\Phi = \Phi_0$ on $\Sigma \subseteq \Gamma_T$ for the inverse solution both with and without 10% added Gaussian noise.

The application of the local regularization technique using a generalized singular value decomposition (GSVD) recovered the voltages to within 6.1% relative root-mean-squared (RMS) error of the true solution. This compares with a previous best from our work of 12.1% relative error [7]. Previous studies have reported the recovery of epicardial potentials with errors in the range of

20–100%, [14,17], although in these cases investigators used three-dimensional geometric models and in some cases measured potentials on both epicardial and torso surfaces. Fig. 3 shows the inverse solution calculated using the local regularization technique compared with the recorded heart voltages as a function of position on the epicardium. The global solution tended to be smoother, not able to follow the extrema as well as the local solution could. The local regularized solution also showed a few areas of local error which suggest that a different partitioning of the sub-matrices might provide even better accuracy.

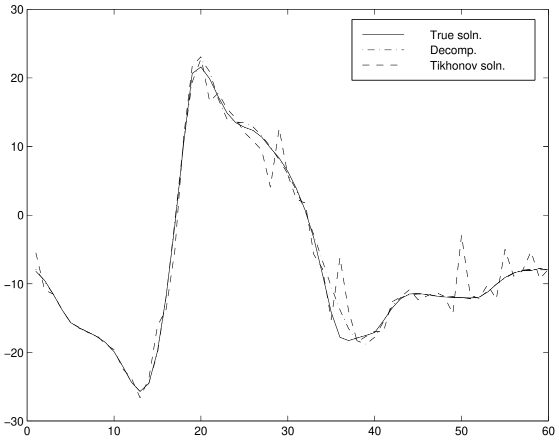

Fig. 3. Local GSVD-based regularization technique. The figure shows the GSVD local regularization solution (marked Decomp), the best Tikhonov solution, and true solution.

AMG Solutions for the Inverse ECG Problem. In this section we report on preliminary work in using AMG for solving the inverse ECG problem. The algorithm used for coarsening in the experiment was based on the graph of the finite element stiffness matrix [9]. The coarsening results in fast computation of the coarse grid matrices, which is a benefit for large-scale problems if the stiffness matrix is known.

For our test problem we used the same 2D slice of the patient torso as used with the hierarchical finite element with local regularization application. Several approaches to coarsening were tested. One of the most valuable was based on the following algorithm:

1. First of all, matrix graph G was created: $G = (v, E)$ where $v = \{i\}$ and $E = \{(i, j) : A[i, j] \neq 0\}$, where $i, j = 1, \ldots, n$

2. Then, graph G was subdivided into trees by the following method: Starting with the first vertex k, not marked as visited, put the vertex into the interpolation graph I, make it a root of the new tree and mark it as visited. Then, for each vertex m in G connected to k and not yet visited, put (k, m) into I and mark m as visited. Repeat these steps until all the nodes are marked as visited.
3. Given interpolation graph I, build a form of a Galerkin interpolation operator: For every sub-tree in I, make its root a coarse grid point, and then build a Galerkin interpolation operator:

$$I_g(i,j) = \begin{cases} 0 & \text{node is not in sub-tree of j-node} \\ 1 & \text{for a coarse grid point} \\ 1/n & \text{otherwise, where n = the number of} \\ & \quad \text{subtrees this node belongs to and} \\ & \text{where } j \text{ belongs to the coarse grid point set.} \end{cases}$$

After construction of the interpolation operator, the restriction operator is defined by $R = I_g^T$, and the final system on the coarse grid will be $A_H x_H = b_H$ where $A_H = RAI_g, b_H = Rb$ and H denotes the original (heart) linear system.

Having constructed the interpolation and restriction operations and coarse grid matrices, we can apply our AMG scheme to the 2D inverse ECG problem. After only four iterations of the AMG method, the reconstruction yields results similar to an exact finite element solution which is obtained at much higher cost, see the comparison in in Fig. 4. This seems to verify an observation made in [15] that multigrid posseses an inherent regularization property. It seems to provide fast convergence for the well-conditioned modes, but is converging poorly for the ill-conditioned modes.

8 Conclusions

In this paper we have described research applying multilevel methods to solve large-scale ill-posed bioelectric field problems using realistic geometries. 3D models based upon patient MRI scans result in meshes that may contain millions of elements, such that even a single forward solution is very compute intensive. Classical multigrid methods can be used as very efficient solvers for the associated forward problem and can then be used as either a preconditioner or for constructing the lead field matrix for the inverse problem.

Finite elements and in particular finite elements on unstructured grids on complicated geometries lead to less sparse, unstructured linear systems, for which a natural grid hierarchy, as needed for classical multigrid, does not exist or would be difficult to construct. Multigrid will in principle be able to solve the forward problem with less than 100 operations per unknown, provided all details within the multigrid algorithm are handled in exactly the right way, but this will not be easy to achieve. Therefore we have used AMG as a black box sparse system solver. As the results indicate, there is still room

Fig. 4. Results from applying AMG to the ill-posed inverse ECG problem. After one iteration the RMS error is 40.1% and after four iterations the RMS error was reduced to 12.7%.

for improvement, so that we plan to implement geometric multigrid methods to further improve the method.

Even more intriguing is the use of multilevel methods to apply multiple levels to regularize the ill-posed problem, thus resulting in a more *local* regularization technique that provides more control and tuning of the regularization process. Current methods for solving inverse bioelectric field problems are primarily aimed at global Tikhonov regularization methods. Compared to these currently used methods, the application of multilevel methods can provide significant improvement that can make many presently intractable problems solvable.

Acknowledgments

This work was supported in part by awards from the NIH NCRR and the NSF. The authors are also in debt to Bart Vanrumste and the MEDISIP group at the University of Ghent for providing the head data sets and Stefan Reitzinger for the Pebbles AMG implementation.

References

1. *Pebbles – User's Guide.* www.sfb013.uni-linz.ac.at/~reitz/pebbles.html.
2. R. E. Alcouffe, A. Brandt, J. E. Dendy, and J. W. Painter. The multi–grid methods for the diffusion equation with strongly discontinuous coefficients. *SIAM J. Sci. Stat. Comput.*, 2:430–454, 1981.

3. A. Brandt. *Multigrid techniques: 1984 guide with applications to fluid dynamics.* GMD–Studien Nr. 85. Gesellschaft für Mathematik und Datenverarbeitung, St. Augustin, 1984.
4. K.R. Foster and H.P. Schwan. Dielectric properties of tissues and biological materials: A critical review. *Critical Reviews in Biomed. Eng.*, 17:25–104, 1989.
5. A. Greenbaum. *Iterative Methods for Solving Linear Systems.* Frontiers in Applied Mathematics. SIAM, 1997.
6. P.C. Hansen. Analysis of discrete ill-posed problems by means of the L-curve. *SIAM Review*, 34(4):561–580, 1992.
7. C.R. Johnson and R.S. MacLeod. Inverse solutions for electric and potential field imaging. In R.L. Barbour and M.J. Carvlin, editors, *Physiological Imaging, Spectroscopy, and Early Detection Diagnostic Methods*, volume 1887, pages 130–139. SPIE, 1993.
8. C.R. Johnson, R.S. MacLeod, and M.A. Matheson. Computer simulations reveal complexity of electrical activity in the human thorax. *Comp. in Physics*, 6(3):230–237, May/June 1992.
9. F. Kickinger. Algebraic multigrid for discrete elliptic second order problems. Technical report, MGNet, November 1997.
10. H. Korab. Mapping the body's electrical fields. *Access: High-performance computing magazine*, pages 14–15, 1995.
11. C.L. Lawson and R.J. Hanson. *Solving Least Squares Problems.* Prentice-Hall, Englewood Cliffs, NJ, 1974.
12. R.S. MacLeod and D.H. Brooks. Recent progress in inverse problems in electrocardiology. *IEEE Eng. in Med. & Biol. Soc. Magazine*, 17(1):73–83, January 1998.
13. R.S. MacLeod and D.H. Brooks. Validation approaches for electrocardiographic inverse problems. In *Advances in Computational Biomedicine.* WIT Press, 2000.
14. B.J. Messinger-Rapport and Y. Rudy. Regularization of the inverse problem in electrocardiography: A model study. *Math. Biosci.*, 89:79–118, 1988.
15. M. Mohr and U. Rüde. Multilevel Techniques for the Solution of the Inverse Problem of Electrocardiography. In E. Dick, K. Riemslagh, and J. Vierendeels, editors, *Multigrid Methods VI, Proceedings of the Sixth European Multigrid Conference*, volume 14 of *Lecture Note in Computational Science and Engineering*, pages 186–192. Springer, 1999.
16. V.A. Morozov. *Methods for Solving Incorrectly Posed Problems.* Springer-Verlag, New York, 1984.
17. H.S. Oster and Y. Rudy. Regional regularization of the electrocardiographic inverse problem: A model study using spherical geometry. *IEEE Trans Biomed. Eng.*, 44(2):188–199, 1997.
18. U. Rüde. *Mathematical and Computational Techniques for Multilevel Adaptive Methods*, volume 13 of *Frontiers in Applied Mathematics.* SIAM, Philadelphia, 1993.
19. J. W. Ruge and K. Stüben. Algebraic multigrid (AMG). In S. F. McCormick, editor, *Multigrid Methods*, volume 3 of *Frontiers in Applied Mathematics*, pages 73–130. SIAM, Philadelphia, PA, 1987.
20. J.A. Schmidt, C.R. Johnson, J.C. Eason, and R.S. MacLeod. Applications of automatic mesh generation and adaptive methods in computational medicine. In I. Babuska, J.E. Flaherty, W.D. Henshaw, J.E. Hopcroft, J.E. Oliger, and T. Tezduyar, editors, *Modeling, Mesh Generation, and Adaptive Methods for Partial Differential Equations*, pages 367–390. Springer-Verlag, 1995.

21. B. Vanrumste, G. Van Hoey, P. Boon, M. D'Havé, and I. Lemahieu. Inverse calculations in EEG source analysis applying the finite difference method, reciprocity and lead fields. In *Proceedings of 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, volume 20, part 4/6, pages 2112–2115, Hong Kong, 1998.
22. R. Verfürth. *A Review of A Posteriori Error Estimation and Adaptive Mesh–Refinement Techniques*. Wiley Teubner, New York, Stuttgart, 1996. ISBN 0-471-96795-5.
23. D.M. Weinstein, L. Zhukov, and C.R. Johnson. Lead-field bases for EEG source imaging. *Annals of Biomedical Engineering*, 28:1–7, 2000.
24. C. Wolters, S. Reitzinger, A. Basermann, S. Burkhardt, U. Hartmann, F. Kruggel, and A. Anwander. Improved tissue modelling and fast solver methods for high-resolution FE-modelling in EEG/MEG-source localization. In J. Nenonen, R.J. Ilmoniemi, and T. Katila, editors, *Biomag 2000, Proc. of the 12th Internat. Conf. on Biomagnetism*.
25. M. Zetlmeisl. Performance Optimization of Numerically Intensive Codes – A Case Study from Biomedical Engineering. Studienarbeit, Lehrstuhl Informatik 10, FAU Erlangen–Nürnberg, 2000.