Published in final edited form as:

Physiol Meas. 2010 August ; 31(8): S1–16. doi:10.1088/0967-3334/31/8/S01.

Electrical impedance tomography reconstruction for three-dimensional imaging of the prostate

A Borsic, R Halter, Y Wan, A Hartov, and K D Paulsen

Thayer School Of Engineering, Dartmouth College, 8000 Cummings Hall Hanover, NH 03755, USA

A Borsic: Andrea.Borsic@Dartmouth.edu

Abstract

Transrectal electrical impedance tomography (TREIT) has been proposed as an adjunct modality for enhancing standard clinical ultrasound (US) imaging of the prostate. The proposed TREIT probe has an array of electrodes adhered to the surface of a cylindrical US probe that is introduced inside of the imaging volume. Reconstructing TREIT images in the open-domain geometry established with this technique poses additional challenges to those encountered with closed-domain geometries, present in more conventional EIT systems, because of the rapidly decaying current densities at increasing distances from the probe surface. We developed a finite element method (FEM)-based dual-mesh reconstruction algorithm which employs an interpolation scheme for linking a fine forward mesh with a coarse grid of pixels, used for conductivity estimation. Simulation studies using the developed algorithm demonstrate the feasibility of imaging moderately contrasting inclusions at distances of three times the probe radius from the probe surface and at multiple angles about the probe's axis. The large, dense FEM meshes used here require significant computational effort. We have optimized our reconstruction algorithm with multi-core processing hardware and efficient parallelized computational software packages to achieve a speedup of 9.3 times when compared to a more traditional Matlab-based, single CPU solution. The simulation findings and computational optimization provide a state-of-the-art reconstruction platform for use in further evaluating transrectal electrical impedance tomography.

Keywords

electrical impedance tomography; ultrasound; finite element method; electrode modeling; three-dimensional; prostate imaging

1. Introduction

In this paper, we present results relative to a reconstruction algorithm developed for a custom transrectal probe that couples electrical impedance tomography (EIT) and ultrasound to image the prostate. Transrectal ultrasound (TRUS) probes are commonly used for image-guided biopsy procedures. Physicians use the US image to delineate the boundaries of the prostate and guide a biopsy needle of typically 6–24 specific sites within the prostate. Unfortunately ultrasound images are not cancer specific and, therefore, are unusable to accurately identify malignant regions for biopsy sampling. Instead systematic sampling patterns, that are not patient specific, are used in current clinical practices.

Based on previous findings demonstrating (Halter *et al* 2009a, 2009b) that normal and malignant prostate tissues present different electrical properties, we developed a combined US and EIT probe that enables combined acquisition of acoustic and electrical property images. By overlaying areas of the prostate with 'abnormal electrical properties' onto the

US images, we aim to improve patient-specific biopsy sampling. In particular, multiple tissue samples might be collected from regions of the prostate that raise concerns. This would reduce the chances of missing small volume cancers, since 12-core protocols only sample 1% of the volume of the typical prostate.

The transrectal probe we developed, described in Borsic *et al* (2009) and in Wan *et al* (2010), is based on a commercial transrectal ultrasound probe to which a set of 30 electrodes has been applied, as an EIT sensing array. Electrodes are printed on a flexible substrate, which is attached to the surface of the probe, and are arranged such that an ‘acoustic window’ is open to ultrasonic energy transmission (see figure 1). The ultrasound transducer resides inside the probe’s shaft and is mechanically translated and rotated inside the shaft to acquire a full 3D image.

This application presents an open-domain problem, in which we image the volume outside the electrodes. This is particularly challenging as the current density used for imaging concentrates mainly near the probe surface, resulting in a rapid decay of the sensitivity with increasing distance from the probe, and hence in a particularly ill-posed problem. Two-dimensional reconstruction of a similar geometry has been demonstrated by Jossinet *et al* (2002a), (2002b), and the sensitivity and ill-conditioning of the three-dimensional problem have been studied by Borsic *et al* (2009).

This paper presents details of a reconstruction algorithm specific to this imaging application, demonstrates successful reconstruction results on synthetic and experimental data and discusses aspects relative to speeding up image reconstruction. We use particularly fine finite element method meshes for solving the forward problem, in order to accurately capture the EIT/US system geometry and to achieve a high level of accuracy in the computed potentials. These large meshes require significant computational time for image reconstruction. In the final section of this paper, we discuss computational aspects pertinent to speeding up the reconstruction algorithm, to accommodate for these large, dense meshes.

2. Reconstruction algorithm

2.1. Forward model

The forward problem is modeled with a low-frequency approximation, where the electric field is conservative and conduction currents dominant with respect to their displacement counterparts, which leads to the classic partial differential equation:

$$\nabla \cdot \sigma \nabla u = 0 \quad \text{on } \Omega, \quad (1)$$

where σ is the conductivity or admittivity of the body to be imaged, u is the electric potential and Ω represents the body to be imaged. Electrodes are modeled with boundary conditions that are referred to as the complete electrode model (Somersalo *et al* 1992) in the EIT literature. The model specifies the following boundary condition for each portion of the boundary, $\partial\Omega_\ell$, underneath the electrode ℓ :

$$u + z_c \sigma \frac{\partial u}{\partial \mathbf{n}} = V_\ell \quad \text{on } \partial\Omega_\ell \quad \ell = 1, \dots, L, \quad (2)$$

where z_c is the contact impedance, ℓ is the electrode number, V_ℓ is the potential developed by the electrode ℓ and L is the number of electrodes. An additional condition is that the flux of the normal component of the current density over the surface of each electrode must equal the applied currents:

$$I_\ell = \int_{\partial\Omega_\ell} \sigma \frac{\partial u}{\partial \mathbf{n}} \quad \ell=1, \dots, L, \quad (3)$$

where I_ℓ is the current applied at the electrode ℓ . In the interelectrode gaps, the following boundary condition is applied:

$$\frac{\partial u}{\partial \mathbf{n}} = 0 \quad \text{on} \quad \partial\Omega \setminus \{\partial\Omega_1 \cup \dots \cup \partial\Omega_L\}. \quad (4)$$

Imposing the Neumann boundary condition (4) is common in typical applications of EIT, where the imaging domain is closed. As the current flows in a closed region, the insulating boundary condition (4) is appropriate. In our specific application of EIT, the domain is effectively open. We model therefore the imaging probe as embedded in a cylinder which represents the volume surrounding the probe, as discussed in section 2.3. By numerical experiments we choose the diameter of the cylinder to be 24 cm; this guarantees a significant decay of the electric field set by the electrodes before it reaches the surface of the cylinder, allowing to impose a Neumann boundary condition (4), without affecting the forward measurements. In our experiments the applied imaging field decays to 1×10^{-4} of the original value at the electrodes, when the field reaches the surface of the embedding cylinder.

2.2. Parameter estimation

The reconstruction algorithm we use is based on a standard nonlinear least-squares, Tikhonov-regularized inverse formulation, using a FEM implementation of the forward model (Holder 2004). In our data acquisition system, we perform tetrapolar measurements, applying a current through pairs of electrodes and measuring potentials through a different pair of sensing electrodes. The reconstruction is stated as

$$\sigma_{\text{rec}} = \arg \min_{\sigma} \frac{1}{2} \|\mathbf{V}(\sigma) - \mathbf{V}_{\text{meas}}\|^2 + \alpha \frac{1}{2} \|L(\sigma - \sigma^*)\|^2, \quad (5)$$

where σ is the vector of conductivities to be estimated (admittivities for the complex-valued case). Specifically, σ_{rec} is the vector of reconstructed conductivities, $\mathbf{V}(\sigma)$ is the vector of simulated voltages resulting from the forward solver, $\mathbf{V}_{\text{meas}}$ is the vector of measured electrode voltages, α is the Tikhonov factor, L is a regularization matrix and σ^* is a reference conductivity distribution. Application of the Newton–Raphson method to (5) results in the iterative formula

$$\delta\sigma_n = -(J_n^T J_n + \alpha L^T L)^{-1} [J_n^T (\mathbf{V}(\sigma_n) - \mathbf{V}_{\text{meas}}) - \alpha L^T L(\sigma_n - \sigma^*)], \quad (6)$$

where n is the iteration number, $\delta\sigma_n$ is the conductivity update for iteration n and J_n is the Jacobian of the forward operator $\mathbf{V}(\sigma)$ calculated for $\sigma = \sigma_n$. Given the nonlinearity of the problem, we update the conductivity with a parabolic line search procedure (Nocedal and Wright 1999):

$$\sigma_{n+1} = \sigma_n + \beta \delta\sigma_n, \quad (7)$$

where β is a scalar value in the range $[0, 1]$ determined from the line search process. The use of a parabolic line search procedure results in the evaluation of the objective function, and thus in the solution of the forward problem, at two additional points in the β interval $(0, 1]$. While this procedure does not result in particular convergence gains on noise-free data, where typically a Newton–Raphson step of length 1 is adequate, in our experience the line search procedure is critical in maintaining stability in the presence of strong noise, and we use it therefore in the present clinical application.

Equations (6) and (7) are iterated to minimize the objective function in (5). As a stopping criterion we compare to norm of the current update to the norm of the current estimate, and stop the reconstruction if the change is less than 5%, or $\|\beta\delta\sigma_n\| < 0.05 \|\sigma_n\|$. This criterion results typically in two or three iterations on noisy data, where the step length reduction contributes to having smaller and smaller steps during the iterative process.

2.3. FEM modeling

The combined EIT and US probe we use consists of a 2.28 cm shaft, which houses internally an US head that is mechanically translated and rotated inside the shaft to acquire a three-dimensional imaging volume (Wan *et al* 2010). An array of 30 EIT electrodes is printed on a flexible substrate and attached to the surface of the probe. The electrodes, as illustrated in figure 1, are arranged to form a rectangular pattern, with an internal ‘acoustic window’. The window is sufficiently wide to allow the US head to scan the entire volume of the prostate without clinically significant interference from the electrodes. In a previous arrangement in Borsic *et al* (2009), we considered an array of thin printed electrodes, arranged in a regular rectangular grid that did cover the US head. While it was possible to acquire US images through the electrodes, which proved to be sono-transparent, this arrangement presented slight artifacts caused by the edges of electrodes, and we opted for a new design of the electrode array.

As we are imaging to the exterior of the probe, we use an overall cylindrical FEM of a diameter of 24 cm to model the space around the probe, and the probe itself is embedded in the center of this cylinder, as illustrated in figure 2. The overall mesh is finer in the proximity of the electrodes and in the imaging region in front of the electrode array, to deliver a better accuracy. The mesh is controlled to be significantly coarser past a radius of 7 cm, as currents decay significantly and this does not effect forward modeling. Overall the FEM mesh presents 97 973 nodes and 541 604 tetrahedral elements.

2.4. Coarse representation of conductivity

In this section we discuss an interpolation method that links a coarse representation of the conductivity to the fine FEM mesh used for forward solving. Reconstruction is based on a dual scheme, with a fine mesh used for forward computations, and a coarse representation of conductivity used for inversion. This arrangement permits independent choices regarding the number of elements in the forward mesh and in the coarse conductivity representation. The choice of number of elements in the fine mesh should in fact be based on forward modeling accuracy considerations, while the number of conductivity elements to be estimated should be based on considerations related to the number of available measurements and to the resolution of the imaging method. The use of such a dual scheme allows one to make these two choices independently, contrary to the methods that use the same FEM for computing forward and inverse solutions.

The method we use is based on joining together the conductivity values of several fine FEM elements, to form ‘coarse pixels’. The conductivity value of these coarse pixels is then estimated as a single parameter, reducing the total number of parameters to be estimated,

and allowing control of this number independently from the underlying FEM mesh. Even though the conductivity representation is coarser, the fine mesh still supports a fine description of the electric potential, which is spanned by the FEM shape functions defined on the fine FEM elements. This approach enables specification of a desired forward accuracy, while retaining separate control on the number of inverse parameters to be estimated.

The process we use for generating a coarse conductivity representation, given an FEM mesh, is automated and based on the generation of a number of ‘seed points’—as illustrated in figure 2. The seed points, shown in yellow in figure 2, are used for growing ‘coarse pixels’, by linking together fine FEM elements that are near to each seed point. The conductivity value of each coarse pixel formed around a single seed point is designated as a single parameter in image reconstruction. The number and locations of seed points determine therefore the number and locations of coarse pixels. In our application, we image only a sub-volume of the mesh, since the sensitivity in the open domain decays very rapidly with distance from the probe surface. We image a volume with an aperture of 140° , a radius of 6 cm and a vertical extension of 7 cm. We use cylindrical coordinates to define the imaging domain. Figure 3 shows, in checkerboard colors, the coarse grid of pixels derived by the use of 14 seed points equispaced along the angular direction, 8 points equispaced along the radial direction and 11 points equispaced along the vertical direction. In our software implementation, the user specifies the three resolutions as an input parameter (in this example [8, 11, 14]) and the imaging software generates the seed points and the resulting coarse representation and coarse–fine interpolation scheme. This permits experimentation with different coarse resolutions, for a given fine FEM forward mesh.

2.5. Image reconstruction with the coarse–fine interpolation scheme

The coarse representation of conductivity introduced in the preceding section results in a coarse–fine conductivity interpolation scheme that can be generally expressed in terms of the linear combination $\sigma^f = P\sigma^c$, where P is a matrix that defines the interpolation, which is sparse given the local nature of the scheme that groups elements on the fine mesh, and where σ^f is the fine conductivity and σ^c the coarse conductivity. In the specific case of the proposed scheme, which groups elements on the fine mesh to form coarse pixels, rows of P will be populated with ‘1’ entries for column indices corresponding to fine elements forming a coarse pixel.

In terms of reconstruction, the algorithm internally applies (6) to the coarse conductivity parameters, σ^c , which represents the image to be reconstructed. The interpolation scheme intervenes on two occasions.

- *Forward computation.* In order to perform a forward computation, the coarse mesh conductivity is projected onto the fine mesh as $\sigma_n^f = P\sigma_n^c$, at every iteration n and supplied to the FEM forward solver.

Jacobian computation. Applying (6) to the coarse mesh conductivity requires computation of the Jacobian on the coarse grid. The (i, j) th element of the Jacobian is defined as

$$J(i, j) = \frac{\partial \mathbf{V}(i)}{\partial \sigma^c(j)}. \quad (8)$$

The Jacobian can be expressed in terms of the fine mesh conductivity by using the differentiation chain rule

$$\frac{\partial \mathbf{V}(i)}{\partial \sigma^c(j)} = \sum_k \frac{\partial \mathbf{V}(i)}{\partial \sigma^f(k)} \frac{\partial \sigma^f(k)}{\partial \sigma^c(j)}, \quad (9)$$

where k indexes the elements of σ^f that depend on $\sigma^c(j)$ through the interpolation scheme. As the fine mesh conductivity depends linearly, through P , on the coarse mesh conductivity, we have

$$\frac{\partial \mathbf{V}(i)}{\partial \sigma^c(j)} = \sum_k \frac{\partial \mathbf{V}(i)}{\partial \sigma^f(k)} P(k, j). \quad (10)$$

Recognizing that $\frac{\partial \mathbf{V}(i)}{\partial \sigma^f(k)}$ is the (i, k) th element of the Jacobian on the fine mesh yields

$$J^c(i, j) = \sum_k J^f(i, k) P(k, j) \quad (11)$$

resulting in

$$J^c = J^f P, \quad (12)$$

where we have used again the c and f superscripts to indicate the coarse and fine grids, respectively. Equation (12) projects the fine mesh Jacobian onto the coarse grid, and is used in (6). In practice, the full fine mesh Jacobian matrix is never constructed and stored in memory, rather single rows are computed and immediately projected onto the coarse mesh Jacobian. The Jacobian on the fine mesh, J^f , can be computed efficiently with the commonly used lead field method (Geselowitz 1971, Polydorides and Lionheart 2002) which states that

$$\delta \mathbf{V}(i) = - \int_{\Omega} \delta \sigma^f \mathbf{E}_{\text{applied}}(p) \cdot \mathbf{E}_{\text{lead}}(\ell), \quad (13)$$

where $\delta \mathbf{V}(i)$ is the variation in the i th current measurement on the electrode ℓ , resulting from application of the pattern p , $\mathbf{E}_{\text{applied}}(p)$ is the electric field generated in the body Ω , by applying the voltage pattern p , and $\mathbf{E}_{\text{lead}}(\ell)$ is the electric field that results from application of a unit stimulus to the ℓ th electrode.

3. Image reconstruction results

In this section we report validation results of the reconstruction algorithm described in the preceding sections on synthetic and experimental data. In the present paper we focus on reconstruction of synthetic data, where we use additive noise to simulate real-world conditions, and we show a single reconstruction from experimental data. In a second paper (Wan *et al* 2010), we present a thorough experimental evaluation of the imaging capabilities of the EIT prostate probe/algorithm and describe in detail the acquisition system used for collecting EIT data, and the physical experimental setup used for testing the probe.

In all the imaging experiments we reconstruct data in a sub-volume of the entire mesh. Angularly the imaging volume has an aperture of 140° , 70° per side. This corresponds to the aperture of the ultrasound field, as determined by the electrode arrangement, which has a central window of a given width (see figure 1). Vertically the imaging volume has an extension of 7 cm, which covers the extension of the electrode array, and we image to a maximum radial distance of 6 cm from the axis of the probe, as the sensitivity decays too much at larger distances. A coarse grid of pixels is formed in this space, with 10 pixels along the radial direction, 14 pixels along the angular direction and 14 pixels along the vertical direction.

In the numerical experiments, synthetic data were generated simulating the presence of a spherical inclusion in a homogeneous background. The background is assumed to have a conductivity of 0.1 Sm^{-1} , while the inclusion a conductivity of 0.15 S m^{-1} . This level of contrast is representative of the benign and malignant contrast observed in the prostate (Halter *et al* 2009a, 2009b). The spherical inclusion has a diameter of 1 cm and has been moved in several positions. In figure 4 the left column shows cross sections in the axial plane of the test conductivity profiles, while the right column shows cross sections in the axial plane of corresponding reconstructions. From top to bottom, the spherical inclusion has been moved along the 0° radial direction at three different distances from the axis of the probe: 2.75, 3.65 and 4.55 cm, respectively. In figure 5 the same spherical inclusion is moved along the left 60° radial direction, almost at the edge of the imaging field, to verify the capability of the probe to image laterally. The sphere is placed also in this case at three locations, having a radial distance of 2.75, 3.65 and 4.55 cm, respectively.

A tetrapolar measurement protocol was adopted to simulate data using an optimal set of 405 linearly independent measurements. The protocol, described in Borsic *et al* (2009), selects the optimal current injection and voltage sense electrode pairs such that sensitivity is maximized in regions where sensitivity is poor due to the rapid spatial decay of applied current densities. A noise vector $\mathbf{n}$ of 405 samples uniformly distributed in the range $[-1, +1]$ was generated and normalized to the voltage data $\mathbf{V}$, forming a normalized noise $\hat{\mathbf{n}}$ as

$$\hat{\mathbf{n}} = \mathbf{n} \frac{\text{std}(\mathbf{V})}{\text{std}(\mathbf{n})}, \quad (14)$$

where $\text{std}(\cdot)$ is the standard deviation of a vector. Noise was added to the data as $\mathbf{V} + \alpha \hat{\mathbf{n}}$, where α is the desired noise level; with $\alpha = 0.01$, we indicate a noise level of 1%. While other studies use Gaussian additive noise for simulating instrumentation noise, we use uniformly distributed noise as we found it to be more representative of measured data in terms of being evenly distributed within a certain range, and in terms of being bounded.

Based on our numerical experiments, and as studied in Borsic *et al* (2009) and Jossinet *et al* (2002a), the open-domain nature of this specific EIT arrangement, where electrode sensing is performed on the outer volume of the electrodes, is particularly ill-posed, due to a very rapid spatial decay of sensitivity with increasing distance from the probe. Simulations demonstrate that the maximum level of noise that the reconstruction could tolerate was approximately 0.1% without giving rise to significant image artifacts or without requiring values of the Tikhonov factor so large that the image would result overly smooth. In this respect, this EIT application is more challenging than a typical closed-domain problem, but we are able to show successful experimental reconstructions, at least for EIT difference data. Reconstructions presented in figures 4 and 5 show in blue an increased conductivity and in red a decreased conductivity. The angular and radial position of the spherical inclusion can be correctly identified from the reconstructed image. The red areas of decreased

conductivity are artifacts caused by the noise in the data and in part are intrinsic in the method: in the presence of no simulated noise, these artifacts were significantly reduced but still present.

Figure 6 represents a reconstruction from experimental data. In this case, the probe was immersed in a tank containing a saline solution with a conductivity of 0.1 S m^{-1} . An initial set of measurements were collected as reference data. A stainless steel ball, representing a high-contrast conductor, was then suspended in front of the probe. The ball has a diameter of 1.28 cm and was suspended at a distance of 2.8 cm from the axis of the probe, and vertically in the center of the electrode array. A second set of EIT measurements was collected and used for reconstructing a difference EIT image with respect to the empty tank. Figure 6 shows a vertical cross section and an axial cross section of the three-dimensional reconstruction. The blue regions indicate an increase in conductivity and the red a decrease in conductivity. The reconstruction successfully identifies the position of the metal ball and shows minor artifacts, smaller than those in the simulations, in part because the contrast of the inclusion is stronger. This example, coupled with those presented in Wan *et al* (2010) demonstrate that it is possible to use the developed prostate EIT probe to successfully image the conductivity in a three-dimensional volume in front of the probe.

4. Reconstruction speed optimization

In reconstructing EIT images from the prostate probe we use fine meshes with approximately 100 000 nodes and 500 000 elements. These FEM meshes are quite large in comparison to what is used traditionally in EIT, but are necessary for modeling accurately the details of the electrodes and interelectrode gaps on the probe tip. Also, because this is an open-domain problem, the imaging volume around the probe needs to be modeled, as a cylindrical form, which contributes to the larger meshes. To minimize the impact of this large domain, the size of the mesh elements is controlled by increasing the element size with distance from the probe. Despite this, reconstruction still results in a significant computational effort.

In this section, we discuss optimizations, in terms of increased computing speed, when compared to traditional plain MATLAB implementation of the reconstruction algorithm. These results are general and applicable to other soft-field tomography applications.

The optimizations and tests we report were produced on an 8-core, shared memory, PC architecture. The PC is based on two quad-core Xeon 5355 CPUs, with an internal clock frequency of 2.66 GHz and a front-side bus speed of 1.33 GHz, for a total of eight computational Xeon 64-bit cores. In terms of software environment, all the tests were run under Ubuntu Linux, version 7.05 64-bit, using MATLAB 2009a 64-bit with eight threads enabled.

In table 1 we report timing results for a single iteration of the reconstruction algorithm, and for the three main functions of the algorithm: computing the forward solution, computing the Jacobian, with (12) and (13), and computing the conductivity update with (6). The reconstruction benchmarking is relative to a mesh with 97 973 nodes and 541 604 tetrahedral elements, and using a coarse grid with $11 \times 14 \times 14$ conductivity elements. The reconstruction is complex valued, estimating conductivity and permittivity distributions.

A single iteration of the algorithm, with a parabolic line search on the conductivity update, takes approximately 25 min, resulting in reconstruction times on the order of 1/2 h if few iterations are performed in order to solve the nonlinear inverse problem. These lengthy computing times make the process of reconstructing EIT images noninteractive, and negatively impact the opportunity of exploring different data sets and different

reconstruction parameters. On the basis of these considerations, we set out to optimize the execution speed. Specifically, we focused on optimizing the forward solution and the computation of the Jacobian, which jointly account for 93.5% of the total execution time.

4.1. Forward solving

In terms of forward solving, the computational time is determined primarily by the sparse system solver used for solving the linear system derived from the FEM method. In this regard we conducted a literature search and selected a few fast solvers for benchmarking our problems. A survey published in Gould *et al* (2005) reports performance benchmarking results for several public domain and commercial sparse linear solvers. In this survey UMFPACK (<http://www.cise.ufl.edu/research/sparse/umfpack/n.d.>), the sparse linear solver used by MATLAB emerges as one of the poorest performers, while in the same survey, PARDISO (PARallel Direct Solver) (<http://www.pardiso-project.org/n.d.>) emerges as one of the best performing solvers in the survey group. We decided to test PARDISO. The solver is available publicly online, and also in an optimized implementation, through Intel's Math Kernel Library. PARDISO can solve real- and complex-valued, symmetric and unsymmetric systems, and is based on a left-right looking LU factorization technique optimized for shared memory systems (Schenk *et al* 2000). PARDISO accepts a sparse matrix format different from MATLAB's compressed sparse column (CSC) format. PARDISO uses a compressed sparse row (CSR) format and requires the first entry of each row to be the diagonal value (if not zero). We therefore developed a MATLAB mex file that performs this sparse matrix conversion and calls PARDISO.

Table 2 reports benchmarking results for PARDISO, from Intel's Math Kernel Library 10.1, and MATLAB's backslash operator, for release R2009a. PARDISO offers significant speedups compared to the use of the backslash operator in MATLAB, resulting in 7–10 speed increases, depending on the test case. The gains are better when multiple threads are used, as PARDISO offers a better scalability. For example, in the complex case the MATLAB's backslash operator has a speedup of 2.2 going from the use of one thread to the use of eight threads. PARDISO has a speedup of 2.9 in going from one thread to eight threads, so the relative performance gap between the two solvers becomes larger when a large number of threads are used.

Other efforts to speed up EIT forward solving are the ones of Horesh *et al* (2006) and of Soleimani *et al* (2005). PARDISO performance compares favorably with these two works. In their study on multilevel preconditioners, Horesh *et al* used a single-threaded solver on a dual Xeon workstation with a CPU frequency of 2.8 GHz, operating under Linux. Given that both their architecture and ours are based on Xeon processors with similar CPU frequencies, the timing results for single-threaded execution on the two architectures are comparable. For complex-valued problems with 100K unknowns, Horesh *et al* (2006) show graphs where the solution time for a single right-hand side (RHS) is about 10 s, resulting in approximately 300 s for 30 RHSs, plus approximately ten additional seconds for the computation of the preconditioner. PARDISO, on a single thread, solved a complex-valued system with 100K unknowns and 30 RHSs in 58 s, about 5.3 times faster. Soleimani *et al* (2005), in comparison between an un-preconditioned conjugate-gradient method and an algebraic multigrid preconditioned conjugate-gradient algorithm, report a solution time of 13 s per RHS for the latter, but they do not indicate the specifics of the computing platform on which the tests were performed.

On the basis of our actual tests and comparisons with available literature results, we therefore adopted PARDISO as a forward solver, as it proved to be significantly faster than other known solvers.

4.2. Jacobian computation

Computation of the Jacobian consumes the largest amount of CPU time in our reconstruction algorithm, accounting for 56% of the total cost. The Jacobian is computed on the fine mesh and projected onto the coarse grid/mesh via the interpolation matrix in (12), $J^c = J^f P$. In our implementation, we construct a single row of J^f at a time and immediately

project it onto the corresponding i th row of the coarse mesh Jacobian, $J_i^c = J_i^f P$, retaining in memory only one row of the fine mesh Jacobian at any time.

Inspection of the computations involved reveals that approximately 6 GFlops are required for the Jacobian construction, when the 97 973-node cylindrical mesh is used, and when 405 tetrapolar measurements are used. The large number of computations accounts, at least in part, for the long computing times necessary to form the Jacobian. In order to optimize this computation, we implemented this function as a C language mex file, making calls to functions of Intel's Math Kernel Library (MKL) and Intel's Integrated Performance Primitives (IPP) library. The implementation is multi-threaded, but almost the same performance was obtained using one or eight threads. A low-level performance analysis revealed that speed gains by the use of multiple processors were limited, in our system and for this algorithm, by the bandwidth between main memory and the processors. The time spent computing the Jacobian is ultimately bound by the speed at which the computed applied and lead fields in (13) can be fetched from memory (since these vectors do not fit in local cache memory of the CPU). Under these circumstances, even a single processor is spending part of its computing time waiting to fetch data. Multiple processors compete for access to data which does not improve performance. Performance gains of the optimized routine, compared to traditional MATLAB implementation, are however significant and are on the order of 16 times, as reported in table 3.

4.3. Overall gains

The combined gains resulting from optimization of the forward solution and of the Jacobian computation of our reconstruction algorithm are significant with respect to a standard MATLAB implementation as illustrated in table 4. For the mesh in use, with 97 973 nodes and 541 604 tetrahedral elements, the speed optimizations result in a relative gain of 7.4 times, bringing the computing time for a single iteration of the algorithm from approximately 25 min down to 3.4 min. This gain is particularly significant in absolute terms when 2/3 iterations of the algorithm are run, to perform a nonlinear reconstruction. In this case, of three iterations, the optimized algorithm terminates in about 10 min compared to almost 1 h and a half with the traditional algorithm implementation. These gains allow significantly better interaction, and therefore better insight in data by experimenting with different reconstruction parameters.

5. Conclusions and future work

In this paper we present an algorithm for tomographic reconstruction for a custom transrectal probe that couples electrical impedance tomography (EIT) and ultrasound to image the prostate. Modeling of the probe with a FEM mesh and an arrangement that allows a coarse representation of the conductivity on the underlying fine forward mesh are discussed. The overall algorithm is shown to produce successful reconstructions on synthetic data with additive noise and on phantom experimental data. This specific application of EIT image reconstruction is particularly challenging as applied currents densities decrease very rapidly with increasing distance from the probe, making it difficult to reconstruct sources of contrast past a certain distance, and rendering the problem particularly ill-posed. We show however that it is possible to image sources of moderate contrast at a distance of 4.55 cm from the probe, which is compatible with the clinical use of the probe. In the clinical

application that we envision, where impedance information will be used as an overlay to ultrasound information, we will use structural US information to guide EIT reconstruction, possibly overcoming part of the effects of the pronounced ill-conditioning, and potentially extending further the volume that can be imaged. Work in this direction is currently undergoing in our research group.

In the present paper, we show results relative to speeding up EIT reconstruction. These results are significant as we use particularly fine meshes in this application, and we can obtain an overall gain of 7.4 times by implementing the most computationally intensive functions of the algorithm as C-mex files that are called from the MATLAB environment. We envision, in the clinical application, the need of a near-realtime reconstruction capability, as EIT information should be used to guide the biopsy procedure. In order to further speed up reconstruction, we are currently implementing the reconstruction algorithm on graphic processing units (GPU), which from preliminary analysis should result in further significant reductions in computing time, thanks to higher computing power and larger bandwidth to memory.

Acknowledgments

This work is partially supported by NIH grant 5R01CA124925.

References

- Borsic A, Halter R, Wan Y, Hartov A, Paulsen KD. Sensitivity study and optimization of a 3D electric impedance tomography prostate probe. *Physiol Meas*. 2009; 30:S1–18. [PubMed: 19491445]
- Geselowitz DB. An application of electrocardiographic lead theory to impedance plethysmography. *IEEE Trans Biomed Eng*. 1971; 18:38–41. [PubMed: 5545046]
- Gould, NIM.; Hu, Y.; Scott, JA. Technical Report RAL-TR-2005-005. CCLRC Rutherford Appleton Laboratory; Oxfordshire, UK: 2005. A numerical evaluation of sparse direct solvers for the solution of large sparse, symmetric linear systems of equations.
- Halter R, Schned A, Heaney J, Hartov A, Paulsen K. Electrical properties of prostatic tissues: I. Single frequency admittivity properties. *J Urol*. 2009; 182:1600–7. [PubMed: 19683745]
- Halter R, Schned A, Heaney J, Hartov A, Paulsen K. Electrical properties of prostatic tissues: II. Spectral admittivity properties. *J Urol*. 2009; 182:1608–13. [PubMed: 19683743]
- Holder, D. Electrical Impedance Tomography: Methods, History and Applications. Bristol: Institute of Physics Publishing; 2004.
- Horesh L, Schweiger M, Bollhofer M, Douiri A, Holder DS, Arrige SR. Multilevel preconditioning for 3D large-scale soft-field medical applications modelling. *Int J Inf Sys Sci*. 2006; 2:532–56. <http://www.cise.ufl.edu/research/sparse/umfpack/n.d>
- <http://www.pardiso-project.org/n.d>
- Jossinet J, Marry E, Matias A. Electrical impedance endotomography. *Phys Med Biol*. 2002; 47:2189–202. [PubMed: 12164581]
- Jossinet J, Marry E, Montalibet A. Electrical impedance endotomography: imaging tissue from inside. *IEEE Trans Med Imaging*. 2002; 21:560–5. [PubMed: 12166851]
- Nocedal, J.; Wright, SJ. Numerical Optimization. Berlin: Springer; 1999.
- Polydorides N, Lionheart WRB. A MATLAB based toolkit for three-dimensional electrical impedance tomography: a contribution to the EIDORS project. *Meas Sci Technol*. 2002; 13:1871–83.
- Schenk O, Gartner K, Fitcher W. Efficient sparse LU factorization with left–right looking strategy on shared memory multiprocessors. *BIT*. 2000; 40:158–76.
- Soleimani M, Powell CE, Polydorides N. Improving the forward solver for the complete electrode model in EIT using algebraic multigrid. *IEEE Trans Med Imaging*. 2005; 24:577–82. [PubMed: 15889545]

- Somersalo E, Cheney M, Isaacson D. Existence and uniqueness for electrode models for electric current computed tomography. *SIAM J Appl Math.* 1992; 52:1023–40.
- Wan Y, Halter R, Borsic A, Manwaring P, Hartov A, Paulsen K. Sensitivity study of an ultrasound coupled transrectal electrical impedance tomography system for prostate imaging. *Physiol Meas.* 2010; 31:S17–29. [PubMed: 20647618]

Figure 1.

Arrangement of the electrodes used in the combined EIT/ultrasound (US) probe. A total of 30 electrodes are arranged in a rectangular pattern and printed on a flexible substrate. This EIT sensing array is applied to the US imaging probe, and the central 'window', not covered by electrodes, allows the US scanning element, housed inside the probe, to scan the volume of the prostate with no interference. (a) Electrode layout. (b) FEM model of the probe tip and electrodes.

Figure 2.

Top and side views of the FEM mesh used for forward solving. The mesh consists of a 24 cm diameter cylinder with the probe embedded—imaging occurs in the space outside of the probe. The yellow dots show ‘seed points’ that are used for generating a coarse conductivity representation, as detailed in section 2.4. (a) Top view of the mesh used for forward modeling, and in yellow ‘seed points’ for the formation of the coarse conductivity grid. (b) Side view of the mesh used for forward modeling, and in yellow ‘seed points’ for the formation of the coarse conductivity grid.

Figure 3.

Coarse conductivity grid: the checkerboard shows the coarse conductivity grid used for representing conductivity values to be estimated in the image reconstruction process. The grid results from the grouping of tetrahedra from the finer mesh around ‘seed points’ to form coarse ‘pixels’. This grid was generated using 14 ‘seed points’ on the angular direction, 8 points on the radial direction and 11 points on the vertical direction. The jagged appearance of each coarse pixel is due to the grouping of several tetrahedra to form the pixel itself.

Figure 4.

Synthetic data. Cross sections of three-dimensional conductivity test profiles (left column) and corresponding reconstructions (right column). The test conductivity profiles consist in a homogeneous background of 0.1 S m^{-1} with a spherical inclusion having a conductivity of 0.15 S m^{-1} , a diameter of 1 cm and a variable position. The inclusion is moved outward along the 0° radial direction, which is centered with the electrode array. All the simulations have 0.1% additive noise and a single common value for the Tikhonov factor, which has been chosen visually. Test conductivity profiles are plotted on a color scale in the range of $0.1\text{--}0.15 \text{ S m}^{-1}$, while reconstructed images are plotted on a color scale covering the range $0.08\text{--}0.12 \text{ S m}^{-1}$. (a) Test profile with a 1 cm spherical contrast at 2.75 cm from the axis of the probe. (b) Reconstruction corresponding to the test profile on the left. (c) Test profile with a 1 cm spherical contrast at 3.65 cm from the axis of the probe. (d) Reconstruction corresponding to the test profile on the left. (e) Test profile with a 1 cm spherical contrast at 4.55 cm from the axis of the probe. (f) Reconstruction corresponding to the test profile on the left.

Figure 5.

Synthetic data. Cross sections of three-dimensional conductivity test profiles (left column) and corresponding reconstructions (right column). The test conductivity profiles consist in a homogeneous background of 0.1 S m^{-1} with a spherical inclusion having a conductivity of 0.15 S m^{-1} , a diameter of 1 cm and a variable position. The inclusion is moved outward along the left 60° radial direction, which corresponds almost to the leftmost edge of the imaging volume (aperture 70° per side). All the simulations have 0.1% additive noise and a single common value for the Tikhonov factor, which has been chosen visually. Test conductivity profiles are plotted on a color scale in the range of $0.1\text{--}0.15 \text{ S m}^{-1}$, while reconstructed images are plotted on a color scale covering the range $0.08\text{--}0.12 \text{ S m}^{-1}$. (a) Test profile with a 1 cm spherical contrast at 2.75 cm from the axis of the probe. (b) Reconstruction corresponding to the test profile on the left. (c) Test profile with a 1 cm spherical contrast at 3.65 cm from the axis of the probe. (d) Reconstruction corresponding to the test profile on the left. (e) Test profile with a 1 cm spherical contrast at 4.55 cm from the axis of the probe. (f) Reconstruction corresponding to the test profile on the left.

Figure 6.

Experimental data. Vertical (left) and axial (right) cross sections of three-dimensional reconstructions of a phantom consisting in a uniform saline bath of a conductivity of 0.1 S m^{-1} and of a metallic ball, of a diameter of 1.28 cm, suspended in front of the electrode array at a distance of 2.88 cm from the axis of the probe. Images are relative to different reconstructions, where the empty tank has been used as reference data. The blue spot in the reconstructed images, which corresponds to an increased conductivity, indicates correctly the location of the metal ball. (a) Vertical cross section of a three-dimensional reconstruction of phantom data. (b) Axial cross section of a three-dimensional reconstruction of phantom data.

Table 1

Timing of a single plain MATLAB-based reconstruction iteration—eight computing threads.

Function	Time (s)	Time (%)
Forward solve ($\times 3$ times) ^a	570	37.5
Jacobian computation	853	56.0
Conductivity update	93	6.5
Total	1516	100.0

^aThe forward solution is computed three times in a single iteration of the algorithm, as a parabolic line search is performed on the conductivity update.

Table 2

Forward solve timing results.

Solver/problem	CPU time 30 RHSs (s)			
	One thread	Two threads	Four threads	Eight threads
MATLAB real	138.74	102.24	85.43	83.74
PARDISO real	18.46	11.19	7.66	6.91
MATLAB complex	418.84	277.22	207.75	190.18
PARDISO complex	58.60	34.21	21.05	19.56

Table 3

Timing of Jacobian computation.

MATLAB routine (s)	Optimized routine (s)
853	52

Table 4

Timing comparison of single reconstruction step.

Implementation	Fwd. Solve $\times 3$ (s)	Jacobian (s)	Inv. Solve (s)	Total (s)	Gain
Plain MATLAB	570	853	93	1516	–
Optimized	59	52	93	204	7.4