Accelerated Markov Chain Monte Carlo Sampling in Electrical Capacitance Tomography

Daniel Watzenig*, Markus Neumayer* and Colin Fox†

*Institute of Electrical Measurement and Measurement Signal Processing, Graz University of Technology, Kopernikusgasse 24/4, A-8010 Graz, Austria

† Department of Physics, University of Otago, 730 Cumberland Street, Dunedin 9016, New Zealand

E-mail: daniel.watzenig@TUGraz.at

***Abstract*—Electrical Capacitance Tomography is an ill-posed inverse problem that aims at recovering the spatial permittivity distribution of an inhomogeneous medium from capacitance measurements at the boundary. We consider the problem of fast robust estimation of inclusion shape and position in binary mixtures. The boundary of the inclusion is represented implicitly using a radial basis function representation. The inverse problem is formulated as Bayesian inference, with Markov chain Monte Carlo sampling used to explore the posterior distribution. An affine approximation to the forward map built over the state space significantly reduces reconstruction time, while introducing minimal extra error. Numerical examples are presented for synthetic data sets, avoiding all inverse crimes.**

***Index Terms*—Electrical capacitance tomography, model reduction, MCMC sampling, statistical inversion**

I. Introduction

Electrical Capacitance Tomography (ECT) is a noninvasive method that aims at recovering the spatial permittivity distribution within an object from measurements of the capacitances between electrodes at the boundary of the object [1]. This measurement process can be modeled as

$$d = F(\theta) + n \tag{1}$$

where d denotes the measured data, θ are parameters representing the the permittivity, and F denotes the forward map that is the deterministic mapping from parameters θ to noise-free data. Measurement noise is denoted by the random variable n.

Deterministic solution methods vary the parameter θ to minimize some norm of the difference between the model output $F(\theta)$ and measured data d. Due to the ill-posed nature of the problem, the norm has to be augmented by a regularization term to obtain a stable solution. The balance between the data misfit term and the regularization term is controlled by a regularization parameter. Typically gradient or even Hessian information about the map F are used to accelerate the convergence of the optimization procedure.

Bayesian statistical inversion follows a different approach. Uncertainties are included as probabilities, to give a 'posterior' probability distribution over parameters that are consistent with measured data, thereby quantifying uncertainty in the solution. Solutions to an ill-posed inverse problem are then well determined problems of statistical inference over the posterior distribution. This paradigm has many advantages over deterministic approaches, such as giving robust estimates, and the ability to treat arbitrary forward maps and error distributions. One non-obvious advantage is the ability to use a wide range of representations of the unknown permittivity, including parameter spaces that are discrete, discontinuous, or even variable dimension.

Whereas regularization gives point estimates, Bayesian methods present averages over all solutions consistent with the data. This leads to a marked improvement in robustness of properties calculated from solutions, since the regularized solution is typically unrepresentative of the bulk of feasible solutions in high dimensional nonlinear problems. However, the better results achieved by Bayesian methods come at the cost of greater modeling and computational effort.

The primary cost in computational implementations of Bayesian inference is the repeated evaluation of the forward map F when using Markov chain Monte Carlo (MCMC) algorithms, that we discuss in section III. Hence, an efficient implementation of the forward map is mandatory to obtain results within an acceptable time. In this paper we use a fast, affine, approximation to F to accelerate the MCMC sampling.

Recently, two approaches have been developed to allow the use of approximate forward maps within statistical inversion, while maintaining the accuracy of calculated uncertainties. The *enhanced error model* augments the noise model by a Gaussian model for the difference between accurate and approximate forward maps [2], and has been shown to drastically reduce computation time at the expense of a small increase in uncertainty. The second is the delayed-acceptance MCMC [3] that modifies the MCMC algorithm, to sample from the correct posterior distribution while avoiding computation of the exact forward map for most steps.

The paper is structured as follows: The next section briefly introduces the physical forward problem of ECT and the modeling of closed contours used to represent material inclusions. Section III addresses the formulation of ECT in a Bayesian framework, and the MCMC algorithm. Section IV introduces two strategies for accelerating the MCMC algorithm. Reconstruction results and a comparison of the achieved improvement in terms of computational effort are presented in section V.

II. Framework for electrical capacitance tomography

A. The forward problem

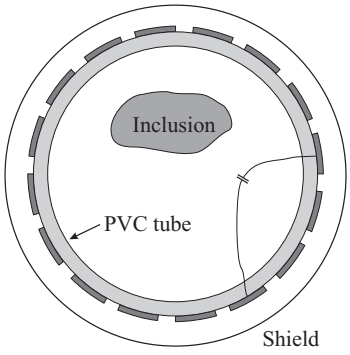

Fig. 1. Schematic of an ECT-sensor.

Figure 1 shows a schematic of a typical ECT sensor for process tomography applications. N_{elec} electrodes are mounted on the exterior of a pipe. The whole problem domain is bounded by an electrical shield. The aim is the reconstruction of the material distribution (i.e. the shape and position of an inclusion) in the pipe using 120 independent measurements of the coupling capacitances between the $N_{\text{elec}} = 16$ electrodes.

The electrostatic field problem is modeled by the governing equation

$$\nabla \cdot (\varepsilon_0 \varepsilon_r \nabla \Phi) = 0 \tag{2}$$

where ε_0 and ε_r denote the absolute and the relative permittivity, and Φ denotes the electric scalar potential. The boundary conditions are of Dirichlet type and given by

$$\Phi = V_{\text{T}} \quad \text{on } \Gamma_{\text{T}} \tag{3}$$

$$\Phi = 0 \quad \text{on } \Gamma_{\text{R}} \text{ and } \Gamma_{\text{shield}} \tag{4}$$

where Γ_{T} represents the boundary of the transmitting electrode and Γ_{R} refers to the boundaries of the receiver electrodes. The coupling capacitances between the active electrode and each receiver electrode can be obtained by evaluating Gauss's law,

$$C_i = -\frac{1}{V_{\text{T}}} \int_{\Gamma_{R,i}} \vec{n} \cdot \varepsilon \nabla \Phi d\Gamma \tag{5}$$

for each receiver electrode.

Equation (2) is solved by means of the finite element method (FEM). For the computation of the forward problem in this paper, a mesh with about 560 finite elements is used.

B. Shape description using RBFs

The representation of unknowns within a reconstruction algorithm forms a fundamental classification between different algorithms. In general one can distinguish between algorithms which either provide pixel or voxel images, that aim at recovering the shapes of objects. Since we are interested in reconstructing shape and position of piecewise constant material inclusions in an otherwise uniform background material, we use a contour model to describe closed boundaries. In general one can distinguish between non-parametric models (i.e. level sets) and parameterized shape models (e.g. spline models, Fourier descriptions). In this work we use a shape description based on radial basis functions (RBFs) [5]. With the RBF approach an object is represented in implicit form given by $f(\boldsymbol{x_i}) = 0$, where $\boldsymbol{x}_j$ represents the Cartesian coordinates of a point on the boundary of the shape. The function f is of form

$$f(\boldsymbol{x}_i) = \sum_j^N \lambda_i \phi(||\boldsymbol{x}_i - \boldsymbol{c}_j||) \tag{6}$$

where $\boldsymbol{c}_i$ represent given locations on the boundary of the object (i.e. scatter data) and the function ϕ includes the individual RBFs. We use the thin plate splines

$$\phi(r) = r^2 \cdot \log(r), \tag{7}$$

where r is the Euclidean distance between two points. The weights λ_i can be found by forming the constraint $f(\boldsymbol{x}_i = \boldsymbol{c}_i) = h_i$, leading to the linear system of equations

$$\boldsymbol{A}\boldsymbol{\lambda} = \boldsymbol{h}, \tag{8}$$

which can be solved for $\boldsymbol{\lambda}$. With thin plate splines, the set of basis functions needs to be augmented by a linear function $P((x, y)) = ax + by + c$, with equation (6) is extended to

$$f(\boldsymbol{x}_i) = \sum_j^N \lambda_i \phi(||\boldsymbol{x}_i - \boldsymbol{c}_j||) + P(\boldsymbol{x}_i). \tag{9}$$

Equation (8) is extended to

$$\begin{bmatrix} \boldsymbol{A} & \boldsymbol{Q} \\ \boldsymbol{Q}^T & \boldsymbol{0} \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda} \\ \boldsymbol{c}_p \end{bmatrix} = \begin{bmatrix} \boldsymbol{h} \\ \boldsymbol{0} \end{bmatrix}, \tag{10}$$

where $\boldsymbol{Q} = \begin{bmatrix} \boldsymbol{x}_p & 1 \end{bmatrix}$ and $\boldsymbol{c}_p = \begin{bmatrix} a & b & c \end{bmatrix}^T$. The contour of an object can be determined by solving (10) for $f(\boldsymbol{x_i}) = 0$ given the locations $\boldsymbol{x}_p$, and asserting a nonzero value interior to the object.

The reconstruction algorithm manipulates the boundary of the inclusion through locations $\boldsymbol{x}_p$. The permittivity is then mapped onto the finite element grid for computation of the forward problem.

III. Statistical Inversion

The presence of uncertainties, such as ubiquitous measurement noise, means that the measurement and imaging processes are probabilistic and the inverse problem is naturally a stated as statistical inference. In the Bayesian formulation, inference about θ is based on the posterior distribution

$$\pi(\theta|d) = \frac{\pi(d|\theta)\pi(\theta)}{\pi(d)} \propto \pi(d|\theta)\pi(\theta) \tag{11}$$

where $\pi(\theta)$ is the prior density expressing the information about θ independent of the measurements d, and $\pi(d|\theta)$ is the likelihood function that shows how measurements affect knowledge about θ. The posterior distribution $\pi(\theta|d)$ is the probability distribution over θ given the prior information and the measurements. The solution of the inverse problem usually involves summarizing the posterior distribution, though in exploratory analyses we often simply display several reconstructions drawn from the posterior distribution. The denominator

$\pi(d)$ is a finite normalizing constant once measurements d are made.

Given the posterior distribution, any statistics of interest can be computed. Commonly used summary statistics are the maximum *a posteriori* (MAP) estimate

$$\theta_{MAP} = \arg\max \pi(\theta|d) \tag{12}$$

and the conditional mean

$$\theta_{CM} = \int \theta \pi(\theta|d) d\theta. \tag{13}$$

In addition, statements on parameter variability and reliability of solutions can be given. Another appealing feature is the the ease of taking into account prior information about parameters.

In the case of Gaussian measurement noise the likelihood can be written as

$$\pi(d|\theta) \propto \exp\left(-\frac{1}{2}\left(d - F(\theta)\right)^T \Sigma^{-1} \left(d - F(\theta)\right)\right), \tag{14}$$

where Σ denotes the covariance of the measurement noise. For a detailed review about this topic we we refer to [6].

Analytic evaluation of integrals such as (13) is intractable in practical cases because the posterior distribution is complex and a function of many variables. Hence, numerical methods that draw samples from $\pi(d_m|\theta)$ are used to explore the feasible solution space, and to evaluate integrals using Monte Carlo integration.

A. Markov Chain Monte Carlo sampling with Metropolis Hastings Green kernel

The efficient exploration of the posterior distribution $\pi(d_m|\theta)$ becomes the major computational concern when implementing Bayesian inference. Grid based exploration and quadrature are only applicable if the space of θ is of low dimension. A very general class of algorithms for performing sample-based Bayesian inference are the Markov Chain Monte Carlo (MCMC) sampling methods, see e.g. [4].

We use MCMC with Metropolis Hastings Green (MHG) kernel, summarized in algorithm 1. Starting from a candidate θ, the MHG-MCMC algorithm proposes a new candidate θ' and computes the acceptance probability α (code line 5). The candidate θ' is accepted with probability α, and is otherwise rejected. The acceptance ratio includes the Jacobian determinant

$$|J_m| = \left|\frac{\partial(\theta', \gamma')}{\partial(\theta, \gamma)}\right| \tag{15}$$

for the mapping from composite parameter (θ, γ) to composite parameter (θ', γ').

Although the algorithm is simple to code, the computational burden is given by line 5, as the computation of $\pi(\theta'|d)$ requires the computation of the forward problem. As the number of sampling steps N_{MCMC} is typically several thousand [4], [6].

The generation of a new candidate requires varying the parameters θ in a random way. We typically use several 'moves' to make up the proposal distribution. A simple random walk of a single randomly chosen control point $\boldsymbol{x}_p$ is used to guarantee convergence, though leads to inefficient sampling by itself. We make further specific manipulations to increase the efficiency of the MCMC sampler. As the shape of an object is reconstructed, suitable moves are rotational and translational movements as well as scaling of the boundary [6].

Algorithm 1 Basic Structure of MH MCMC

1: Pick a valid initial state θ and the prior $\pi(\theta)$ and likelihood $\pi(d|\theta)$
2: **for** $i = 1$ to N_{MCMC} **do**
3: Generate a proposal candidate θ'
4: **if** $\pi(\theta') > 0$ **then**
5: Compute $d' = F(\theta')$ and the acceptance ratio $\alpha(\theta, \theta') = \min\left(1, \frac{\pi(d|\theta')\pi(\theta')}{\pi(d|\theta)\pi(\theta)}|J_m|\right)$
6: Draw $u \sim \mathcal{U}(0, 1)$
7: **if** $u < \alpha(\theta, \theta')$ **then**
8: $\theta = \theta'$
9: $\pi(\theta) = \pi(\theta')$
10: $\pi(d|\theta) = \pi(d|\theta')$
11: **end if**
12: **end if**
13: **end for**

Choice of the prior distribution $\pi(\theta)$ is critical, as it has a direct impact on the results. If the true material distribution has a low probability with respect to the prior distribution $\pi(\theta)$, the bulk of feasible probability will be biased away from the true distribution. For a review on prior models we refer to [4] and [6]. In the present work we use a prior which penalizes small contours, only allows contours wholly inside the pipe, and penalizes local curvature in the boundary contour.

IV. Accelerating MCMC methods

To accelerate the used sampling algorithm a strategy has to be found to decrease the number of costly evaluations of the forward map F in line 5 of the MCMC algorithm given in 1. Hence, it seems obvious that an approximation of F has to be introduced for the time-critical computation of the likelihood of the proposal. However, the use of an approximation $\tilde{F}$ causes a model error

$$\tilde{n} = F(\theta) - \tilde{F}(\tilde{\theta}). \tag{16}$$

In deterministic inversion theory this error may become critical. In the framework of Bayesian inversion the model error can be incorporated as an additional additive noise [2]. Assuming independence, the probability density function of the sum of the two noises is given by

$$\pi(n_{sum}) = \pi(n) * \pi(\tilde{n}), \tag{17}$$

which is the convolution of the model error and the measurement noise. The statistics $\pi(\tilde{n})$ of $\tilde{n}$ can be obtained by sampling, as we show later.

A. Delayed acceptance algorithm

The delayed acceptance Metropolis Hastings (DAMH) algorithm was introduced by A. Christen and C. Fox in 2005 [3], to gain advantage from a cheap state-dependent approximation of F for computing the likelihood function. In case of an

accepted state the accurate forward map is computed and a modified acceptance probability used to ensure convergence to the correct posterior distribution. Hence, the DAMH algorithm shows an excellent performance gain in the case of a low acceptance rate. A simple approximation is the first-order series expansion of F

$$F(\theta') \approx F(\theta) + \left.\frac{\partial F}{\partial \theta}\right|_{\theta} \Delta\theta, \tag{18}$$

where $\Delta\theta = \theta' - \theta$. Only a matrix operation has to be computed to evaluate the likelihood function of a new proposal. The second term in (18) denotes the Jacobian of F with respect to the elements of θ. We calculate the Jacobian with respect to the material values ε_r of the finite elements, and use mapping from contour θ to the finite element grid via the chain rule to evaluate desired Jacobian. In this work we compute the Jacobian matrix using an adjoint calculation. The time for this calculation can be removed by maintaining the Green's functions of a self-adjoint version of the forward map, as outlined in [3]. The local approximation given by (18) restricts proposal states θ' to be close to the current stat θ, to keep the approximation error low.

B. Approximated forward model

The use of the Jacobian for a local linear approximation is reasonable from the physical point of view. However, any approximation to the forward map may be used, even if not physically motivated. Typical examples for such approximations are neural networks or polynomial approaches, while coarse grid approximations are suggested by the numerical computation. In the following we propose using an affine transformation to approximate the accurate ECT forward map. This is motivated by the procedure that we use to calibrate the offset and gain of the ECT sensor [6].

Hence we approximate,

$$\boldsymbol{C} = \boldsymbol{P}\boldsymbol{\varepsilon}'_r, \tag{19}$$

where $\boldsymbol{\varepsilon}'_r$ is the augmented permittivity vector given by

$$\boldsymbol{\varepsilon}'_r = \begin{bmatrix} 1 & \boldsymbol{\varepsilon}_r^T \end{bmatrix}^T \tag{20}$$

where $\boldsymbol{\varepsilon}_r$ is the vector of relative permittivities on the finite element mesh. The augmentation is required to obtain an affine mapping.

To obtain the elements of the matrix $\boldsymbol{P}$, we use a least squares fit of form

$$\begin{bmatrix} 1 & \boldsymbol{\varepsilon}_{r,j}^T \\ 1 & \cdots \\ \vdots & \vdots \end{bmatrix} \boldsymbol{p}_i = \boldsymbol{c}_i \tag{21}$$

where $\boldsymbol{p}_i$ denotes the i-th row of the matrix $\boldsymbol{P}$ and $\boldsymbol{c}_i$ is the corresponding capacitance. The vectors $\boldsymbol{\varepsilon}_{r,j}$ are representitive samples of of expected permittivity distributions. These samples could be drawn from the prior distribution $\pi(\theta)$, however, we decided to generate samples of $\boldsymbol{\varepsilon}_{r,j}$ using the procedure in algorithm 2.

Using this approach to draw samples $\boldsymbol{\varepsilon}_{r,j}$ of the material distribution, a map $\boldsymbol{P}$ can be established by solving (21) in a least squares sense using the pseudo inverse. To evaluate the performance of the map $\boldsymbol{P}$ another set of samples of $\boldsymbol{\varepsilon}_{r,j}$ has been generated, and the error $\tilde{n}$ was calculated.

Algorithm 2 Algorithm to draw samples $\boldsymbol{\varepsilon}_{r,j}$.

```
Set all elements in the FEM grid to permittivity value one
Draw #_incl ~ U({1, 2, 3})
for i = 1 to #_incl do
    Draw (x_center, y_center) ~ U(in ROI)
    Draw ε_r ~ U(ε_r,min, ε_r,max)
    Draw r_incl ~ U(r_min, r_max)

    Map circular inclusion with parameters to FEM grid
end for
Compute the forward problem
```

Figure 2 shows the trend of the mean μ and the standard deviation σ of $\tilde{n}$ for all capacitances. Two relevant aspects can be observed from the results depicted in figure 2. Firstly, the mean μ is almost zero and, hence, the map $\boldsymbol{P}$ is unbiased. Secondly, the standard deviation σ is in the range of the typical measurement error of our ECT sensor, which has a standard deviation of $\sigma = 5 \times 10^{-3}$ [7]. The determined standard deviation is even smaller than that of a typical measurement error.

To verify the quality of the estimator, the signal to noise ratio (SNR) of our ECT sensor were performed. Figure 3 illustrates the measured SNR and the SNR of the approximation error. The SNR from the sensor was measured for an empty (air filled) sensor. One can see that the approximation has almost the same quality as the measurements.

Figure 4 depicts a specific distribution of $\tilde{n}$ for a particular measurement and an equivalent Gaussian distribution. The distribution $\pi(\tilde{n})$ is slightly non symmetric and has a small tail towards positive values of $\tilde{n}$. For simplicity reasons, we used an equivalent Gaussian distribution given the measured mean and the standard deviation. Hence, the distribution of

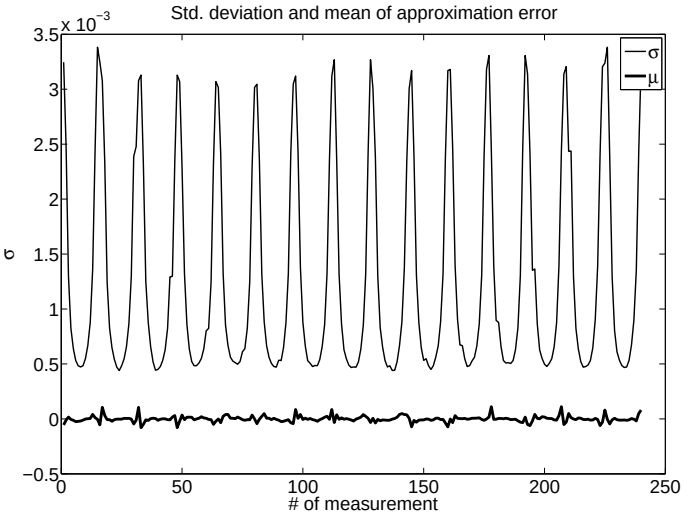

Fig. 2. Standard deviation and mean of the approximation error for the 15×16=240 measurements.

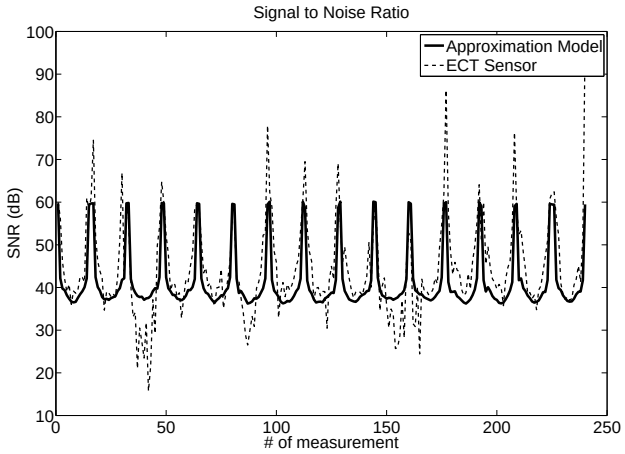

Fig. 3. Signal to noise ratio (SNR) of the approximation model and of the real ECT sensor.

$\pi(n_{sum})$ remains Gaussian, as the convolution of two Gaussian functions remains a Gaussian function.

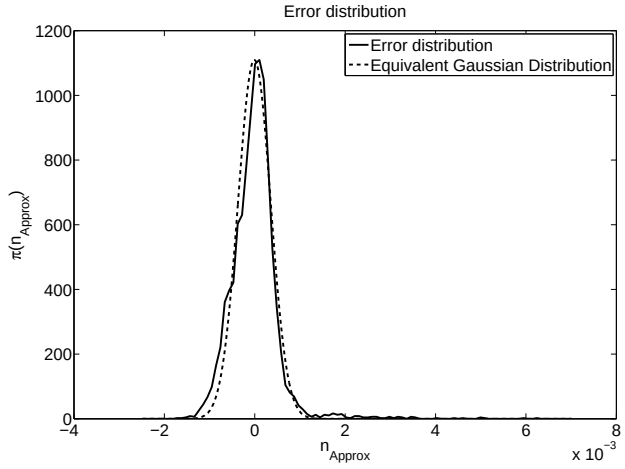

Fig. 4. Probability density function of the approximation error and approximated Gaussian fit.

V. Evaluation and results

To evaluate the performance of the different algorithms, we performed two reconstruction tasks. In the first case a circular object is reconstructed, in the second experiment a more complex contour consisting of two merged circular bubbles is reconstructed. In order to avoid an inverse crime, the reconstruction data was generated using a different finite element mesh (2662 finite elements). In addition, the data was corrupted with zero mean Gaussian noise with a variance of $\sigma^2 = 1 \times 10^{-5}$. As initializing contour for all reconstructions, we chose a circular inclusion in the center of the pipe.

A. Reconstruction of a circular contour

Figure 5 depicts the reconstruction results of the first example using the affine transformation to approximate the accurate forward map. In figure 5(a) the MAP state and the CM estimate are illustrated, while figure 5(b) depicts the

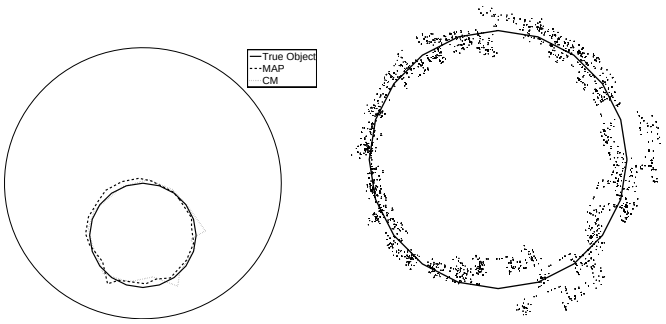

(a) MAP and CM estimates of a circular contour.

(b) Scatter Plot.

Fig. 5. Results from the reconstruction of a circular inclusion ($\varepsilon_{r,inclusion}$=2 in ε_r=1).

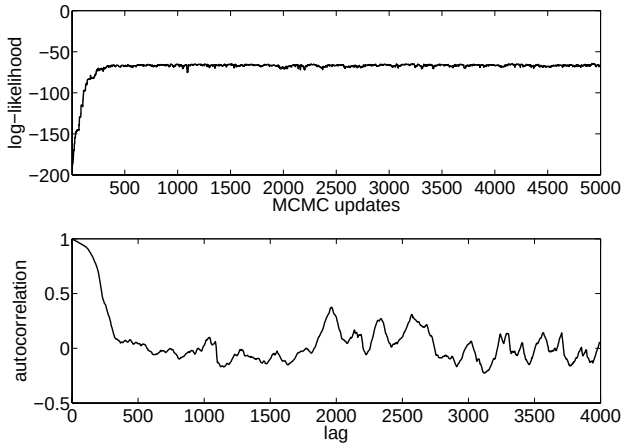

Fig. 6. Results from the reconstruction of a circular inclusion ($\varepsilon_{r,inclusion}$=2 in ε_r=1).

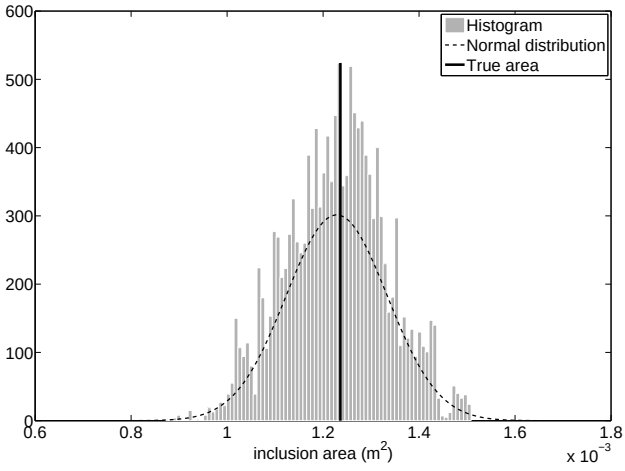

Fig. 7. Results from the reconstruction of a circular inclusion.

scatter plot highlighting the posterior variability of inclusion shape. MAP and CM are in good accordance indicating a distinct mode in the posterior distribution.

Figure 6 depicts the MCMC output trace of the log-likelihood function and the autocorrelation function, which is a measure of the statistical efficiency. The faster the autocorrelation function decays the less correlation is between

consecutive states of the Markov chain and consequently more reliable estimates are obtainable. In particular, the autocorrelation function should be after falling off smoothly to zero distributed with some noise about the x-axis. Figure 7 shows the distribution of inclusion area of the reconstructed inclusion for 15000 MCMC steps.

B. Speed comparison

The identification of the test distribution in subsection V-A was carried out with the delayed acceptance algorithm and the full (accurate) forward model. The reconstruction results are similar to the results presented in subsection V-A. Table I summarizes the computation times for the first 1500 MCMC steps. $t_{\boldsymbol{C}}$ denotes the cumulated computation time for solving the forward problem, $t_{\boldsymbol{J}}$ is the cumulated computation time for calculating the Jacobian $\boldsymbol{J}$, which is necessary in the delayed acceptance algorithm. As we noted previously, this cost can be avoided using more sophisticated schemes for operating by the Jacobian, rather than forming it. r_{Accept} represents the acceptance ratio, which is comparatively high, i.e. in [6] an acceptance ratio of $2-5\%$ is stated, when the Markov chain reaches its equilibrium. However, in the starting phase of the MCMC procedure, such a high ratio can occur causing a high computation time for the two stage algorithm. Comparing $t_{\boldsymbol{C}}$ for all three methods the advantage of the fast approximation is obvious.

TABLE I
COMPARISON OF THE COMPUTATION TIMES FOR 1500 MCMC STEPS.

Method	r_{Accept}	$t_{\boldsymbol{C}}$	$t_{\boldsymbol{J}}$
	$\%$	s	s
Accurate forward model	67.1	318	x
Two stage algorithm	61.1	244	397
Affine transformation	64.8	0.93	x

C. Reconstruction of a more complex contour

In a second experiment, a more complex contour consisting of two partially merged circular contours is reconstructed. The shape offers several difficulties for the reconstruction approach. First, the rotation move is an almost useless move, for this shape. Second, the shape offers a characteristic contraction placed in the center of the pipe, which is the region where the sensitivity is lowest.

Figure 8(a) shows the reconstruction results which were obtained after 20000 MCMC steps by the proposed algorithm using the fast approximation model. We limited the spatial resolution of the contour model to 20 control points (corners of contour). Given the challenging contour of the true object, the posterior variability indicated by the scatter plot in figure 8(b) shows a reasonable variance. Calculated MAP state and CM estimate match the true object quite well. The zigzag behavior of the point estimates could be decreased by using more control points or a stricter prior model regarding the angles between the adjacent elements of contour. However, we note that these represent preliminary results that could be improved upon by more comprehensive sampling.

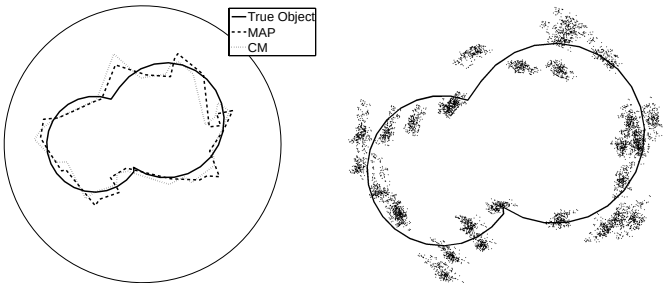

(a) MAP and CM estimates. (b) Scatter Plot.

Fig. 8. Results after 20000 MCMC steps using the fast approximation model (affine transformation).

VI. CONCLUSION

We have presented an accelerated Markov chain Monte Carlo methods for performing sample-based inference in ECT, using an appropriate approximation of the forward model. The approximation consisted of an affine model to the forward map, best fit to sampled permittivity distributions that represent the type of inclusion we expect to find. Compared to an MCMC using the accurate electrical capacitance tomography forward problem, a reduction in computation time by a factor of 340 can be obtained using the affine transformation. We have demonstrated that a tolerably small increase in posterior uncertainty of relevant parameters (inclusion area, contour shape) is traded for a huge reduction in computing time without introducing bias in estimates. The proposed cheap approximation indicates that accurate real-time inversion of capacitance data using statistical inversion is possible.

REFERENCES

[1] B. Brandstätter, G. Holler, and D. Watzenig, "Reconstruction of inhomogeneities in fluids by means of capacitance tomography, in Jnl. for Comp. and Math. in Electrical and Electronic Eng. (COMPEL), vol. 22, 2003, pp. 508-519.

[2] J. M. J. Huttunen and J. P. Kaipio, "Approximation errors in nonstationary inverse problems", in Jnl. of Inverse Problems and Imaging, vol. 1, no. 1, 2007, pp. 77-93.

[3] J. A. Christen and C. Fox, "MCMC using an approximation", in Jnl. of Comp. and Graphical Statistics. December 1, vol. 14, no. 4, 2005, pp. 795-810.

[4] J. P. Kaipio and E. Somersalo, "Statistical and computational inverse problems", New York: Applied Mathematical Sciences, Springer, 2004.

[5] K. Uhliir, J. Patera, and V. Skala, "Radial basis function method for isoline extraction", in Elect. Comp. and Informatics, 2004, pp. 439-444.

[6] D. Watzenig and C. Fox, "A review of statistical modelling and inference for electrical capacitance tomography", in Jnl. of Meas. Sci. and Techn., vol. 20, no. 5, 2009, 052002.

[7] D. Watzenig, G. Steiner, A. Fuchs, H. Zangl, and B. Brandstätter, "Influence of the discretization error on the reconstruction accuracy in ECT", in Jnl. for Comp. and Math. in Electrical and Electronic Eng. (COMPEL), vol. 26, no. 3, 2007, pp. 661-676.