X-ray Tomography of Large Objects with Limited Measurement Angle

Mikko Vepsäläinen, Markku Markkanen, and Pauli Sundberg

Eigenor Corporation, R&D Center, 00100 Helsinki, Finland

Abstract. In this paper we present an efficient implementation of an algorithm for reconstructing a 3D volume from limited angle projection data, based on statistical inversion theory. We demonstrate the strength of the method for detecting structural defects in large composite aerospace components, whose dimensions prevent acquiring measurements over the full circle. In comparison with a number of other tomographic reconstruction methods that can be applied to the limited angle case, such as tomosynthesis or simultaneous algebraic reconstruction technique (SART), we achieve superior depth resolution with reduced noise and artifacts.

Keywords: Computed Tomography, Inverse Problems
PACS: 81.70.Tx, 02.30.Zz

INTRODUCTION

X-ray computed tomography has proven a valuable tool for determining the three-dimensional inner structure of complex objects. The presence of various internal defects, together with their size, shape and other characteristics, can be reliably detected from computer-generated numerical reconstructions. With the advent of sensitive digital detector arrays, fast cone-beam measurement setups and ever increasing computing power, nondestructive evaluation using a full 3D reconstruction has become feasible for a large class of objects.

Mobile measurement setups enable CT imaging of targets that are either too large to transport into laboratory or already installed in place (e.g. airplane wings, pipe tomography). Because of the large size or supporting structures, the object in many cases is not accessible for measurements over full 180° angle. This poses severe challenges for the standard reconstructing technique based on filtered backprojection (FBP). The problem becomes underdetermined due to limited angular information and the premises of the FBP algorithm are not fulfilled.

A viable alternative to FBP in the case of limited angle or noisy measurements are algebraic methods such as ART, SIRT and their derivatives. Iterative by nature, these algorithms only aim to produce an approximate solution to the linear problem at hand, and are therefore less sensitive to uniqueness and existence issues as well as measurement noise. Even then, when the problem is severely underdetermined, one has very little control over where the iteration will converge, and visible reconstruction artifacts will be produced.

The approach we employ in this paper is to supplement the measured data with any prior information we may have on the measured object, within a Bayesian framework. An explicit three-dimensional CT reconstruction algorithm using the *statistical inversion* and Markov Chain Monte Carlo methods (MCMC) was proposed already in 2003 by Siltanen et. al. [1], but an efficient implementation for 3D cone-beam setup with realistic detector and grid resolutions has not been available until now.

This paper is organized as follows. We first briefly review the FBP and SART algorithms as they are the most commonly used algorithms in computed tomography. In the following section we describe the theory behind our statistical inversion algorithm, the model for prior information we use as well as some implementation details. We then proceed to apply the described methods to a simulated case of airplane wing ribs measured from moderate and small opening angles and show that our approah compares favorably to other methods.

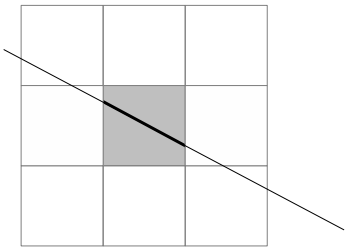

FIGURE 1. Nearest-neighbor, or box-line integrated, modeling of theory matrix.

RECONSTRUCTION ALGORITHMS

Analytical Methods

When radiographic projections from all directions (over 180°) are available at dense angular spacing, the reconstruction method of choice is usually filtered backprojection (for details, see e.g. [2]). Used widely in both medical and industrial applications, FBP is simple, fast and can produce excellent results as long as a complete set of low noise data can be provided. Filtered backprojection is an analytical method based on inverting the Radon transform (projection data) via the Fourier slice theorem, and basically boils down to multiplying the data with given geometrical factors, convolving with a ramp filter and backprojecting over all angles:

$$f(\mathbf{r}) = \int_0^{2\pi} \mathrm{d}\theta \, G_2(\mathbf{r},\theta) \Big(h * [P_\theta \cdot G_1(\mathbf{r},\theta)] \, (p(\mathbf{r},\theta)) \Big), \tag{1}$$

where $G_{1,2}$ are geometry factors and $h(p)$ is the filter. Theoretically, it corresponds to a pseudo-inverse of the projection data together with a residual image resulting from noise and model errors [3], which makes it very sensitive to measurement imperfections.

In practice, the data as well as the reconstruction volume are discrete, and the integrals and convolutions in eq. (1) will be replaced with finite sums. When the angular spacing is small, the quality of resulting images remains good, but as the angles become sparser, visible streak artifacts begin to show in the image.

Algebraic Methods

Another common approach to the tomography reconstruction problem is to start directly with the discretized problem

$$Ax = m, \tag{2}$$

where x is a (finite) column vector of the unknown grid values (mass attenuation coefficient), m contains the measured pixel values in all projections and A is *theory matrix* representing the measurement. The theory matrix elemens A_{ij} model the attenuation of the X-ray beam due to a given voxel, a typical example would be using the line length and a constant attenuation inside the voxel [4], see fig. 1.

Because the theory matrix is very large and generally non-square, the linear system is solved iteratively. A version of Kaczmarz iteration known as the ART algorithm proceeds via multiple projection-backprojection steps by projecting the current grid values to detector, comparing with the actual measurements and backprojecting the difference to update the grid values. Other versions of the algorithm (SIRT, SART, etc.) differ mostly in the way they choose the projections/pixels to consider between successive updates.

The iterative algorithms are less sensitive to noise and sparse data than FBP, but without using prior information or regularization (in the form of filtering between updates, see e.g. [5] for comparison) can not reconstruct images from limited angle data without artifacts.

STATISTICAL INVERSION

In order to reconstruct a volume from limited angle (insufficient) data without generating artifacts, the measurements have to be complemented with some external, or prior, information. A natural framework for this is Bayesian inference, where all variables are considered random variables. The "reconstruction" process then corresponds to determining the conditional (given the measurement) *a posteriori* distribution for the unknown grid values, and further computing a point estimate from that distribution to produce an image. Below we outline these steps. More details can be found in [1, 6].

Theory

To quantify our uncertainty in the measured data, we augment the linear measurement model of eq. (2) with additive noise:

$$Ax = m + \varepsilon, \tag{3}$$

and we model the noise with a Gaussian distribution, $\varepsilon \sim N(0, \sigma)$ (the choice of Gaussian is not crucial to the application, but presents a good trade-off between speed, generality and image quality). For simplicity, we assume here that the measurement error in each pixel is identically distributed, with variance σ^2, and independent of data and errors in other pixels. More realistically, the Gaussian error model could be improved by letting the variance depend on the measured pixel value, at the cost of additional complexity.

The prior information is quantified in terms of prior distribution $P_{\mathrm{pr}}(x)$, which assigns a high probability to the kind of images we expect to see, before the actual measurement is taken into account. In the implementation we are discussing in this paper we use a smoothness prior for suppressing noise and artifacts (see below).

With the measurement model at hand, we can write the conditional probability for the given measurement as

$$P(m|x) \propto \exp\left[-\frac{1}{2\sigma^2}(Ax-m)^{\mathrm{T}}(Ax-m)\right] \tag{4}$$

and with the help of Bayes' formula we finally arrive at the conditional posterior probability

$$P(x|m) \propto P(m|x)P_{\mathrm{pr}}(x). \tag{5}$$

In theory, this is the final answer. The posterior probability, defined (up to a normalization constant) by eqs. (4-5) and our specified prior distribution, describes the likelihood of a given set of grid values, given the measured data and the prior expectation. Ideally, one would like to have enough measurements to constrain the result without depending too much on the prior, but in cases with limited angle or sparse data the prior distribution has a central role in reducing undesired artifacts.

In order to produce an image, a point estimate from the distribution defined in eq. (5) needs to be computed. The two standard choices here are the conditional mean (CM) value and the point of maximum a posteriori probability (MAP), but we should point out here that having the whole distribution at hand (and being able to sample it, as described below) enables computing many other kinds of images as well, for example plotting the probability of a voxel's value belonging to a specified interval.

In our implementation of the algorithm, we have chosen to compute the conditional mean

$$\bar{x} = \int \mathrm{d}x\, x P(x|m)\,, \tag{6}$$

which gives an optimal estimate with respect to the mean square error. We compute the integral above via Markov Chain Monte Carlo methods, generating samples of the distribution using the so-called single component Gibbs sampler and averaging over the samples (again, for more details see [1]).

It should be pointed out here that explicitly allowing noise in the measurement, eq. (3), enables the algorithm to construct a solution which satisfies $Ax = m$ only approximately. We are thus not reconstructing the noise, a critical requirement in all inverse problems.

Prior Models

The prior distribution encodes any knowledge or expectations we might have on the measured object. Among the popular choices are simple white noise models (leading to Tikhonov regularization) and various smoothing priors, but in principle one could craft very refined models for specific use cases. For example, in pipe tomography we are expecting circular cross section of nearly homogeneous matter.

In order to keep our software general and fast enough, and to avoid losing detail due to a dominating prior, we have chosen to use a total variation (TV) smoothness prior to suppress noise, oscillations and other reconstruction artifacts. On a discrete three-dimensional grid, the prior is defined as

$$P_{\mathrm{pr}}(x) \propto \exp(-\alpha \,\mathrm{TV}(x)), \qquad \text{where} \quad \mathrm{TV}(x) = \sum_{i=1}^{N} \sum_{j=1}^{3} \left| x(\mathbf{r}_i) - x(\mathbf{r}_i \pm \hat{e}_j) \right|, \tag{7}$$

or the sum over all neighboring voxel differences. This prior clearly favors smooth volumes, but does not penalize large steps over many small ones, thus preserving sharp edges where data suggests them.

In addition to the total variation prior, we typically constrain the grid attenuation values to be non-negative with the positivity prior $x_i \geq 0$. This is not strictly necessary, but usually helps the algorithm converge faster.

Implementation Details

We have implemented the Markov Chain Monte Carlo algorithm with the Gibbs sampler in pure C99, optimized for parallel processing. As in all three-dimensional reconstruction problems with realistic volume and resolution, the main consideration is the massive size of the problem. With typical grid dimensions of the order 512^3, detector sizes around 600^2 and 360 projections we have a linear system of size $1.3 \cdot 10^8$ unknowns and equations. Despite the theory matrix A being very sparse, storing it in the computer main memory would require hundreds of gigabytes and usually it is not feasible. This means that the relevant parts of the theory matrix need to be (at least partially) recomputed for each update.

In ART-type algorithms one considers a single detector pixel together with all voxels on the ray to that pixel (in a given projection) at a time when updating the volume values. This enables fast computation of theory weights using ray-tracing techniques. On the contrary, the statistical inversion algorithm accesses the theory matrix column-wise, considering information from *all* projections for updating a single voxel. Combined with the smoothness prior, which removes the independence of neighboring voxels, this makes the efficient parallelization of the algorithm a delicate issue. A central part of our algorithm is indeed finding variable sized sets of voxels, each of which can be updated independently of other sets.

As the sparsity of the theory matrix is crucial to resource usage, care has to be taken to avoid computations leading to dense matrices, such as explicitly forming the ubiquitous normal matrix $Q \equiv A^{\mathrm{T}}A$ to compute Qx. Instead, we keep track of the vector Ax and multiply it with (sparse) A. To effectively calculate Ax, we use another optimization related to the Gibbs sampler: as the grid is updated one voxel at a time, the forward projection can be computed as an update to the previous value,

$$Ax^{(i)} = Ax^{(i-1)} + A[x^{(i)} - x^{(i-1)}], \tag{8}$$

where the difference $x^{(i)} - x^{(i-1)}$ is nonzero only at a single point.

A common issue with most regularization approaches is the introduction of new parameters, whose values need to be tuned for obtaining successful results. To alleviate the additional computational burden, we have developed discretization-invariant derived parameters corresponding to α in eq. (7) and σ in eq. (4), whose optimal values can be sought on a coarser grid before the final reconstruction. Similarly, as the MCMC algorithm requires quite many sampling rounds before converging to the desired posterior distribution (known as *burn-in period*), we can initialize the volume on a coarser grid and only do the final sampling with full resolution to cut down the computation time.

SIMULATION EXAMPLE

To study the performance of the statistical inversion algorithm on limited angle tomography problems, we present a study of a crude model simulating an airplane wing rib, with a superimposed test pattern of circular defects with varying radius and density, see figs. 2 and 3.

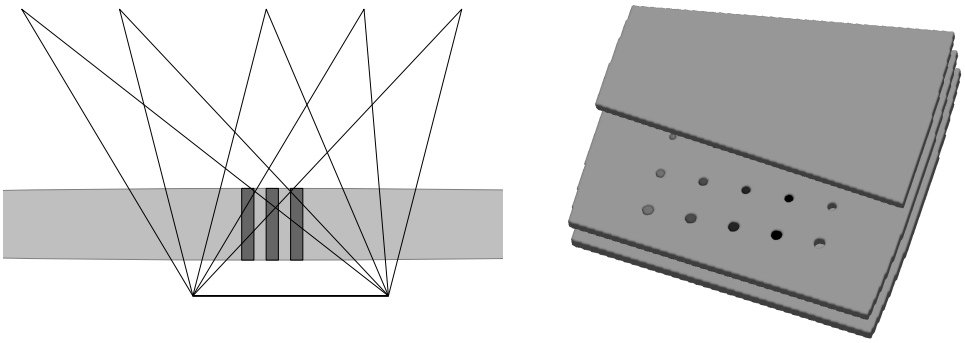

FIGURE 2. Simulation model.

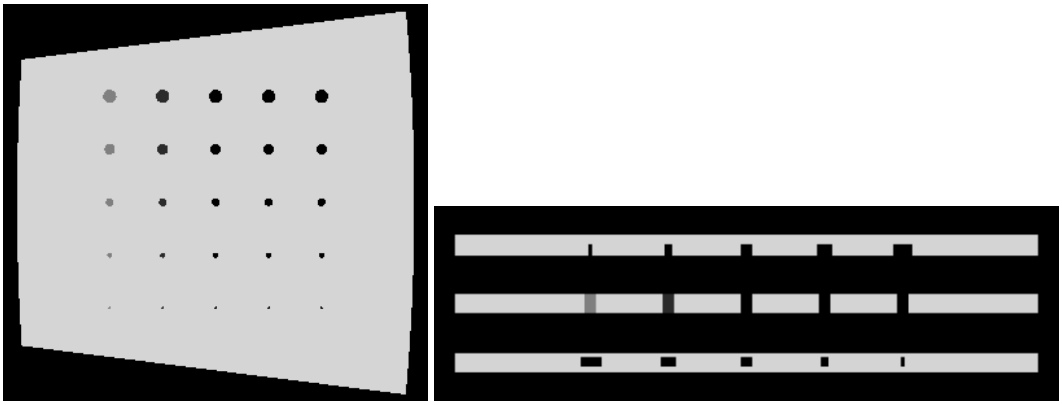

FIGURE 3. Model central slice seen from side (left) and above (right).

We have created simulated projections of 300×300 pixels, with the X-ray source moving on a linear path above the target and covering a 110° opening angle, with projections spaced one degree apart. The problem is thus the limited angle, whereas the measurement angles are not particularly sparse.

The result of reconstructions from noiseless measurements using filtered backprojection, an ART-type algorithm and our statistical inversion method (called Psig) are shown in fig. 4. As expected, FBP fails as it's underlying assumptions are violated by the limited angle. The algrebraic method performs better, but there are oscillatory artifacts and the shape of the defects is elongated in the direction of measurement. The statistical inversion algorithm performs best, although artifacts can be seen near the edges, which are only visible in a small fraction of projections.

Effect of Geometry Error

The precise knowledge of the measurement geometry is usually essential for reconstruction quality. Nevertheless, the position of the X-ray source or the detector can be hard to determine accurately, in particular with mobile setups, and the related error can not be reduced simply by increasing the exposure time as with the quantum noise. To study this effect, we added 1% offset to the source position while generating the simulated projections while using the non-offset geometry for reconstruction. The results are shown in fig. 5. As can be seen in the pictures, the artifacts expectedly increase with both algorithms, while the shape and edges of the defects are somewhat better defined in Psig reconstruction due to regularization.

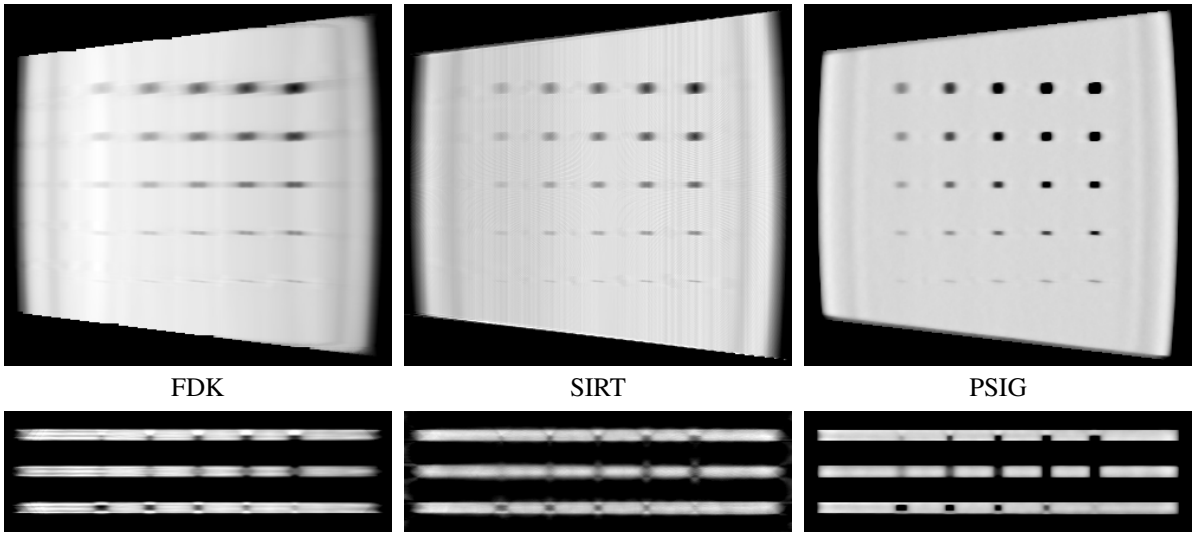

FIGURE 4. Reconstruction from noiseless measurements over 110° opening angle.

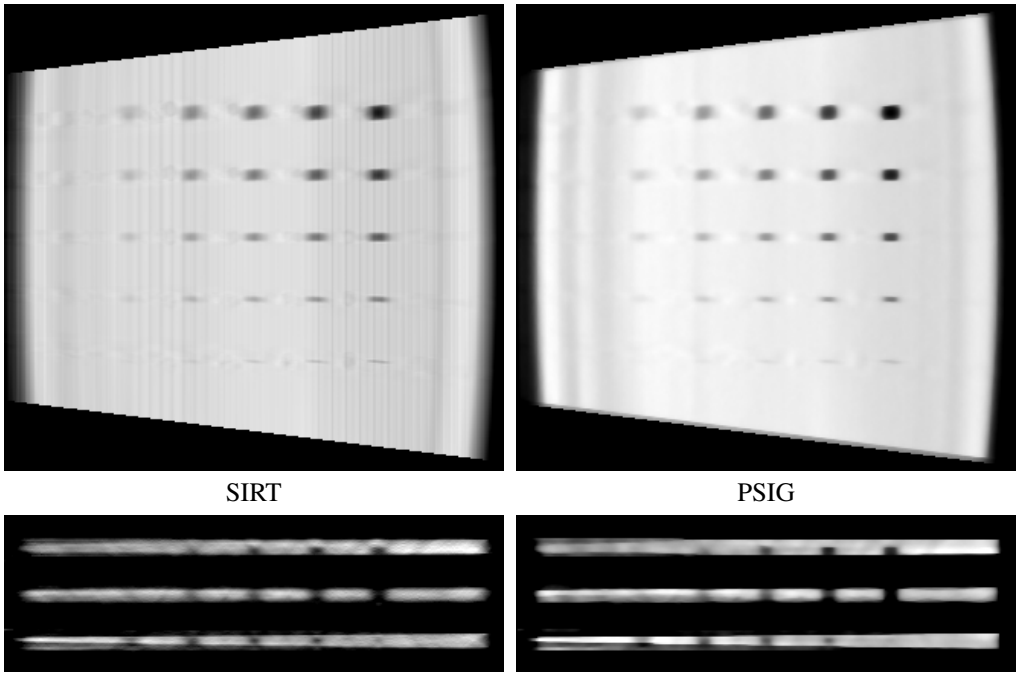

FIGURE 5. Reconstruction from measurements with 1% geometry error.

Small Angle Results

In both industrial and medical imaging situations arise when measuring over large enough opening angle for accurate 3D reconstruction is not possible. In these situations one may still study images parallel to the detector plane to gain information on the target, conventionally using the simple tomosynthesis algorithm. More involved algorithms can improve the contrast and depth resolution also in these situations, at the cost of increased computational burden. To study the performance of our algorithm at very small angles, we limited the opening angle to 30° ($\pm 15°$ around the detector normal). The central reconstruction planes parallel to the detector with different algorithms are shown in fig. 6, together with a profile plot through the center line of the picture.

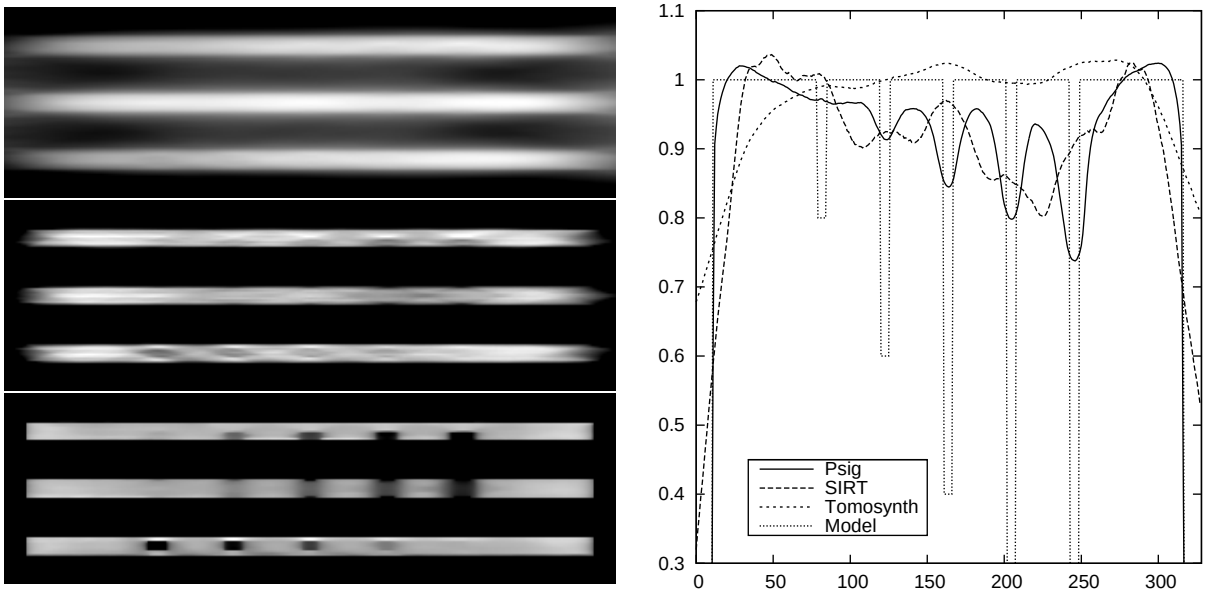

FIGURE 6. Reconstruction from 30° opening angle with tomosynthesis (top), SIRT (middle) and Psig (bottom). The right panel shows a profile plot through the center line.

As shown in the figure, the defects are clearly visible only in the reconstruction with statistical inversion algorithm. The profile plot shows that the Psig reconstruction follows the general shape of true values, but it is strongly smoothed and the numerical levels are not reliable. Other algorithms fail to bring out the location and shape of the defects, and the values often move into wrong direction altogether.

CONCLUSIONS

We have implemented the statistical X-ray tomography inversion algorithm in the form suitable for parallel computing, written in C language. The algorithm is based on a Bayesian model, complementing the measurement data with prior information, which in our case is formulated as a total variation smoothness prior together with a positivity constraint. The model results in a posterior distribution for the unknown attenuation values inside the reconstruction volume, and we compute the (conditional) mean of this distribution using Markov Chain Monte Carlo methods to produce the final 3D image.

In this paper we have studied the performance of our Psig algorithm in the limited angle CT reconstruction problem, where the standard filtered backprojection method fails due to incomplete measurement data. Total variation regularization together with built-in assumption of noise in the measurement model make the resulting reconstruction less sensitive to lacking and erroneous data. As a result, the limited angle reconstructions show less noise and artifacts together with more clearly defined features. We also demonstrated that Psig can provide improved contrast and depth resolution when the opening angle is very small and usually tomosynthesis or SART are used, although a full, accurate 3D reconstruction can not be obtained.

The simulation model can be argued to give an overly optimistic impression of our method, as the reconstruction target is piecewise constant, with very small total variation. On the other hand, many industrial, synthetic objects fall into this class, mostly consisting of homogeneous matter with little variation. Nevertheless, it is clear that for highly irregular objects the prior information is different and should be reflected in the model. In the future we plan to experiment with other kind of priors as well, possibly also including structural information such as the generic shape of the object (if known in advance) into the reconstruction process.

Currently the main restriction on using our code with even larger problems is the rapidly growing memory consumption from partially storing the theory matrix values. To circumvent this problem we plan to move most of the theory matrix computation on GPU to enable on-fly calculation with smaller memory footprint.

REFERENCES

1. S. Siltanen, V. Kolehmainen, S. Järvenpää, J. P. Kaipio, P. Koistinen, M. Lassas, J. Pirttilä, and E. Somersalo, *Physics in Medicine and Biology* **48**, 1437 (2003).
2. A. C. Kak, and M. Slaney, *Principles of Computerized Tomographic Imaging*, IEEE Press, 1988.
3. M. A. Anastasio, X. Pan, and E. Clarkson, *IEEE Transactions on Medical Imaging* **20**, 539 (2001).
4. R. L. Siddon, *Medical Physics* **12**, 252 (1985).
5. W. Xu, and K. Mueller, "Evaluating popular non-linear image processing filters for their use in regularized iterative CT," in *Nuclear Science Symposium Conference Record (NSS/MIC), 2010 IEEE*, 2010, pp. 2864–2865.
6. J. Kaipio, and E. Somersalo, *Statistical and computational inverse problems*, Springer, New York, 2005.