DOI: 10.2352/ISSN.2470-1173.2016.19.COIMG-155

Making Advanced Scientific Algorithms and Big Scientific Data Management More Accessible

S. V. Venkatakrishnan [a,b], K. Aditya Mohan [c], Keith Beattie [d], Joaquin Correa [e], Eli Dart [f], Jack R. Deslippe [e], Alexander Hexemer [a,b], Harinarayan Krishnan [a,d], Alastair A. MacDowell [b], Stefano Marchesini [a,b,d], Simon J. Patton [d], Talita Perciano [a,d], James A. Sethian [a,d,g], Rune Stromsness [d], Brian L. Tierney [d], Craig E. Tull [d], Daniela Ushizima [a,d], and Dilworth Y. Parkinson [a,b]

[a] **Center for Advanced Mathematics for Energy Research Applications, Lawrence Berkeley National Lab, Berkeley, CA 94720**
[b] **Advanced Light Source, Lawrence Berkeley National Lab, Berkeley, CA 94720**
[c] **School of Electrical and Computer Engineering, Purdue University, Purdue, IN 47907**
[d] **Computational Research Division, Lawrence Berkeley National Lab, Berkeley, CA 94720**
[e] **National Energy Research Scientific Computing Center, Berkeley, CA 94720**
[f] **Energy Sciences Network, Berkeley, CA 94720**
[g] **Department of Mathematics, UC Berkeley, Berkeley, CA 94720**

Abstract

Synchrotrons such as the Advanced Light Source (ALS) at Lawrence Berkeley National Laboratory are user facilities - they are sources of extremely bright X-ray beams, and scientists come from all over the world to perform experiments that require these beams. As the complexity of experiments has increased, and the size and rates of data sets has exploded, managing, analyzing and presenting the data collected at synchrotrons has been an increasing challenge. The ALS has partnered with high performance computing, fast networking, and applied mathematics groups to create a"super-facility", giving users simultaneous access to the experimental, computational, and algorithmic resources to overcome this challenge. This combination forms an efficient closed loop, where data despite its high rate and volume is transferred and processed, in many cases immediately and automatically, on appropriate compute resources, and results are extracted, visualized, and presented to users or to the experimental control system, both to provide immediate insight and to guide decisions about subsequent experiments during beam-time. In this paper, We will present work done on advanced tomographic reconstruction algorithms to support users of the 3D micron-scale imaging instrument (Beamline 8.3.2, hard X-ray micro-tomography).

Introduction

User facilities like the Advanced Light Source (ALS) at Lawrence Berkeley National Laboratory host scientists and researchers from all over the world who come because of the extremely bright X-ray beams that enable unique science. The main expertise of most of these users is not directly in computing and data management but in their field of science. However, the size and rate of data produced at synhcrotrons has exploded over the last few years. This volume and velocity can easily overwhelm many users–with their existing tools, they are not able to fully analyze or take advantage of all of the data. The ALS has partnered with high performance computing (NERSC), fast networking (LBLnet and Energy Sciences Network), computer scientists (SPOT Suite), and applied mathematics groups (CAMERA), to give users simultaneous access to the experimental, computational, and algorithmic resources to overcome the data challenge.

In this contribution, we will focus on new work at the X-ray Tomography Beamline (8.3.2) of the ALS. We have previously described SPOT Suite [1, 2], which detects and transfers data to the NERSC high performance computing facility, where processing jobs are automatically launched. Raw and processed data and parameters are search-able by users through a web portal. One of the processing steps is tomographic reconstruction by a filtered back-projection method. This approach is fast and straightforward and is useful for giving users real-time feedback. But one of the original goals in building an infrastructure that included the NERSC high perfomance computing center was to make it possible to launch more computationally expensive algorithms on data sets that merit it, yielding improved results.

While the co-location of data and compute resources is required, it is not sufficient to allow most ALS users of the tomography instrument to take advantage of advanced tomographic reconstruction algorithms: in general they lack the required expertise in navigating the command line and queue system of the high performance computer, and also lack the required expertise in advanced tomographic reconstruction algorithms, which in general require detailed adjustment of parameters.

In this manuscript, we describe the work we have done to change this, making computationally intensive Model-Based Iterative Reconstruction (MBIR) more readily available to all users of the beamline. While this has involved some work on the SPOT Suite interface to enable users to choose to launch MBIR reconstructions on selected data sets, in this manuscript we focus on describing the work done to improve the ability of MBIR to autotune, and to minimize the number of parameters users must adjust. We show that this has allowed MBIR to give excellent results over a range of samples.

Advanced Algorithms for Tomography - Model-Based Iterative Reconstruction

MBIR algorithms represent the state-of-the-art in attaining a balance between resolution and noise for a variety of tomographic applications [3–9]. These methods cast the reconstruction as a maximum a posteriori estimation problem involving the formulation of a probabilistic forward model for the data and a prior

DOI: 10.2352/ISSN.2470-1173.2016.19.COIMG-155

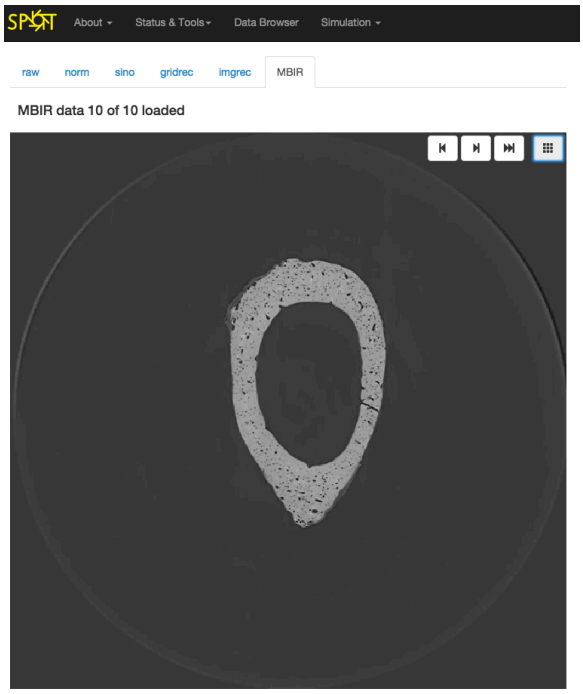

Figure 1. *A screen-shot of the result of applying the MBIR algorithm to a data set via the SPOT interface. Users can reconstruct the data using multiple approaches and compare the results.*

model for the 3D volume. The estimation is typically equivalent to minimizing a cost-function of the form,

$$(\hat{x},\hat{\phi}) = \underset{x,\phi}{\operatorname{argmin}}\{-\log p(y|x,\phi) - \log p(x)\}$$

where x is a vector containing the unknown voxels, ϕ represents unknown calibration parameters associated with the measurement, y is a vector containing all the measurements, $p(y|x,\phi)$ is the likelihood and $p(x)$ is the prior probability of x. The central challenges in the MBIR approach are the formulation of the likelihood function, a suitable prior model and the subsequent design of an optimization method to find the minimum of the resulting cost function.

We have used a specific version of this algorithm [10–12] that is designed to automatically account for the presence of miscalibrations that account for rings and the presence of streaks due to outliers; artifacts that are common in synchrotron X-ray tomography. In particular, our implementation optimizes the cost-function [10]

$$c(x,d,\sigma) = \frac{1}{2}\sum_{n=1}^{N}\sum_{i=1}^{M}\beta_{T,\delta}\left((y_{n,i} - A_{n,i,*}x - d_i)\frac{\sqrt{\Lambda_{n,i,i}}}{\sigma}\right) + MN\log(\sigma) + \sum_{\{j,k\}\in\mathcal{N}} w_{jk}\rho(x_j - x_k) \quad (1)$$

where, $\beta_{T,\delta}$ is the generalized Huber function [9] with parameters T and δ, N is the total number of views, M is the number of detector pixels, $y_{n,i}$ is the log normalized measurement at view n and pixel location i, A_n is the forward projection matrix at view n, d_i is a offset calibration parameter associated with detector pixel i (that models the presence of rings), Λ_n is a diagonal weight matrix with entries set to be proportional to the inverse variance of the noise at each measurement [3, 13], and σ is a scaling parameter for the noise associated with the measurements. The function $\rho(.)$ is set to a q-Generalized Markov Random Field [14] be of the form

$$\rho(\Delta) = \frac{\Delta_s^3\left|\frac{\Delta}{\Delta_s\sigma_s}\right|^2}{c_s + \left|\frac{\Delta}{\Delta_s\sigma_s}\right|^{2-p}}$$

where c_s is typically set to a small number (in the range of $1/100$), Δ_s is the length of each side of a voxel, σ_s and p are model parameters. Notice that since c_s is small, this term is approximately corresponds to a generalized Markov Random field [15]. One advantage of this model is that when $p = 1$ we obtain a total-variation like prior that corresponds to sharp edges and when $p = 2$ we obtain smooth reconstructions corresponding to a Gaussian prior. In our experience with several data sets, values of $p = 1.2$ provide a good balance between preserving texture and edges in the reconstructions.

One of the challenges in making MBIR accessible to experimental scientists, some of whom may not be performing quantitative imaging (i.e. do not have an estimate for the values of their reconstruction in units of inverse distance), is to be be able to automatically set the parameters to obtain a reasonable reconstruction. Furthermore, we need to factor the cost function so there are "knobs" that can be adjusted in a intuitive manner so the users can obtain the desired image quality. In particular, we re-write the prior so that $\sigma_s = \frac{\sigma_{s,0}}{\kappa}$. Here κ is a dimensionless quantity while $\sigma_{s,0}$ is a value having units of inverse distance. In the code, we automatically initialize $\sigma_{s,0}$ using the following heuristic reasoning. For a GGMRF, the maximum likelihood estimate of this quantity is equal to be the "standard-deviation" like analogue for the GMRF [16]. However, we do not have the reconstruction to be able to estimate this value. Instead, we obtain an estimate of this quantity by finding the standard deviation in each projection image y_n and dividing this quantity by the projected thickness of the sample (units of distance). We average the estimate from all the views n to set the value of $\sigma_{s,0}$ after adjusting it for the variance computation of the noise. Given this setting, the user only has to pick the value of κ which is a dimensionless smoothness value. Typically, κ on values in the range of $0-1$, with a value of 0.5 producing a significantly more noisy reconstruction than a value of 1. The parameter T is set so that if the fitting error is greater than T times the noise standard deviation then that measurement is classified as an outlier [9]. Hence, values in the range of as 3 to 5 are usually sufficient to reject outliers in typical data sets. The value of δ (in the range [0,1] with a typical value of 0.1) is set to reject the bad measurements. In summary, the principal parameter the user is left to adjust to perform the reconstruction is the dimensionless smoothness value - κ. However, all the other parameters (p, T and δ) can also be adjusted via the input interface.

The MBIR algorithm uses a multi-CPU version of NHICD [12, 17] algorithm to optimize the cost function (1). The division of the 3-D volume across cores follows the scheme in [12] i.e. the volume is divided into chunks of slices along the axial direction,

DOI: 10.2352/ISSN.2470-1173.2016.19.COIMG-155

a fraction of the voxels in each chunk are updated by a single compute node in parallel, and each chunk communicates its "end" voxels to the adjacent chunks after every iteration to ensure we lower the value of the overall cost function with each update. The chunk sizes are optimized to ensure we fully utilize the processing power of each node in the National Energy Research Scientific Computing clusters.

Finally, in order to make the MBIR code more accessible to users, we have built a simple GUI based tool that is integrated into the existing infrastructure for the collection, processing and analysis of tomography data sets at the ALS, termed SPOT suite [1,2]. Using this interface, the user can set the basic MBIR parameters, reconstruct the data and compare the output of multiple algorithms (see Fig. 1).

Results

In this section, we present results using the MBIR method to process several data sets acquired at the ALS and demonstrate how the algorithm can help drive new scientific experiments at the synchrotron. In each case, we compare the MBIR method ($p = 1.2$, $\delta = 0.1$) to an analytic reconstruction algorithm (gridrec [18] or filtered back projection - FBP) that incorporates pre/post-processing to remove ring artifacts. For each result, we present a single slice from the 3D reconstructions to compare the different methods. All the reconstructions are performed using the NERSC HPC compute clusters.

Sparse View Sampling

One of the strengths of the MBIR method, is the ability of the algorithm to reconstruct a sample from a sparse set of views [10]. Fig. 2 demonstrates three such examples where the MBIR reconstruction has a higher quality than the the analytic approach, albeit using only a fraction of the total number of views compared to the analytic techniques (Data Set (c) courtesy - Andrew McElrone; Data Set (e) courtesy- TE Connectivity). While the amount of sub-sampling that can be used without a significant loss of detail depends on the sample, in general we have observed it is possible to obtain a factor of 2 to 4 compared to analytic approaches across a wide range of scientific data sets.

Limited View Sampling

As a part of novel scientific experiments at a synchrotron based beamline, there are cases when the sample holders used for the tomography experiment are such that they block (or highly attenuate) a collection of views. Fig. 3 (a) shows one such case, where a diamond anvil cell blocks the measurements for about 60° out of the 180° rotation resulting in a limited view data set. As a result, the analytic reconstruction approach produces a reconstruction with significant streak artifacts as well as loss of detail (Fig. 3 (c)). In contrast, the MBIR approach produces a reconstruction that suppresses the streak artifacts and preserves edge detail better. This example illustrate how using advanced algorithmic approaches can help reconstruct data from non-traditional sampling geometries.

Robustness to noise

In general, MBIR style approaches have been shown to be very useful to handle low-dose and noisy data sets for a variety of applications [17, 19, 20]. Fig. 4 shows the reconstruction from the two algorithms for a noisy data set containing 512 views over a 180 degree rotation (data courtesy : Dilworth Y. Parkinson). Notice that for the same data set the MBIR method produces a lower noise reconstruction compared to the analytic approach while preserving the edges in the reconstruction. MBIR algorithms are set to play a crucial role in high speed synchrotron tomography [12, 21] as well as cases when the samples are sensitive to beam damage.

Other modalities - Two-step in-line phase contrast tomography

The MBIR algorithm can also be adapted to other parallel beam tomography applications. Fig. 5 (a) and (b) shows the reconstruction of phase-contrast tomography data (carbon/resin sample) in which the phase is first retrieved using a Bronikov algorithm [22]. The MBIR approach produces a sharp reconstruction compared to the FBP method while preserving the details in the edges. We note that a more principled approach to the phase contrast tomography problem would involve a single unified cost-function that directly operates on the raw measurements [23].

Conclusion

In this paper, we have introduced the ALS "super-facility" and the implementation of an advanced tomographic reconstruction algorithm - MBIR into this framework - demonstrating how the algorithm can be used for reconstructions in several challenging scenarios. We foresee such advanced algorithms being a key driver for new scientific experiments which are not possible using current approaches. In addition to better modeling, accelerating MBIR methods by developing new mathematical optimization techniques that best utilize modern parallel computing platforms [24–26] will play a vital role in driving the next generation of scientific experiments.

Acknowledgements

This work was supported by the Office of Science, Office of Basic Energy Sciences, and the Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This work used resources of the National Energy Research Scientific Computing Center (DOE-ASCR) and the Advanced Light Source at Lawrence Berkeley National Laboratory (DOE-BES, Scientific User Facilities Division). Research was supported by Lawrence Berkeley National Laboratory's Laboratory Directed Research and Development Program and the SciDAC Institute of Scalable Data Management, Analysis, and Visualization (SDAV, DOE-ASCR). CAMERA is supported by DOE-ASCR and DOE-BES). S.V. Venkatakrishnan and Alexander Hexemer were supported by Alexander Hexemer's DOE Early Career Award. K. A. Mohan was supported by an AFOSR/MURI grant #FA9550-12-1-0458. We thank the users of Advanced Light Source Beamline 8.3.2 for use of their data, as cited in the figure captions. We thank Prof. Charles A. Bouman at Purdue University for many helpful discussions and suggestions.

DOI: 10.2352/ISSN.2470-1173.2016.19.COIMG-155

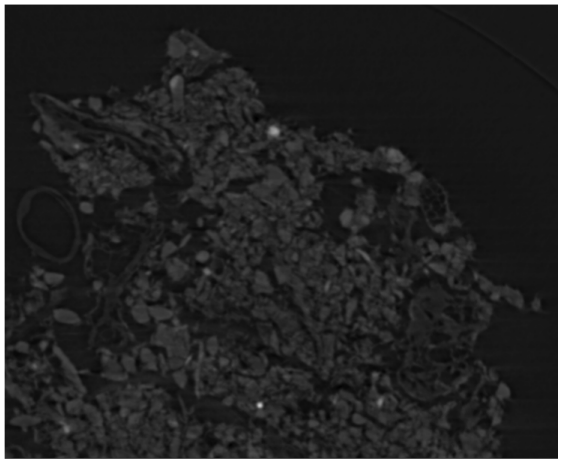

(a) Gridrec with 1024 views

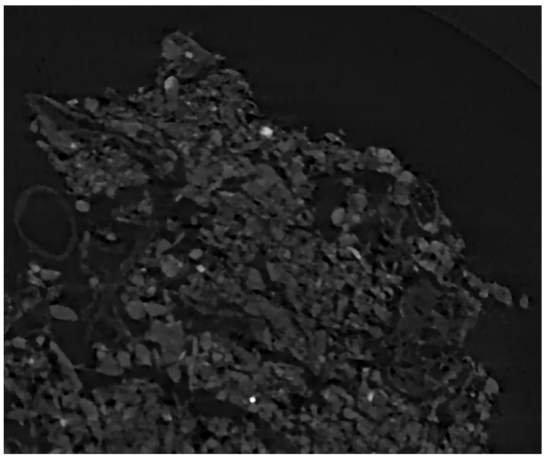

(b) MBIR with 256 views

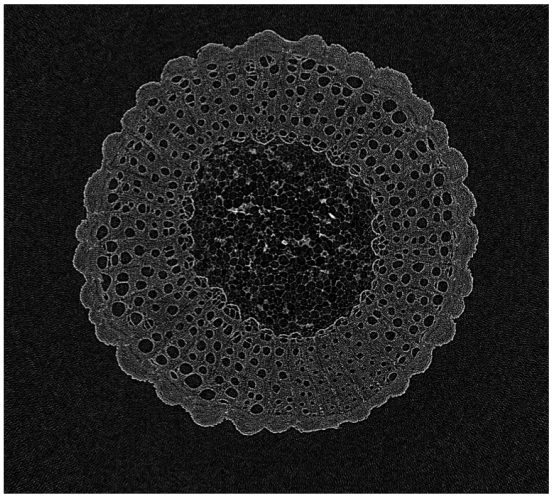

(c) Gridrec with 1024 views

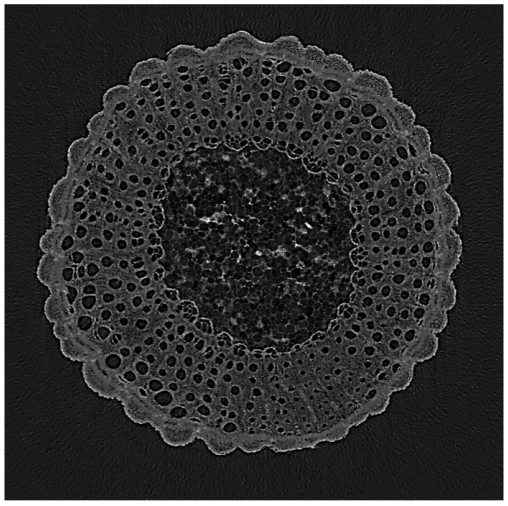

(d) MBIR with 512 views

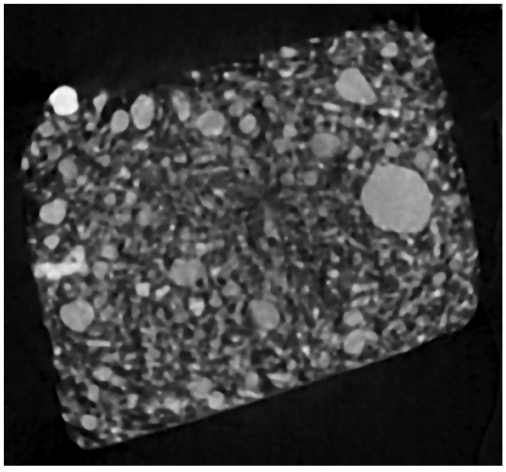

(e) Gridrec with 2048 views

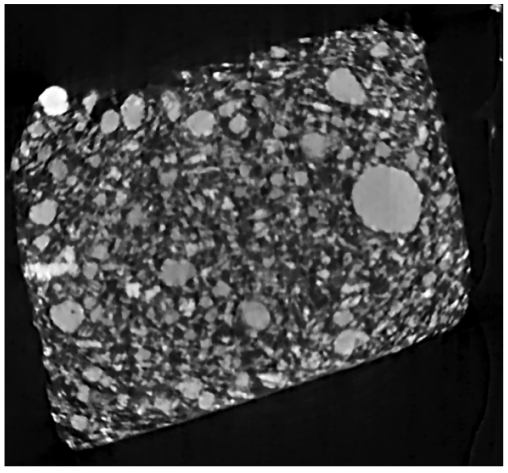

(f) MBIR with 1024 views

Figure 2. *Illustration of how the MBIR algorithm (a,d,f) can provide superior quality reconstructions than the conventional approach (a,c,e) with a fraction of the typically measured views. (a) and (b) are from a soil data set. (c) and (d) are reconstructions of a plant stem data set (courtesy Andrew McElrone, UC Davis). (e) and (f) are from a conductive composite sample (courtesy TE Connectivity). The use of MBIR algorithms can drive new scientific experiments by reducing the amount of data required for a given application.*

DOI: 10.2352/ISSN.2470-1173.2016.19.COIMG-155

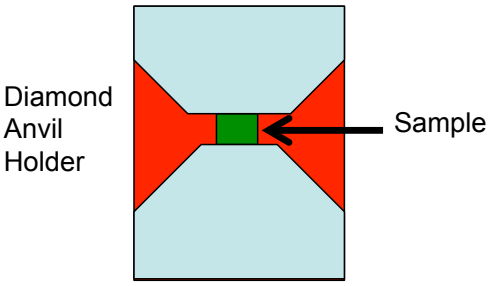

(a) Schematic of the diamond anvil cell.

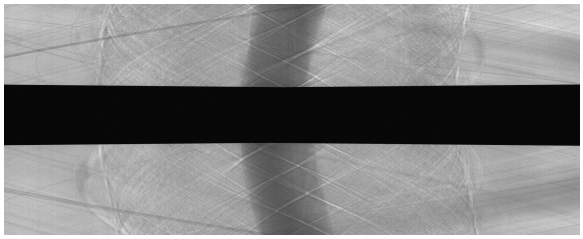

(b) Sinogram from a single slice.

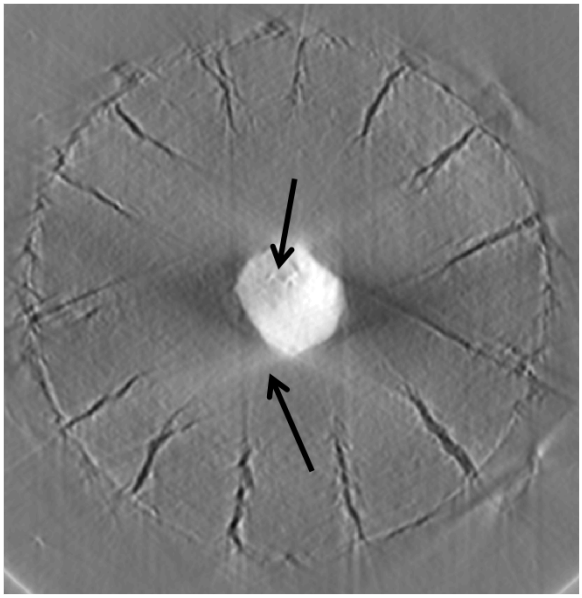

(c) FBP

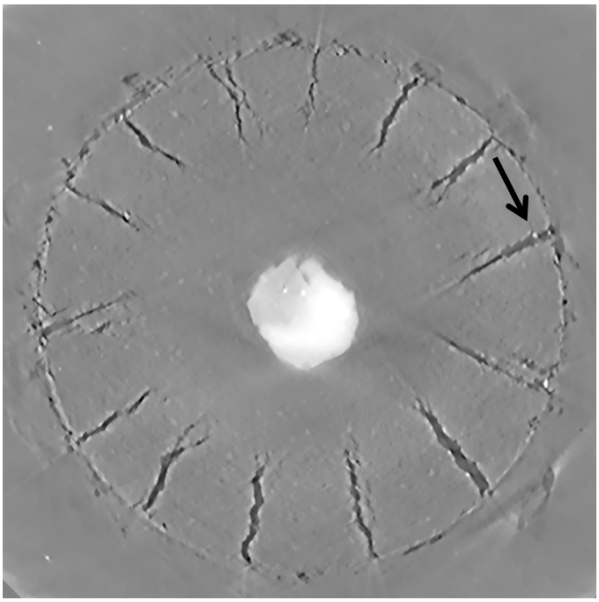

(d) MBIR

Figure 3. *Illustration of how the MBIR algorithm is robust to limited angular sampling. In this case, a diamond anvil cell (a) blocks the measurements (very high attenuation) of an obsidian sample for about 60°, resulting in a sinogram pattern of the type seen in (b) (Data : courtesy Alastair McDowell and Martin Kunz). Hence, there are streaks in the FBP reconstruction due to the limited angular sampling (c). In contrast, the MBIR method (d) reduces streak artifacts and preserves detail, thereby enabling experiments using sample holders that may block a fraction of the views during acquisition.*

DOI: 10.2352/ISSN.2470-1173.2016.19.COIMG-155

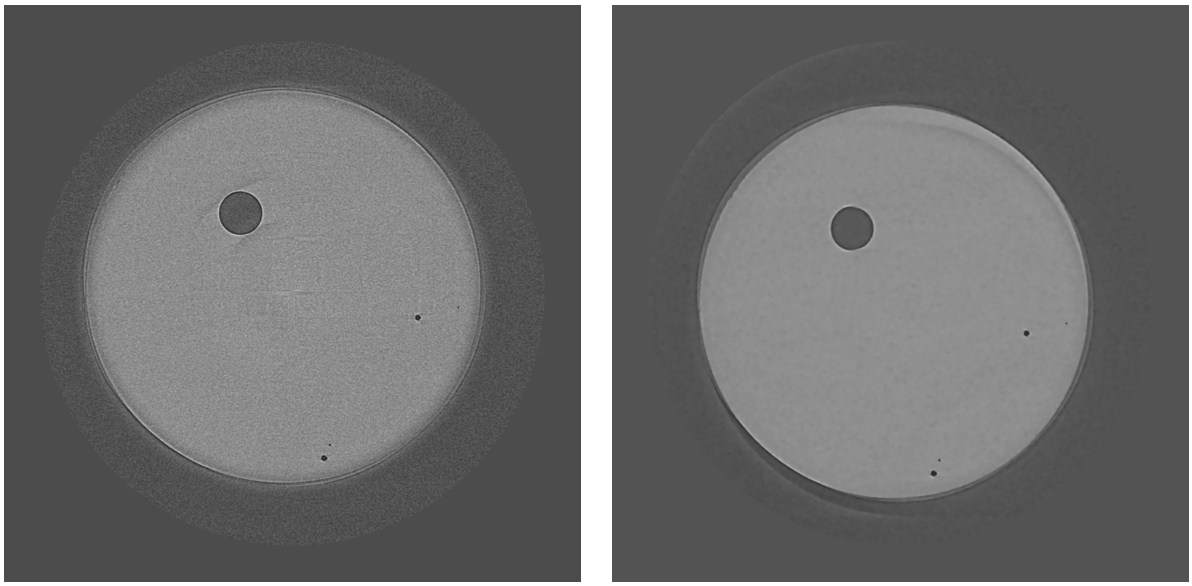

(a) Gridrec with 512 views (b) MBIR with 512 views

Figure 4. *Illustration of how the MBIR algorithm can produce a lower noise reconstruction while preserving edge detail compared to the typically used algorithm (Gridrec) while operating on the same data set. The above reconstruction is from a single slice of a bead data set.*

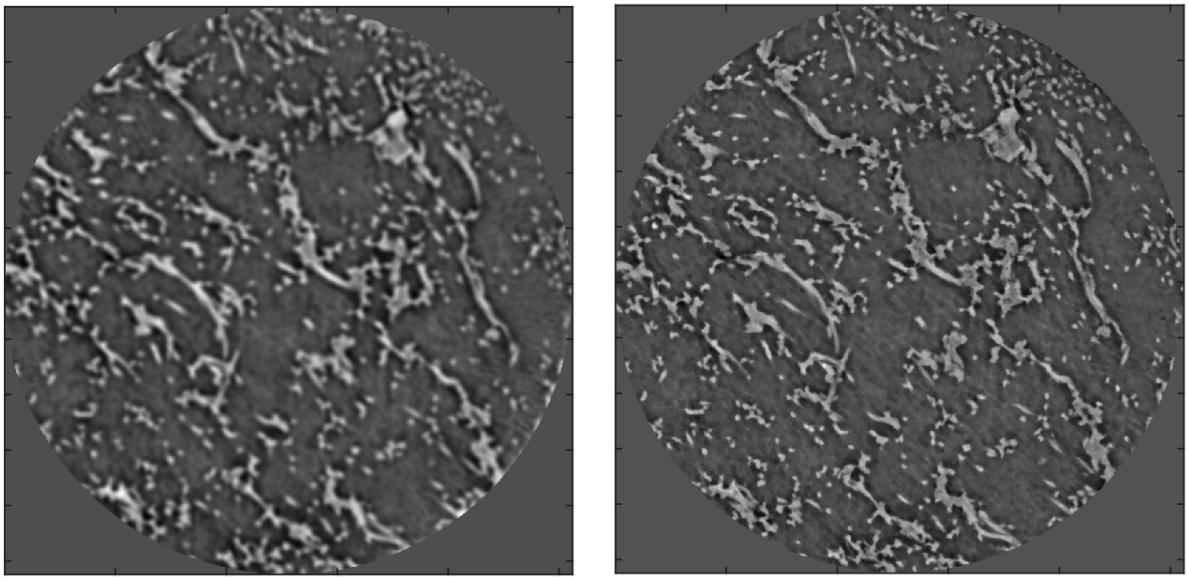

(a) FBP with 512 views (b) MBIR with 512 views

Figure 5. *The MBIR algorithm adapted for in-line phase contrast tomography data of a carbon/resin sample. The reconstructions are performed after retrieving the phase using the Bronikov algorithm [22]. The MBIR method produces a sharper reconstruction while preserving the edges.*

DOI: 10.2352/ISSN.2470-1173.2016.19.COIMG-155

References

[1] Blair, J., Canon, R. S., Deslippe, J., Essiari, A., Hexemer, A., MacDowell, A. A., Parkinson, D. Y., Patton, S. J., Ramakrishnan, L., Tamura, N., Tierney, B. L., and Tull, C. E., "High performance data management and analysis for tomography," *Proc. of SPIE* **9212**, 92121G–92121G–9 (2014).

[2] Deslippe, J., Essiari, A., Patton, S. J., Samak, T., Tull, C. E., Hexemer, A., Kumar, D., Parkinson, D., and Stewart, P., "Workflow management for real-time analysis of lightsource experiments," in [*Proceedings of the 9th Workshop on Workflows in Support of Large-Scale Science*], *WORKS '14*, 31–40, IEEE Press, Piscataway, NJ, USA (2014).

[3] Sauer, K. and Bouman, C., "Bayesian Estimation of Transmission Tomograms Using Segmentation Based Optimization," *IEEE Trans. on Nuclear Science* **39**, 1144–1152 (1992).

[4] Fessler, J., "Penalized weighted least-squares image reconstruction for positron emission tomography," *IEEE Trans. on Medical Imaging* **13**, 290 –300 (June 1994).

[5] Qi, J. and Leahy, R. M., "Iterative reconstruction techniques in emission computed tomography," *Physics in Medicine and Biology* **51**(15), R541 (2006).

[6] Stojanovic, I., Pien, H., Do, S., and Karl, W. C., "Low-dose X-ray CT reconstruction based on joint sinogram smoothing and learned dictionary-based representation," in [*9th IEEE International Symposium on Biomedical Imaging (ISBI)*], 1012 –1015 (May 2012).

[7] Boas, D., Brooks, D., Miller, E., DiMarzio, C., Kilmer, M., Gaudette, R., and Zhang, Q., "Imaging the body with diffuse optical tomography," *IEEE Signal Proc. Magazine* **18**, 57 – 75 (Nov. 2001).

[8] Venkatakrishnan, S., Drummy, L., Jackson, M., De Graef, M., Simmons, J., and Bouman, C., "A model based iterative reconstruction algorithm for high angle annular dark field-scanning transmission electron microscope (HAADF-STEM) tomography," *Image Processing, IEEE Transactions on* **22**, 4532–4544 (Nov 2013).

[9] Venkatakrishnan, S., Drummy, L., Jackson, M., De Graef, M., Simmons, J., and Bouman, C., "Model-based iterative reconstruction for bright-field electron tomography," *Computational Imaging, IEEE Transactions on* **1**, 1–15 (March 2015).

[10] Mohan, K., Venkatakrishnan, S., Drummy, L., Simmons, J., Parkinson, D., and Bouman, C., "Model-based iterative reconstruction for synchrotron X-ray tomography," in [*Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*], 6909–6913 (May 2014).

[11] Mohan, K., Venkatakrishnan, S., Gibbs, J., Gulsoy, E., Xiao, X., De Graef, M., Voorhees, P., and Bouman, C., "4D model-based iterative reconstruction from interlaced views," in [*Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*], 783–787 (April 2015).

[12] Aditya Mohan, K., Venkatakrishnan, S., Gibbs, J., Gulsoy, E., Xiao, X., De Graef, M., Voorhees, P., and Bouman, C., "TIMBIR: A method for time-space reconstruction from interlaced views," *Computational Imaging, IEEE Transactions on* **1**, 96–111 (June 2015).

[13] Sauer, K. and Bouman, C. A., "A local update strategy for iterative reconstruction from projections," *IEEE Trans. on Signal Processing* **41**, 534–548 (Feb. 1993).

[14] Thibault, J.-B., Sauer, K. D., Bouman, C. A., and Hsieh, J., "A three-dimensional statistical approach to improved image quality for multislice helical CT," *Medical Physics* **34**(11), 4526–4544 (2007).

[15] Bouman, C. and Sauer, K., "A generalized Gaussian image model for edge-preserving MAP estimation," *IEEE Trans. on Image Processing* **2**, 296 –310 (July 1993).

[16] Saquib, S., Bouman, C., and Sauer, K., "ML parameter estimation for Markov random fields with applications to Bayesian tomography," *Image Processing, IEEE Transactions on* **7**, 1029–1044 (Jul 1998).

[17] Yu, Z., Thibault, J., Bouman, C., Sauer, K., and Hsieh, J., "Fast model-based X-ray CT reconstruction using spatially nonhomogeneous ICD optimization," *IEEE Trans. on Image Processing* **20**, 161 –175 (Jan. 2011).

[18] Marone, F. and Stampanoni, M., "Regridding reconstruction algorithm for real-time tomographic imaging," *Journal of Synchrotron Radiation* **19**, 1029–1037 (Nov 2012).

[19] Yamada, Y., Jinzaki, M., Tanami, Y., Shiomi, E., Sugiura, H., Abe, T., and Kuribayashi, S., "Model-based iterative reconstruction technique for ultralow-dose computed tomography of the lung: a pilot study.," *Investigative radiology* **47**(8), 482 (2012).

[20] Venkatakrishnan, S., Ming-Siao, H., Garvin, N., Jackson, M., De Graef, M., Simmons, J., Bouman, C., and Drummy, L., "Model-based iterative reconstruction for low-dose electron tomography," in [*Microscopy and Microanalysis*], (2015).

[21] Gibbs, J. W., Mohan, K. A., Gulsoy, E. B., Shahani, A. J., Xiao, X., Bouman, C. A., De Graef, M., and Voorhees, P. W., "The three-dimensional morphology of growing dendrites," *Scientific Reports* **5**, 11824 (Jul 2015). Article.

[22] Bronnikov, A. V., "Phase-contrast CT: fundamental theorem and fast image reconstruction algorithms," (2006).

[23] Tian, L., Petruccelli, J. C., Miao, Q., Kudrolli, H., Nagarkar, V., and Barbastathis, G., "Compressive X-ray phase tomography based on the transport of intensity equation," *Opt. Lett.* **38**, 3418–3421 (Sep 2013).

[24] Liu, J., Venkataramani, S., Venkatakrishnan, S., Bouman, C. A., and Ragunathan, A., "EMBIRA: An accelerator for model-based iterative reconstruction," *Very Large Scale Integration Systems, to appear IEEE Transactions on* (2016).

[25] Wang, X., Bouman, C. A., and Midkiff, S., "High performance model based image reconstruction," in [*International Conference for High Performance Computing, Networking, Storage and Analysis*], (2015).

[26] McGaffin, M. and Fessler, J., "Alternating dual updates algorithm for X-ray CT reconstruction on the GPU," *Computational Imaging, IEEE Transactions on* **1**, 186–199 (Sept 2015).