https://doi.org/10.2352/EI.2026.38.15.COIMG-168

Fast Large-Scale Model-Based Iterative Tomography via Exploiting Mathematical Structure, Hierarchical Optimization, Smart Initialization, and Distributed GPU Computing

Dinesh Kumar and Jeffrey J. Donatelli,
Center for Advanced Mathematics for Energy Research Applications,
Applied Mathematics and Computational Research Division,
Lawrence Berkeley National Laboratory, Berkeley, CA

Abstract

Model-Based Iterative Reconstruction (MBIR) is important because direct methods, such as Filtered Back-Projection (FBP) can introduce significant noise and artifacts in sparse-angle tomography, especially for time-evolving samples. Although MBIR produces high-quality reconstructions through prior-informed optimization, its computational cost has traditionally limited its broader adoption. In previous work, we addressed this limitation by expressing the Radon transform and its adjoint using non-uniform fast Fourier transforms (NUFFTs), reducing computational complexity relative to conventional projection-based methods. We further accelerated computation by employing a multi-GPU system for parallel processing.

In this work, we further accelerate our Fourier-domain framework, by introducing four main strategies: (1) a reformulation of the MBIR forward and adjoint operators that exploits their multi-level Toeplitz structure for efficient Fourier-domain computation; (2) an improved initialization strategy that uses back-projected data filtered with a standard ramp filter as the starting estimate; (3) a hierarchical multi-resolution reconstruction approach that first solves the problem on coarse grids and progressively transitions to finer grids using Lanczos interpolation; and (4) a distributed-memory implementation using MPI that enables near-linear scaling on large high-performance computing (HPC) systems.

Together, these innovations significantly reduce iteration counts, improve parallel efficiency, and make high-quality MBIR reconstruction practical for large-scale tomographic imaging. These advances open the door to near-real-time MBIR for applications such as in situ, in operando, and time-evolving experiments.

Introduction

Synchrotron micro- and nano-computed tomography (micro-CT and nano-CT) [14, 18, 6] have become indispensable tools across a wide range of scientific fields, including materials science, geoscience, and energy research. These techniques enable non-destructive three-dimensional imaging of the internal structure of materials with high spatial resolution. Recent advances in synchrotron instrumentation and detector technology have significantly increased acquisition speeds, enabling new experimental modalities such as time-resolved tomography.

Time-resolved tomography allows researchers to observe the evolution of microstructural processes within materials. Examples include crack propagation under mechanical stress [13], degradation of polymer electrolytes during battery operation [11], and fluid transport through porous media [1]. Capturing these dynamic processes requires rapid acquisition of projection data, often under strict dose or time constraints. As a result, the collected datasets frequently contain limited angular sampling and increased measurement noise, making accurate image reconstruction significantly more challenging. Another application that would benefit from MBIR is laminography which is inherently limited-angle with a missing-wedge because of sample geometry [15].

Model-based iterative reconstruction (MBIR) has emerged as a powerful framework for addressing these challenges. By incorporating prior knowledge and accurate physical models into the reconstruction process, MBIR can produce higher-quality reconstructions from noisy or incomplete data compared to conventional analytical approaches [20, 16, 12]. However, these benefits come at the cost of substantial computational expense. Traditional MBIR implementations rely on ray-tracing methods for evaluations of forward and adjoint projection operators, which makes reconstruction prohibitively slow for large-scale or time-resolved tomography experiments.

To address this limitation, the TomoCAM framework [12] was developed as a GPU-accelerated platform for tomographic reconstruction. TomoCAM implements the Radon transform and its adjoint using non-uniform fast Fourier transforms (NUFFTs) [9, 4], enabling significantly faster computation compared to conventional ray-tracing methods. In addition, the framework leverages multi-GPU architectures to parallelize the reconstruction process, allowing large datasets to be processed more efficiently.

In this work, we present several strategies that further improve the performance and scalability of the TomoCAM framework. The main contributions of this work are:

- Reformulating the gradient operator using a multi-level Toeplitz structure enables the use of FFT-based fast convolutions
- A hierarchical multi-resolution optimization strategy that accelerates convergence by progressively refining the reconstruction grid.
- A filtered backprojection-based initialization that significantly reduces the number of iterations required for convergence.
- A hybrid distributed MPI and multi-GPU implementation

that enables large-scale tomographic reconstruction on modern HPC systems.

The code is available under an open-source license at `https://github.com/lbl-camera/tomocam`.

Theory and Mathematical Background

Tomographic reconstruction is fundamentally based on the inversion of the Radon transform. We briefly review the mathematical background here and refer the reader to standard texts for additional details [3, 10]. Let $f : \mathbb{R}^2 \to \mathbb{R}$. The Radon transform of f is defined by line-integral

$$R_\theta f(t) = \int_{\ell(t,\theta)} f(x,y) d\ell \tag{1}$$

where $\ell(t,\theta)$ denotes the line parameterized by the offset t and projection angle θ. The Radon transform is closely related to the *Central Slice Theorem*, which states that the one-dimensional Fourier transform of a projection acquired at angle θ corresponds to a slice of the two-dimensional Fourier transform of the object $f(x,y)$ taken at the same angle θ, i.e.,

$$\mathbb{F}Rf(\omega) = \hat{f}(\omega\cos\theta, \omega\sin\theta) \tag{2}$$

where $\hat{f}$ in two-dimensional Fourier transform of f, and $(\omega\cos\theta, \omega\sin\theta) = (k_x, k_y)$ the Fourier-domain coordinates.

Traditional iterative tomographic reconstruction methods often rely on repeatedly applying forward and back-projection operators, which can be interpreted as large implicit linear systems. In previous work [12], we implemented these operators using non-uniform fast Fourier transforms (NUFFTs), reducing the computational complexity from $O(N^2)$ to $O(N\log N + M\log_{10}\varepsilon^{-1})$, where N denotes the number of voxels in target reconstruction, M is the total number of pixels in the projection data, and ε is the desired numerical precision.

The forward projection, i.e., the output of the Radon transform can be calculated as

$$Rf = \mathbb{F}^{-1}\mathscr{F}f, \tag{3}$$

where $\mathscr{F}$ denotes the NUFFT Type-2 operation that maps the Cartesian grid in the object domain to a polar grid in the Fourier domain, and $\mathbb{F}^{-1}$ denotes the one-dimensional inverse FFT applied along the radial direction for a given projection angle. Similarly, the adjoint operation, commonly referred to as back-projection, can be expressed as

$$R^*g = \mathscr{F}^*\mathbb{F}g, \tag{4}$$

where $\mathscr{F}^*$ denotes the NUFFT Type-1 operation that maps the polar grid in the Fourier domain back to the Cartesian grid in the object domain, and $\mathbb{F}$ denotes the one-dimensional FFT along the radial direction.

When the projection data doesn't satisfy angular Shannon-Nyquist criterion, regularization Φ is required. The tomographic reconstruction problem can then be formulated as the following optimization problem:

$$\hat{f} = \arg\min_f \frac{1}{2}\|Rf - g\|_2^2 + \lambda\Phi(|\nabla f|). \tag{5}$$

Here, g denotes the measured projection data, Φ is a regularization function that enforces prior knowledge, and λ is a regularization parameter that balances data fidelity and regularization. TomoCAM uses an edge-preserving qGGMRF prior for Φ; we refer the reader to [20, 16] for details. The formula for the iterative step is obtained by differentiating (5), i.e.,

$$f^{k+1} = f^k - \frac{1}{L}\left(R^*(Rf - g) - \lambda\nabla\Phi(|\nabla f|)\right), \tag{6}$$

where L is the Lipchitz constant of the gradient.

It is worth noting that, for 3D reconstruction, the Radon-based data fidelity term is evaluated slice-by-slice, whereas the qGGMRF prior enforces spatial coupling across the full 3D volume. This optimization problem can be solved using various iterative algorithms. In this work, we employ the Fast Iterative Method with Restart [7], an accelerated variant of Nesterov's method [17], due to its favorable convergence properties for large-scale problems.

In the following sections, we describe several newly implemented strategies that further improve the convergence rate and computational efficiency of this framework.

Toeplitz Structure Exploitation

The primary computational bottleneck in solving Eq. (5) arises from repeated evaluations of $R^*(Rf - g)$ in (6), during gradient computation. Each evaluation requires two NUFFTs. Although NUFFTs have log-linear complexity, the gradient calculation can be further accelerated by exploiting the Toeplitz structure of operator composition R^*R. Specifically, this composition can be expressed as a convolution between f and a *point spread function* (PSF) kernel K, that depends on the coordinates sampled in the Fourier space.

While Toeplitz acceleration has been explored in MRI reconstruction [5, 19] and cryo-EM [21], its application to NUFFT based tomographic formulations has received limited attention. This observation allows the data fidelity gradient to be computed using FFT-based convolution.

The data fidelity gradient is given by

$$\nabla_f e = R^*(Rf - g) = R^*Rf - R^*g. \tag{7}$$

In (7), we can rewrite

$$R^*Rf = K \circledast f, \tag{8}$$

where $\circledast$ refers to convolution.

The PSF kernel K can be precomputed by applying $\mathscr{F}^*$ to a function $\mathbf{1}$, where $\mathbf{1}(x,y) = 1$ for all (x,y) sampled in Fourier space. To accurately recover R^*Rf, the PSF kernel K must be sampled on a larger grid with at least $2N-1$ grid points in each dimension. Accordingly, to compute the gradient, f is zero-padded to size $2N-1$ in each dimension, and the resulting convolution is restricted to an $N \times N$ region. The convolution can then be evaluated efficiently using FFTs:

$$R^*Rf = Res_{N\times N}\left(\mathbb{F}^{-1}\left(\mathbb{F}K \odot \mathbb{F}Pad_{2N-1\times 2N-1}f\right)\right), \tag{9}$$

where $\mathbb{F}$ denotes the two-dimensional fast Fourier transform, *Res* is restriction, *Pad* is extension, and $\odot$ represents element-wise multiplication. The reformulated loss and gradient equations are given as,

$$e = \frac{1}{2} f^{\mathsf{T}} K \circledast f - f^{\mathsf{T}} R^{*} g + \frac{1}{2} g^{\mathsf{T}} g \tag{10}$$

$$\nabla_f e = K \circledast f - R^{*} g \tag{11}$$

Both $\mathbb{F}K$ and $R^{*}g$ are computed once prior to the iterative reconstruction. This reformulation significantly reduces the computational cost of the data fidelity gradient and loss, as it requires only two FFTs and one element-wise multiplication instead of a forward and a backprojection.

Initialization Strategy

The choice of initialization can significantly influence the convergence behavior of iterative reconstruction algorithms. Starting the optimization from a poor initial estimate may require many iterations before the algorithm approaches a meaningful solution. To mitigate this issue, we adopt an informed initialization strategy based on filtered backprojection (FBP), which provides a fast approximate reconstruction from the measured projection data [3, 10].

Let g denote the measured projection data and R the forward projection operator. A quick initial estimate can be obtained by applying the adjoint operator R^{*} to the projection data. However, an unfiltered backprojection produces a blurred object due to the frequency response of the Radon transform. To compensate for this effect, the projection data are first filtered with a ramp filter prior to backprojection [8].

Let $p_\theta(s)$ denote the projection at angle θ and sinogram coordinate s. The filtered projection is computed as

$$\tilde{p}_\theta(s) = (h \circledast p_\theta)(s), \tag{12}$$

where $\circledast$ denotes convolution and $h(s)$ is the ramp filter kernel. In the Fourier domain, the ramp filter is defined as

$$\hat{h}(\omega) = |\omega|, \tag{13}$$

which amplifies high-frequency components to compensate for the smoothing introduced by the projection process.

The initial guess is then obtained using filtered backprojection

$$f_0 = R^{*} \tilde{g}, \tag{14}$$

where $\tilde{g}$ denotes the collection of filtered projections. This initialization provides a computationally inexpensive approximation of the true solution.

Although the filtered backprojection reconstruction may contain artifacts due to limited-angle sampling or measurement noise, it typically captures the dominant low-frequency structure of the underlying image. Consequently, it serves as an effective starting point for the iterative optimization. By initializing the reconstruction with f_0, the algorithm begins closer to the desired solution, reducing the number of iterations required for convergence and lowering the overall computational cost.

Hierarchical Multi-Resolution Reconstruction

To accelerate convergence, we employ a hierarchical multi-resolution reconstruction strategy. Instead of solving the inverse problem directly at the target resolution, the reconstruction is performed progressively across a sequence of grids with increasing spatial resolution. This approach leverages the observation that the low-frequency components of the solution dominate the early stages of reconstruction and can be estimated reliably on coarser grids at significantly lower computational cost.

We define a hierarchy of grids $\Omega_0, \Omega_1, \ldots, \Omega_L$, where Ω_0 corresponds to the coarsest discretization and Ω_L to the desired reconstruction resolution. At each level l, the optimization problem in Eq. (5) becomes

$$f^{(l)} = \arg\min_f \frac{1}{2} \|R_{(l)} f - g_{(l)}\|_2^2 + \lambda \Phi(|\nabla f|), \tag{15}$$

where $R_{(l)}$ denotes the forward operator defined on grid Ω_l, and $g_{(l)}$ represents projection data that have been downsampled to match that resolution. For downsampling, we select coarse-grid points by subsampling the full-resolution projection data.

The optimization is first performed on the coarsest grid Ω_0, producing an estimate $f^{(0)}$. After convergence at this level, the solution is interpolated to the next finer grid using Lanczos interpolation. Lanczos interpolation is based on a windowed sinc kernel that provides high-quality band-limited interpolation while suppressing ringing artifacts.

Let P_l denote the interpolation operator mapping grid Ω_l to Ω_{l+1}. The initialization for the next level is given by $f_0^{(l+1)} = P_l f^{(l)}$. The Lanczos kernel with window parameter a is defined as

$$L(x) = \begin{cases} \frac{\sin(\pi x)\sin(\pi a x)}{a\pi^2 x^2}, & |x| < a, \\ 0, & \text{otherwise}, \end{cases} \tag{16}$$

In practice, the interpolation operator P_l is implemented as a convolution with a separable Lanczos kernel. For a discrete signal $f[i]$, the interpolated value at location x is given by

$$\tilde{f}(x) = \sum_i f[i]\, L(x - i). \tag{17}$$

For multi-dimensional reconstructions, the interpolation is applied separably along each spatial dimension. The Lanczos window truncates the sinc kernel to a finite support $|x| < a$, limiting computational cost while preserving most of the spectral properties of ideal sinc interpolation. In our implementation, we use a fixed window size $a = 3$, which provides a good compromise between interpolation accuracy and computational efficiency.

The reconstruction is then refined by solving the optimization problem on Ω_{l+1} using $f_0^{(l+1)}$ as the initial estimate. This coarse-to-fine procedure is repeated until the finest grid Ω_L is reached.

This hierarchical strategy accelerates convergence because large-scale structures are recovered on coarse grids where the number of unknowns is small and the optimization landscape is smoother. The interpolated solution therefore provides a high-quality initialization for subsequent levels, allowing the optimization on finer grids to focus primarily on recovering high-frequency details. As a result, high-quality reconstructions can be obtained with significantly fewer iterations at the finest resolution.

Distributed Multi-GPU Implementation

To handle large-scale tomographic datasets and accelerate computation, the reconstruction framework is implemented using a hybrid parallel architecture that combines distributed-memory parallelism with GPU acceleration. Inter-node communication is handled using the Message Passing Interface (MPI), while GPUs within each node are coordinated using shared-memory parallelism. This design distributes both the computational workload and memory footprint across multiple GPUs while maintaining efficient evaluation of the gradient and regularization terms required by the iterative reconstruction algorithm.

The reconstruction volume is partitioned along the rotation axis of the tomography geometry. The domain is first decomposed across MPI processes (compute nodes), and each node further partitions its assigned sub-volume among the GPUs available on that node. This hierarchical decomposition enables efficient scaling across both nodes and GPUs.

The data fidelity gradient depends only on local volume partitions and no point-to-point communication between neighboring processes is required. In contrast, the qGGMRF prior used in this work relies on interactions between each voxel and its 26 nearest neighbors in three dimensions. Thus, boundary information must be exchanged between adjacent sub-volumes during each iteration. To support this, a halo region is maintained for each partition that stores boundary voxels from neighboring processes. After every iteration, halo regions are updated through MPI communication to ensure that the regularization gradient is computed consistently across partition boundaries.

Within each compute node, GPU execution is orchestrated using a lightweight scheduling system designed to overlap data transfers and GPU computation. The scheduler consists of three cooperating threads connected through thread-safe work queues. The first thread transfers input data from host memory to GPU memory and inserts the corresponding work items into a bounded work queue. The second thread retrieves tasks from this queue and executes the associated GPU kernels. Completed tasks are placed into an outgoing queue, from which a third thread asynchronously transfers the results from GPU memory back to host memory.

The resulting hybrid MPI–GPU framework scales efficiently on modern multi-GPU HPC systems, enabling reconstruction of large volumetric datasets that would otherwise exceed the memory capacity of a single device. By combining distributed computation, shared-memory parallelism, and GPU acceleration, the proposed implementation makes high-quality model-based iterative reconstruction (MBIR) practical for large-scale tomographic imaging applications.

Numerical Benchmarks

We used a sparsely sampled nano-CT experimental dataset, published on the Tomobank data repository [2], to evaluate the performance of the proposed methods. The dataset consists of 212 projections with detector size 2048×2448 pixels; see Figure 1. The benchmarks are designed to assess four key aspects of the method: (i) the computational efficiency of the forward and adjoint operators, (ii) the impact of the initialization strategy, (iii) the effectiveness of the multi-resolution reconstruction scheme, and (iv) the scalability of the distributed HPC implementation.

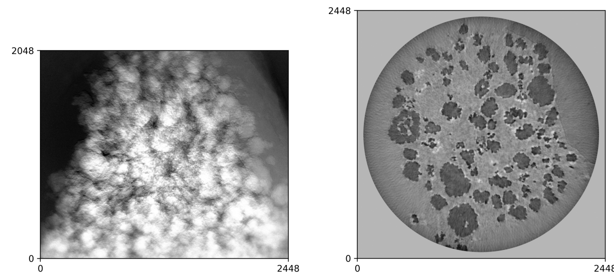

(a) (b)

Figure 1. A Nano-CT projection data from the Tomobank was used for benchmarking. (a) A projection image from the tomogram. (b) A slice of reconstructed volume using MBIR.

Toeplitz Acceleration of Forward and Adjoint Operators

The performance of the forward and adjoint operators is critical for model-based iterative reconstruction since these operators dominate the computational cost of each iteration. We compare two implementations of the operator, i.e., the direct method that uses $\|Rf - g\|_2$ and $R^*(Rf - g)$ to calculate loss and gradient respectively, against the reformulated operators that leverage multi-level Toeplitz structure (10) and (11), across increasing image widths.

Figures 2 report the runtime for evaluating the gradient and loss. The optimized implementation consistently reduces computation time while maintaining high numerical accuracy. The relative difference between the two implementations remains small across all tested image sizes, indicating that the accelerated operator preserves the accuracy of the original method. These results demonstrate that the optimized operator provides a significant performance improvement without compromising reconstruction fidelity.

Effect of Initialization

We next evaluate the impact of the proposed filtered backprojection (FBP) initialization compared with a constant initialization. Figure 3 shows the data fidelity term $\|Rf - g\|_2$ as a function of iteration for both approaches.

The FBP initialization begins with a substantially lower residual and maintains a consistent advantage throughout the optimization. This occurs because the FBP estimate captures the dominant low-frequency structure of the object, placing the iterative solver loser to the optimal solution. Consequently, the reconstruction algorithm requires fewer iterations to reach the same residual level compared with a constant initialization. FBP-based initialization significantly reduces the initial residual.

Hierarchical Multi-Resolution Reconstruction

We now evaluate the impact of the hierarchical multi-resolution reconstruction strategy. While the filtered backprojection (FBP) initialization described in the previous subsection naturally serves as an effective starting point for the coarse grid, it is

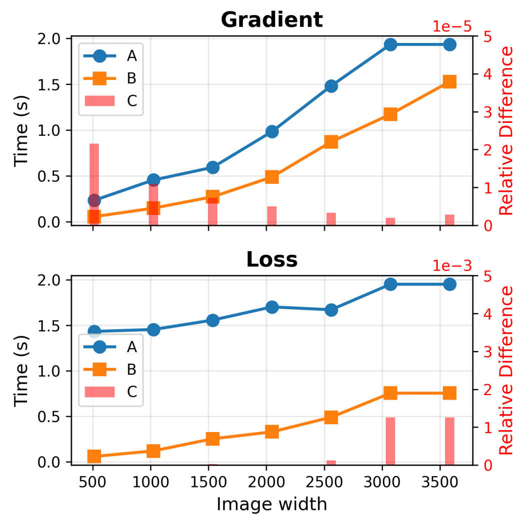

Figure 2. *Runtime comparison between the direct method (A) and the Toeplitz-based method (B) for gradient and loss, as well as the relative difference (C) between the two methods, as a function of image width. The optimized Toeplitz implementation consistently loss, as well as the relative difference (C) between the two methods, as a function of image width. The optimized Toeplitz implementation consistently reduces computation time while maintaining high numerical accuracy. The left y-axis represents runtime, and the right y-axis represents the relative difference between the direct and Toeplitz-based methods.*

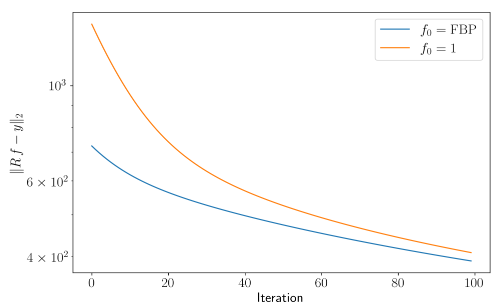

Figure 3. *The initial residual is approximately **a factor of two lower** compared to constant initialization, effectively reducing the optimization effort by **20–25 iterations**.*

intentionally disabled in this experiment to study the effect of the hierarchical reconstruction strategy independently.

Figure 4 compares the convergence behavior of the proposed coarse-to-fine approach against a single-resolution optimization. The reconstruction begins on a coarse grid of (611×611), where large-scale structures are recovered efficiently at low computational cost. The solution is then interpolated onto progressively finer grids, with the resolution increased by a factor of 2 at each stage, providing improved initialization for subsequent optimization. As shown in Figure 4, this strategy substantially reduces the number of iterations required at the highest resolution. While the exact performance gains may vary across datasets, in this case the total runtime is reduced by approximately a factor of 5. The coarse-to-fine scheme therefore improves convergence while reducing the overall computational effort.

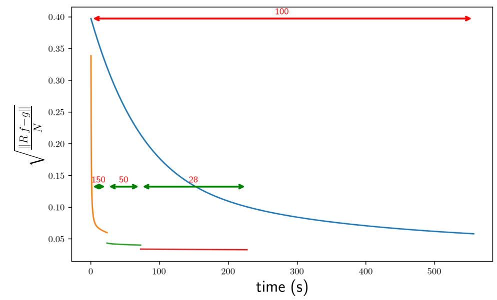

Figure 4. *Convergence behavior of the multi-resolution reconstruction strategy. The coarse-to-fine approach reduces computational cost at the highest resolution by providing improved initial estimates at each refinement stage. In this example, three hierarchical grid levels are used, starting from a grid of size $(N/4 \times N/4)$, doubling the resolution at each stage. Numbers next to arrows indicate iterations at each resolution.*

Distributed HPC Performance

Finally, we evaluate the scalability of the distributed implementation on the Perlmutter supercomputer at the National Energy Research Scientific Computing Center (NERSC). The reconstruction volume size in this experiment is $2048 \times 2447 \times 2447$ voxels.

Figure 5 shows the reconstruction time as a function of the number of compute nodes. The results demonstrate scaling as the number of nodes increases from 16 to 128. The total reconstruction time decreases from approximately 30 minutes on 16 nodes to less than 9 minutes using 128 nodes. These results confirm that the distributed MPI implementation effectively utilizes multiple GPUs while keeping communication overhead manageable. The proposed framework therefore enables efficient reconstruction of large volumetric datasets on modern HPC systems.

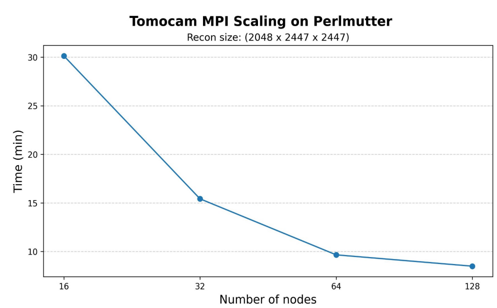

Figure 5. *Performance of the distributed MPI implementation on the National Energy Research Scientific Computing Center's Perlmutter supercomputer for a reconstruction volume of size $2048 \times 2447 \times 2447$.*

Conclusion

We presented an efficient framework for large-scale model-based iterative reconstruction (MBIR). By leveraging the multi-level Toeplitz structure of the Radon transform and its adjoint, we reformulated the data fidelity term, enabling significant computational performance improvements. To further accelerate convergence, we introduced a smart initialization strategy based on filtered backprojection and a hierarchical multi-resolution strategy. This approach captures the low-frequency structure of the object early in the optimization process, leading to faster and more stable convergence. We also developed a distributed multi-GPU implementation that partitions the volume along the rotation axis and efficiently manages inter-node inter-GPU communication through halo exchanges.

Numerical experimental results demonstrate that the proposed strategies achieve significantly faster reconstruction while maintaining quality for large volumetric datasets. These results highlight the practical viability of combining operator-efficient formulations, hierarchical optimization, and distributed computing for next-generation tomographic reconstruction. The presented framework opens the door to near-real-time MBIR for applications such as in situ, in operando, and time-evolving experiments.

Acknowledgments

This work was supported by the Center for Advanced Mathematics for Energy Research Applications (CAMERA), funded by the Advanced Scientific Computing Research (ASCR) and Basic Energy Sciences (BES) programs of the Office of Science of the U.S. Department of Energy (DOE), as well as by the DOE ASCR Early Career Research Program, under Contract No. DE-AC02-05CH11231.

References

[1] R. Boutchko, V. L. Rayz, N. T. Vandehey, J. P. O'Neil, T. F. Budinger, P. S. Nico, J. L. Druhan, D. A. Saloner, G. T. Gullberg, and W. W. Moses. Imaging and modeling of flow in porous media using clinical nuclear emission tomography systems and computational fluid dynamics. *J Appl Geophy*, 76:74–81, Jan. 2012.

[2] F. D. Carlo, D. Gürsoy, D. J. Ching, K. J. Batenburg, W. Ludwig, L. Mancini, F. Marone, R. Mokso, D. M. Pelt, J. Sijbers, and M. Rivers. TomoBank: a tomographic data repository for computational x-ray science. *Measurement Science and Technology*, 29(3):034004, feb 2018.

[3] C. L. Epstein. *Introduction to the Mathematics of Medical Imaging, Second Edition*. second edition, 2007.

[4] J. Fessler and B. Sutton. Nonuniform fast fourier transforms using min-max interpolation. *Signal Processing, IEEE Transactions on*, 51:560 – 574, 03 2003.

[5] J. A. Fessler. On nufft-based gridding for non-cartesian mri. *Journal of Magnetic Resonance*, 188(2):191–195, 2007.

[6] M. Ge, D. S. Coburn, E. Nazaretski, W. Xu, K. Gofron, H. Xu, Z. Yin, and W.-K. Lee. One-minute nano-tomography using hard x-ray full-field transmission microscope. *Applied Physics Letters*, 113(8):083109, 2018.

[7] P. Giselsson and S. P. Boyd. Monotonicity and restart in fast gradient methods. In *53rd IEEE Conference on Decision and Control, CDC 2014, Los Angeles, CA, USA, December 15-17, 2014*, pages 5058–5063. IEEE, 2014.

[8] R. G.N. and L. A.V. Three-dimensional reconstruction from radiographs and electron micrographs: application of convolutions instead of Fourier transforms. *Proc Natl Acad Sci U S A*, 1971.

[9] L. Greengard and J.-Y. Lee. Accelerating the nonuniform fast fourier transform. *SIAM Review*, 46(3):443–454, 2004.

[10] A. C. Kak and M. Slaney. *Principles of computerized tomographic imaging*. Society for Industrial and Applied Mathematics, Philadelphia, 2001.

[11] D. Kulkarni, S. J. Normile, L. G. Connolly, and I. V. Zenyuk. Development of low temperature fuel cell holders for operando x-ray micro and nano computed tomography to visualize water distribution. *Journal of Physics: Energy*, 2(4):044005, oct 2020.

[12] D. Kumar, D. Y. Parkinson, and J. J. Donatelli. tomoCAM: fast model-based iterative reconstruction via GPU acceleration and non-uniform fast Fourier transforms. *Journal of Synchrotron Radiation*, 31(1):85–94, Jan. 2024.

[13] N. M. Larson and F. W. Zok. Insights from in-situ x-ray computed tomography during axial impregnation of unidirectional fiber beds. *Composites Part A: Applied Science and Manufacturing*, 107:124–134, 2018.

[14] A. A. MacDowell, D. Y. Parkinson, A. Haboub, E. Schaible, J. R. Nasiatka, C. A. Yee, J. R. Jameson, J. B. Ajo-Franklin, C. R. Brodersen, and A. J. McElrone. X-ray micro-tomography at the Advanced Light Source. In S. R. Stock, editor, *Developments in X-Ray Tomography VIII*, volume 8506, page 850618. International Society for Optics and Photonics, SPIE, 2012.

[15] M. A. Marcus. Information content of and the ability to reconstruct dichroic x-ray tomography and laminography. *Opt. Express*, 30(22):39445–39465, Oct 2022.

[16] K. A. Mohan, S. V. Venkatakrishnan, L. F. Drummy, J. Simmons, D. Y. Parkinson, and C. A. Bouman. Model-based iterative reconstruction for synchrotron x-ray tomography. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6909–6913, 2014.

[17] Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Publishing Company, Incorporated, 1 edition, 2014.

[18] V. Nikitin, A. Tekawade, A. Duchkov, P. Shevchenko, and F. De Carlo. Real-time streaming tomographic reconstruction with on-demand data capturing and 3D zooming to regions of interest. *Journal of Synchrotron Radiation*, 29(3):816–828, May 2022.

[19] T. Ou. gnufftw: Auto-tuning for high-performance gpu-accelerated non-uniform fast fourier transforms. Master's thesis, EECS Department, University of California, Berkeley, May 2017.

[20] S. V. Venkatakrishnan, L. F. Drummy, M. D. Graef, J. P. Simmons, and C. A. Bouman. Model based iterative reconstruction for Bright Field electron tomography. In C. A. Bouman, I. Pollak, and P. J. Wolfe, editors, *Computational Imaging XI*, volume 8657, pages 75 – 86. International Society for Optics and Photonics, SPIE, 2013.

[21] L. Wang, Y. Shkolnisky, and A. Singer. A fourier-based approach for iterative 3d reconstruction from cryo-em images. *arXiv preprint arXiv:1307.5824*, 2013.

Author Biography

Dinesh Kumar is a Research Scientist within the Center for Advanced Mathematics for Energy Research Applications (CAMERA) at Lawrence Berkeley National Laboratory. He received his Ph.D. in Mechanical Engineering from the University at Buffalo in 2012 and was awarded an Advanced Light Source Postdoctoral Fellowship in 2014. His work involves new techniques for analyzing synchrotron data for X-ray Scattering and Tomography.

Jeffrey Donatelli is a Staff Scientist at Lawrence Berkeley National Laboratory, where he leads the Mathematics for Experimental Data Analysis Group and serves as the Deputy Director of the CAMERA center. He received his Ph.D. in Applied Mathematics from the University of California, Berkeley in 2013. His research focuses on inverse problems and computational harmonic analysis, with applications to reconstruction and data analysis for light, electron, and neutron source experiments.