Research Statement

Xiao Wang

1 My Research Topic and Significance

My research interest is to apply image processing, medical physics, high performance computing, and machine learning to computed tomographic (CT) reconstruction. In a CT scan, an X-ray source and a detector rotate around a sample object and produces signals at every view angle. The goal of CT reconstruction is then to process signals from all view angles and produce images for 3D interior structure of the object.

The performance and quality of CT reconstruction are significant for many applications. In medical imaging, 26 million CT scans are performed each year in US alone, and a CT reconstruction with a high spatial resolution and little noise allows radiologists to identify lesions and early-stage cancer nodules easily. Furthermore, a high-quality CT reconstruction with little image noise allows radiologists to reduce radiation doses used in a CT scan while maintaining image quality, and thereby reducing radiation risk for patients. In scientific and biology imaging, scientists use synchrotron imaging and electron microscopy, which are both CT imaging in nature, to get an interior view of biology and material samples. An advanced CT reconstruction with super-resolution allows scientists to see details at unprecedented image resolution and discover new bio-structure beyond the hardware limit of CT detector resolution. In security imaging, industrial CT is used for checked baggage and cargo scanning at airports and seaports. High-resolution and real-time reconstruction enables airport and seaport workers to tell home-made explosives from other liquid with more confidence, reduce the need to open checked baggage manually, and allows faster security screening at airports and seaports. To support the above applications, **my research concerns how to significantly increase CT reconstruction spatial resolution while reducing artifacts and noise. At the same time, my research develops high performance and parallel algorithms to achieve real-time reconstructions for time-sensitive applications.**

2 Background and Summary of My Contributions

The mainstream approach to reconstruct from CT signals to images is Fourier methods. Fourier methods are the standard method for almost every commercial CT and they use Fourier and inverse Fourier transform to interpolate discretely sampled CT signals acquired from detectors back to a reconstruction in the continuous spatial domain. Because of the extensive amount of signal interpolation and geometry approximation in Fourier methods, Fourier methods fail to take advantage of the high detector resolution from the state-of-the-art CT scanner, and have unsatisfactory reconstruction quality with blurry image details and significant aliasing artifacts [3].

To achieve a high reconstruction spatial resolution with few artifacts and little noise, my research work, known as Model-Based Iterative Reconstruction (MBIR), uses medical physics to construct a forward model to represent the CT scanner geometry, X-ray photon acquisition physics as well as the signal noise, and fit the CT signals with the forward model. At the same time, my work constructs an image prior model to represent the reconstruction spatial property, such as image noise or texture, through statistical distribution. As the forward and prior models realistically retain the CT scan's physics-geometry, preserve X-ray acquisition information and image spatial property, the final reconstruction can often achieve better image quality with higher spatial resolution, less noise and fewer artifacts than the Fourier methods [2, 3].

MBIR, however, has several key technical issues that prevent it from a wide use, and in summary these issues are: (1) The amount of computations for MBIR is several magnitudes of order more than the Fourier methods. Furthermore, MBIR exhibits poor cache locality as the data access layout for each 3D pixel, also called a voxel, requires an array of memory to be accessed in sinusoidal patterns, and thus leads to a poor per-core efficiency. (2) Various numerical algorithms to compute MBIR has a trade-off between parallelism, convergence rate and cache locality. Therefore, it is difficult to have a massively parallel implementation without worsening its convergence rate or cache locality. (3) The realistic but complex forward model has a very large memory requirement to store its physics-geometry information, whose size grows linearly with image size. Limited by its memory footprint, MBIR, therefore, has a poor super-resolution capability. Any naive way to shrink or distribute the forward model among nodes corrupts the physics-geometry information. (4) The conventional MBIR is successful in removing high-frequency noise and preserving high-contrast edges, but cannot effectively denoise at low-frequency. Therefore, the image appearance is quite different from what vendors offer on their clinical CT, and radiologists are not used to MBIR's noise texture. (5) The conventional MBIR has only been implemented on the single-source CT scanner. For the state-of-the-art CT scanner (dual-source and flying-focal spot CT), the forward model must be reinvented to accommodate the new geometry and acquisition. To address the above issues, my research work makes the following contributions:

1. Introduces the idea of super-voxel (SV) to significantly increase cache locality, prefetching, and completely regularize data access for improved SIMD operations [7, 4] (Project 1).
2. Proposes a parallel SV algorithm for fully 3D reconstruction, and achieves massive parallelism while achieving a convergence rate faster than other numerical algorithms [8] (Project 2).
3. Proposes an Asynchronous Consensus MBIR (AC-MBIR) that significantly reduces memory footprint and increases image super-resolution while keeping all the physics-geometry information intact [9] (Project 3).
4. Develops a Consensus Equilibrium framework that allows MBIR to use advanced but non-differentiable denoising filters, such as Non-Local Mean (NLM) and Block-Matching 3D Filtering (BM3D), as a prior model. With these advanced prior models, MBIR can better denoise at low-frequency with an image texture similar to that of the Fourier methods [6] (Project 4).
5. Develops the first MBIR solution for Dual-Source Flying Focal Spot CT scanner. The proposed solution has a new forward model for such CT geometry, and outperforms the state-of-the-art clinical reconstruction (Siemens ADMIRE) in terms of spatial resolution and with much fewer image artifacts [10] (Project 5).

On the practical side, I have co-developed a software for MBIR for training and education purpose, and can be downloaded at [5]. In addition, my collaborators and I have developed a cloud-based graphical interface solution, Inversion Engine, that allow synchrotron beamline scientists to use high performance MBIR without the need to access the source code. Furthermore, the high-performance MBIR solution is patented [1] and is being commercialized through High Performance Imaging LLC for airport rapid security screening. As to the MBIR clinical solution I developed in project 5, the image quality has brought attention from major CT vendors, such as Siemens and General Electric, to investigate how to port my MBIR solution to their next-generation commercial CT products. Currently, A formal research agreement has formed between Siemens, my collaborators and me to explore a commercial MBIR for Siemens newest scanner under Siemens Contract Loan No. 40001591.

For each contribution and the related project mentioned above, here is the detailed introduction:

Project 1: High Performance and Cache-Friendly MBIR [7, 4].

MBIR produces higher quality images and allows for the use of more general CT scanner geometries than is possible with more commonly used methods. The high computational cost of MBIR, however, often makes it impractical in applications for which it would otherwise be ideal. This project describes a new MBIR implementation that significantly reduces the computational cost of MBIR while retaining its benefits. It proposes a novel organization of the scanner data into super-voxels (SV) that, combined with a super-voxel buffer (SVB), dramatically increase cache locality, prefetching, SIMD operations, and enable parallelism across SVs with an average speedup of 187 times on 20 CPU cores compared to the conventional MBIR, on 3200 security screening datasets.

A separate endeavor of this project develops the first GPU algorithm for Iterative Coordinate Descent based MBIR. The algorithm leverages the SV design, and efficiently exploits three levels of parallelism available in MBIR to better utilize the GPU hardware resources. This project also transforms the CT signal data layout to obtain

more coalesced accesses with several GPU-specific optimizations to boost performance. Across 3200 security screening datasets, our GPU implementation obtains a mean speedup of 4.43 times over the iso-power CPU cores.
Project 2: Fast Converging and Massively Parallel MBIR [8]. Despite of its high performance, project 1 is a single node implementation and restricts its parallelism to shared-memory CPU or GPU cores, and cannot scale to a large cluster with distributed memory. Furthermore, implementations for MBIR have a trade-off among parallelism, cache locality and algorithmic convergence. This project proposes a massively parallel MBIR that ensures a fast convergence through non-homogeneous updates, but also keeps the high cache locality benefits from the SV design. The massively parallel MBIR is compared with two state-of-the-art MBIR implementations on a 69632-core distributed system for both synchrotron and security screening datasets. Results indicate that the massively parallel MBIR has an average speedup of 1665 compared to the fastest state-of-the-art implementations. The huge speedup achieved in this project has received attention from ACM award committee and this work was selected as a finalist for 2017 ACM Gordon Bell Prize.
Project 3: Super-Resolution MBIR with Minimal Memory Footprint [9]. MBIR has a high memory requirement that limits the achievable image resolution. Because of the memory requirement, MBIR often has a limited super-resolution capability. This project proposes Asynchronous Consensus MBIR (AC-MBIR) that uses Consensus Equilibrium (CE) to provide a super-resolution algorithm with a small memory footprint, low communication overhead and a high scalability. Super-resolution experiments on large synchrotron datasets show that AC-MBIR has a 6.8 times smaller memory footprint and 16 times more scalability, compared with project 2's massively parallel MBIR, and maintains a 100% strong scaling efficiency at 146880 cores. In addition, AC-MBIR achieves an average bandwidth of 3.5 petabytes per second at 587520 cores.
Project 4: Plug-and-Play MBIR with Non-Differentiable Prior Model [6]. The conventional MBIR prior model estimates a voxel's value based on its neighboring voxels. The local-neighbor prior model, however, is a low-pass filter and does not denoise effectively at low frequency and significantly changes the noise texture. In contrast, advanced denoisers such as Non-Local Mean (NLM) or Block Matching 3D Filtering (BM3D) not only learn from local neighborhoods, but also from remote voxels, and with much better low frequency denoising capability. Unfortunately, these denoisers are non-differentiable and cannot be easily fused together with the forward model and solved as an optimization problem. This project uses a Consensus Equilibrium (CE) Method to fuse non-differnetiable denoiser as a prior model for MBIR. Experiments on security screening datasets show that MBIR with a BM3D prior model improves image spatial resolution, with less noise and fewer artifacts.
Project 5: MBIR for Dual-Source and Flying-Focal Spot CT [10]. The existing MBIR implementation is for single-source CT reconstruction only, and has no implementation for the newest CT geometry, dual-source flying-focal spot CT (DS-FFS CT), which features a much shorter patient scan time and a much higher sampling rate. This project presents the first MBIR reconstruction method for DS-FFS CT that realistically models its geometry and introduces a new noise model that more accurately estimates CT signal noise than the existing one. We compare image quality between MBIR and the state-of-the-art Siemens reconstruction, ADMIRE, on multiple clinical and phantom datasets. Results show that MBIR has a much higher spatial resolution in terms of Signal-Noise Ratio and Contrast-Noise Ratio while having fewer image artifacts.

3 My Vision for Future Work

I will apply deep learning to MBIR for CT reconstruction, and my deep learning MBIR project have two phases:
Phase 1: Adaptive Priors for Deep Learning MBIR. The existing deep learning method for CT reconstruction adds synthetic noise to noiseless training images, with the noise variance chosen to match that in clinical images. Then a convolution neural network (CNN) learns the image noise from the noiseless and noisy training images, and is applied to test images to remove noise. Despite that deep learning is great for image denoising and shows promise for high quality reconstruction, deep learning has a fatal error: the training data is limited and CNN cannot generalize its learning beyond the training data. With finite training data, a small training data pertubation, such as highly attenuated metal or noise fluctuation in CT signal acquisition, can cause over-denoising and loss of image details. For the existing method, it is necessary to train on very large numbers of cases for every

indication, every patient population, and every CT acquisition parameters to guarantee a stable performance.

To address the above issue, I propose to use a set of image priors for MBIR, including statistical distribution prior (generalized markov random field), non-differentiable denoising prior (NLM and BM3D), and multiple CNNs with each CNN trained with different synthetic noise variances. In addition, each image prior is weighted. When image noise is significant, more prior model weight is given to a CNN trained with a similarly high noise variance, and a low weight is given to others. Vice versa with little image noise, more weight is given to a CNN trained on a similarly low noise variance. When the CT signal noise is large or when anomaly, such as metal, is detected, more weight is given to statistical distribution prior and NLM denoising prior, and less weight is given to CNNs to prevent data pertubation. With a diverse but adaptive choices of prior models, over-denoising and loss of image details can be prevented. As far as I know, this will be the first deep learning CT reconstruction that uses a diverse but adaptively chosen image priors to prevent adversarial pertubation.

Phase 2: High Performance for Deep Learning MBIR. With a diverse choices of image priors, the deep learning MBIR faces several issues for achieving real-time reconstruction. One issue is the CNN training activation function. The convention for training is to use the ReLu activation function as ReLu is more stable and faster to train than others. However, the ReLu function is susceptible to data pertubation and anomaly input values. Another issue is the slow convergence rate for deep learning MBIR. To compute the solution to deep learning MBIR, a popular method is to use linear algebra to fuse the set of image priors with the forward model as a non-convex optimization problem, and use the alternating minimization to compute such problem. With so many prior models, however, the alternating minimization is bound to have a slow convergence. In this project, I propose to use the Consensus Equilibrium (CE) Method to compute such optimization problem and CE can achieve fast convergence even with many image prior models (refer to CE convergence experiments in Figure 7 in [9]). More specifically, the CE Method computes an individual solution for each prior and forward model, and then combines individual solutions into a consensus one in a lossless way. The CE Method, however, tends to have convergence overshoot issue when computed in parallel. To address the two issues mentioned above, I will propose a new activation function that is equally or more trainable than ReLu, but less susceptible to data pertubation. In addition, I will develop a new parallel CE algorithm that is fast converging without convergence overshoot.

References

[1] C. A. Bouman, S. P. Midkiff, S. J. Kisner, and X. Wang. Tomography system. U.S. Patent US20180025514A1, 2015.

[2] K. Jensen, A. C. Martinsen, A. Tingberg, T. M. Aalokken, and E. Fosse. Comparing Five Different Iterative Reconstruction Algorithms for Computed Tomography in an ROC Study. *European Radiology*, 24(12):2989–3002, Dec 2014.

[3] P. J. Pickhardt, M. G. Lubner, D. H. Kim, J. Tang, J. A. Ruma, A. M. D. Rio, and G. H. Chen. Abdominal CT With Model-Based Iterative Reconstruction (MBIR): Initial Results of A Prospective Trial Comparing Ultralow-Dose With Standard-Dose Imaging. *American Journal of Roentgenology*, 199(6):1266–1274, Dec 2012.

[4] A. Sabne, X. Wang, S. J. Kisner, C. A. Bouman, A. Raghunathan, and S. P. Midkiff. Model-based iterative ct image reconstruction on gpus. In *Proceedings of the 22Nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '17, pages 207–220, New York, NY, USA, 2017. ACM.

[5] V. Sridhar, X. Wang, P. Jin, and C. A. Bouman. Baseline mbir reconstruction for 2d and 3d parallel-beam ct, 2019. Download at `https://github.com/cabouman/OpenMBIR`.

[6] V. Sridhar, X.Wang, G. Buzzard, and C. A. Bouman. Distributed memory framework for fast iterative ct reconstruction using multi-agent consensus equilibrium. Under review at *IEEE Transaction on Computational Imaging*. Preprint availabe at ArXiv, November 2019.

[7] X. Wang, A. Sabne, S. J. Kisner, A. Raghunathan, C. A. Bouman, and S. P. Midkiff. High performance model based image reconstruction. In *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '16, pages 2:1–2:12, New York, NY, USA, 2016. ACM.

[8] X. Wang, A. Sabne, P. Sakdhnagool, S. J. Kisner, C. A. Bouman, and S. P. Midkiff. Massively parallel 3d image reconstruction. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '17, pages 3:1–3:12, New York, NY, USA, 2017. ACM.

[9] X. Wang, V. Sridhar, Z.Ronaghi, R. Thomas, J. Deslippe, D. Dilworth, G. Buzzard, S. Midkiff, C. Bouman, and S. Warfield. Consensus equilibrium framework for super-resolution and extreme-scale ct reconstruction. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '19, New York, NY, USA, 2019. ACM.

[10] X.Wang, R. MacDougall, C. A. Bouman, and S. K. Warfield. Model-based iterative reconstruction for dual-source and flying-focal spot computed tomography. Under review at *AAPM Medical Physics Journal*. Preprint availabe at ArXiv, November 2019.