High Performance Model Based Image Reconstruction

Student: Xiao Wang, Advisors: Charles Bouman, Samuel Midkiff

Imaging Lab

wang1698@purdue.edu, bouman@purdue.edu, smidkiff@ecn.purdue.edu

Contribution

In Computed Tomography (CT) methods, Model Based Iterative Reconstruction (MBIR) produces higher quality images than commonly used Filtered Backprojection (FBP) but at a very high computational cost.

We describe a new MBIR implementation, PSV-ICD, which significantly reduces the computational cost of MBIR while retaining its benefits. It describes a novel organization of the scanner data into *super-voxels* (SV) that, combined with a *super-voxel buffer* (SVB), dramatically increases locality and prefetching while enabling parallelism across SVs. Experimental data is presented showing an average speedup of $187X$ on 20 cores.

Model

MBIR is based on the numerical solution of an optimization problem described by

$$\hat{u} = \underset{u \geq 0}{\operatorname{argmin}} \left\{ \frac{1}{2}(v - Au)^T D(v - Au) + S(u) \right\} \tag{1}$$

We consider the image u as a N dimension vector whose elements are called voxels and the data v as a vector of size M containing the CT measurement data. A voxel represents a single data point on the final image, whose value corresponds to the data being reconstructed. The matrix A is a $M \times N$ forward system matrix of the scanner geometry, D is a $M \times M$ diagonal weighing matrix containing the inverse variance of the scanner noise and $S(u)$ is the regularizing prior function which depends upon voxels only. To reduce the amount of computation, iterative coordinate descent(ICD) [2] is used and each individual voxel in the image is updated sequentially. Because of the geometry of the CT tomography system shown in Fig. 1a, the major computation cost comes from the poor locality in the data layout in each voxel update shown in Fig. 1b.

The Super-Voxel (SV)

A SV is a contiguous group of voxels in the shape of a square in the image space and whose measurement data (from scanners) is likely to be close together in memory. Thus, operating on the data for these voxels increases temporal and spatial locality. The data access to a SV follows a sinusoidal band pattern, shown in Fig. 2a. In addition, data in this sinusoidal band is copied to a super-voxel buffer (SVB), shown as a straight band in Fig. 2b, to further improve hardware prefetching and cache locality. Parallelism across SVs (PSV-ICD) also shows good speedup with other proposed parallelism such as grouped coordinate descent (GCD) [1, 2].

References

[1] J. Fessler, E.P. Ficaro, N.H. Clinthorne and K. Lange. Grouped-Coordinate Ascent Algorithms for Penalized-Likelihood Transmission Image Reconstruction. In *IEEE Trans. Med. Img. 1997*

[2] J. Zheng, S.Squib, K.Sauer and C.Bouman . Parallelizable Bayesian Tomography Algorithms With Rapid, Guaranteed Convergence. In *IEEE Trans. Img. Proc. 2000*

CT Projections and the Sinogram

1a)

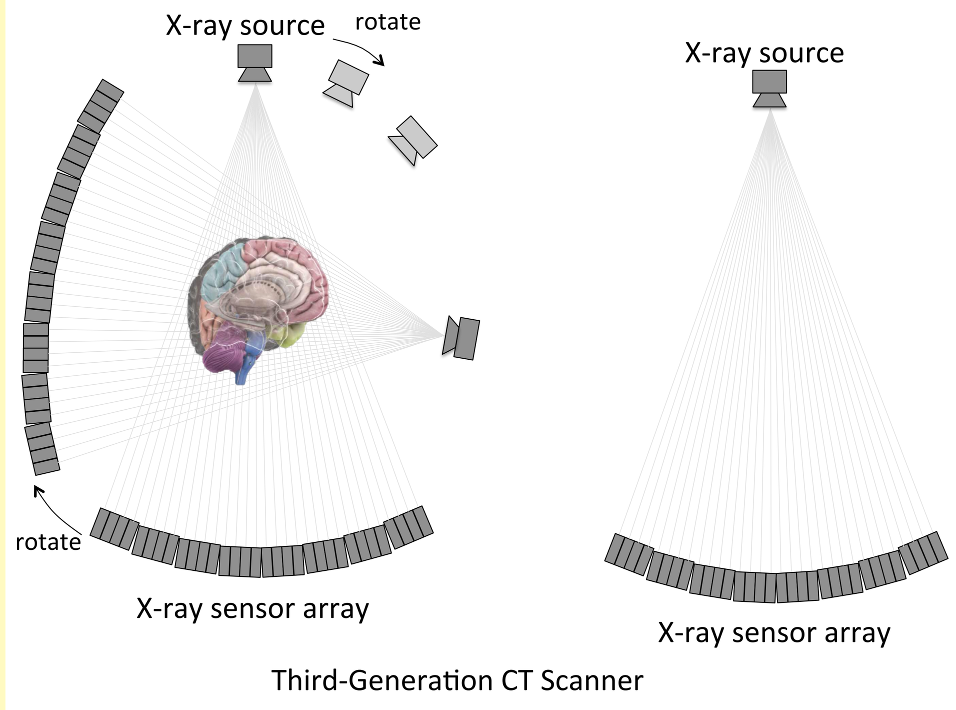

1b)

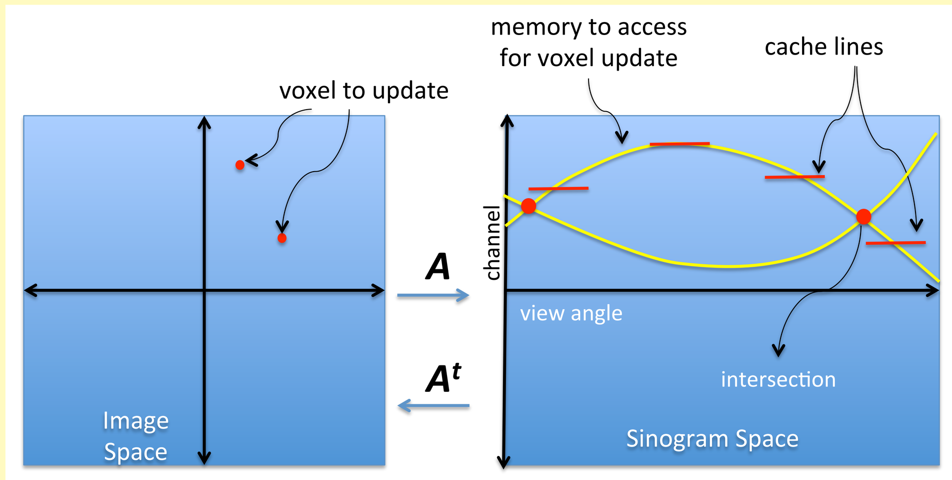

Fig. 1a Illustration of CT system operations. Fig. 1b This figure illustrates the image and measurement data space (or sinogram space) for a typical 2D projection geometry. Notice that each single voxel in the image domain traces out a sinusoidal pattern in the sinogram domain.

PSV-ICD

2a)

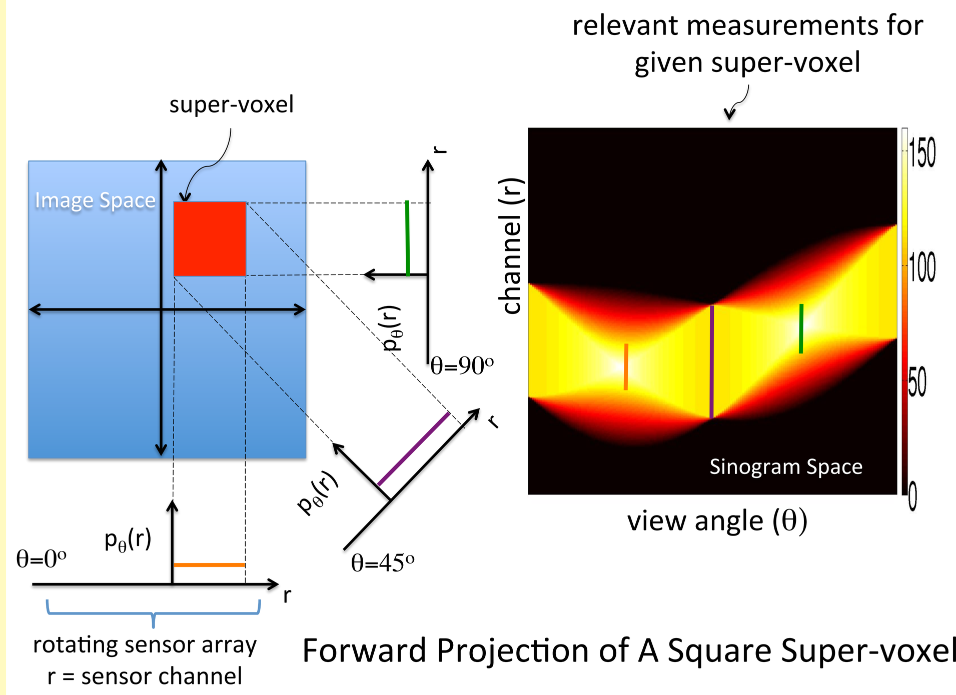

2b)

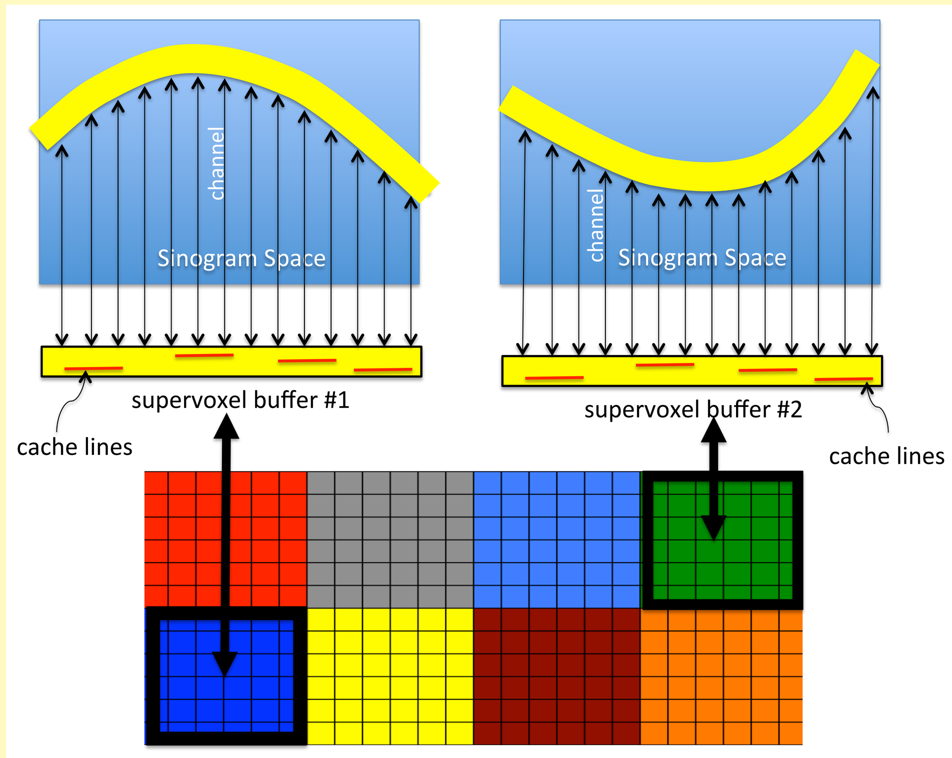

Fig. 2a The radon transformation of a square super-voxel. Fig. 2b PSV-ICD with SVBs. SVBs are shown as straight yellow bands.

Results

PSV-ICD is evaluated on a benchmark data set obtained from a set of Imatron C-300 test scans containing 3200 test cases. Each slice in this data set has the following specification: (1) 720 views uniformly distributed between 0 and 180 degrees. (2) 1024 channels uniformly sampled over the region of interest (ROI). (3) 512 × 512 reconstructioned image size with an embedded circular ROI. All data in this section was collected on 2 standard 2.6 GHZ clock rate Intel Processors Xeon-E5 2660 v3 with 10 cores in each processor.

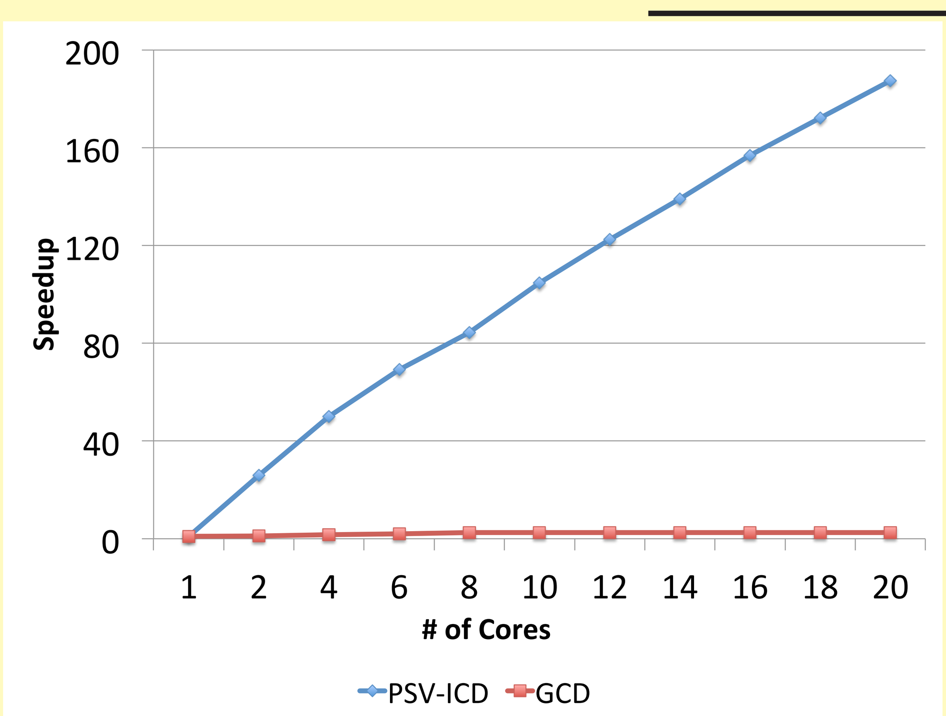

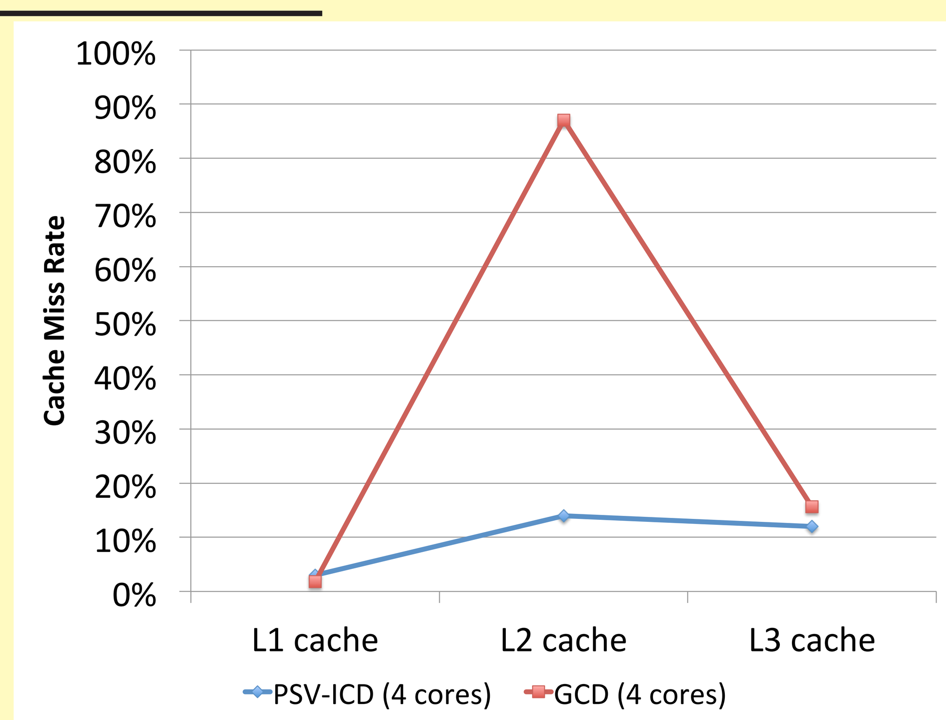

We achieve over $187X$ speedup of the baseline algorithm by using the PSV-ICD algorithm on 2 Intel Xeon-E5 processors with 20 computing cores in total. The existing parallel algorithm, GCD, has only a $2.57X$ speedup. The graph on the right shows the PSV-ICD and GCD's data cache miss rate at different levels of cache.

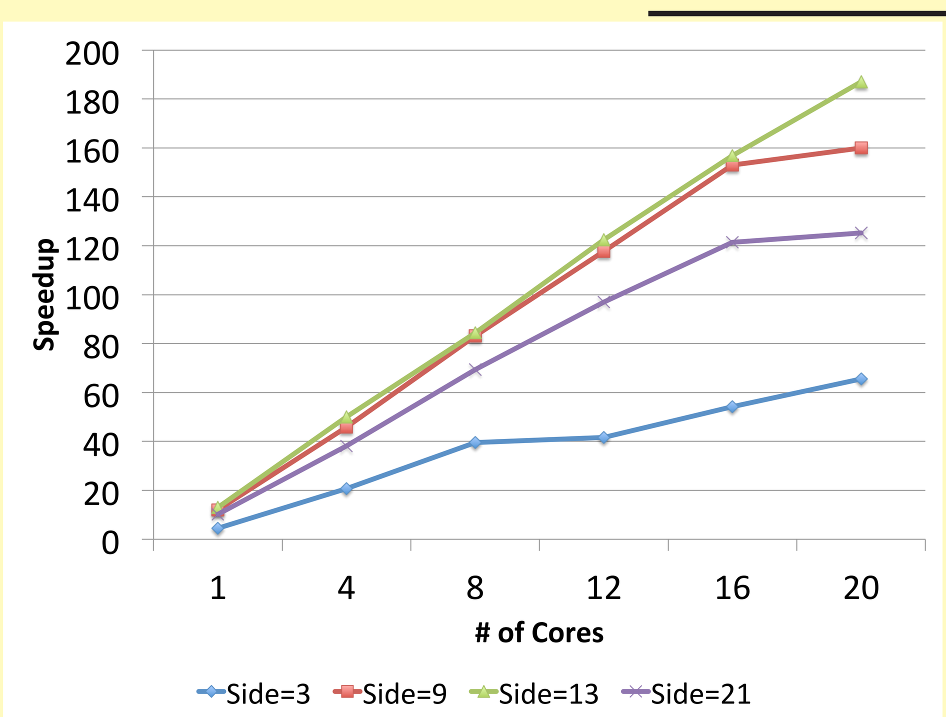

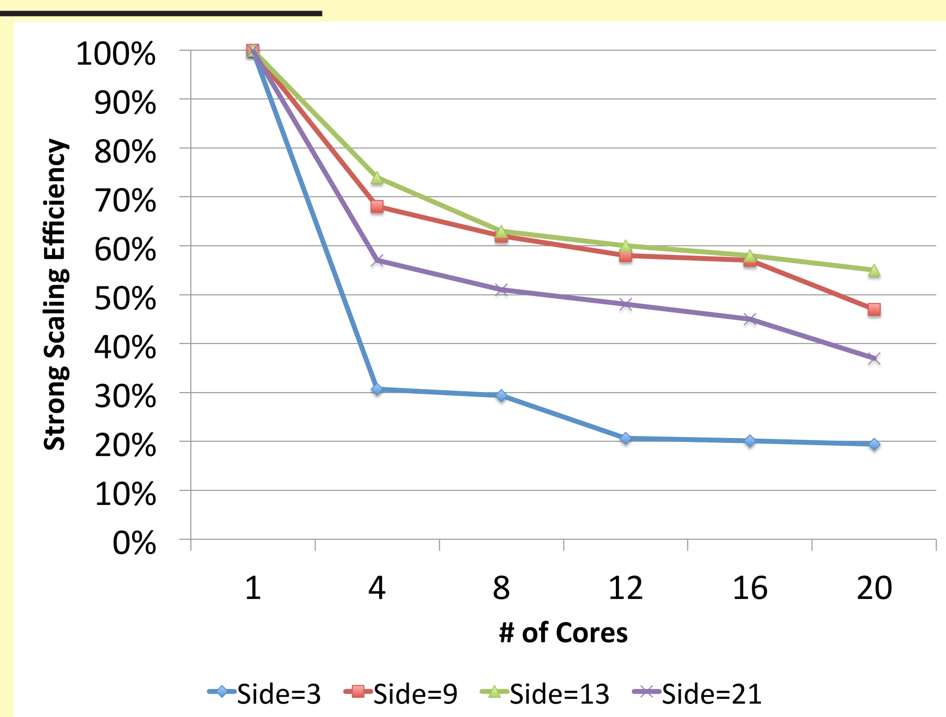

The graph on the left shows PSV-ICD strong scaling speedup at different square SV side lengths. The graph on the right shows the PSV-ICD strong scaling efficiency at different square SV sides lengths.

Funding

This research was supported by the U.S. Department of Homeland Security under SBIR Award HSHQDC-14-C-00058, subcontracted through High Performance Imaging LLC, and by the Indiana Economic Development Corporation (IEDC).