

High-Performance Iterative CT Reconstruction using Super-Voxel Technology

High Performance Imaging (HPI)

Charles A. Bouman, HPI/Purdue University

Sherman Jordan Kisner, HPI

Sam Midkiff, HPI/Purdue University

Anand Raghunathan, HPI/Purdue University

Support by DHS SBIR (HSHQDC-14-C-00058)

and ALERT DHS Center

“This material is partially based upon work supported by the U.S. Department of Homeland Security, Science and Technology Directorate, Office of University Programs, under Grant Award 2013-ST-061-ED0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

Overview

Opportunity

Model-based iterative reconstruction (MBIR) can improve detection probability (P_D) and reduce false alarms (P_{FA}) by reducing artifacts in X-ray CT reconstructions relative to FBP and direct Fourier methods. (ALERT TO3/TT)

Problem

MBIR is too computationally expensive to be practical for security CTX scanners (~1000x more computation than traditional methods)

Solution:

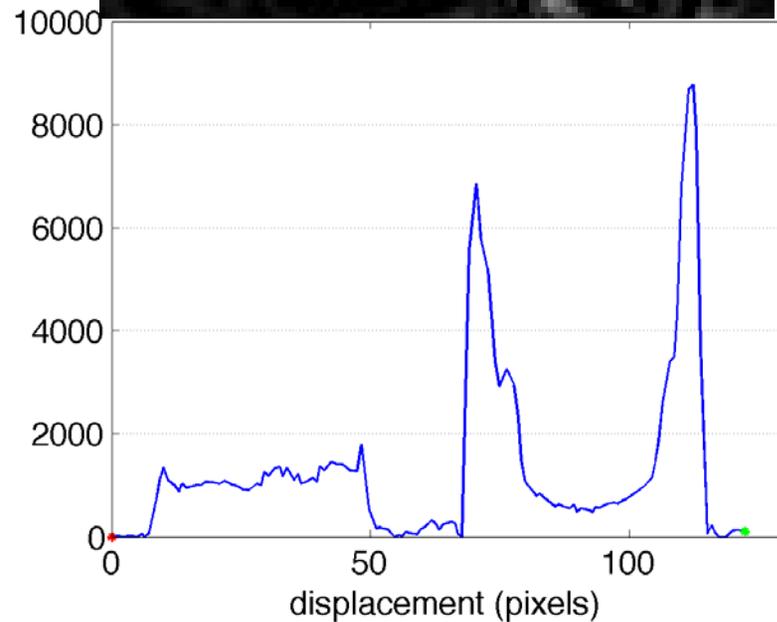
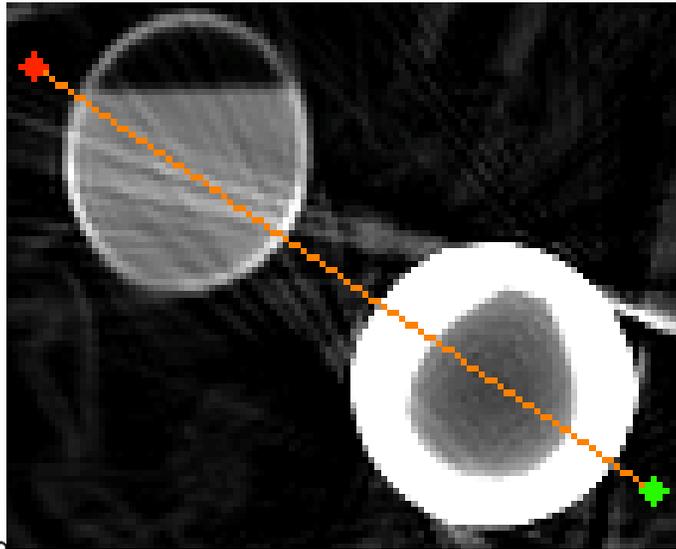
- PSV-ICD: Parallel super-voxel optimization for MBIR
- Unlocks hardware potential based on algorithmic/computer architecture co-design
- Increases processor efficiency by factor of ~20
- Allows for efficient parallelization on GPU and CPU architectures

Support

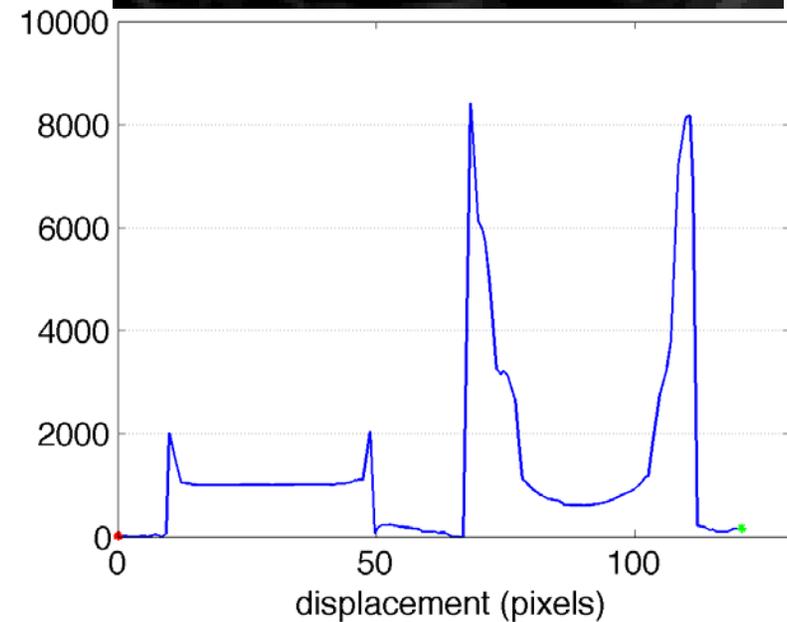
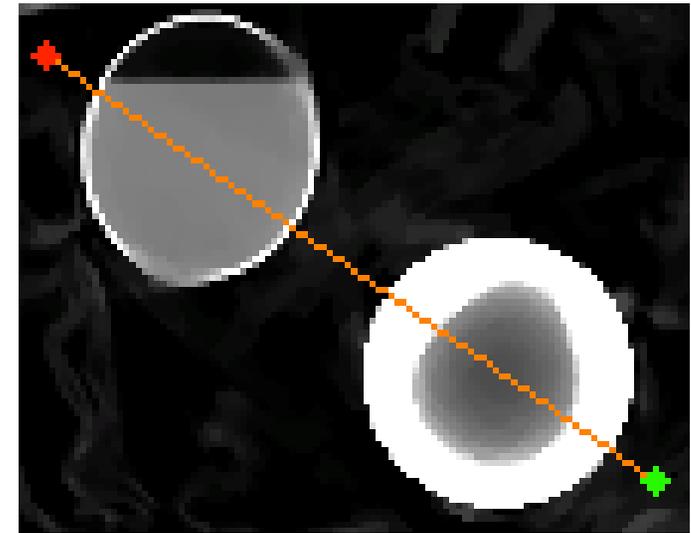
- ALERT
- DHS SBIR (HSHQDC-14-C-00058)

Results: artifact reduction

DFM



MBIR



S.J. Kisner, C.A. Bouman, et al., "Innovative Data Weighting for Iterative Reconstruction in a Helical CT Security Baggage Scanner", *The 47th IEEE Int'l Carnahan Conf on Security Technology (ICCST)*, Medellin, Colombia, 2013.

MBIR ALERT Transition Task

- In collaboration with Morpho Detection, we investigated the application of fully 3D MBIR in an EDS system for aviation security
- Demonstrated significant potential for,
 - Improved segmentation
 - Reduction of false alarms
 - Improved operator experience
 - Reduced cost of additional detection
- First study to evaluate IQ from iterative reconstruction in baggage screening

What is limiting MBIR reconstruction time?

- Rough example
 - Benchmark: 512 channels; 512 views; 400 slices; 512x512 resolution.
 - Processor: 20 core, 2.6 GHz Intel Xeon-E5 => 1.64 TFLOPS
 - Baseline reconstruction time on single core: **9035 sec**
 - Theoretical 3D reconstruction time: **1.65 sec** ← **Wow!**
- Conclusion:
 - Hardware is fast enough, but we need to use it more efficiently
- What are the limits to hardware efficiency?
 - Operations per second => Not really the problem
 - Memory access speed => Yeah, that's a problem
 - Parallelization => Need to keep all those cores busy

Efficient Computation of MBIR

- **Memory bandwidth**

- $(L1 \text{ cache speed}) = 10 * (L2 \text{ cache speed}) = 100 * (\text{Main memory speed})$
- Reduce cache misses
- Keep data local
- Increase **data reuse** ← **Very important**

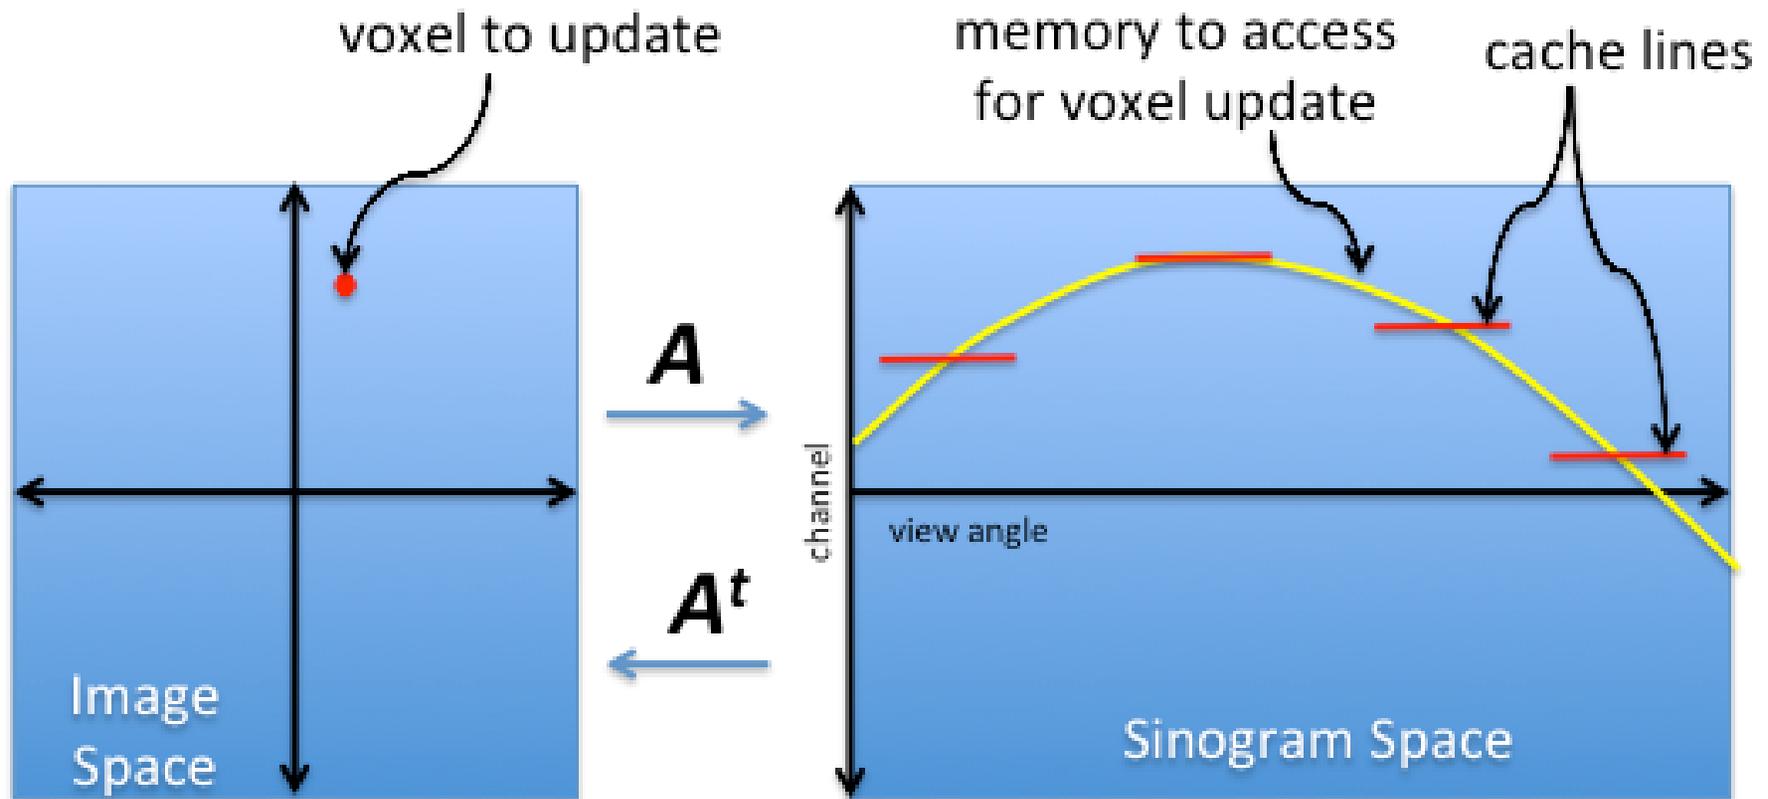
- **Parallelization**

- 1,000 to 4000 cores on a modern processor (or vector units)
- Must keep these busy
- Increase **independent operations** ← **Very important**

- How do we increase **data reuse** and **independent operations**?

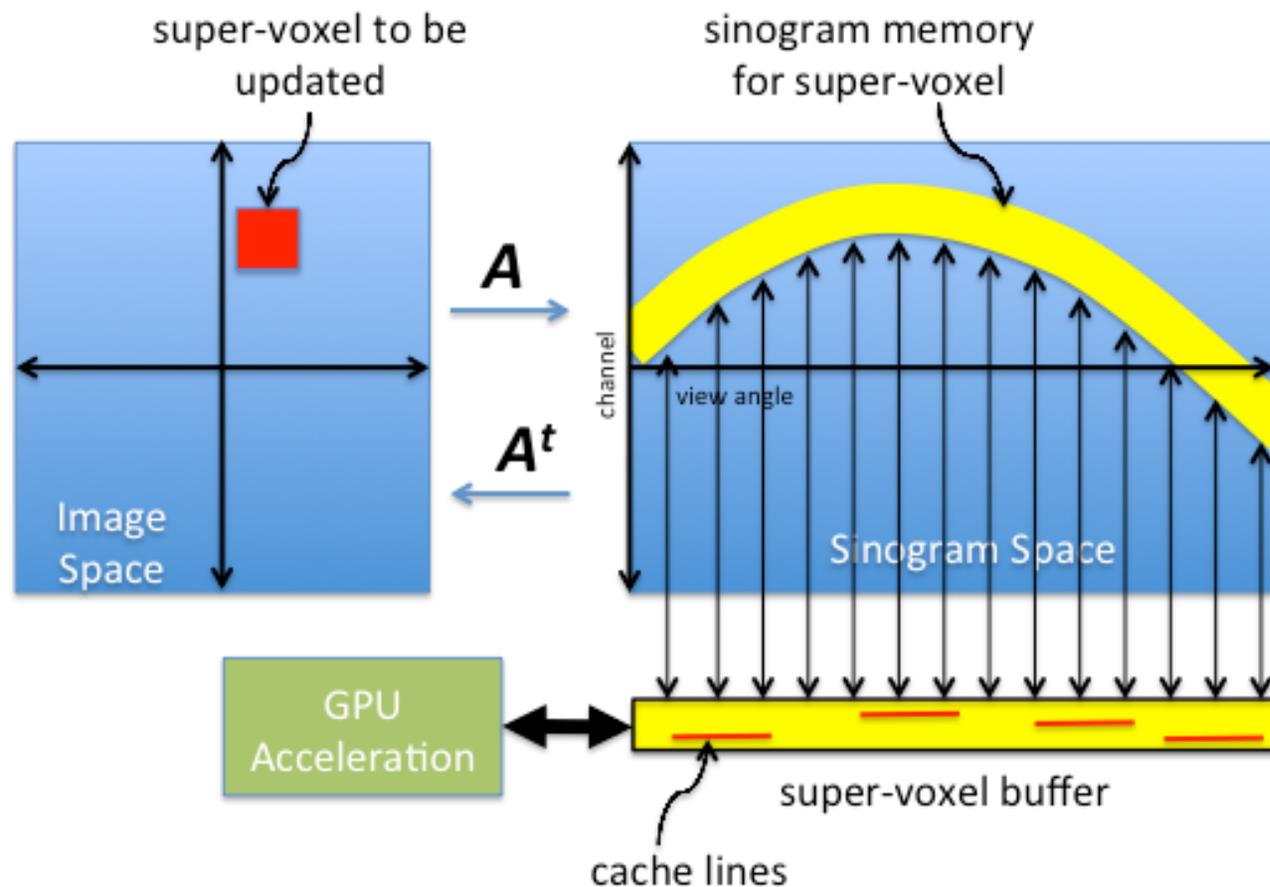
MBIR Cache Access Patterns

- Traditional ICD memory access patterns
 - Wasted cache
 - Little memory reuse



Super-Voxel-ICD Cache Access Patterns

- Super-voxel (SV): rectangular array of voxels to be updated
- Super-voxel buffer (SVB): cache buffer containing sinogram data
 - *No cache waste*
 - *Great deal of memory reuse*



But what about parallelization?

Parallelization Hierarchy

- **Intra-voxel parallelism:**
 - Parallelism within update of single voxel
- **Intra-SV parallelism:****
 - Parallelism across multiple voxels in an SV
- **Inter-SV parallelism:****
 - Parallelism across different SVs
- **Inter-slice parallelism:**
 - Parallelism across different slices of the 3D volume

Comments:

- Parallelisms are orthogonal
- Listed in fine to cores grain order

Reconstruction Performance Goals

- Benchmark 3D recon problem
 - 4 GPUs (<\$10,000)
 - Number of channels = 512
 - Number of views = 512
 - Number of slices = 400
 - Spatial Resolution = 512x512
 - Reconstruction time < 15 sec
- Equivalent benchmark Single Slice TO3 recon problem
 - 1 GPU
 - Number of channels = 1024
 - Number of views = 720
 - Number of slices = 1
 - Spatial Resolution = 512x512
 - Reconstruction time < 420 msec

$$\begin{aligned} \text{TO3 Recon time} &< 15 \text{ sec} * \frac{1 \text{ slice}}{400 \text{ slices}} * \frac{720 \text{ views}}{512 \text{ views}} * \frac{1024 \text{ channels}}{512 \text{ channels}} * \frac{4 \text{ GPUs}}{1 \text{ GPU}} \\ &= 420 \text{ msec} \end{aligned}$$

Multicore performance evaluation

Data

- Imatron-300 (ALERT TO3)
- 720 views, 1024 channels
- single slice recon, 512x512 image size
- ~3200 test slices

Multicore CPU Hardware

- two Intel Xeon-E5 2660 (2.6 GHz), each with 10 cores
- cache size per core: L1 32KB, L2 256 KB, L3 25 MB shared

GPU benchmark hardware

- Nvidia Tesla K40
- 15 streaming multiprocessors (SMs)
- For each SM, private L1 cache 64KB
- L2 cache 1.5MB shared across all SMs

Preliminary speed-up for multicore

ALERT Task Order #3 test data set

- 512x512 image size; 720 views; 1024 channels; 1 slice
- T_r = total reconstruction time
- N_F = no. of floating points operations (FLOP) per equit
- N_e = number of equivalent iterations
- O_F = Theoretical FLOPS of CPU/GPU
- E_F = Processing efficiency

Factor	Baseline ICD	SV-ICD(1)	PSV-ICD(4)	PSV-ICD(16)	PSV-ICD(20)
O_F (GFLOP)	83	83	332	1331	1664
N_F (GFLOP)	18.52	18.38	18.38	18.38	18.38
N_e (equits)	4.6	4.0	4.1	4.1	4.2
E_F (efficiency)	0.41%	5.95%	4.24%	3.54%	3.92%
T_r (sec)	253	15.0	5.28	1.69	1.27
T_{opt} (sec)	1.03	0.89	0.22	0.06	0.05
Speedup		16.9x	45.2x	150x	199x

*Intel Xeon-E5 2660 (2.6 GHz), 1-20 cores

Summary

- MBIR offers great potential in baggage screening applications and improvement in EDS performance
 - Improved image quality and resolution
 - Reduced artifacts
 - Increased design flexibility
- PSV-ICD provides a cost-effective solution for MBIR implementation
 - ~20x efficiency increase (memory reuse)
 - Linear parallelization (parallelization)
 - *More to come ...*