REAL TIME RECONSTRUCTION OF SYNCHROTRON TOMOGRAPHIC DATA USING GPUs

Kumar, D. *[1], Venkatakrishnan, S.V. †[2], Parkinson, D. ‡[1], and Marchesini, S. §[1]

[1]Lawrence Berkeley National Lab, Berkeley, CA, USA
[2]Oak Ridge National Lab, Oak Ridge, TN, USA

Keywords: x-ray tomography, microct, synchrotron, tomographic reconstruction, GPU

Summary: We'll discuss a GPU implementation of non-uniform fast Fourier transforms-based volume reconstruction of the synchrotron tomographic data. A Python interface was developed to manage the work-flow using either a GUI or CLI.

1. INTRODUCTION

Scientists are drawn to synchrotrons and accelerator based light sources for X-ray experiments because they provide unmatched brightness, coherence, and flux. Synchrotron-based micro- and nano-tomography play a vital role in deciphering the internal structure of modern materials, in particular in dynamic situations. Among many other applications, they have been used to study cracking and failure of ceramic matrix composites–new lightweight materials now being used in jet engines–under load at high temperature; flows of oil, brine, and carbon dioxide through rocks; and dendrites formation in batteries that lead to loss of capacity and failure. The rate of improvement in brightness and detector technology has outpaced Moore's law growth seen for computers, networks, and storage, and is enabling novel observations and discoveries with faster frame rates, larger fields of view, higher resolution, and higher dimensionality. Most synchrotron microCT facilities now have cameras that can acquire 12bit images with 2000x2000 pixel resolution at over 2000 fps (producing $>$ 300 GB/hour), and some facilities have recently deployed cameras that can go orders of magnitude faster. For many dynamic experiments, it is critical to "see" what is happening in 3D in real time, so that the experiment can be controlled.

2. RECONSTRUCTION ALGORITHM

Radon and inverse Radon transforms constitute the computational crux of many tomographic reconstruction algorithms. The 2D Radon transform of an image can be implemented in a number of ways. One of the most efficient ways to perform Radon is to first perform a 2D FFT on a regular grid, and then interpolate the transform onto a polar grid before 1D inverse Fourier transforms are applied to interpolated points along the same radial lines [1]. The interpolation between the Cartesian and the polar grid is the key step in this procedure. It can be carried out using a "gridding" algorithm that maintains the desired accuracy with low computational complexity. The gridding operation requires the convolution between irregular samples and a kernel calculated at regular sample position and vice versa. The advantage of this approach is that it scales well with a large volume of data. Our implementation takes advantage of open source libraries, such as Arrayfire [2], Afnumpy [3] to minimize data exchanges between host and the device over the slow PCIe. This is a major bottleneck in state of the art GPU based open-source solutions available at the brightest light sources around the world, e.g. ASTRA. Most of the code in our work is written in python using Afnumpy, which is a Numpy like wrapper for Arrayfire. On a single machine with 4 K80s, we were able to achieve reconstruction of a 2560 X 2560 X 2000 voxel volume in about 3.5 minutes. We are working towards extending the capability of the code to perform Model Based Iterative Reconstructions (MBIR). A CPU only version of MBIR can take 1000s

*dkumar@lbl.gov
†venkatakrisv@ornl.gov
‡dyparkinson@lbl.gov
§smarchesini@lbl.gov

of CPU hours. Because of the iterative nature of MBIR, the speedup provided by our algorithm is especially important because it is used repeatedly over each iteration.

3. RESULTS

We are able to achieve considerable faster reconstruction time with GPU-NUFFT. It takes about 192.7s on a single GPU, to run 15 iteration of SIRT algorithm on a single Nvidia K80 fig. 1. The turn around times for NUFFT based MBIR are expected to be below 10 minutes for a 2560 X 2560 X 2000 data, which is typical for 8.3.2 (micro tomography) beamline at the Advanced Light Source.

References

[1] Samuel Matej, Jeffrey A. Fessler,and Ivan G. Kazantsev, Iterative Tomographic Image Reconstruction Using Fourier-Based Forward and Back-Projectors, *IEEE Transactions on Medical Imaging*, vol. 23, no. 4, pp. 401-412, April 2004. doi: 10.1109/TMI.2004.824233

[2] Yalamanchili, P., Arshad, U., Mohammed, Z., Garigipati, P., Entschev, P., Kloppenborg, B., Malcolm, J. and Melonakos, J. (2015). ArrayFire - A high performance software library for parallel computing with an easy-to-use API. Atlanta: AccelerEyes. Retrieved from *https://github.com/arrayfire/arrayfire*

[3] F. Maia, https://github.com/FilipeMaia/afnumpy

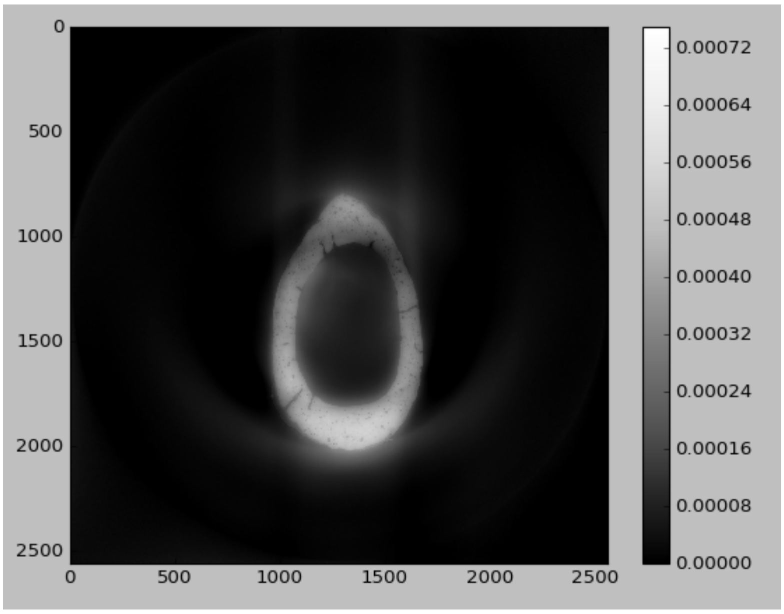

Figure 1: Single slice of a bone sample reconstructed using NUFFT