

Submission to Geoffrey C. Fox

Proposal for a chapter in the book to be published by Morgan Kaufman, 2000.

Exact title and publication information will be submitted later.

Date 20 Dec, 1999

Version 1.1

Formatting: This document has been kept on purpose in a very simple format, as to allow the editor to choose modifications for the proper page style. Since I am not familiar with Word, the placement of the figures is fixed. Reformatting the text may lead to a wrong page layout. Please, correct if necessary.

A **cross-reference** to the chapter about computational Grids is made in the text. (cf. Chapter ..).

All pictures can be obtained from the author in encapsulated postscript or PowerPoint format. If necessary, please, request them from gregor@mcs.anl.gov.

Corresponding Author:

Gregor von Laszewski
9700 S. Cass Ave
Argonne, IL 60439
gregor@mcs.anl.gov

o: (630) 252 0472

h: (630) 378 0837

During the time from 12/23/99-12/27/99 I can be reached under the number (716) 225 5278.

Please, do not hesitate to contact me if I should reformat or improve the chapter in any way.

Gregor

Quasi Real-Time Microtomography Experiments at Photon Sources

Gregor von Laszewski, Mei-Hui Su, Joe Insley, Ian Foster, Carl Kesselman

1. Introduction

Computed microtomography (CMT) is a powerful tool for obtaining nondestructively a three-dimensional view of the internal structure of opaque objects. In contrast to the widespread use of this technique in the millimeter scale as part of diagnostic procedures in hospitals, we are interested in the investigation of objects on the micrometer scale.

An application of this method is, for example, the quality control processes during the production of three-dimensional semiconductor wafers. Being able to visualize the details of chip wafers in all three dimensions allows engineers to improve the chip design before production. Other examples can be found in the field of earth science, where common tasks include investigation of the interior of a very small meteorite and study of the enclosures of very tiny materials in opaque diamonds formed 100,000 years ago, in order to determine more about the origin and development of the earth.

The energy and the infrastructure necessary to conduct such experiments can be provided by using x-ray beams at synchrotrons. The use of x-rays for investigating the internal structure of materials at the micron scale has grown rapidly over the past decade as a result of the availability of synchrotron radiation sources. One such facility is the Advanced Photon Source (APS) at Argonne National Laboratory.

A typical computed microtomography experiment at the APS proceeds as follows. A sample is mounted in the experiment station, parameters are adjusted, and the sample is illuminated by a collimated beam of x-rays. Data is collected for multiple sample orientations by using a charge-coupled device. A time-consuming reconstruction process is then used to obtain a three-dimensional representation of the raw data with spatial resolution of as little as 1

μm . The 3-D image contains quantitative information about the x-ray attenuation coefficient at a particular x-ray energy.

The many orders of magnitude increase in brilliance now available at third-generation sources such as the APS allows dramatic improvements in temporal resolution and makes it feasible to record fully 3-D, high-resolution tomographic data on time scales of less than a second per image.

Nevertheless, a major difficulty with the current practice is the turnaround time between the data acquisition and the reconstruction, often due to lack of available computing power. This is especially problematic for all synchrotron-based experiments because only a limited amount of beam time is available for a user. The use of distributed supercomputing power can reduce this turnaround time to a few hours or minutes, enabling the users to view the results in quasi-real time and to alter experiment conditions on the fly. This capability can greatly improve the usefulness of a synchrotron radiation facility.

2. The Computational Processing Pipeline Framework

We have provided a general framework for CMT applications that is based on the concept of a processing pipeline. The pipeline consists of a series of components that communicate with each other via input and output channels. Each of the components can be mapped, in principle, onto different computational resources. Thus the framework ideally can be mapped onto computational Grids (cf. chapter about Grids). The CMT pipeline has the following additional properties:

Data format: HDF is used to guarantee portability of the data across a variety of diverse compute platforms. In addition, it provides the ability for self-describing data, which will enable the organization of large subject-related data archives in near future.

Data interchange: The Globus Nexus communication library is used to allow the support of multiple protocols as part of the message exchange mechanism. The protocol is selected based on its availability and performance characteristics between the computational processing units connected via the channels.

Preprocessing algorithm: To improve the quality of the reconstruction, images should be preprocessed with appropriate filters. The preprocessing algorithm usually varies from experiment to experiment. A set of predefined standard preprocessing algorithms is available that can be easily used without recompiling the code. The user can extend the available preprocessors.

Reconstruction algorithm: Currently, we use a high-performance parallel implementation of reconstruction algorithms for microtomography datasets, based on a filtered backprojection technique.

Interleaved reconstruction and visualization: Resulting images are shipped in real time to a

visualization unit to gradually update a 3-D rendered image during the experiment. This gradual update is important to allow decisions as to whether the experiment should continue. If the experiment does not perform as expected, it is terminated.

Integrated visualization and collaboration engine: One goal of this project is to enable researchers to participate in an experiment from their home institution rather than to travel to the APS. A beamline scientist will be able to handle the experiment locally while communicating with the scientist designing the experiment. A remote video conferencing tool allows such collaboration. In addition, the 3-D image analysis tool contains a control component that enables collaborative visualization on a variety of output devices including graphics workstations, ImmersaDesks (Figure 1), and CAVEs. This allows the resulting 3-D image to be rendered rapidly. Moreover, it provides a shared control among the users participating in a collaborative session, enabling computational steering of the experiment. This general visualization framework is currently reused also by scientists from different disciplines with similar visualization requirements (e.g., electron microscopy and astrophysics).




Figure 1: A screen shot of the ImmersaDesk taken during a collaborative session with

two users. The control panel allows the user to modify parameters such as the color tables used to control the position, the size, and the orientation of the object. The inner markup of the object can be analyzed further with the help of cross-sections and transparent masking of uninteresting features.

3. Scientific Challenges

The computational framework based on a processing pipeline presents several scientific and computational challenges.

I Reconstruction

As noted, we currently use a reconstruction algorithm based on filtered backprojection. This algorithm parallelizes nicely in that each slice in a dataset can be processed independently. Hence, the principal challenge is to develop efficient techniques for moving data between detector, secondary storage, supercomputers, and workstations for visualizing results.

3.2 New Acquisition and Reconstruction Strategies

Access to a large amount of computational power allows the use of new acquisition and reconstruction strategies. Traditionally, data is collected at microtomographic beamlines at constant angle offsets: for example, 0, 1, 2, ... degrees if 360 samples are to be taken. In an interactive environment such as we describe here, it becomes attractive instead to use an interleaved angle list. For example, we may first gather images at 60-degree offsets (0, 60, 120, 180, 240, 300), then collect additional images to provide a 30-degree sampling, and so on until a complete 1-degree dataset is obtained. The advantage of this strategy is that the reconstruction algorithm can be run repeatedly, once for each more detailed set of data; hence, the scientist obtains a series of more refined images and may be able to detect a flawed experimental setup early in the data collection process.

Another interesting direction that is enabled by the availability of supercomputer resources is the following. In principle, reconstruction quality can be improved by performing multiple reconstructions with different algorithms and parameter settings. We are hopeful that the enhanced compute power made accessible by grid environments will initiate a new area in the development of reconstruction algorithms for computed microtomography and other disciplines.

3.3 Computational Requirements

The data rates and compute power required to address a CMT problem are prodigious, easily reaching one gigabit per second and one teraflop per second. We illustrate this statement with a scenario. A 3-D raw dataset generated by a typical detector will comprise 1000 1024×1500 two-byte slices (3 GB); detectors with significantly higher resolutions will soon be available. If we assume current reconstruction techniques and make fairly optimistic scaling assumptions,

reconstruction of this dataset requires about 10^{13} floating-point operations (10 Tflops). On a 100 Mflop/sec workstation, this translates to 32 hours; on a 1 Tflop/sec computer, it would take 10 seconds. With current detector technologies, this dataset might take 1500 seconds to acquire; however, new detectors will improve readout times considerably. Besides the computational demand resulting from the reconstruction of the 3-D object, the display of the rendered final result is also a problem with the current state-of-the-art imaging hardware. The size of the datasets generated in these experiments can be quite large, typically on the order of 1024^3 floating-point values. Currently, even the accelerated graphics hardware used by our application has trouble keeping up with volumes of this size, and the dataset needs to be subsampled down to 256^3 floating-point values in order to maintain its frame rates for interactive usage.

4. Benefits of Real-Time X-Ray Microtomography Experiments

The framework described in this chapter offers several benefits to the end user. First, a fast reconstruction algorithm can be used to help decide whether the current experiment has to be interrupted prematurely because of an error in the setup. This will allow for an increase in the number of experiments to be conducted per hour. In order to handle the complicated and diverse supercomputing environments, it is essential to provide a simple interface giving the beamline experimentalist control over the parameter set, as well as the possibility to terminate the current calculation at any time.

Besides the requirements driven by the computational aspect of the application, organizational aspects benefit from the framework described in this chapter. Because of the hazardous and often unpleasant environment at the beamline, remote operation is desirable. With remote operation, the facility can maintain a small but well-trained team of beamline staff experimentalists. This approach offers several advantages. It reduces the operational and user-specific cost and minimizes travel cost to the unique facility. Furthermore, it increases the access time to the beamline while minimizing the effort required by trained experts to set up experiments. With the availability of a collaborative and remote steering environment, new user communities in commercial and educational facilities are likely to use the supercomputing-enhanced light sources in remote fashion. During an experiment multiple users using different visualization engines at geographically dispersed locations should be able to collaborate easily with each other. The details of this infrastructure will be hidden from the end users, the microtomography scientists. For these users it is irrelevant where and how the result is achieved, as long as time and computational accuracy requirements are met. Figure 2 shows such a grid-enabled collaborative application.




Figure 2: A The “grid-enabled” CMT application allows researchers to display the same state of the visualized object on all display stations participating in a collaborative session. Remote computation and steering become possible across multiple access points.

5. Summary

In this chapter we have described a grid-enabled real-time analysis, visualization, and steering environment for microtomography experiments. Specifically, we have provided a portable parallel framework that allows different reconstruction algorithms to be ported on the Grid. A standard data format based on HDF is defined to distribute the data among scientists in a meaningful way.

The real-time visualization environment developed fulfills the basic needs of the microtomography scientists. Moreover, with the availability of this environment, we anticipate that scientists will make algorithm improvements, for example, including a priori knowledge of a previous reconstruction in order to increase the quality of the image.

In the future, we will focus on the use of new modalities in real-time reconstruction for interactive use and will explore the collaborative analysis of results. In addition, we will emphasize improvements to the usage of dynamic scheduling of computers, high-speed

networking, and collaboration technologies.

Acknowledgments

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Computational and Technology Research, U.S. Department of Energy, under Contract W-31-109-Eng-38. Globus research and development is supported by DARPA, DOE, and NSF.

We based the grid enabled version of the code from the code provided by Mark Ellisman and Steve Young from the National Center for Microscopy and Imaging Research at San Diego Supercomputing Center. We like to thank Derrick C. Mancini, Steve Wang, Ian McNulty, Mark Rivers, and Francesco DeCarlo for conducting the experiments at the APS beamline. This truly interdisciplinary project would not have been possible without the countless hours of work performed by the Globus team: John Bresnahan, Joe Bester, Steve Fitzgerald, Brian Toonen, Steve Tuecke, Karl Czajkowski, and all the others who have helped us throughout the years.

References

- (Foster1998) I. Foster and C. Kesselman, *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufman, 1998.
- (Foster1999) I. Foster, J. Insley, G. von Laszewski, C. Kesselman, M. Thiebaux, Distance Visualization: Data Exploration on the Grid, *IEEE Computer* **32**, No.12, pp. 36-43. Dec. 1999.
- (Haddad1994) W. S. Haddad, I. McNulty, J. E. Trebes, E. H. Anderson, R. A. Levesque, and L. Yang, "Ultra high resolution x-ray tomography," *Science* **266**, pp. 1213--1215, 1994.
- (HDF) NCSA HDF homepage, <http://hdf.ncsa.uiuc.edu/>, 1999
- (Klosowski1998), P. Klosowski, M. Koennecke, J. Z. Tischler, and R. Osborn, "Nexus: A common format for the exchange of neutron and synchrotron data," *Physica* **B241-243**, pp. 151--153, 1998.
- (Mercurio1992) P. J. Mercurio, T. T. Elvins, S.J. Young, P. S. Cohen, K.R. Fall, M. H. Ellisman, "The Distributed Laboratory: An Interactive Visualization Environment for Electron Microscopy and Three-Dimensional Imaging", *Comm. Assoc. Comp. Mach.*, Vol. 35, No. 6, pp. 54-63., 1992.
- (Wang 1999) Y. Wang, F. De Carlo I. Foster, J. Insley, C. Kesselman, P. Lane, G. von Laszewski, D. Mancini, I. McNulty, M.-H. Su, B. Tieman, "A Quasi-Realtime X-Ray Microtomography System at the Advanced Photon Source", *Proceedings of SPIE99*, Vol.

3772, pp.301-309, 1999.

(vonLaszewski1999) G. von Laszewski, M.-H. Su, J. A. Insley, I. Foster, J. Bresnahan, C. Kesselman, M. Thiebaut, M. L. Rivers, S. Wang, B. Tieman, I. McNulty, Real-Time Analysis, Visualization, and Steering of Microtomography Experiments at Photon Sources, *Ninth SIAM Conference on Parallel Processing for Scientific Computing*, Apr. 1999. Also: <http://www.mcs.anl.gov/xray>