

A high-throughput x-ray microtomography system at the Advanced Photon Source

Yuxin Wang,^{a)} Francesco De Carlo, Derrick C. Mancini, Ian McNulty, and Brian Tieman
Advanced Photon Source, Argonne National Laboratory, Argonne, Illinois

John Bresnahan, Ian Foster, Joseph Insley, Peter Lane, and Gregor von Laszewski
Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Illinois

Carl Kesselman, Mei-Hui Su, and Marcus Thiebaux
Information Sciences Institute, University of Southern California, California

(Received 14 November 2000; accepted for publication 23 January 2001)

A third-generation synchrotron radiation source provides enough brilliance to acquire complete tomographic data sets at 100 nm or better resolution in a few minutes. To take advantage of such high-brilliance sources at the Advanced Photon Source, we have constructed a pipelined data acquisition and reconstruction system that combines a fast detector system, high-speed data networks, and massively parallel computers to rapidly acquire the projection data and perform the reconstruction and rendering calculations. With the current setup, a data set can be obtained and reconstructed in tens of minutes. A specialized visualization computer makes rendered three-dimensional (3D) images available to the beamline users minutes after the data acquisition is completed. This system is capable of examining a large number of samples at sub- μm 3D resolution or studying the full 3D structure of a dynamically evolving sample on a 10 min temporal scale. In the near future, we expect to increase the spatial resolution to below 100 nm by using zone-plate x-ray focusing optics and to improve the time resolution by the use of a broadband x-ray monochromator and a faster detector system. © 2001 American Institute of Physics.
[DOI: 10.1063/1.1355270]

I. INTRODUCTION

X-ray radiation has been widely used for imaging since its discovery. Its large penetration depth makes it well suited for three-dimensional (3D) imaging by tomography, particularly when the internal structure of an object is difficult to study by other means, such as sectioning. As an example, medical computed tomography (CT) systems are used to image the 3D structure of a patient at mm-scale resolution. In the past decade, the resolution of x-ray tomography has been pushed to the 100-nm level with the emergence of second- and third-generation synchrotron radiation sources. One such system typically uses 5–20 keV hard x rays in a straight-ray projection imaging scheme with high-resolution detectors. With this system, samples of a few mm thickness can be imaged in 3D at 1 μm resolution.^{1–3} Another type of system uses soft x rays with a few hundred to a few thousand eV energy and a zone-plate optical element for magnification. With this system, small samples a few micrometers in size can be imaged at 50–100 nm resolution.^{4–6} If photon statistics were the only concern, the third-generation synchrotron radiation facilities would be able to provide high enough brilliance to produce high-resolution images with signal-to-noise ratio of 5:1 at close to video frame rates. In principle, 1 μm resolution images can be acquired at a bending magnet beamline at the rate of 1–10 images per second while 100-nm image resolution can be achieved at an undulator

beamline at a similar rate. At this speed, a complete tomographic data set can be acquired in less than a few minutes if the data acquisition system can sustain the same speed.

To take full advantage of high-speed acquisition, one requires comparably rapid reconstruction and visualization calculations. A typical data set contains 800 images, each with 1024 \times 1024 pixels, at 16-bit dynamic range, adding up to about 1.6 GByte. A high-performance data transfer and computer system is therefore needed to transfer, filter and reconstruct these data at a rate similar to that at which they are acquired. This capability to reconstruct the data rapidly and provide a beamline user with immediate feedback is extremely valuable for user-shared facilities, since each user receives a limited number of days of operation annually. In most current synchrotron-based x-ray tomography systems, a user typically performs the processing and reconstruction calculations after the experimental period is finished. The inability to examine reconstruction from newly acquired data makes it difficult to make crucial decisions during the experimental period, resulting in poor use of beam time. (A common consequence is that most experimenters accumulate a large number of unprocessed data sets.) Our inline reconstruction system therefore presents a significant advantage over existing systems by providing the users with rendered reconstruction results for analysis immediately after the data are acquired.

^{a)}Present address: Xradia, Inc., 4075A Sprig Dr., Concord, CA 94520.




FIG. 1. Illustration of the straight-ray projection scheme used at the APS 2-BM beamline to acquire the tomographic projections. The x-ray beam is monochromatized by a Kohzu double-crystal monochromator. The monochromatic beam passes through the sample to reach the YAG:Ce scintillation screen. The visible light emitted by the scintillator is imaged to the CCD detector by a microscope objective. Both the sample and the detector are mounted on an x - y stage for alignment to the x-ray beam. The sample is fixed on a miniature x - z stage mounted on the rotation stage so that the region of interest can be centered on the rotation axis. The scintillation screen is mounted on a z stage for focusing.

We have constructed a quasireal time x-ray tomography system at the 2-BM (bending magnet) beamline at the Advanced Photon Source (APS) at Argonne National Laboratory. It integrates a fast detector system, high-speed data networks, a massively parallel computer for reconstruction calculations, and a hardware-optimized visualization computer in a pipelined fashion. The imaging system uses 5–20 keV x rays to record projections in either amplitude or phase contrast mode with a straight-ray projection scheme. Depending on the imaging resolution and sample contrast, this system can collect a complete tomographic data set within tens of minutes to a few hours. Immediately after each projection is obtained, it is rapidly transferred to the Mathematics and Computing Sciences (MCS) Division (also located at Argonne National Laboratory) where the preprocessing and reconstruction calculations are performed by the massively parallel computer concurrently with the data acquisition. The reconstruction results, once completed, are transferred to the visualization computer that performs the hardware rendering calculations. A user at the beamline can view the rendered 3D images interactively minutes after the data acquisition is complete. This system also gives users the option to acquire the tomographic data set in several cycles, initially with coarse then with fine angular steps. In effect, the user observes the reconstructed images gradually being refined with the completion of each cycle.

II. IMAGING SYSTEM

The energy range of the x-ray beam at 2-BM is selectable between 5 and 20 keV by the use of a Kohzu double-crystal monochromator.^{3,7} This monochromator has an energy resolving power of over 7000. The imaging system is illustrated in Fig. 1. The sample is placed in the x-ray beam such that the transmitted beam reaches the scintillator screen. We have used two types of scintillators in our experiments: a 0.5 mm thick CdWO_4 crystal and a 0.5 mm thick yttrium-aluminum-garnet (YAG) crystal with a 5 μm layer of

YAG:Ce. The YAG crystal provides relatively higher resolution and contrast at the cost of efficiency, especially at higher energies.

The visible light emitted by the scintillator is imaged to a charge coupled device (CCD) camera by a microscope objective lens. The camera (Princeton Instruments) uses a back-illuminated CCD chip cooled by liquid nitrogen. The CCD has 1024×1024 pixels with a 24 μm pixel size and 25 mm chip size. A series of Nikon plan achromat objectives with magnifications of 4, 10, 20, and 40 are used for different field-of-view and resolution requirements. The field of view and the resolution obtainable with each lens are determined by the ratio of the CCD chip size and pixel size to the magnification, except with the $40\times$ objective, the resolution of which was measured to be 0.8 μm . Figure 2 shows a microfabricated test pattern imaged with the $40\times$ objective. Figure 3 shows a cricket imaged with the $6\times$ objective and a fruit




FIG. 2. Microfabricated test pattern consisting of two-dimensional crossed gold lines with 2 μm width and 2 μm spacing imaged at 5 keV with a $40\times$ objective. The resolution measured from this image was approximately 0.8 μm .




FIG. 3. Projections of biological samples recorded at 5 keV. (a) A cricket imaged at 0° rotation angle. (b) The cricket imaged at 90° . (c) A fruit fly imaged at 0° . (d) The fruit fly imaged at 90° . The cricket was imaged with a $4\times$ objective, which gave a field of approximately 6.25 mm and resolution of $6\ \mu\text{m}$. The fruit fly was imaged with a $20\times$ objective, which gave a field of approximately 1.25 mm and resolution of $1.2\ \mu\text{m}$.

fly imaged with the $20\times$ objective, both at 5 keV. With the combination of this CCD camera and the YAG scintillator, typical integration times range from 1 s with the $6\times$ and 10 s with the $40\times$ objective (see Table I). The readout and transfer time is over 4 s for a 1024×1024 image and approximately 1 s for a 512×512 image, when 2×2 binning is used. At this rate, a data set containing 720 projections, each with 1024×1024 pixels, requires over 1 h to acquire while one containing 360 projections, each of 512×512 size, requires about 10 min. Several “black-field” and “white-field” images are also acquired along with the projection data. A black-field image is a background measurement acquired without x-ray beam illumination, and a white-field image is one acquired with illumination but with the sample removed from the imaging field. These images are essential for normalizing the projection images (see Sec. V). In practice, we find one black and white field for every 20 projections is sufficient.

The data acquisition is controlled by a graphical user

TABLE I. Properties of the detector system for different objectives and the acquisition time for 1024×1024 sized images. The time per data set includes the readout and transfer time.

Objective (magnification)	Numerical aperture	Imaging field (mm)	Effective resolution (μm)	Time per image (s)	Time/data set (h)
4	0.1	6.25	6	1–5	1.5
10	0.25	2.5	2.4	2–10	2
20	0.4	1.25	1.2	5–20	3
40	0.6	0.42	0.8	10–50	6




FIG. 4. Group structure of the tomographic data files. The CMT_ImageType field contains an index of the data files containing black/white-field images, projections, sinograms, and the reconstructions. The information on the experimental conditions, users, sample, and computation parameters is stored in subsequent fields.

interface (GUI) program running on a Sun workstation that allows the user to adjust data-acquisition parameters, such as the integration time, file locations, and computation configurations, etc. This program can be executed either locally on a computer at the beamline or remotely via an x-window interface. It also allows the projections to be acquired in multiple cycles, starting with coarse angular intervals then in increasingly finer steps. The reconstruction results from the coarse cycles are usually available for viewing at the beamline while the next cycle is being acquired, giving the user an opportunity to change acquisition or calculation parameters or to abort the acquisition process. The series of black- and white-field images can also be acquired in various combinations with the projections. After each projection is acquired, the control program launches a background process to transfer the data to a parallel computer at MCS.

III. DATA FILE STRUCTURE

We developed a file structure for the tomographic data so that they can be easily organized and interfaced to the calculation program. A data set is divided into a header file and a series of data files. The header file is generated by the GUI control program at the beginning of data acquisition and functions as a master index of the individual data files. It contains the experimental conditions, sample information, parameters for the preprocessing and reconstruction calculations, and the names of the data files. The data files themselves contain the actual projection and black-field and white-field data. These files are written in the Hierarchical Data Format (HDF) created by the National Center for Supercomputing Applications.⁸ The HDF format is platform in-




FIG. 5. Illustration of the data acquisition and computation pipeline. An operator at the experiment station at the APS controls the data acquisition system on a Sun workstation and views the reconstruction result on the feedback monitor. The data acquired at the beamline are first written to a local hard disk and then transferred via a 100 Mb/s network to the Mathematics and Computer Science Division, where the preprocessing and reconstruction calculations are performed by the parallel computer. The reconstruction results are saved to a hard disk then loaded into the Viz Engine. The Viz Engine provides rendered images interactively to the beamline feedback monitor via the high-speed network and also periodically writes rendered still images to a web server.

dependent, self-describing, and supports nearly all data types that are likely to be used in tomography experiments. Furthermore, we use the hierarchical structure to organize a large number of data fields in each data set, very much like a file directory structure. The structure of the HDF header file is shown in Fig. 4. This structure is an extension of the NeXus convention developed for the exchange of data from neutron- and synchrotron-based experiments.^{9,10} An application programming interface (API) written in C++ was implemented to provide an interface between the data and the reconstruction program. The API has C++ class definitions that correspond to the HDF data structures.

IV. COMPUTATION PIPELINE

The preprocessing and reconstruction computing structure is organized in a pipeline fashion, as shown in Fig. 5. The data files acquired at the beamline are transferred from the beamlines local storage to the parallel computer's local disks via a 100-Mb/s network. The massively parallel computer used in our experiments is a Silicon Graphics Inc. (SGI) Origin 2000 with 96 available nodes; typically, only 80 nodes are used in our calculations. It utilizes a grid-enabled computing environment with the Globus metacomputing toolkit.¹¹ Analogous to electric power grids, the Globus framework dynamically handles resource allocation, resource configuration, program submission, input/output

(I/O) operation, program scheduling, file staging, and uniform program executions within a system of processors or work stations.¹² This computing environment also allows many remote computers around the world within the framework to be utilized for our calculations.

After each new file arrives at MCS, the preprocessing routines running on the parallel computer first perform filtering and alignment calculations on the data (details are described in Sec. V). The filtered and aligned data are then reassembled into a sinogram format for tomographic reconstruction. The initial parameters and specification for the calculations are contained in the HDF header files. The reconstruction program generates volume slices that can be stacked to form a reconstructed volume. They are first written to a local disk and later to a terabyte storage robot for long-term archival. After each cycle of reconstruction is complete, a visualization program running on another SGI computer is activated to load the newly reconstructed data and perform rendering calculations. This computer, called the Viz Engine, uses a hardware-optimized volume-rendering software called Volumizer from SGI to produce 3D renderings at a video rate.¹³ The 3D structure of the reconstruction can be viewed in a virtual-reality environment with stereoscopic goggles. Currently, because of limitations of the graphics hardware, in order to maintain an interactive frame rate the size of the rendered volume is limited to 256^3 . Typically, most of our reconstructions are subsampled to this size before they are loaded into the Viz Engine. Alternatively, we could observe the original reconstructed slices or load a 256^3 subregion of the data at full resolution, enabling us to observe the finer features. In addition to the virtual-reality environment at MCS, a toolkit running on beamline workstations can also receive the live output of the Viz Engine and allows a beamline user to control the rendering and viewing parameter interactively.

The Viz Engine also periodically writes still images of the rendered data to a web server so that its output can be viewed worldwide with a web browser by collaborators. This feature is valuable for experiments conducted by a large collaboration, when some members cannot be physically present at the experiment station.

V. PREPROCESSING AND RECONSTRUCTION

The raw projection data typically contain many artifacts resulting from beam nonuniformity and defects in the scintillator, lens, and the detector. They can be effectively corrected by use of the black- and white-field images with the following method:

$$f = \ln \frac{f_w - f_b}{f_0 - f_b},$$

where f is the filtered image, f_0 is the raw projection, and the f_w and f_b are the white- and black-field images, respectively. This filtering process gives the line integral of the absorption through the sample along the beam direction. The collection of such filtered projections $f(\theta, x, y)$ is then the radon transform of the sample's three-dimensional absorption map $\rho(x, y, z)$. Some phase information is usually present in the

TABLE II. Reconstruction time using the filtered backprojection method for various size data sets. In each case, 80 of the 128 processors of the parallel computer were used.

No. of projections	Projection size (pixels)	Calculation time (min)
721	1024×1024	17
721	512×512	2
361	512×512	1

images but is minimized by reducing the distance between the sample and the scintillator screen (less than 5 mm in our case). We can therefore ignore the phase information in our calculations without introducing observable artifacts.

An outlier filter is sometimes used to remove isolated high-intensity points resulting from pixel defects in the CCD chip or from cosmic rays. It calculates the local median and standard deviation for each pixel in the image, and replaces the pixel by the median if the pixel value is more than a certain number of standard deviations away from the median. A least-square deconvolution (Wiener) filter has also been implemented to restore the images degraded by the optical system.

Before the projections are used in reconstruction calculations, they must be aligned to one another so that the rotation axis is located at the center of the images. We have learned that our rotation stage bearing typically has a radial runout error on the order of 1 μm . Therefore, for images with 1–2 μm resolution, each projection must be aligned individually. For the projections at 6 μm or lower resolution, however, the stage error can be ignored, and we only need to determine the location of the rotation axis in one then collectively shift them all by the same amount. A cross-correlation function is used to identify the rotation axis. We reverse the projection acquired at a 180° angle and compute the cross-correlation function with the 0° image:

$$C_{fg}(x', y') = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f^*(x-x', y-y')g(x, y) dx dy,$$

where f represents the 0° image and g represents the reversed 180° image. Ideally, f and g are the same image but shifted from each other, and the rotation center is located halfway between the shifted features. The peak of the correlation function indicates the amount of shift between the two images and therefore how much each image is to be shifted. In practice, the cross-correlation function is calculated using the property,

$$\mathcal{F}\{C_{fg}\} = \mathcal{F}^*\{f\}\mathcal{F}\{g\},$$

where \mathcal{F} indicates a Fourier transform. For the higher resolution images, such cross-correlation calculations must be performed on each projection to correct both the rotation stage errors and the shift from the rotation center. In the special cases in which the object is of high contrast and completely located inside the imaging field, the image centroid can be computed and can serve as the alignment mark for centering.

The reconstruction programs used in our experiments are based on a code provided by Ellisman and Young of the

National Center for Microscopy and Imaging Research at the San Diego Supercomputing Center.¹⁴ Three commonly used algorithms, filtered backprojection, the algebraic reconstruction technique (ART), and the simultaneous iterative reconstruction technique (SIRT), have been implemented. However, because of the high angular sampling rate used in our experiments, filtered backprojection is almost always used because of its higher speed. The original code was optimized for our parallel computer and acquisition scheme. In our case in which a single rotation axis is used, the reconstruction calculation for each section is independent. Therefore, the problem of parallelizing the reconstructions reduces to sim-




FIG. 6. (a) Reconstructed sections from the cricket data. Four sections of the reconstruction volume, indicated on the projection image, are shown. The fine features, such as individual hairs and details in the skeletal structure, including small cracks, are clearly visible. (b) Reconstructed sections from the fruit fly data. Four sections of the reconstruction volume, indicated on the projection image, are shown.





FIG. 7. (Color) Illustration of the visualization environment. (a) The rendering control and visualization environment displayed on the ImmersaDesk. The user can wear stereo goggles to view the display with 3D perception. A similar control tool set is also available at the beamline experiment station but with only 2D displays. (b)–(d) Three renderings from the cricket reconstruction digitally cut at different planes to show the internal structures of the cricket's head.

ply assigning reconstruction sections to individual processors. The reconstruction time using 80 processors is summarized in Table II. For full-sized images of 1024×1024 pixels at a 0.25° angular sampling rate, the reconstruction calculation can be completed in about 17 min. If the image size is reduced to 512×512 and the angular sampling rate to 0.5° , the reconstruction time can be reduced to under 1 min. Compared with the corresponding acquisition times, this calculation time is sufficient for monitoring the experimental results in real time. Some reconstruction results are shown in Figs. 6 and 7. The computation time scales approximately linearly

with the number of processors and the number of projections in the data set. Hence if the data are acquired in a few cycles, the reconstruction time for each cycle can be determined by the fraction of the number of projections acquired in that cycle. As an example, assume that 721 projections of 1024×1024 were to be acquired in a few cycles, and 37 projections with angular interval of 5° are to be acquired in the first cycle. (This angular interval is adequate for a good-quality preview of most samples.) The first cycle can be acquired in about 3 min and reconstructed in about 1 min. Therefore, the beamline user can preview the reconstruction from the first

cycle about 5 min after the data acquisition started.

VI. DISCUSSION

We have constructed a high-throughput x-ray microtomography system with sub- μm resolution and with combined acquisition and reconstruction times in the minutes to 1 h time scale. Currently, the speed of the system is primarily limited by the readout and transfer time of the CCD camera. In addition, for most tomography experiments, the double-crystal monochromator has too narrow a bandwidth which reduces the x-ray flux. The use of a CCD camera with 10 frame/s readout rate and a multilayer monochromator with an energy resolving power of 100 should reduce the acquisition time by a factor of 10–50 and make it similar to or less than the reconstruction time. Both improvements are currently being implemented at the 2-BM beamline. Further improvement in the preprocessing and reconstruction code along with a network speed upgrade to 1 Gbit/s would allow the speed of this system to be limited only by the photon statistics.

With the rapid increase in the speed of personal computers based on G4 and Pentium 4 microprocessors, it is possible to construct a cluster of microcomputers to perform the reconstruction calculations at the rate of current supercomputers. Digital signal processors can also be used effectively when Fourier-based reconstruction algorithms are applied, as in most medical imaging systems. Such a reconstruction system dedicated to an experiment station can be a cost-effective addition to the beamline instrumentation. It is also more desirable in practice since it removes the need to schedule usage with other computer users, as in the case with supercomputers.

A limitation of the current optical system is that the thickness of the scintillator requires special correction for the objectives with high numerical aperture because they have been designed for use with 0.17 mm thick glass coverslips. Although this correction is not possible with our current objectives, a new series of oil-immersion infinity-corrected Zeiss objectives with variable coverslip thicknesses will be used in the future, and this should be able to further increase

the resolution of our imaging system. Furthermore, the use of high-resolution Fresnel zone plate x-ray lenses can lead to a factor of 10–20 improvement in the spatial resolution. A full-field imaging x-ray microscope being implemented at the 2-ID-B (undulator insertion device) beamline at the APS has demonstrated 130 nm resolution with 1.8 keV soft x rays. Further development is likely to increase the resolution to 50 nm with acquisition times of 10 s per image. We anticipate that these developments will make high-throughput nanotomography not only possible, but routine.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy, Basic Energy Sciences, Office of Science, under Contract No. W-31-109-ENG-38. The authors would like to thank Mark Ellisman and Steve Young of the National Center for Microscopy and Imaging Research at the San Diego Supercomputing Center for providing the basis for their reconstruction codes.

¹Q. C. Johnson, J. H. Kinney, U. Bonse, M. C. Nichols, R. Nusshardt, and J. M. Brase, *Mater. Res. Soc. Symp. Proc.* **69**, 203 (1986).

²T. M. Breunig, J. C. Elliot, S. R. Stock, P. Anderson, G. R. Davis, and A. Guvenilir, *X-ray Microscopy*, edited by A. G. Michett and G. R. Morrison (Springer, Berlin, 1991), Vol. III, pp. 465–468.

³H. R. Lee, B. Lai, W. Yun, D. C. Mancini, and Z. Cai, *Proc. SPIE* **3149**, 257 (1997).

⁴W. S. Haddad, I. McNulty, J. E. Trebes, E. H. Anderson, R. A. Levesque, and L. Yang, *Science* **266**, 1213 (1994).

⁵J. Lehr, *Optik (Stuttgart)* **104**, 166 (1997).

⁶Y. Wang, C. Jacobsen, J. Maser, and A. Osanna, *J. Microsc.* **197**, 80 (2000).

⁷B. Lai, D. C. Mancini, W. Yun, and E. Gluskin, *Proc. SPIE* **2880**, 171 (1996).

⁸NCSA HDF homepage: <http://hdf.ncsa.uiuc.edu/>

⁹P. Klosowski, M. Koennecke, J. Z. Tischler, and R. Osborn, *Physica B* **241–243**, 151 (1998).

¹⁰Nexus homepage: <http://www.neutron.anl.gov/NeXus>

¹¹I. Foster and C. Kesselman, *The Grid: Blueprint for a New Computing Infrastructure* (Kaufman, 1998).

¹²G. von Laszewski *et al.*, *SIAM J. Sci. Comput. (USA)* **April**, 1999 (1999).

¹³R. Drebin, L. Carpenter, and P. Hanrahan, *Proceedings of SIGGRAPH '88*, 1988, pp. 65–74.

¹⁴P. J. Mercurio, T. T. Elvins, S. J. Young, P. S. Cohen, K. R. Fall, and M. H. Ellisman, *Commun. ACM* **35**, 54 (1992).