

MASSIVE: AN HPC COLLABORATION TO UNDERPIN SYNCHROTRON SCIENCE

Wojtek James Goscinski, Chris Hines, Paul McIntosh,
Monash University, Clayton, Australia

Keith Bambery, Ulrich Felzmann, Chris Hall, Anton Maksimenko,
Santosh Panjekar, David Paterson, Mark Tobin,
Australian Synchrotron, Clayton, Australia

Chris Ryan, Darren Thompson,
CSIRO, Clayton, Australia

Abstract

MASSIVE is the Australian specialised High Performance Computing facility for imaging and visualisation. The project is a collaboration between Monash University (lead), Australian Synchrotron (AS) and CSIRO, and it underpins a range of advanced instruments, including AS beamlines. This paper will report on the outcomes of the MASSIVE project since 2012, in particular focusing on instrument integration, and interactive access for analysis of synchrotron data. MASSIVE has developed a unique capability that supports an increasing number of researchers, including an instrument integration program to help facilities move data to an HPC environment and provide in-experiment data processing. This capability is best demonstrated at the AS Imaging and Medical Beamline (IMBL) where fast CT reconstruction and visualisation is essential to performing effective experiments. A workflow has been developed to integrate beamline allocations with HPC allocation providing visitors with access to a dedicated project space and a CT reconstruction service using a remote desktop environment. The work herein describes the processing and analysis workflows developed for processing AS data at the IMBL, Macromolecular crystallography beamlines (MX), X-ray Fluorescence Microscopy (XFM), and the Infrared Microspectroscopy Beamline (IR).

THE MASSIVE FACILITY

MASSIVE was established to provide data analysis and visualisation services to the imaging and instrumentation community. The facility provides computer hardware, software and expertise to drive research in the biomedical science, materials research, engineering, and neuroscience communities, and it underpins advanced imaging research in synchrotron X-ray and infrared imaging, functional and structural magnetic resonance imaging, X-ray computed tomography (CT), electron and optical microscopy. MASSIVE offers Australian scientists access to two specialized computing facilities at Monash University and AS, plus specialised services offered through the cloud.

MASSIVE underpins the processing and analysis

requirements of beamlines at AS. The IMBL is able to produce data at over 500 Mbytes/s and data sets over a terabyte in size. This introduces obvious challenges for researchers to capture, process, analyze and visualise data in a timely and effective manner. Researchers are also increasingly eager to perform data analysis “in-experiment” to make quick decisions.

Hardware

MASSIVE provides two computers, M1 and M2, that operate at over 5 and 30 teraflops respectively, using CPU processing, and over 50 and 120 teraflops, respectively, using Nvidia and Intel co-processors.. The computers are connected using a dedicated network for fast file transfer and management*.

GPUs have proved an important part of the MASSIVE environment. A number of key applications, including the X-TRACT [1] CT reconstruction software, have been parallelized to take advantage of GPUs and optimized on MASSIVE which is critical to fast processing of data in near real-time at the IMBL. Furthermore, GPUs underpin interactive visualisation through the MASSIVE Desktop and through parallel rendering tools such as Paraview [2].

Both systems have a GPFS [3] file system that is capable of a combined 5GB+ per second write speed. Fast file system performance has proved essential to support both the fast capture of data from instruments, and file system intensive image processing.

Instrument Integration

MASSIVE runs a program for the integration of imaging instruments with high performance computing capability. This program helps develop workflows for scientists to process and visualise data as an experiment progresses or immediately after it completes. A number of instruments have been integrated with MASSIVE [4]. The IMBL CT reconstruction workflow (described within) is a major output of this program.

The MASSIVE Desktop

MASSIVE provides users with an easily accessible scientific desktop environment configured for analysis and visualisation of research data. It provides researchers with access to a range of existing tools and software, including commercial and open-source applications used for synchrotron data processing such as X-TRACT [1],

* Detailed specifications available at: <https://www.massive.org.au/high-performance-computing/resources>

GeoPIXE [5], Drishti [6], and Avizo[†]. Many tools are integrated with the HPC scheduling environment [4].

The desktop environment is an effective way for MASSIVE to provide access to tools, without rewriting or wrapping those tools, underpinned by a high performance file system and running on high-end CPUs and GPUs.

The remote desktop uses the CentOS operating system running the KDE or Gnome environments. For remote access, the desktop uses TurboVNC/VirtualGL[‡], an open source VNC client, which supports remote hardware rendering and clients on Windows, Mac and Linux.

To make access as easy as possible the MASSIVE team develops a tool called Strudel[§] (short for Scientific Desktop Launcher) that automates the steps to access a desktop session. Strudel launches an interactive visualisation job through a visualisation queue, and connects using TurboVNC over a secure SSH connection. It is open source, configurable and also being applied at other Australian HPC facilities.

MASSIVE AND APPLICATIONS TO SYNCHROTRON SCIENCE

A number of workflows have been deployed by AS to leverage the M1 system. Five beamlines use MASSIVE to process user data during or shortly after data capture: the IMBL, MX1 and MX2 beamlines, XFM, and IR.




Figure 1: HPC usage by beamline 2012 to 2014 (MX1 and MX2 are shown combined).




Figure 2: Research areas supported by the M1 and M2 computers during or shortly after beamtime.

During the period 2011 to 2014, approximately 446 beamtime sessions have used M1. Figure 1 shows the number of beamtime allocations that used MASSIVE. Figure 2 characterizes the sciences using the facility.

IMAGING AND MEDICAL BEAMLINE

In order to control the data collection and optimise the experimental conditions at the IMBL, scientists must be able to visualise collected data in near real-time as the experiment is in progress. In particular, groups that are imaging anesthetized animals need to establish whether a previous scan has successfully produced the desired data before proceeding with the next step of an experiment.

The most popular workflow on the IMBL is for CT, which requires collection of a set of X-ray projection images with the specimen being rotated between each. To start a data collection the user sets up the specimen on a rotary table then runs acquisition from a custom GUI on a control computer. The GUI allows specimen information to be recorded and checks the CT parameters and instruments. IMBL motion control hardware produces a train of pulses between a defined start and finish angle of the rotation stage. Each pulse edge occurs at the defined projection angle for the image. The imaging detector responds to pulses by capturing an image at a set exposure time. Image data is rapidly recorded to an image file on the file store. With the IMBL standard 5.5 megapixel X-ray imager the full frame rate is around 30 frames per second. Tagged Image File Format (TIFF) files are currently used but a HDF5 format is being developed.

The computer used to control and collect data from the detector is connected via a high-speed fibre network link to the file system on M1. Following a single CT collection a set of projection images will be saved for either 180 degree or 360 degree rotation of the sample. Since the synchrotron beam has a high aspect ratio, it is often the case that a single specimen will require several segments collected through its height.

The preferred reconstruction software for CT data sets collected on IMBL is X-TRACT from CSIRO [1]. It runs on the MASSIVE desktop and provides a simple and easy to use interface for reconstruction. Each user group which uses IMBL is provided with an individual MASSIVE project, along with both archived data storage space, and scratch space for intermediate calculations (details described in following section). To stitch projection images prior to reconstruction IMBL provides a script which uses ImageMagick^{**} and CTAS [7].

The first imaging workflow step is to remove noise and pedestals using standard flat and dark field correction. Sinograms are compiled in memory, then reconstructed into image slices and then written to disk. 32 bit TIFF files are used since the popular output from X-TRACT is the calculated linear attenuation coefficient map. Visualisation of reconstructed data is done using the MASSIVE Desktop running Drishti, ParaView, or Avizo.

[†] <http://www.fei.com/software/avizo3d/>

[‡] <http://www.virtualgl.org/>

[§] <https://www.massive.org.au/userguide/cluster-instructions/strudel>

^{**} <http://www.imagemagick.org/script/index.php>

Annually IMBL provides around 100 days of CT imaging experiment time. On average a group will collect 5 to 10 TB of raw data. The scratch space required when processing this can be double that.

HPC Integration at IMBL

In 2015 AS and MASSIVE initiated a project to tightly integrate IMBL beamline allocations and project accounts on MASSIVE. The intent is to provide IMBL users with automatic access to desktop sessions for data processing and collaboration during and after an experiment. The technology to support this project was developed under the NeCTAR Characterisation Virtual Laboratory^{††}.

The integration workflow is as follows (Figure 3): (1) A user creates an ID in the AS user database and submits a proposal for access using the AS portal; (2) If beamtime is approved a booking is created and collaborators can be added to the booking; (3) The MASSIVE ID system periodically queries the AS user database using a RESTful web service; (4) A new MASSIVE user matching each collaborator and a new project matching each experiment is created in the MASSIVE ID system. Passwords on MASSIVE are not set (as they will never be used) and user names are not published (as they will be determined automatically); (5) At the beginning of a beamline session, an operator resets the system. This archives data captured during the previous experiment and defines the project for new user data on MASSIVE; (6) User visits the beamline; (7) Data captured is written to the defined location on MASSIVE and file ownership set; (8) For processing, the user logs into MASSIVE using Strudel providing their AS user credentials; (9) Strudel authenticates against the AS ID system. On successful authentication a token is returned; and (10) Strudel uses the token to authenticate the user to a MASSIVE web service and to establish a secure desktop instance.

This workflow has a number of attributes: First, it provides AS the ability to automatically create (and discontinue) projects and user identities on a shared HPC system; Second, it provides AS users and their collaborators with a dedicated project space for each experiment; and third AS users are provided with direct access to the MASSIVE Desktop while visiting the IMBL and the same environment after their visit.




Figure 3: IMBL HPC project and allocation workflow.

This workflow is deployed at the IMBL and since April 2015 it has been used by 35 experiments. There are plans to deploy the same workflow to XFM.

MACROMOLECULAR CRYSTALLOGRAHY BEAMLINES

The AS hosts two MX beamlines: the bending-magnet high-throughput MX1 beamline [8] and the micro-focus MX2 beamline powered by a small-gap in-vacuum undulator. Both beamlines deliver tuneable X-rays in the energy range from 6 to 18 keV for the standard MX and anomalous diffraction experiments for experimental phasing.

The beamlines use Auto-Rickshaw (AR) [9, 10], a data evaluation system, which provides computer coded decision-makers for successive and automated execution of a number of existing macromolecular crystallographic computer programs thus forming a software pipeline for automated and efficient crystal structure determination. AR is triggered by the collection of a diffraction image.

The main AR server (<http://www.embl-hamburg.de/Auto-Rickshaw>) runs at EMBL-Hamburg and is freely accessible publically. A mirror server has been installed on MASSIVE for automatic data processing for users. The system is triggered from the beamline computers when the auto-processed X-ray data indicates an anomalous signal. The results from AR are automatically directed to the same directory from where the jobs were triggered. The automatic triggering from the beamline computer are available as an opt-in mode for users. Apart from the automatic system, AR can be triggered from the beamline computers for structure solutions using various phasing methods over a web-based user interface by selecting appropriate parameters and by selecting the processed and scaled datasets. This workflow on MASSIVE has helped to produce the structure of the BC component of ABC toxins that explains the general mechanism for protein encapsulation and delivery [11]. Additionally, it helped to determine the crystal structure of the Extracellular Adhesion Domain of the Sialic Acid-binding Adhesin SabA from *Helicobacter pylori* [12] the Mg²⁺ transporter CorA from the Archaea *Methanocaldococcus jannaschii* [13], an Indole-3-Acetic Acid Amido Synthetase from *Grapevine* [14] and cyanuric acid hydrolase [15].

X-RAY FLUORESCENCE MICROSCOPY

Remote login to MASSIVE via the portal provides users at the XFM beamline [16] a persistent desktop environment to run the GeoPIXE package from CSIRO [17] for the reconstruction of list-mode data-sets, acquired at XFM using the Maia detector and fluorescence imaging system [18], and the exploration and analysis of these rich data-sets. Maia is a 384 element detector array with dedicated FPGA-based real-time processor, which can acquire data at event rates of up to around 10 MHz, with each event tagged by XY position in a scan, dwell time and incident and transmitted beam flux count for each

^{††} <https://www.massive.org.au/cvl>

pixel. First-pass processing is performed in real-time using the Dynamic Analysis method (DA) [5] implemented in the FPGA.

Off-line processing of the list-mode stream and deconvolution of its rich spectral content is performed in GeoPIXE exploiting 10-12 cores on M1. GeoPIXE also generates the DA transform matrix used for the real-time deconvolution. Using the matrix approach and parallel processing on MASSIVE, DA achieves the result at rates of around 10^8 events into 10^4 - 10^5 pixels per second. Further tools in GeoPIXE permit the reconstruction of XANES image stacks for speciation analysis and their correction, for example for spatial drift. Users then explore this multi-element or speciation image space using GeoPIXE and extract spectra from specific features for further quantitative analysis or to verify the accuracy of imaging. The resulting image data-sets, which can be as large as 1 G pixel showing detail in 20-30 elements, are explored interactively using GeoPIXE or using statistical analysis of the hyperspectral data.

This capability allows super high-definition X-ray fluorescence element images to be obtained with detail at $\sim 2 \mu\text{m}$ spatial resolution in areas of several square cm. The high pixel count opens up a suite of 3D imaging techniques, which map over 3 dimensions to provide high-resolution X-ray absorption near-edge structure (XANES) images of chemical speciation (scan in XYE) and fluorescence tomography reconstructions of 3D multi-element distribution (scan in XY θ).

GeoPIXE on MASSIVE is used for processing 95% of experiments at XFM and usage after a visit is increasing.

INFRARED MICROSPECTROSCOPY BEAMLINE

The IR beamline manages and maintains a project on MASSIVE to aid in the analysis of infrared spectra recorded from biological samples. Fourier transform infrared (FTIR) microspectroscopy has been applied to diverse areas of biological and medical research. The technique provides a direct and rapid probe of chemical composition without the need for sample pre-treatment with stains and has the potential to identify and validate important biomarkers. However, light scattering effects can hamper the analysis of these FTIR spectra and this is especially problematic for the case of Mie scattering from intact cell samples. MASSIVE is being used to computationally model and correct for these light scattering effects in synchrotron FTIR microscopy data. The algorithm used is Resonant Mie Scattering modelling through Extended Multiplicative Signal Correction (RMieS-EMSC) [19]. Recent studies on cells performed at the beamline include work on developing a diagnosis for malaria [20] and understanding the biological effects of radiation [21]. Parallel computing allows RMieS correction to be performed around 100 times faster than is possible on desktop computers. MASSIVE allows data characterisation to be performed directly on the spectra stored on the synchrotron facilities archive, so results can

be shared with collaborators conveniently and securely.

CONCLUSION

The MASSIVE collaboration has significantly enhanced data processing at AS beamlines. The analysis workflows and desktop described within are important AS services. This is most strongly demonstrated when MASSIVE is being used "in-experiment" allowing the visitor make quick decisions, increasing AS productivity. Flexibility in scheduling and HPC management is important and it is essential that capability can be reserved for instrument and interactive use. During these times, the system maybe underutilised however the impact is gained when scientists have access to HPC capability on demand.

ACKNOWLEDGMENT

This work is funded by project partners, the NCI Specialised Facilities, and Victorian Government and the NeCTAR CVL. The authors acknowledge contributions from staff at AS, Monash and CSIRO.

REFERENCES

- [1] T. E. Gureyev, *et al.*, "Toolbox for advanced X-ray image processing," *Proc. SPIE8141 B*, vol. 81410, pp. 81410B-14, 2011.
- [2] A. Henderson, *et al.*, *The ParaView Guide: Kitware*, NY, 2004.
- [3] F. B. Schmuck and R. L. Haskin, "GPFS: A Shared-Disk File System for Large Computing Clusters," in *FAST*, 2002, p. 19.
- [4] W. J. Gosinski, *et al.*, "The Multi-modal Australian ScienceS Imaging and Visualisation Environment (MASSIVE) high performance computing infrastructure: applications in neuroscience and neuroinformatics research," *Frontiers in Neuroinformatics*, vol. 8, 2014-March-27 2014.
- [5] C. Ryan, "Quantitative trace element imaging using PIXE and the nuclear microprobe," *International Journal of Imaging Systems and Technology*, vol. 11, pp. 219-230, 2000.
- [6] A. Limaye, "Drishti: a volume exploration and presentation tool," in *SPIE Optical Engineering+ Applications*, 2012,
- [7] A. Maksimenko, M. Ando, S. Hiroshi, and T. Yuasa, "Computed tomographic reconstruction based on x-ray refraction contrast," *Applied Physics Letters*, vol. 86, p. 124105, 2005.
- [8] N. P. Cowieson, *et al.*, "MX1: a bending-magnet crystallography beamline serving both chemical and macromolecular crystallography communities at the Australian Synchrotron," *Journal of Synchrotron Radiation*, vol. 22, pp. 187-190, 2015.
- [9] S. Panjekar, *et al.*, "Auto-Rickshaw: an automated crystal structure determination platform as an efficient tool for the validation of an X-ray diffraction experiment," *Acta Crystallographica Section D: Biological Crystallography*, vol. 61, pp. 449-457, 2005.

- [10] S. Panjikar, et al, "On the combination of molecular replacement and single-wavelength anomalous diffraction phasing for automated structure determination," *Acta Crystallographica Section D: Biological Crystallography*, vol. 65, pp. 1089-1097, 2009.
- [11] J. N. Busby, et al, "The BC component of ABC toxins is an RHS-repeat-containing protein encapsulation device," *Nature*, vol. 501, pp. 547-550, 2013.
- [12] S. S. Pang, et al., "The three-dimensional structure of the extracellular adhesion domain of the sialic acid-binding adhesin SabA from *Helicobacter pylori*," *Journal of Biological Chemistry*, vol. 289, pp. 6332-6340, 2014.
- [13] A. Guskov, et al., "Structural insights into the mechanisms of Mg²⁺ uptake, transport, and gating by CorA," *Proceedings of the National Academy of Sciences*, vol. 109, pp. 18459-18464, 2012.
- [14] T. S. Peat, et al., "Cyanuric acid hydrolase: evolutionary innovation by structural concatenation," *Molecular microbiology*, vol. 88, 2013.
- [15] T. S. Peat, et al, "Crystal structure of an indole-3-acetic acid amido synthetase from grapevine involved in auxin homeostasis," *The Plant Cell*, vol. 24, pp. 4525-4538, 2012.
- [16] D. Paterson, et al., "The X-ray Fluorescence Microscopy Beamline at the Australian Synchrotron," in THE 10TH INTERNATIONAL CONFERENCE ON X-RAY MICROSCOPY, 2011, pp. 219-222.
- [17] C. Ryan, et al., "MAIA X-ray fluorescence imaging: capturing detail in complex natural samples," in *Journal of Physics: Conference Series*, 2014, p. 012002.
- [18] R. Kirkham, et al., "The Maia Spectroscopy Detector System: Engineering for Integrated Pulse Capture, Low-Latency Scanning and Real-Time Processing," *SRI 2009, 10th International Conference on Radiation Instrumentation. AIP Conference Proceedings*, vol. 1234, pp. 240-243, 2010.
- [19] P. Bassan, et al., "Resonant Mie scattering (RMieS) correction of infrared spectra from highly scattering biological samples," *Analyst*, vol. 135, pp. 268-277, 2010.
- [20] B. R. Wood, et al., "Diagnosing malaria infected cells at the single cell level using focal plane array Fourier transform infrared imaging spectroscopy," *Analyst*, vol. 139, pp. 4769-4774, 2014.
- [21] E. Lipiec, et al., "SR-FTIR Coupled with Principal Component Analysis Shows Evidence for the Cellular Bystander Effect," *Radiation research*, vol. 184, pp. 73-82, 2015.