DATA LIFECYCLE IN LARGE EXPERIMENTAL PHYSICS FACILITIES: THE APPROACH OF THE SYNCHROTRON ELETTRA AND THE FREE ELECTRON LASER FERMI

Fulvio Billè[#], Roberto Borghes, Francesco Brun, Valentina Chenda, Alessio Curri, Venicio Duic, Daniele Favretto, Georgios Kourousias, Marco Lonza, Milan Prica, Roberto Pugliese, Martin Scarcia, Michele Turcinovich, Elettra Sincrotrone Trieste, Basovizza, Trieste, Italy

Abstract

Often the producers of Big Data face the emerging problem of Data Deluge. The literature includes experiences and proposals for dealing with this problem. Nevertheless experimental facilities, such as synchrotrons and free electron lasers, may have additional and specialised requirements. This is partially due to the plethora of cutting-edge custom instrumentation but also to the necessity for a complex system that manages the access to the facility for thousands of scientists. A complete data lifecycle describes the seamless path that joins distinct IT tasks such as experiment proposal management, user accounts, beamline software, data acquisition, reduction and analysis, archiving, cataloguing, and remote access. This short paper presents the data lifecycle of the synchrotron Elettra and the free electron laser FERMI. It attempts to go beyond the overgeneralised binary distinction of information between data and metadata since in practice there are many heterogeneous but related data sources. With the focus on the broad concept of data access, the Virtual Unified Office is presented. It is a core element involved in scientific proposal management, user information database, scientific data oversight, and remote access. Eventually recent choices of advanced beamline software are here discussed as well as approaches to distributed control systems. The latter holds the key role to data and metadata acquisition but it also requires integration with the rest of the system components in order to facilitate additional services such as cataloguing, archiving and remote access. With the experience of the past years and the forecasting of future requirements, the expected system upgrades are also here outlined. The scope of this paper is to disseminate the current status of a complete data lifecycle, discuss key issues and hints for future directions. This is potentially useful for similar facilities especially for those under construction or upgrade.

INTRODUCTION

Elettra and FERMI

The research centre Elettra consists of two advanced light sources, the electron storage ring Elettra and the free-electron laser (FEL) FERMI hosting 32 experimental stations.

Elettra is a third-generation synchrotron serving the scientific and industrial community since 1993. It is involved in a broad spectrum of research in physics, chemistry, biology, life sciences, environmental science, medicine, forensic science, and cultural heritage. Currently 28 beamlines, including a storage-ring free-electron laser, utilize the radiation generated by the Elettra source. All of the most important x-ray based techniques in the areas of spectroscopy, spectromicroscopy, diffraction, scattering and lithography are present, together with facilities for infrared microscopy and spectroscopy, ultraviolet inelastic scattering, and band mapping. In 2010 the third-generation electron storage ring source Elettra was upgraded to operate in top-up mode. It is the only third-generation synchrotron radiation source in the world that operates routinely at two different electron energies; i) 2.0 GeV for enhanced extended ultraviolet performance and spectroscopic applications, and ii) 2.4 GeV for enhanced x-ray emission and diffraction applications. A plan for a substantial upgrade named Elettra 2.0 is in motion [1].

FERMI is a new seeded free electron laser (FEL) facility in operation next to the third-generation synchrotron radiation facility. Unique among the FEL sources currently operating in the ultraviolet and soft x-ray range worldwide, FERMI has been developed to provide fully coherent ultrashort 10-100 femtosecond pulses with a peak brightness ten billion times higher than that made available by third-generation light sources. It is a single-pass FEL that covers the wavelength range from 100 nm to 4 nm in the first harmonic. It comprises two separate coherent radiation sources, FEL-1 and FEL-2, that are being brought online sequentially. FEL-1 operates in the wavelength range between 100 and 20 nm via a single cascade harmonic generation, while the latest FEL-2 is designed to operate at shorter wavelengths (20-4 nm) via a double cascade configuration.

Data Deluge and Discontinuity

The era of Big Data has raised relevant problems like that of Data Deluge [2]. Elettra and similar facilities can be seen as information generators. The processes that generate such information flows are rather specialised since the facilities themselves have particular workflows. Such workflows include proposal submission, evaluation, beamline access, execution of experiments with custom instrumentation, data analysis, dissemination of scientific results etc. At this stage the issue of Data Deluge is studied in the focused context of a synchrotron facility.

[#]fulvio.bille@elettra.eu

ISBN 978-3-95450-148-9

This brief manuscript suggests that the elements with the biggest impact in Elettra are: advances in detectors, upgrades of the sources, demanding data analysis, and archiving requirements. These four elements contribute to a very demanding information flow which is a challenge to manage. The next sections provide the reader with an overview on how it is done at Elettra.

PROPOSAL MANAGEMENT

An experiment in a synchrotron facility requires a very specific procedure. While the core remains the actual experiment, the complete process includes multiple subprocesses like proposal submission and evaluation, access request, data access and processing, remote operations and general accounting. In 1997 Elettra has been among the first institutes of its kind, to develop a complete system named Virtual Unified Office which serves as the backbone for the above-mentioned processes. This document refers to proposal management as the collective of all operation related to it. Even if this section is focused on proposal management for reasons of connecting it with the broader data flow, the VUO extends much further acting as a sort of ERP system.

Virtual Unified Office

The web portal of the VUO manages the whole lifecycle of an experiment: from the first approach of the researchers to the publication of the results. It is based on Oracle and mod_owa while each webpage corresponds to a PL/SQL stored procedure.

The first phase is the registration: it is of paramount importance since the username/password combination is used by the user not only on the web portal but also on the acquisition workstation, on the data storage devices, on the computational cluster for data analysis, on the instruments used for remote operations and Wi-Fi access. It is a AAA system different from the company one and it is dedicated to the scientific users. An alternative option is the Umbrella authentication system [https://www.umbrellaid.org] which is also available.

The proposal workflow is fully managed from the portal and depending on the result of its evaluation it may end up in the beamtime schedule of the portal calendar.

The researcher submitting the proposal can use a 3-level storage: Scratch, Online, and Offline. Scratch and Online are high reliability storages using the Gluster distributed file system. Offline is managed by an external storage server farm at the CINECA supercomputing center and is based on iRODS. Scratch is the working disk, where the acquisition and the data reduction and pre-screening take place. Initially, access is allowed only from the beamline workstations. Eventually when the data are finalised they can be transferred to Online by means of simple operations from the VUO portal. The data store in Online cannot be further edited and serves as the main data deposit.

Even if the main interaction is through a simple web portal, the actual operations take place on different servers and they are performed by agents written in Python that are communicating between each other through TANGO [3]. Such operations include the generation of the unix users, creation of the directory tree, permission setting for POSIX Access Control.

The transfer from Scratch to Online happens through the use of encapsulated rsync commands coordinated by Tango agents. There are cases where this transfer occurs in a scheduled and automatic manner as a task of the acquisition applications; these tasks can also deal with possible data preprocessing (i.e. lossless compression of HDF files).

The whole team of participants of the proposal can access the data through WebDAV and though web browsers. In special cases other systems and protocols can be utilised (e.g. GridFTP) even if are not regularly supported. There are also in place systems for remote access allowing for collaboration between the researchers who stay in their home institutes. These are tunnel systems that allow safe access to the resources through RDP, NX, VNC. Naturally on-site access is faster than off-site not only due to the bandwidth but also due the available protocols.

Once the experiment is finished, the user must provide 3 levels of feedback at different time schedule and with increasing detail. The first is requested immediately after the experiment and it mainly concerns the functionality of the machine. The second one regards the achievements, while the final one is the experimental report. Such procedures may seem purely logistical but in practice may serve a real scientific purpose like associating a specific dataset with a peer-reviewed publication.

The same portal is also used for managing the users' reimbursement of expenses therefore it is integrated with the company's ERP. The VUO portal is also multi-facility since it manages not only Elettra and FERMI but also those related to the European CERIC-ERIC consortium [http://www.ceric-eric.eu/].

Another application of the VUO is the registration of all the publications related to experiments performed at the facility. The publication database is also used for evaluating the future research proposals.

DATA ACQUISITION

The actual acquisition of data is a process of paramount importance. The beamline end-stations of synchrotrons and free electron lasers are complex instruments of interconnected components like motors, pumps, and detectors. The data acquisition software is an advanced system capable of commanding all of these components

ISBN 978-3-95450-148-9

while at the same time provides the user with a graphical interface [4,5].

Endstation Control Systems

The Scientific Computing team of Elettra has designed a modern, flexible and extensible endstation control based on technologies compatible with the TANGO distributed control system. The user interface is based on QT while the core language is Python. The target data format is custom structured HDF5. In order to avoid multiple developments, the system has a common core that aims at easy adaption to each beamline. Special attention has been paid for allowing easy integration with processing pipelines. Due to the brief nature of this manuscript, the reader is encouraged to refer to the manuscript "*A flexible system for end-user data visualisation, analysis prototyping and experiment logbook*" by Borghes et al. present at the ICALEPCS2015 proceedings.

DATA ANALYSIS

In most experiments, the data produced require various stages of processing prior reaching a stage where they can be interpreted. In their initial form they are often referred to as RAW data but this is an oversimplification. During an experiment there are multiple detectors that acquire data (e.g. images, spectra, beam characteristics, machine parameters etc). As a unique set they get assembled and processed. This processing may include multiple steps and later be separated in different operation pipelines. It soon becomes clear that the workflow is neither standard nor static thus it requires dynamic approaches that often make optimized solutions really difficult to be achieved.

Multiple Approaches

Considering that data analysis through computational processing does always take place even in non ideal situation (i.e. lack of a specialised team/resources), in Elettra presently exist multiple approaches. The minimal is the one where the beamline scientist tries to process their own data with the use of personal computers. On the other extreme there is a specialised team for scientific computing where, in collaboration with the beamline personnel, develop solutions that use high end enterprise hardware according to standardized approaches. In real case scenarios both approaches coexist but also all the in-between options as well. While trying to enforce the more advanced solutions, the scientific computing team is still providing support for a variety of options. The core development language is Python but in most cases it relies on external modules often developed in C or Fortran. For GPGPU computing the current developments are mostly in CUDA but past projects included OpenCL. The target operating system is any modern distribution of Linux but without forcing a specific one. Many systems are command-line but where a GUI is necessary it is developed in QT. The TANGO Python binding is often necessary for collecting metadata useful to the analysis. Nevertheless there are still plenty of data analysis software written in IDL and Matlab for Windows.

Data Formats

Decades ago, before the Big Data era, communities like that of crystallography had identified the need for standard and common formats [6]. It is well established why this is a necessity. Nowadays reproducibility, repeatability, exchange, size, speed, standardization of data analysis are even more relevant thus data formats require special attention. Elettra has participated in relevant projects (PaNData Europe, PaNData-ODI) [7] and has specialized studies leading to implementations for formats for X-ray florescence microscopy (XRF), Ptychography and Computed Tomography. All of the beamlines of the FEL FERMI do use HDF5 as the container of metadata-rich structures [8]. This experience provided the scientific computing team with a good insight regarding the advantages and disadvantages of these formats.

Hardware

The hardware resources for data analysis that is non uniform is a complex issue. They require non only HPC and server class hardware but also high end work stations often equipped with professional monitors (i.e. imaging applications). The advance use of GPGPU requires systems like NVIDIA TESLA. In Elettra their main use is in applications for x-ray microscopy imaging. Such applications often require GPU memory boards with >10GB of memory. For the in-house CDI/ptychography capacity these resources still pose the main limitation regarding the size of scans. Finally, the analysis requires accessing the data but this access, if it is through a slow connection, may become the bottleneck of the process. For this reason, the network and relevant infrastructure should always be adequate enough and this is often a difficult task especially with the modern detectors.

ACCESS

Accessing the data happens in two ways: supervised, where a user access them on will, and unsupervised when it is part of an automated process like a data analysis workflow. The access may allow for read and write operation and have specific requirements of performance and user rights. The facility requires a specific structure to describe in-house and visiting research teams. This is done by suitable mapping to advanced ACL on filesystem. These may require user accounts for instruments and automated group hierarchy based on beamtime proposals. Even if this in some cases requires adhoc system administration operations, the majority of such operations are taking place through automatic and semi-automatic procedures in the Virtual User Office. Finally there is a special class of access requirements, that of remote access.

ISBN 978-3-95450-148-9

Remote Operations

Any form of off-site data access or operation is considered remote. It usually takes place after a beamtime experiment and in its simplest form regards data download. This is already included as a service in the VUO. Still more complex forms of remote access may be necessary. Such cases go beyond the remote desktop solutions like NX and cloud computing resources for running analysis programs. Among the most demanding is that of remote beamline operation. The latest challenge in Elettra is remote operation of a beamline in an unattended mode (Xpress beamline) meaning in absence of local personnel. Past projects have successfully implemented scientific data analysis systems based on the model of software and hardware as a service [8,9].

INTERCONNECTION

Among the trends of the computational community, there is the focus on metadata and their importance. A beamline dataset usually includes a set of instrument parameters as metadata for the core raw data acquired by the principal detectors. It often becomes apparent that the conceptual distinction of metadata and data in such specific cases is rather difficult. In real world, acquisition and analysis scenarios, this becomes even more complex when a complete dataset is not enough for a complete processing due to the fact that additional pieces of information (like experiment information) are stored in other systems like the VUO. At this point at Elettra a dataset is only a node of a network that all together provides the necessary information. This set of interconnected nodes uses the RDBMS of the VUO, the HDF5 dataset's data and metadata, realtime information from the TANGO control system, user accounting from the file system, etc. The full analysis of this approach is not the argument of this brief overview paper but the core idea is that of abandoning the single file as a dataset and represent it as a set of interconnected heterogeneous information providers. Defining the systems as such should allow us for optimisations and future upgrades.

FUTURE UPGRADES

The synchrotron Elettra was built in 1993 so most of its beamlines have well established operation modes that permit only gradual upgrades and maintenance. For the past years the Elettra scientific computing team which is responsible of all the activities mentioned in this manuscript, had the opportunity to implement new technologies in the new Free Electron Laser FERMI and its beamlines. To extend this activities there is a plan for an upgrade of the synchrotron to a new source: Elettra 2.0 that should include a substantial upgrade and modernization of all the existing beamlines and procedures [1]. The aforementioned axis of the i) set of systems on the Virtual Unified Office, ii) the cloud based distributed and modular data acquisition model, and iii) the systematic data analysis will be among the core elements of the upgrade plan. At the same time there will be a detailed investigation aiming at the optimization of concepts such that of the dataset as an interconnection of heterogeneous information sources. Eventually this will required advances in both infrastructure and scientific software.

CONCLUSION

Synchrotron and free electron laser facilities like Elettra and FERMI have special needs regarding their data flow. In order to face problems related to data deluge, Elettra implements a complete data lifecycle. This covers all the experimental stages; from proposal submission to beamline access and data analysis. In order to do so key elements like proposal management, acquisition software, and data analysis need to be seamlessly interconnected. This brings the introduction to the concept of data as a set of connected information nodes from heterogeneous sources in contrast to the traditional file centric view. Future work will extend and build on this idea while the planned upgrades to Elettra 2.0 may be the ground for implementing the potential advances.

ACKNOWLEDGEMENT

The authors, members of the Scientific Computing team of Elettra, thank the IT Group and the beamlines Groups of Elettra Sincrotrone Trieste. Moreover they acknowledge the importance of advanced technologies and open source software like that of the distributed control system TANGO.

REFERENCES

[1] E. Karantzoulis, “Elettra 2.0 – The next machine”, Proceedings of IPAC2015, Richmond, VA, USA
[2] Anderson, Chris. The End of Theory: The Data Deluge Makes the Scientific Method Obsolete. Wired, 2008.
[3] Götz, A., E. Taurel, J. L. Pons, P. Verdier, J. M. Chaize, J. Meyer, F. Poncet, G. Heunen, E. Götz, and A. Buteau. “TANGO a CORBA Based Control System.” ICALEPCS2003, Gyeongju, October, 2003.
[4] Borghes, R., V. Chenda, A. Curri, G. Gaio, G. Kourousias, M. Lonza, G. Passos, R. Passuello, L. Pivetta, and M. Prica. “Control and Data Acquisition Systems for the FERMI@ Elettra Experimental Stations.” Proceedings of ICALEPCS2011, Grenoble, France, MOPMU015, 2012, 462.
[5] Borghes, R., V. Chenda, A. Curri, G. Kourousias, M. Lonza, M. Prica, and R. Pugliese. “A Common Software Framework for FEL Data Acquisition and Experiment Management at FERMI.” Proceedings of ICALEPCS2013, San Francisco, CA, USA, 2013, 6–11.
[6] Hall SR, Allen FH, Brown ID. “The Crystallographic Information File (CIF): a new standard archive file for crystallography”. Acta Crystallographica A47 (6): 655–685 (1991)

ISBN 978-3-95450-148-9

[7] Prica, Milan, George Kourousias, Alistair Mills, and Brian Matthews. “Requirements for Data Catalogues within Facilities.” Work 501 (n.d.): 63688.
[8] Brun, Francesco, Serena Pacilè, Agostino Accardo, George Kourousias, Diego Dreossi, Lucia Mancini, Giuliana Tromba, Roberto Pugliese, “Enhanced and Flexible Software Tools for X-ray Computed Tomography at the Italian Synchrotron Radiation Facility Elettra”, Fundamenta Informaticae 141 (2015) 233–243
[9] Brun, Francesco, Lucia Mancini, Parnian Kasae, Stefano Favretto, Diego Dreossi, and Giuliana Tromba. “Pore3D: A Software Library for Quantitative Analysis of Porous Media.” Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 615, no. 3 (2010): 326–32.

ISBN 978-3-95450-148-9