

Data management and database framework for the MICE experiment

J Martyniak¹, JJ Nebrensky² and D Rajaram³ on behalf of the MICE collaboration

¹ Imperial College London, Physics Dept. London SW7 2BW, UK

² Brunel University, Uxbridge UB8 3PH, UK

³ Illinois Institute of Technology, Chicago, IL 60616, USA

janusz.martyniak@imperial.ac.uk

Henry.Nebrensky@physics.org

durga@fnal.gov

Abstract. The international Muon Ionization Cooling Experiment (MICE) currently operating at the Rutherford Appleton Laboratory in the UK, is designed to demonstrate the principle of muon ionization cooling for application to a future Neutrino Factory or Muon Collider. We present the status of the framework for the movement and curation of both raw and reconstructed data. We also review the implementation of a robust database system that has been designed for MICE. A raw data-mover has been designed to safely upload data files onto permanent tape storage as soon as they have been written out. The process has been automated, and checks have been built in to ensure the integrity of data at every stage of the transfer. The data processing framework has been recently redesigned in order to provide fast turnaround of reconstructed data for analysis. The automated reconstruction is performed on a dedicated machine in the MICE control room and any reprocessing is done at Tier-2 Grid sites. In conjunction with this redesign, a new reconstructed-data-mover has been designed and implemented. The processing of data, whether raw or Monte Carlo, requires accurate knowledge of the experimental conditions. MICE has several complex elements ranging from beamline magnets to particle identification detectors to superconducting magnets. A Configuration Database, which contains information about the experimental conditions (magnet currents, absorber material, detector calibrations, etc.) at any given time has been developed to ensure accurate and reproducible simulation and reconstruction. A fully replicated, hot-standby database system has been implemented with a firewall-protected read-write master running in the control room, and a read-only slave running at a different location. The actual database is hidden from end users by a Web Service layer, which provides platform and programming language-independent access to the data.

1. Introduction

Ionization cooling is a technique which allows reducing the emittance of a beam of charged particles. The momentum of the particles is reduced by the mechanism of energy loss as they pass through

Interim placeholder PDF for BURA:

Presented at: 22nd *International Conference on Computing in High Energy and Nuclear Physics (CHEP2016)*, San Francisco, USA (10th – 14th October 2016); contribution I.D. 540

To be published in: *IoP Journal of Physics: Conference Series (JPCS)*

absorber material, and then subsequent reacceleration restores the lost longitudinal momentum thereby “cooling” the beam.

The International Muon Ionization Cooling Experiment (MICE) [1] located at the Rutherford Appleton Laboratory, UK, (RAL) is designed to demonstrate the principle of muon ionization cooling for the first time. The main component of MICE is the cooling channel which consists of a series of absorbers and high precision scintillating fibre trackers placed in a strong 4T solenoidal field created by a set of superconducting magnets. To discriminate between the muons of interest and the background of other particles MICE is equipped with a set of particle identification detectors (time-of-flight detectors, Cherenkov counters, and calorimeters) [2].

Figure 1 shows the current setup of MICE with absorbers, particle identification detectors and high precision scintillating-fibre trackers [3].

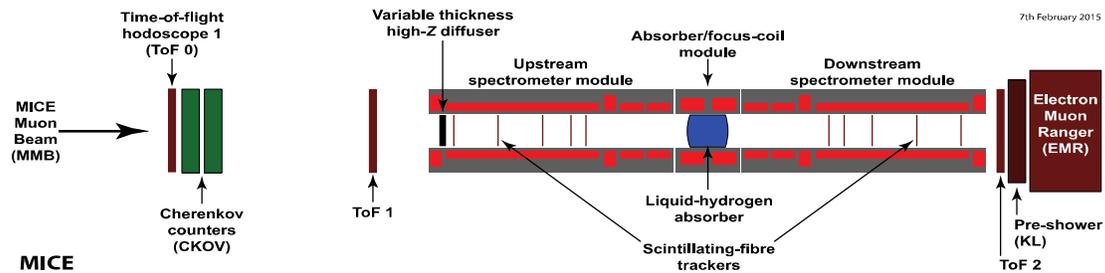


Figure 1. Current MICE configuration.

2. Raw and reconstructed data handling

2.1. Raw data handling

The MICE data acquisition (DAQ) [4] system, developed on the basis of the DATE package from the ALICE experiment is responsible for recording data consistent with experimental triggers. Data are taken in “runs” of varying configurations and a run typically corresponds to ~ 2 hours of operation. The DAQ writes raw data to disk in 250 MB chunks during each run. When the run has finished a File Compactor process collects the recorded chunks along with essential logs and outputs from the online monitoring and online reconstruction programs and collates them into a single tarball. The compactor also adds the list of files and their MD5 checksums to allow integrity checking later in the processing chain. The data-mover process actively monitors the creation of new tarballs and transfers data to the CASTOR Storage Element (SE) at the GridPP Tier-1 facility at RAL [5] which maintains 2 tape copies for added security. The process is described in detail in [6].

The raw data movement task has 2 steps: 1) an initial, temporary, copy with self integrity check and 2) a final copy to the Grid. The first step is intended to avoid bottlenecks on the DAQ disk access as the DAQ writing must take priority during the run.

At the end of each run the compacted tarball is placed in a directory actively watched by the data movement process and the initial copy is started only when the tarball is ready. This is indicated by a presence of a run-dependent semaphore. After the initial copy is made and the file’s integrity is verified, the CASTOR upload process is started. The CASTOR upload process makes a checksummed transfer, registers the file with the LCG File Catalog (LFC) and stores basic file attributes in the MICE Metadata Catalogue which stores essential attributes about the files. We have since upgraded the system for robustness and now use GFAL2 Python API to transfer data and perform LFC registration. This has proved to be more robust than the binary LCG tools used in the original data mover.

2.2. Handling of reconstructed data

Interim placeholder PDF for BURA:

Presented at: 22nd International Conference on Computing in High Energy and Nuclear Physics (CHEP2016), San Francisco, USA (10th – 14th October 2016); contribution I.D. 540

To be published in: *IoP Journal of Physics: Conference Series (JPCS)*

In our paper [6] we described a distributed Grid MICE data reconstruction system. It involved submitting a series of reconstruction jobs to the Grid and copying data back to the RAL CASTOR storage. The reconstruction process was triggered by raw data uploaded to CASTOR and a relevant run record present in the MICE Metadata Catalogue. This method was prone to delays and unnecessary bottlenecks associated with the processing and submission chain. A lot of effort was made to improve the reconstruction speed of the software which has allowed us to perform offline reconstruction on a dedicated machine in the Control Room at a speed consistent with the data-taking rate. The reconstruction mover was then accordingly modified and uses a single stage data movement process, similar to the raw data handling task. The reconstructed data transfer task is now triggered by the same method as the raw data transfer – namely, a semaphore indicating that the reconstructed tarball is ready to be moved. This made the raw and reconstruction movement tasks independent of each other. The Grid-based distributed reconstruction system is still available and is planned to be used for large scale reprocessing.

3. MICE Configuration Database

Since MICE takes data under a variety of configurations – different magnetic fields, triggers, absorber materials, etc – information about experimental conditions is vital for accurate simulation and reconstruction. MICE uses a Configuration Database (CDB) – a fully replicated, hot-standby database system running on a master-slave PostgreSQL DBMS. Access to the database is provided by a Java EE JAX-WS Web Service (WS) layer deployed on Apache Tomcat. We have set up a firewall protected read-write master node located in the experiment Control Room and a public read-only server hosted within the RAL PPD Tier 2 [5]. A JAX-WS Web Service, which is a SOAP service, allows platform and programming language independent access to data.

3.1. CDP API

We provide following CDB APIs:

- *C*, to allow the MICE Run Control system to read various parameters from the hardware and store them in the CDB. Run Control is based on EPICS, which is a toolkit with C APIs, so it was natural to design a WS client interface in the same programming language. We provide a selection of read-write operations used on a run-by-run basis to define magnet currents, absorber and target settings, beam parameters, and any other information that would have to be read from or written to the CDB. We also supply functions to define run condition templates (tags) to be read and written to the CDB. We use a *gSOAP* based client API to contact the CDB.
- *C++*, to be natively used by the reconstruction software, especially in areas where speed is critical. We have designed a selection of read-only operations to access detector calibrations, electronics and detector channel maps. We utilise *gSOAP C++* bindings to access the CDB.
- *Python*, to store certain predefined conditions such as detector geometry definitions, detector calibration, cabling and other non run-based information. This is the most complete set of operations widely used both by experts and end-users for their analyses. The API is regularly released and is bundled with the reconstruction software distribution but can also be installed stand-alone. It uses *Python suds* module which is the only external dependency.
- *Java*, used by a Google Toolkit based CDB-viewer WebApp. Only a slave database can be contacted this way.

3.2. CDB replication and recovery

Since the CDB is actively used during data taking it is very important to ensure that the master server is available at all times. Running PostgreSQL DBMS in a hot-standby mode allows a slave server to be promoted to a master in the event of the original master being unusable. This is a built-in feature of PostgreSQL. We use the following procedure at MICE:

Interim placeholder PDF for BURA:

Presented at: 22nd International Conference on Computing in High Energy and Nuclear Physics (CHEP2016), San Francisco, USA (10th – 14th October 2016); contribution I.D. 540

To be published in: *IoP Journal of Physics: Conference Series (JPCS)*

- Reconfigure (promote) a slave server to become a new master,
- Start a firewalled read-write WS, so the Control Room can now write to the new master. We still maintain the original read-only slave WS as before, so the read access is not affected,
- Synchronise remaining slaves, if applicable. We have an operational master-slave system at this point,
- When the old master is available again, swap the servers back and stop read-write access started above.

In addition to running a master slave(s) system in hot-standby mode we perform regular CDB back-ups for safety.

4. Conclusions

We have presented updated, automated versions of the raw and reconstructed data transfer tasks for MICE which allow files to be reliably and efficiently uploaded to permanent tape storage. Both the raw and reconstruction movers have been performing reliably during data-taking. We have also presented a robust, fully replicated MICE Configuration Database holding vital information about experiment settings. An extensive set of APIs have been developed supporting C, C++, and Python clients thus enabling end-users to easily interact with the database.

References

- [1] M. Bogomilov *et al.*, “The MICE Muon Beam on ISIS and the Beam-line Instrumentation of the Muon Ionization Cooling Experiment” *Journal of Instrumentation* **7** P05009 (2012)
G. Gregoire *et al.*, “MICE and International Muon Ionization Cooling Experiment Technical Reference Document” http://www.mice.iit.edu/trd/MICE_Tech_ref.html (October 2005).
- [2] Bertoni *et al.*, *Nucl. Instrum. Methods A* **615**(1), 14 (2010)
- [3] Ellis *et al.*, *Nucl. Instrum. Methods A* **659**(1), 136 (2011)
- [4] Y. Karadzhov, “MICE Experiment Data Acquisition System” #307 CHEP2013
- [5] David Britton *at al.*, “GridPP: the UK Grid for particle physics” *Phil. Trans. R. Soc. A* June 28, 2009
- [6] J. Martyniak “MICE data handling on the Grid”, *Journal of Physics: Conference Series* **513**, track 3

Interim placeholder PDF for BURA:

Presented at: 22nd *International Conference on Computing in High Energy and Nuclear Physics (CHEP2016)*, San Francisco, USA (10th – 14th October 2016); contribution I.D. 540

To be published in: *IoP Journal of Physics: Conference Series (JPCS)*