

Small experiment, big data: the data production of the Muon $g - 2$ Experiment

Paolo Girotti^{a,b,†,*}

^a*INFN, Sezione di Pisa, Pisa, Italy*

^b*Università di Pisa, Pisa, Italy*

E-mail: paolo.girotti@pi.infn.it, paolo.girotti@phd.unipi.it

The Muon $g - 2$ Experiment at Fermilab aims to measure the muon anomalous magnetic moment with the unprecedented precision of 140 parts per billion (ppb). In April 2021 the collaboration published the first measurement, based on the first year of data taking. The result confirmed the previous experiment at Brookhaven National Laboratory (BNL), and increased the long-standing tension with the Standard Model prediction to 4.2σ . The experiment is now running the sixth year of data acquisition of positive muon data, having accumulated a total of ~ 19 times the statistics of the BNL experiment. A collaboration-wide effort is now in place to help produce the multi-petabyte-sized data sets, a challenge typically faced by much bigger experiments. Having a quick production turnaround time is of critical importance in order to achieve a timely analysis and publication schedule. In this paper I will describe the production workflow, the former and current challenges, the resources, tools, and the future prospects of the Muon $g - 2$ Experiment.

*41st International Conference on High Energy physics - ICHEP2022
6-13 July, 2022
Bologna, Italy*

[†]on behalf of the Muon $g - 2$ collaboration.

*Speaker

1. Introduction

The Muon $g - 2$ Experiment at Fermilab is a precision physics experiment aiming to measure the anomalous magnetic moment of the muon with a precision of 140 parts per billion (ppb). The value of the muon anomaly $a_\mu \equiv \frac{g-2}{2}$ encodes all the possible interactions between the lepton and virtual particles. A precise measurement of it represents a strong test of the Standard Model, as well as a possible indicator of new physics. In April 2021, the Muon $g - 2$ Collaboration released their Run-1 measurement of the muon anomaly with a precision of 460 ppb [1]. This result brings the world average to a discrepancy of 4.2σ with respect to the current consensus of the Standard Model prediction [2]. If the analysis of the complete datasets of the Muon $g - 2$ Experiment will confirm the current central value and if the SM prediction stays unchanged [3], then the discrepancy would represent a strong indication of new physics.

1.1 Small experiment

A measurement of the muon anomalous magnetic moment can be obtained through a spin-polarized beam of muons, an external and measurable magnetic field, and a technique to determine the spin precession through time. The muon anomaly can be expressed as:

$$a_\mu = \frac{\omega_a}{\tilde{\omega}_p} \frac{g_e}{2} \frac{m_\mu}{m_e} \frac{\mu_p}{\mu_e}, \quad (1)$$

where ω_a is the muon anomalous precession frequency and $\tilde{\omega}_p$ is the Larmor precession frequency of the proton (ω_p) convolved with the beam distribution, representing the average field intensity experienced by the muons. The remaining factors are taken from external data. The Muon $g - 2$ Experiment in operation at Fermilab employs a 15-meter diameter superconducting storage ring and a 3.1 GeV polarized positive muon beam. As the muons circulate and decay, emitted positrons curl inward and hit a set of 24 calorimeters [4]. The beam distribution is measured with a set of trackers [5], and the magnetic field is measured with Nuclear Magnetic Resonance (NMR) probes [6]. With a group of ~ 200 collaborators and 5 years of accumulated physics runs, this *small* experiment is now approaching its final years of operation, while most of the data is now getting analyzed.

1.2 Big data

The 24 homogeneous calorimeters are segmented as matrices of 54 crystals, each one coupled to a SiPM, totaling 1296 channels sampled at 800 MHz. The two tracker stations are composed of 8 modules each, which in turn consist of 64+64 straw tubes in a U-V plane configuration, totaling 2048 channels read out at 400 MHz. Each $g - 2$ event is defined as a stored bunch, which lasts for 700 μ s. In this period, muons circulate roughly 4500 times while decaying, producing ~ 2000 hits on the calorimeters. Raw files are called subruns, defined to be ~ 2 GB in size, each one containing ~ 8 seconds of data and roughly 100 events, as the Fermilab muon accelerator chain provides 16 bunches every 1.4 seconds. Laser pulses fired to calorimeters between one bunch and the next are saved as well. When the experiment is running (typically from autumn to spring), data is collected around the clock, producing ~ 12000 subruns and 24 TB of data every day. Subruns are grouped into runs every ~ 500 files, summing to 1 TB size. Finally, a dataset is defined to be a contiguous set of runs with the same running conditions. A typical dataset contains $\sim 50k$ files, and there are 70 datasets from Run-1 to Run-5.

The current data collected in the first 5 years sums up to ~ 7 PB. While some data-quality selection is performed after production, most of the data is actually used for the determination of a_μ , as the goal of the experiment is to collect more than 10^{11} high energy positrons ($E > 1700$ MeV and $t > 30 \mu\text{s}$) for a final statistical uncertainty of 100 ppb. In order to achieve this, roughly 10 PB of raw data (including simulation) must be handled, produced, and analyzed entirely, granting the experiment a place in the *Big Data* category.

2. Offline production workflow

The reconstruction of a dataset is a process that requires multiple sequential steps. A single uninterrupted procedure is not possible as the calorimeters need multiple calibrations both at the subrun (a single raw file) and at the dataset level, and the calibrations can be extracted from reconstructed data only. For this reason, the production is split into a pre-production phase and a full-production phase. Both phases start from the raw files, but the pre-production only performs the minimum reconstruction needed for the extraction of the calorimeter calibrations. For convenience, a series of per-fill Data Quality Checks (DQC) are also performed at this stage.

The full reconstruction can be broken down into 8 steps as illustrated in Figure 1: *Pre-staging*, *Pre-production*, *Calibration and In-Fill DQC*, *Database upload and migration*, *Test production*, *Test DQC and validation*, *Full production*, and *Final DQC*. Starting from Run-2 the production was performed in a rolling scheme: each of the 8 steps happening simultaneously on different datasets, like a typical manufacturing production chain.




Figure 1: The 8 steps of the Muon $g - 2$ production.

2.1 Pre-production

The first step of pre-production involves copying the data from the tape system to the disk cache. This pre-staging of the data is necessary as the amount of raw data is too big to be always stored on disks. Moreover, "on-the-fly" direct copying from tape to the grid nodes would be extremely inefficient. As the amount of data to be reconstructed substantially increased from Run-1 to Run-5, this step became increasingly delicate. Careful planning is needed to ensure that resources are optimally distributed among experiments and that the next datasets are ready for production as soon as the previous ones are completed.

After a dataset has been fully pre-staged, pre-production is performed. It consists of an unpacking step followed by a light reconstruction of the positron data. At this stage, only pulse fitting and minimal calibration is executed. After pre-production, analyzer jobs automatically run to extract gain calibration data and quality checks of individual muon fills. The calibration constants and quality cuts are then manually verified by calibration responsables, before proceeding to upload them to a database.

The final step in the pre-production phase consists of uploading the extracted constants and conditions to a *development* database and then migrating them to a *production* database. These constants are stored as *data* and *status* tables, each marked with an Interval of Validity (IoV) which can be a range of runs or subruns. The database is accessed via a PostgreSQL interface, and when a subrun is reconstructed a `https` query extracts the corresponding table from the production database.

2.2 Full-production

When pre-production and calibration is done, the dataset is ready to be fully produced. Full production includes the unpacking of raw data and the full reconstruction of all detectors, applying the calibrations extracted in the previous phase. Three independent reconstructions for the calorimeters are executed.

The first step in this phase is a test production of a small portion ($\sim 10\%$) of the entire dataset. The subruns that undergo this test production are sampled uniformly across the dataset. The test production does not differ from the full production, and the reason for performing such test is to detect possible problems belonging to a dataset before committing with the full resources. Examples are mis-calibrations, detector faults, software bugs, or changes in the experiment configuration. For this, a dedicated validation step is performed on the produced files where plots are generated to check for physical quantities, such as calorimeter energy spectra, and anomalies in the calibration constants. The different reconstructions are compared with each other and long-term trends on the beam storage efficiency are checked. Once the data passes the inspection from production shifters and experts, full production resumes on the remainder of the dataset.

2.2.1 Data quality checks

Finally, when all the files of a dataset have been produced, a second round of validation is performed and then cuts are applied to select high-quality files only. Quality tests include the number of positrons observed in a fill, the number of muons lost from storage before decaying, the intensity of beam at injection, and the average pass rate in a subrun. After this process, typically $\sim 10\%$ of the files are discarded while $\sim 99\%$ of the positrons are actually kept. The selection of good files is then delivered to the analyzer teams.

3. Resources

The production of many petabytes of raw data requires a non-trivial amount of resources and a careful management of the whole process. This section will summarize the resources needed for the

production of the $g - 2$ data at Fermilab, the challenges faced and the lesson learned. The majority of the tools for managing the data, monitoring, and submitting the jobs are provided by the *FabrIc for Frontier Experiment* (FIFE) toolkit developed by Fermilab [7].

3.1 Tapes and disks

As already mentioned in section 2.1, an important step is the handling of the data itself. Both the raw and produced files are stored in tape libraries containing thousands of 12 TB tape cassettes. Each file belongs to a *file family*, and different families are not mixed in the same tapes. This ensures that, when a file family can be discarded (i.e. obsolete data), the relevant tapes can be recycled.




Figure 2: Data movement from DAQ to production. The pre-staging of raw files from tape to the dedicated GM2 read pool is highlighted in red.

As mentioned in section 2.1, the data has to be copied from tape to disk before being able to process it. The Muon $g - 2$ Experiment has a dedicated ~ 1400 TB disk pool for temporary storage of the raw data to be produced (Figure 2). All the output from production goes to a general ~ 7 PB pool and is then copied to tape. A new *migration* mode involving the dump of an entire tape at once recently helped us increase the efficiency of the pre-staging step.

3.2 Computing nodes

The offline production of the data runs both on the Fermilab computing nodes and on the Open Science Grid (OSG) distributed around the world. The Muon $g - 2$ Experiment has currently 5700 reserved slots in the Fermilab grid, where each slot consists of a CPU core and 2 GB of memory.

3.3 Shifters

While the multi-step *rolling* production scheme mentioned in section 2 proved to greatly increase the efficiency of production, the ever-increasing number of datasets of Run-4 and Run-5 resulted in a unmanageable amount of work to be done by the small production team. In late 2020, a group of *production shifters* was asked to contribute to the production of the data, and the effort was extended in late 2021 to the whole collaboration, as was already the case for the *online shifts*. Every week, 8 shifters guided by the production managers help monitor and execute the various steps of production. A substantial effort has been carried out to automate the process and simplify the checklists, and the number of needed shifters is now decreased to 3 per week.

3.4 Simulation

The simulation is an important piece for the physics analysis deployed mainly for studying the beam motion inside the storage ring. Many simulation packages are used to simulate various parts

of both the beamline starting from the proton target and the storage ring with all the detectors. The simulation has to consider electric and magnetic fields, the beam dynamics, the muon decays, and the interaction with the detectors. This means precisely tracking the particles over 220 km, which is 5000 turns inside the ring, during the 700 μ s bunch duration. The most CPU-intensive parts of the simulation are performed with the help of High-Performance-Computing (HPC) jobs at the National Energy Research Scientific Computing Center (NERSC).

4. Conclusions

The data production of the Muon $g - 2$ Experiment is a challenging but required task for achieving a new world-best measurement of the muon anomalous magnetic moment. The data grew almost 20-fold from the Run-1 publication to the current amount, which will allow a final statistical uncertainty of 100 ppb on a_μ . During the last year, great efforts managed to increase the resources, the production speed, and to improve the managing and the efficiency of the data production workflow. Analysis of Run-2 and Run-3 is now being finalized. Run-4 is now fully reconstructed and Run-5 is halfway pre-produced. Run-6 is now starting as of November 2022 and we are planning to attempt the pre-production step automatically as soon as the data gets collected.

Acknowledgments

This work was supported in part by the US DOE, Fermilab, the Istituto Nazionale di Fisica Nucleare (Italy), and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreements No. 101006726 (aMUSE) and No. 734303 (NEWS).

References

- [1] B. Abi et al (Muon $g - 2$ Collaboration), *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm*, Phys. Rev. Lett. **126**, 141801 (2021)
- [2] T. Aoyama et al, *The anomalous magnetic moment of the muon in the Standard Model*, Phys. Rep. **887** (2020)
- [3] L. Di Luzio et al, *New physics behind the new muon $g - 2$ puzzle?*, Phys. Lett. B **829**, 137037 (2022)
- [4] L. Cotrozzi, *Measurement of the anomalous spin precession frequency in the Muon $g - 2$ Experiment at Fermilab*, Proceedings of this conference, PoS(ICHEP2022)749
- [5] A. Driutti, *Beam Dynamics Effects in the Muon $g - 2$ Experiment*, Proceedings of this conference, PoS(ICHEP2022)053
- [6] R. Reimann, *Tracking the magnetic field in the Fermilab Muon $g - 2$ Experiment*, Proceedings of this conference, PoS(ICHEP2022)054
- [7] K.R. Herner et al, *The FIFE Project at Fermilab: Computing for Experiments*, PoS(ICHEP2016)176