Gminus2 computing requests

Gminus2

(Dated: 1 September 2019)

Introduction

The E989 collaboration goal is to measure $a_\mu = (g-2)_\mu/2$ with a precision of 140 ppb, determined by a statistical precision of 100 ppb on the muon precession frequency, ω_a, and systematics precisions of 70 ppb for ω_a and the proton precession frequency in the average magnetic field, ω_p. In terms of required computing resources and Italian involvement it is appropriate to consider only the resources related to the ω_a measurement, which requires the dominant amount of computing resources also for the entire collaboration.

The BNL experiment analyzed $N_e = 7.5{\cdot}10^9$ detected electrons and positrons with energy larger than 1.86 GeV to measure ω_a with a precision of 460 ppb, and would have required a 21× larger data sample ($N_e = 160{\cdot}10^9$) to measure ω_a to 100 ppb. Improvements on the data analysis, which have been tested on the Run 1 data sample, indicate that E989 can measure ω_a with a statistical precision of 100 ppb with 18.5× the BNL statistics.

The E989 data-taking has begun in 2018 and Table 1 reports the present status and prospects on the recorded data samples for the ω_a measurement. Run 2 showed that 1×BNL data can be collected in 25 days (4% BNL per day).

Table 1. Present and expected future data samples for the ω_a measurement.

Run	period	collected data [×BNL]	good data [×BNL]
Run 1	Apr 2018 - Jul 2018	1.9	1.0
Run 2	Apr 2019 - Jul 2019	2.2	
Run 3	Oct 2019 - May 2020	9.0	
Run 4	Oct 2020 - May 2021	6.0	

Computing model and computing requirements

The Muon $g-2$ experiment uses the MIDAS software framework for data acquisition, a package developed at PSI and TRIUMF. Raw data in MIDAS format are recorded from the E989 detector, which includes the electromagnetic calorimeter (24 modules with a total of 1296 lead fluoride crystals), the laser calorimeter gain monitoring system, the tracker (composed of 2 stations of straw chamber modules), and the beam and magnetic field monitoring subsystems.

At the front-end level, the data rate is 20 GB/s, mostly from the SIPM readout samples of the calorimeter crystals over 700 μs muon fills occurring with a 12 Hz frequency. A threshold-based algorithm ("island chopping") running on twelve GPUs substantially reduces the calorimeter data rate. The resulting raw data are stored and later processed by the offline reconstruction software (clustering, pulse template fitting, tracking) to produce reconstructed samples of comparable size as the raw data. The reconstruction is performed on the FermiGrid. The reconstructed data are processed to perform 6 ω_a analyses in parallel, producing 6 significantly smaller data samples.

In the February 2019 Fermilab Scientific Computing Portfolio Mangement Team meeting (SCPMT) Review the E989 collaboration assessed the amount of computing resources and storage that have been used to process and analyze the Run 1 dataset in 2018 and estimated the usage of computing resources for 2019, 2020 and 2021, assuming some expected improvements in CPU efficiency and data reduction, and in particular assuming that 6×BNL data would be collected and processed in 2019. No significant resources have been used so far or are planned to the used for simulation, despite our anticipations at the beginning of 2018. The expected CPU usage at Fermilab in 2019 is 21.0M core hours (wall). In 2018, a total of 15.1M core hours have been used, 13M at Fermilab and 2.1M (13.9%) off-site. Jobs submitted to the Pisa grid contributed 12.5% of the total off-site CPU in 2018 and 34.8% in 2019. If in 2019 the off-site contribution remains the same in proportion, the total expected CPU usage in 2019 is 24.4M core hours, of which 23.2M core hours are to be used to process 6×BNL of data, which corresponds to 3.9M core hours per 1×BNL of data.

According to the experience of the Italian group in the Italy-UK ω_a analysis of the Run 1 data (containing 1×BNL of good data), the reduced analysis data sample size this analysis is 10 TB. Adequate storage requirements for one analysis of 1×BNL of data should include space for two versions of analysis data (20 TB) and an additional 10 TB for work space, corresponding to a total of 30 TB per 1×BNL.

Using the above mentioned information, we summarize in Table 2 how much CPU, disk storage and tape storage is required per 1×BNL of data, with all storage and CPU improvements implemented since 2019. We follow the collaboration computing model, which estimated that all analysis activities use about the same CPU time as the reconstruction during data-taking. Since the 6 ω_a analyses dominate the analysis computing requirements, we estimate that each of them needs 1/12 of the total estimated CPU requirements by E989. We include for comparison the corresponding estimates we used in March 2018. The March 2018 estimates per 1×BNL are obtained dividing our requests by the amount of expected data at that time (3×BNL).

Table 2. CPU and disk space requirements per 1×BNL of data: present estimates compared with the estimates we did in March 2018.

estimates	activity	CPU [core hours]	disk	tape
now	raw + reco + analysis	3.90 M	330.0 TB	1100 TB
	Italian-UK ω_a analysis	0.32 M	30.0 TB	
March 2018	raw + reco (disk+tape)		470.0 TB	
	Italian analysis activities		20.0 TB	
	Italian simulation		6.7 TB	
	detector studies		6.7 TB	

Past funding requests for computing resources

Table 3 summarizes the past Gminus2 funding requests for computing. These request were related to the expected data samples to be collected in 2018 and 2019, which are also reported.

Table 3. Computing resources funds obtained for the years 2018 and 2019, based on the estimates in March 2018.

year	estimated datasets	disk per 1×BNL	disk	cost per TB	cost	note
2018	3×BNL	33.3 TB	100 TB	0.17 k€	17 k€	(used)
2019	15×BNL	16.7 TB	250 TB	0.17 k€	43 k€	(SJ)

We purchased in 2018 100 TB of disk space for the Pisa Tier 2 computing center. We have used so far only about 6 TB, for several reasons:

- E989 collected 1×BNL of good data rather than the expected 3×BNL;
- E989 did not produce significant amount of simulation;
- the Italian group collaborated to the "Europe" ω_a analysis using UK computing resources at FNAL.

More in detail, Italian and English collaborators have joined forces to perform the "Europe" ω_a analysis, and have found convenient and effective to use FNAL-located UK computing resources, which are dedicated for the UK tracked detectors studies, but have been underused so far and available for analysis jobs.

Funding requests for computing resources

In 2019 and 2020 we would like to continue to do analysis on computing resources located at FNAL, as this arrangement has several advantages:

- we used as guests the UK computing resources and found their setup convenient and effective;
- one gets fast access to the reconstructed data stored at FNAL;

- one has faster and better support for the system administration and for the required software frameworks installation and maintenance;
- we estimate that FNAL-located RAID disk storage costs 0.08 k€ per TB, less than one half the cost we paid for the disk storage we contributed to the Pisa Tier 2 site (0.17 k€).

We plan to continue to use the Pisa Tier 2 grid for data reconstruction and possibly simulation in an opportunistic way, without additional funds other than the ones we already used to buy disk storage.

Therefore,

- we drop our original request for 2019 to purchase 250 TB of disk space for the Pisa Tier 2 site, for which 43 k€ s.j. funds are presently reserved;
- we ask to use part of the above s.j. funds (16 k€, see below) to buy computing resources for doing analysis at FNAL.

We need to buy resources for analysis at FNAL since the UK workstation that we have used so far as guests will be mostly reserved for detector studies starting with Run 3.

The UK computing facilities that we used for analysis consist of:

- a workstation Dell Precision T7810 with
  - 2× Xeon E5-2630 v4, each CPU has 10 cores, each core has 2 threads and is rated 10.3 HS06
  - 64 GB RAM
  - 1 year of operation delivers ∼1.8M core hours ≃ 1.8 kHS06.years
- a RAID disk storage server composed of:
  - Synology RackStation RS2818RP+ 16-Bay NAS Server with 4×1Gb-ethernet interfaces
  - 16×6 TB disks

For 2019, we estimate that 100 TB of disk storage are appropriate to analyze the data collected so far (Run 1 + Run 2, about 3×BNL, requiring ∼3×30 TB), for a total request of 16kE, as reported in the first row of Table 4.

In 2020 we expect to analyze 9×BNL of data collected in Run 3. Our original estimate (done in July 2019) was based on the cost of 300 TB for storage in the Pisa Tier 2 site in 2018 (0.17 k€/TB). Based on the considerations so far, we would like to change our request to buy 300 TB (≈ 9×30 TB) of RAID disk space at FNAL. Since there are no detailed plans on simulation production, we request for 2020 150 TB of disk storage for simulation, to be updated when we will have more reliable estimates. Therefore our request for 2020 is 24 k€ for data and 12 k€ for simulation.

We summarize our requests in Table 4.

Table 4. Funding requests for computing resources.

year	requests	notes
2019	16 k€	from the release of part of the 43 k€ s.j. funds
		- 8 k€ for ∼250 HS06 workstation at FNAL
		- 8 k€ for 100 TB RAID disk space at FNAL, to analyze 3×BNL
2020	36 k€	
		- 24 k€ for 300 TB RAID disk space at FNAL, to analyze 9×BNL
		- 12 k€ for 150 TB RAID disk space at FNAL, for simulation production