Data Acquisition for the New Muon g-2 Experiment at Fermilab

Wesley Gohn

University of Kentucky

July 15, 2015

Outline

1. Introduction
   - Physics of Muon g-2
   - Project Status
   - Detectors and Backend Electronics
2. Data Acquisition System
   - System Requirements
   - MIDAS
   - GPU Processing
   - Prototyping
   - Tracker
   - Event Display
3. Conclusion

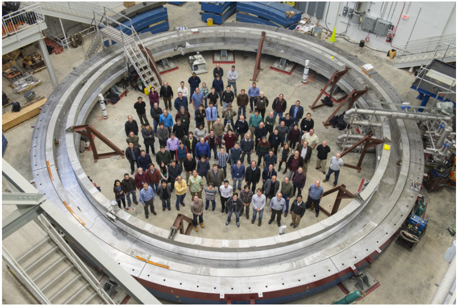

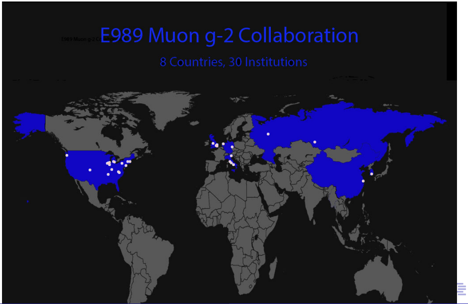

Physics of Muon g-2

- In the standard model, the muon is a spin 1/2 pointlike particle.
- It has a magnetic dipole moment of $\vec{\mu} = g\frac{q}{2m}\vec{s}$, with $g = 2$ for a pointlike particle (Dirac)
- Additional effects from QED, electroweak theory, and hadronic factors move the standard model prediction of g away from 2. It has become customary to measure this discrepancy, g-2.
- If a discrepancy with the standard model value is found, beyond standard model contributions to g-2 could come from SUSY, dark photons, or other new physics (NP).

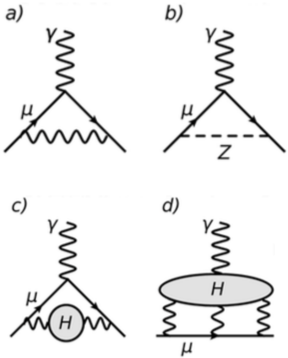

Example diagrams that contribute to $g-2$

- a: Leading-order QED Schwinger term
- b: Electroweak Z exchange diagram
- c: Lowest-order hadronic vacuum polarization
- d: Hadronic light-by-light scattering.

$$a_\mu = a_\mu^{QED} + a_\mu^{EW} + a_\mu^{QCD} + a_\mu^{NP}$$

Measurements of g-2

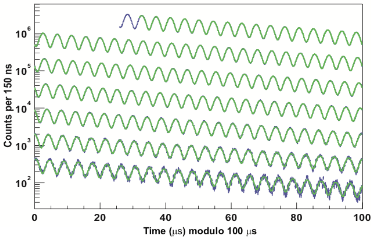

$$a_\mu \equiv \frac{g-2}{2}$$

$$\vec{\omega}_a = -\frac{Qe}{m}[a_\mu \vec{B} - (a_\mu - (\frac{mc}{p})^2)\frac{\vec{\beta}\times\vec{E}}{c}]$$

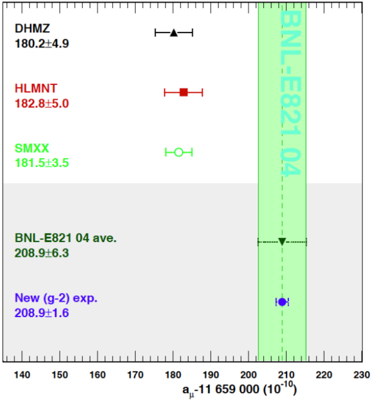

- BNL E821 measured g-2 to have a 3.3σ discrepancy from the standard model (2006).
- Fermilab E989 will measure 20 times the number of muons, reducing the uncertainty on this measurement by a factor of 4.
- Without theory improvements, discrepancy could reach $> 5\sigma$.

Measurements of g-2

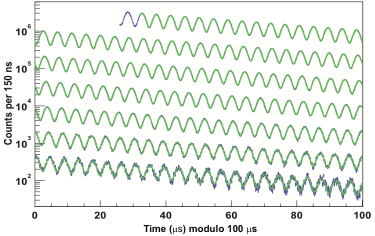

$$a_\mu \equiv \frac{g-2}{2}$$

$$\vec{\omega}_a = -\frac{Qe}{m}\left[a_\mu \vec{B} - \left(a_\mu - \left(\frac{mc}{p}\right)^2\right)\frac{\vec{\beta}\times\vec{E}}{c}\right]$$

- BNL E821 measured g-2 to have a 3.3σ discrepancy from the standard model (2006).
- Fermilab E989 will measure 20 times the number of muons, reducing the uncertainty on this measurement by a factor of 4.
- Without theory improvements, discrepancy could reach $> 5\sigma$.

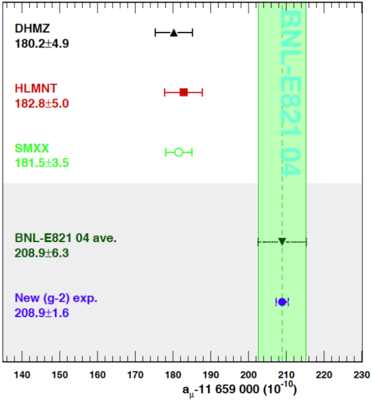

Project Status

- The ring was moved from BNL to FNAL in 2013.
- It has been installed in our new MC1 building and cooled to 5K. Currently we are working to bring the magnetic field to full strength.
- Plan for data taking to begin in early 2017.

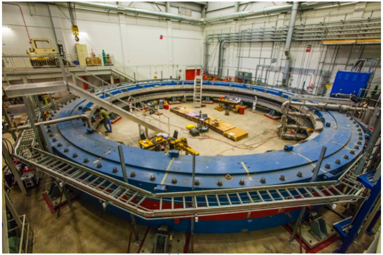

Detectors and Backend Electronics

- Measurement will utilize 24 calorimeters (each composed of 54 PbF_2 crystals read out by SiPMs), 3 straw trackers, and several auxiliary detectors.
- Each calorimeter will be readout by a custom WFD in a μTCA crate with an AMC13 control module controlled via IPBus.

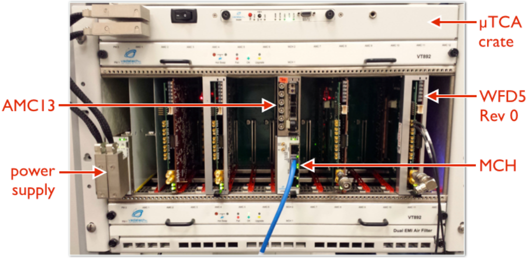

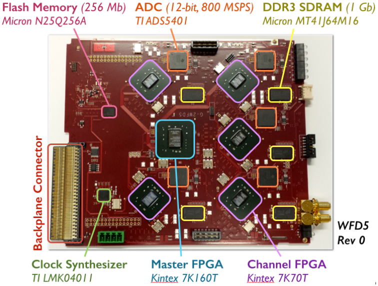

Images courtesy of David Sweigart

Requirements

- Accomodate a 12 Hz average rate of muon spills that consists of sequences of four successive 700 μs spills with 11 ms spill-separations
- Handle the readout, processing, monitoring and storage of the data obtained from the twenty-four electromagnetic calorimeters, each comprising 9×6 arrays of PbF_2 crystals read out by SiPMs.
- The signals derived from individual crystals are read out by 1296 channels of custom 800 MHz, 12-bit, waveform digitizers.
- Provide both the readout of the raw ADC samples and the derivation of T-method, Q-method, and other calibration, diagnostic and systematic datasets.
- For a 12 Hz spill rate the time-averaged rate of raw ADC samples is 18 GB/s in total.

DAQ Schematic

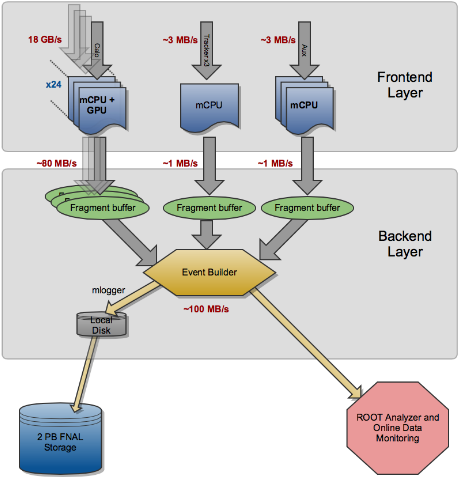

- Layered array of commodity, networked processors
- FE layer for readout of digitizer (calo), MHTDCs (straws)
- BE layer for assembly of event fragments, storage
- Slow control layer for setting, monitoring of HVs, etc.
- Online analysis layer using ROME for monitoring the integrity of raw data, physics data.

MIDAS

- Much experience with MIDAS in the collaboration from MuLan, MuCap and MuSun.
- Frontend acquisition code written in C/C++ with IPBus libraries for communication with μTCA electronics and CUDA libraries for data processing in GPUs.
- ROME analyzer for online data monitoring.
- Data will be written to tape as MIDAS datafiles, and later converted to the Fermilab *art* framework.
- MSCB will be used for slow controls.

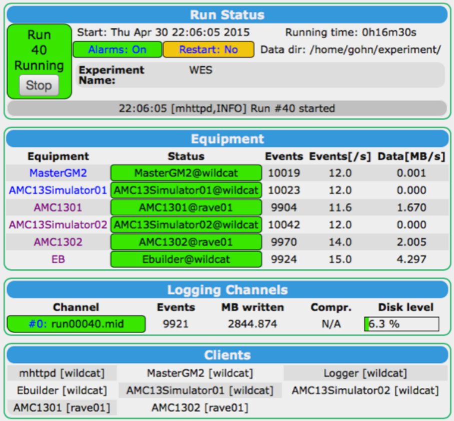

DAQ Frontends

MIDAS frontends and applications currently in use:

- **MasterGM2:** Receives trigger signal and controls other frontends via remote procedure calls.
- **CaloReadoutAMC13:** Reads data from AMC13 via 10 GbE, processes data in GPU using CUDA routines for T and Q methods. Also configures AMC13 via MIDAS online database and IPBus.
- **RiderTrigger:** Sends IPBus trigger to AMC13, which is relayed to Rider via TTC signal.
- **AMC13Simulator:** Simulates AMC13 data format and sends data via TCP/IP to CaloReadoutAMC13.
- **EventBuilder**: Combines data fragments from multiple buffers into MIDAS events.
- **MLogger:** Writes data to midas file.
- **Analyzer:** Reads data from TCP socket or from MIDAS file and creates ROOT histograms (to be replaced by ROME analyzer).

AMC13 Readout Frontend

- Collects data asynchronously from waveform digitizers and AMC13.
- Configures AMC13 and WFDs using MIDAS ODB and IPBus.
- Mutual Exclusion locks are used to break CPU processing into three threads.
- CUDA is used for multithreading of data processing in GPU.

Multithreading with mutex-locks.

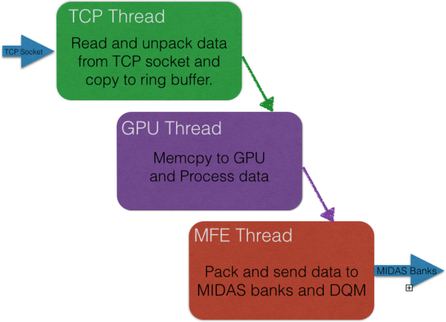

GPU Processing

- Data will be processed in an array of 24 GPUs (One GPU per calorimeter)
- Utilizing NVIDIA TESLA K40 GPUS
  - Peak double precision floating point performance: 1.43 Tflops
  - Peak single precision floating point performance: 4.29 Tflops
  - Memory bandwidth 288 GB/sec
  - Memory size (GDDR5): 12 GB
  - CUDA cores: 2880
- Data processing code is written using CUDA libraries in C++.

Results of bandwidth tests:

PCIe Version	GPU	Host to device, Pageable	Host to device, Pinned
2	K20	3326.6 MB/s	5028.3 MB/s
3	K20	5628.6 MB/s	6003.6 MB/s
3	K40	6647.8 MB/s	10044.3 MB/s

T and Q Methods

CUDA routines process data with two complimentary methods.

- T-method
  - Positron events in the calorimeter are individually identified, sorted and fit to obtain time and energy.
  - All events above an energy threshold are included.
  - $\vec{\omega}_a$ is determined from a fit to a pileup-subtracted histogram.
  - This was the method used in BNL E821.
- Q-method
  - Individual positron events are not identified.
  - Detector current is integrated as a proxy for event energy.
  - No pileup correction is necessary!

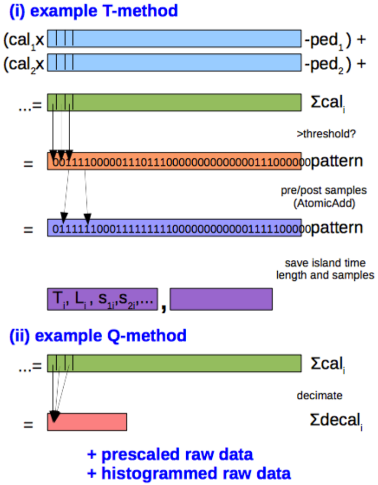

Test Stands

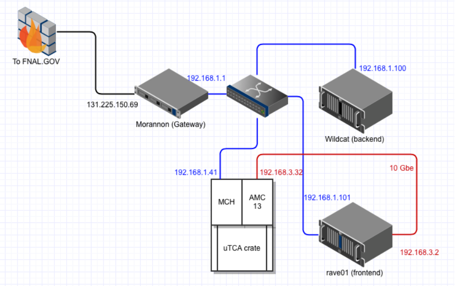

- Prototype of the DAQ is operational at Fermilab.
- Currently includes backend, frontend, gateway, and μTCA crate with WFD and AMC13
- Three more frontend machines have been purchased and will be installed in the next few weeks.

WFD readout

- We have been performing tests by reading digitized noise from a preliminary version of our 5-channel WFDs.
- Triggers are sent via TTC signal via the AMC13.
- We hope to have ≈5 WFDs for testing later this Summer, and a full crate of 12 by the end of the year.

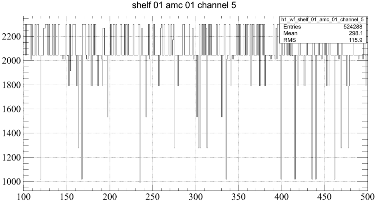

AMC13 Simulator

- Generates realistic waveforms and packs the data in the AMC13 data format.
- Allows us to exersize the DAQ without dependence on hardware.
- Plan to develop this into a tool that will recreate the full spectrum of DAQ input, which will be used for testing the complete data acquisition system.

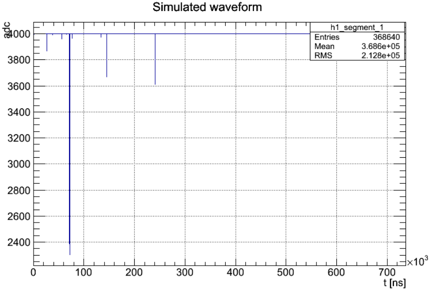

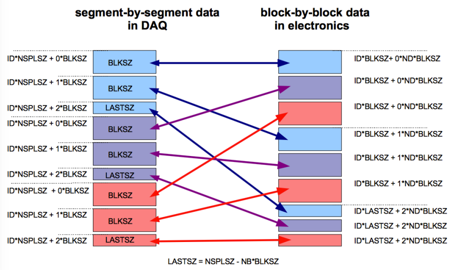

Event building test

MIDAS experiment "UKY"	Wed May 29 20:08:06 2013 Refr:60

Stop | ODB | Messages | Alarms | Programs | Config

RunLog | Logbook | Elog | Doc

Run #2071 | Running | Alarms: On | Restart: No | Data dir: /data/UKY/mid

Start: Wed May 29 20:07:16 2013 | Running time: 0h00m50s

Equipment	Status	Events	Events[/s]	Data[MB/s]
MagicBox	magic_box@mb	0	0.0	0.000
VMEcrate	(frontend stopped)	0	0.0	0.000
masterMT	(frontend stopped)	365	0.0	0.000
EB	Ebuilder@be	0	0.0	0.000
ATS9870	(frontend stopped)	0	0.0	0.000
EMC	(frontend stopped)	5	0.0	0.000
master	master@fe02	574	11.9	0.001
FakeCalo01	(frontend stopped)	0	0.0	0.000
FakeData01	FakeData01@fe01	564	12.0	2.058
FakeData02	FakeData02@fe01	578	12.0	2.057
FakeData03	FakeData03@fe01	555	12.0	2.053
FakeData04	FakeData04@fe01	566	12.0	2.059
FakeData05	FakeData05@fe01	575	11.7	2.002
FakeData06	FakeData06@fe01	551	12.0	2.058
FakeData07	FakeData07@fe01	564	12.0	2.059
FakeData08	FakeData08@fe01	576	12.0	2.059
FakeData09	FakeData09@fe01	551	11.6	1.999
FakeData10	FakeData10@fe01	563	12.0	2.059
FakeData11	FakeData11@fe01	573	12.0	2.059
FakeData12	FakeData12@fe01	551	11.9	2.052
FakeData13	FakeData13@fe01	561	12.0	2.058
FakeData14	FakeData14@fe01	571	12.0	2.057
FakeData15	FakeData15@fe01	547	12.0	2.060
FakeData16	FakeData16@fe01	558	12.0	2.059
FakeData17	FakeData17@fe01	570	12.0	2.057
FakeData18	FakeData18@fe01	544	12.0	2.057
FakeData19	FakeData19@fe01	555	12.0	2.057
FakeData20	FakeData20@fe01	567	11.6	1.997
FakeData21	FakeData21@fe01	578	12.0	2.059
FakeData22	FakeData22@fe01	555	12.0	2.060
FakeData23	FakeData23@fe01	567	12.0	2.059
FakeData24	FakeData24@fe01	578	12.0	2.059
FakeCaloNewQ01	(frontend stopped)	0	0.0	0.000
CaloSimulatorTCPIP01	(frontend stopped)	0	0.0	0.000
CaloReadoutTCPIP01	(frontend stopped)	0	0.0	0.000

Channel	Events	MB written	Compression	Disk level
#0: run02071.mid	579	99.383	N/A	63.9 %

20:07:17[mtransition,INFO] Run #2071 started

```
---------------------- Event# 1 ------------------------------
Evid:0001- Mask:0000- Serial:0- Time:0x512b8c40- Dsize:2881352/0x2bf748
#banks:48 - Bank list:-FD01SR01FD02SR02FD03SR03FD04SR04FD05SR05FD06SR06FD07SR07FD08SR08FD09SR09FD10SR10FD11S
R11FD12SR12FD13SR13FD14SR14FD15SR15FD16SR16FD17SR17FD18SR18FD19SR19FD20SR20FD21SR21FD22SR22FD23SR23FD24SR24-
```

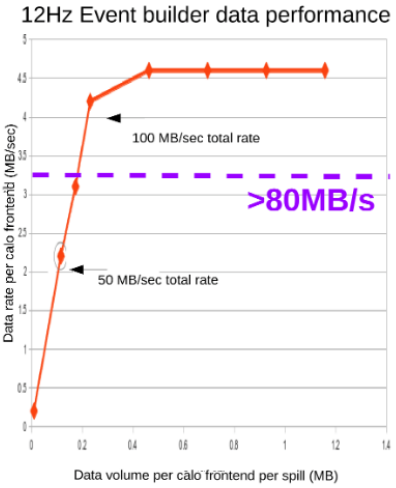

GPU Processing Time

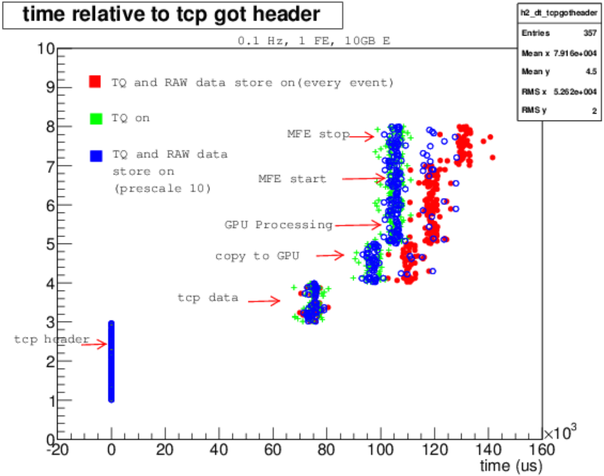

Comparison of processing time needed with processed vs raw data.

Dual GPU Test

Question: Can we run two MIDAS frontends per frontend machine?

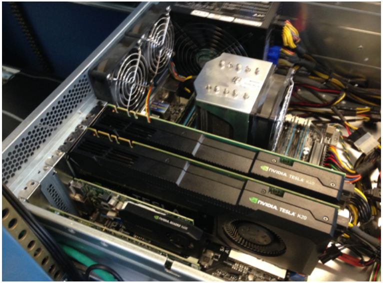

This was tested in three steps:

- Run two AMC13Simulator - AMC13Readout pairs on one machine over localhost
- Run two AMC13Simulator on backend, send data to AMC13Readout over dual 10 GbE.
- Process data from both AMC13s on 2 GPUs.

Running 2 Frontends Per Machine

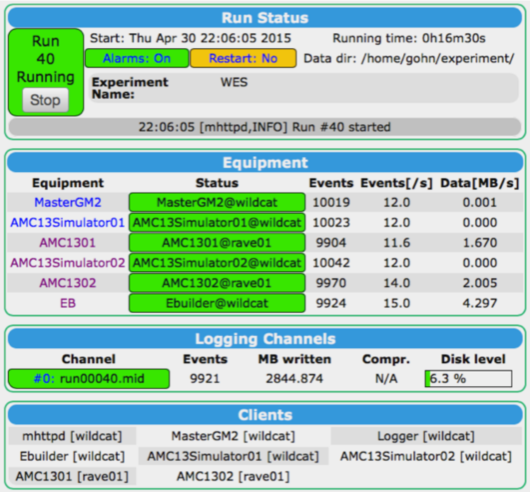

Run Status

Run 40 Running — Stop

Start: Thu Apr 30 22:06:05 2015 — Running time: 0h16m30s

Alarms: On — Restart: No — Data dir: /home/gohn/experiment/

Experiment Name: WES

22:06:05 [mhttpd,INFO] Run #40 started

Equipment

Equipment	Status	Events	Events[/s]	Data[MB/s]
MasterGM2	MasterGM2@wildcat	10019	12.0	0.001
AMC13Simulator01	AMC13Simulator01@wildcat	10023	12.0	0.000
AMC1301	AMC1301@rave01	9904	11.6	1.670
AMC13Simulator02	AMC13Simulator02@wildcat	10042	12.0	0.000
AMC1302	AMC1302@rave01	9970	14.0	2.005
EB	Ebuilder@wildcat	9924	15.0	4.297

Logging Channels

Channel	Events	MB written	Compr.	Disk level
#0: run00040.mid	9921	2844.874	N/A	6.3 %

Clients

mhttpd [wildcat]	MasterGM2 [wildcat]	Logger [wildcat]
Ebuilder [wildcat]	AMC13Simulator01 [wildcat]	AMC13Simulator02 [wildcat]
AMC1301 [rave01]	AMC1302 [rave01]	

Answer: YES

- Test completed at full rate over 10 GbE at MC-1.
- TCP/IP tuned to achieve maximum rate.
- Data were processed using the dual GPUs in the MC-1 frontend machine.

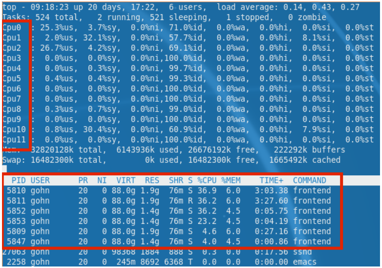

```
top - 09:18:23 up 20 days, 17:22,  6 users,  load average: 0.14, 0.43, 0.27
Tasks: 524 total,   2 running, 521 sleeping,   1 stopped,   0 zombie
Cpu0  : 25.3%us,  3.7%sy,  0.0%ni, 71.0%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu1  :  2.0%us, 32.1%sy,  0.0%ni, 57.7%id,  0.0%wa,  0.0%hi,  8.1%si,  0.0%st
Cpu2  : 26.7%us,  4.2%sy,  0.0%ni, 69.1%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu3  :  0.0%us,  0.0%sy,  0.0%ni,100.0%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu4  :  0.0%us,  0.3%sy,  0.0%ni, 99.7%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu5  :  0.4%us,  0.4%sy,  0.0%ni, 99.3%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu6  :  0.0%us,  0.0%sy,  0.0%ni,100.0%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu7  :  0.0%us,  0.0%sy,  0.0%ni,100.0%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu8  :  0.3%us,  0.7%sy,  0.0%ni, 99.0%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu9  :  0.0%us,  0.0%sy,  0.0%ni,100.0%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Cpu10 :  0.8%us, 30.4%sy,  0.0%ni, 60.9%id,  0.0%wa,  0.0%hi,  7.9%si,  0.0%st
Cpu11 :  0.0%us,  0.0%sy,  0.0%ni,100.0%id,  0.0%wa,  0.0%hi,  0.0%si,  0.0%st
Mem:  32820128k total,  6143936k used, 26676192k free,   222292k buffers
Swap: 16482300k total,        0k used, 16482300k free,  1665492k cached

  PID USER      PR  NI  VIRT  RES  SHR S %CPU %MEM    TIME+  COMMAND
 5810 gohn      20   0 88.0g 1.9g  76m S 36.9  6.0   3:03.38 frontend
 5811 gohn      20   0 88.0g 1.9g  76m R 36.2  6.0   3:27.60 frontend
 5852 gohn      20   0 88.0g 1.4g  76m S 36.2  4.5   0:05.75 frontend
 5853 gohn      20   0 88.0g 1.4g  76m S 23.2  4.5   0:04.19 frontend
 5809 gohn      20   0 88.0g 1.9g  76m S  4.6  6.0   0:27.16 frontend
 5847 gohn      20   0 88.0g 1.4g  76m S  4.0  4.5   0:00.86 frontend
27063 gohn      20   0 98368 1884  888 S  0.3  0.0   0:17.56 sshd
 2258 gohn      20   0  245m 8692 6368 T  0.0  0.0   0:00.00 emacs
```

Tracker Test Beam

The tracker DAQ just finished the first end-to-end MIDAS to *art* to processed data run during the test beam in June.

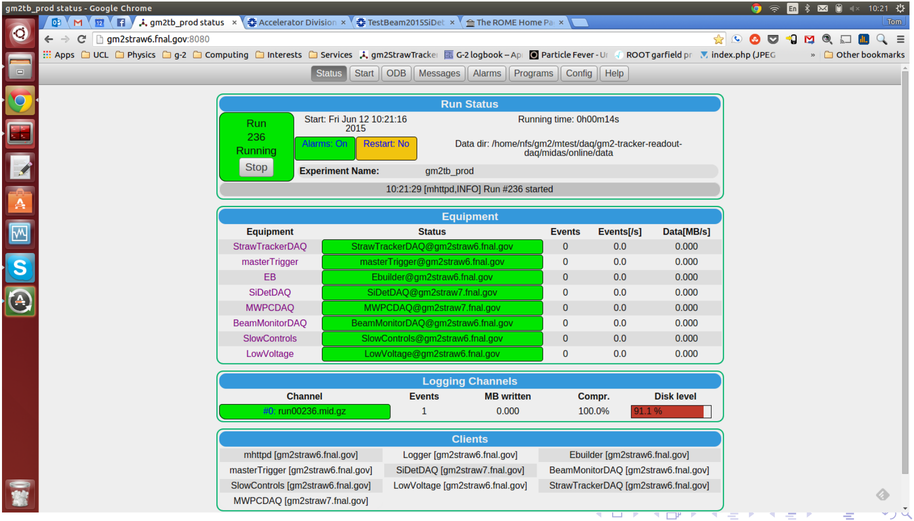

Tracker Test Beam

The tracker DAQ just finished the first end-to-end MIDAS to *art* to processed data run during the test beam in June.

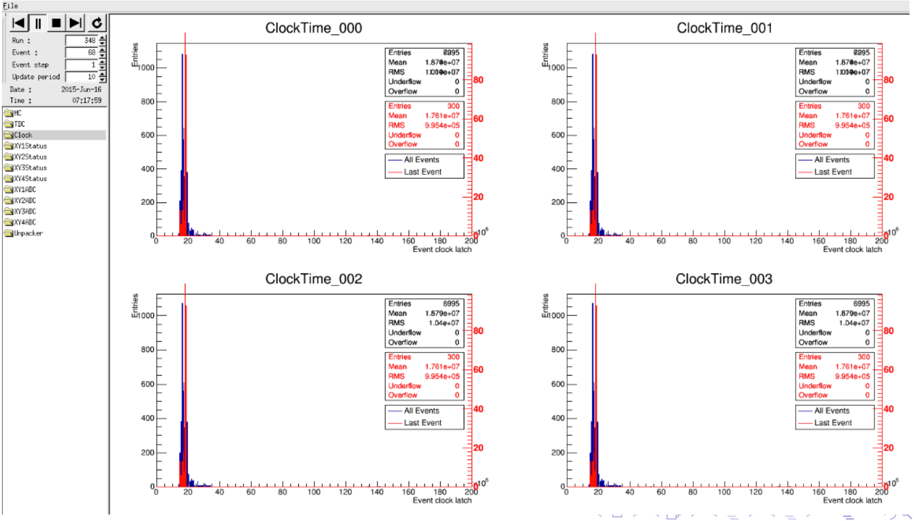

Tracker Test Beam

The tracker DAQ just finished the first end-to-end MIDAS to *art* to processed data run during the test beam in June.

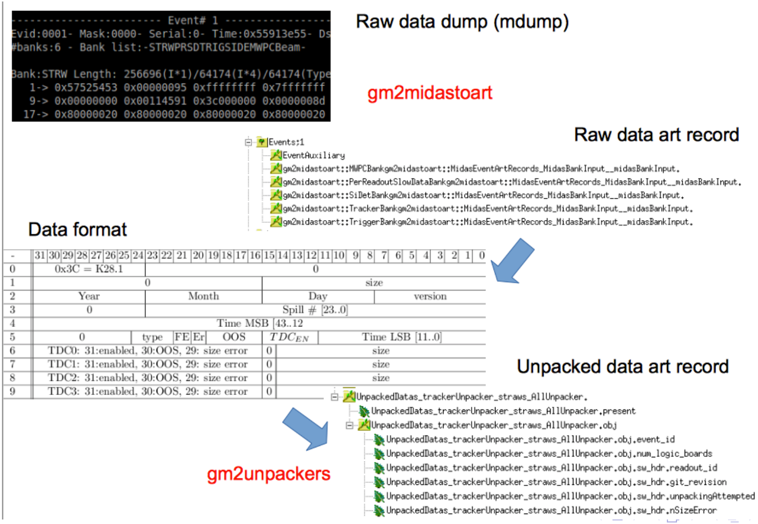

-	bits 31–24	bits 23–16	bits 15–8	bits 7–0
0	0x3C = K28.1	0		
1	0		size	
2	Year	Month	Day	version
3	0	Spill # [23..0]		
4	Time MSB [43..12			
5	0 / type / FE / Er / OOS		TDC_{EN} / Time LSB [11..0]	
6	TDC0: 31:enabled, 30:OOS, 29: size error		0 / size	
7	TDC1: 31:enabled, 30:OOS, 29: size error		0 / size	
8	TDC2: 31:enabled, 30:OOS, 29: size error		0 / size	
9	TDC3: 31:enabled, 30:OOS, 29: size error		0	

Field DAQ

- The magnetic field measurement uses an independent MIDAS implementation.
- Currently in use for magnetic field shimming.
- The field is measured during shimming with 25 NMR probes.
- Laser tracker gives position of magnetic field measuring trolly.

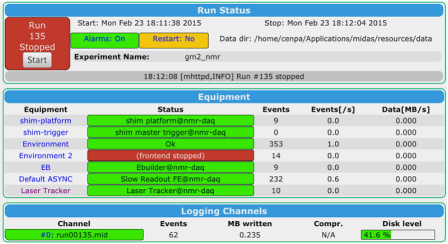

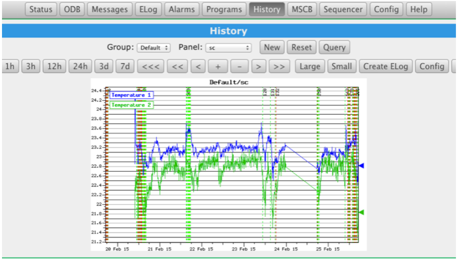

Slow Controls

- Slow Controls will use MSCB with SCS2000 modules.
- Must write to PostgreSQL database

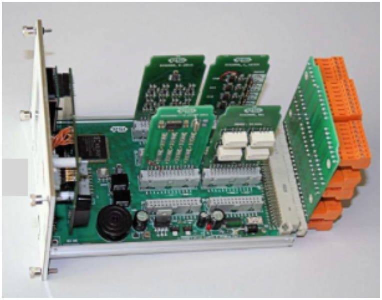

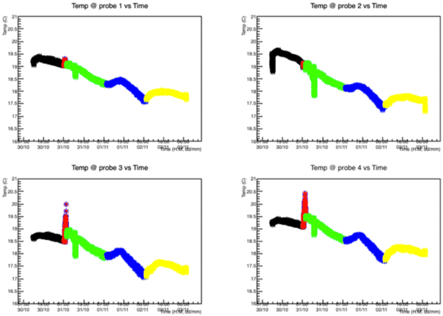

Event Display in Paraview

Planned supplement to ROME

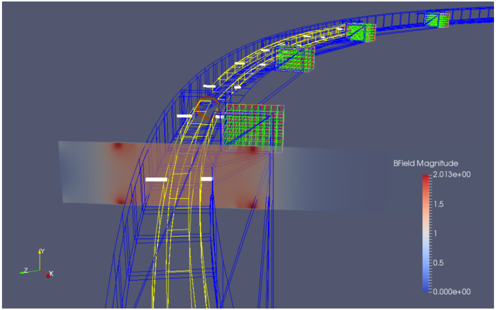

Simulation in Ring

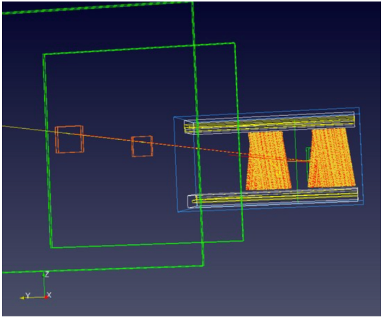

Offline Tracker data

Conclusion

- The new muon g-2 experiment will run at Fermilab beginning in 2017 with the goal of reaching $20\times$ the BNL statistics.
- A new state-of-the-art data acquisition system utilizing parallel data processing in a hybrid system of multi-core CPUs and GPUs is required to achieve the necessary data rates.
- The DAQ will acquire data from 1296 channels of custom μTCA waveform digitizers, as well as straw trackers and auxiliary detectors at a rate of 18 GB/s.
- Prototyping of the DAQ is underway, and construction will be complete by mid-2016.