



Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

Mu2e Computing Status and Outlook


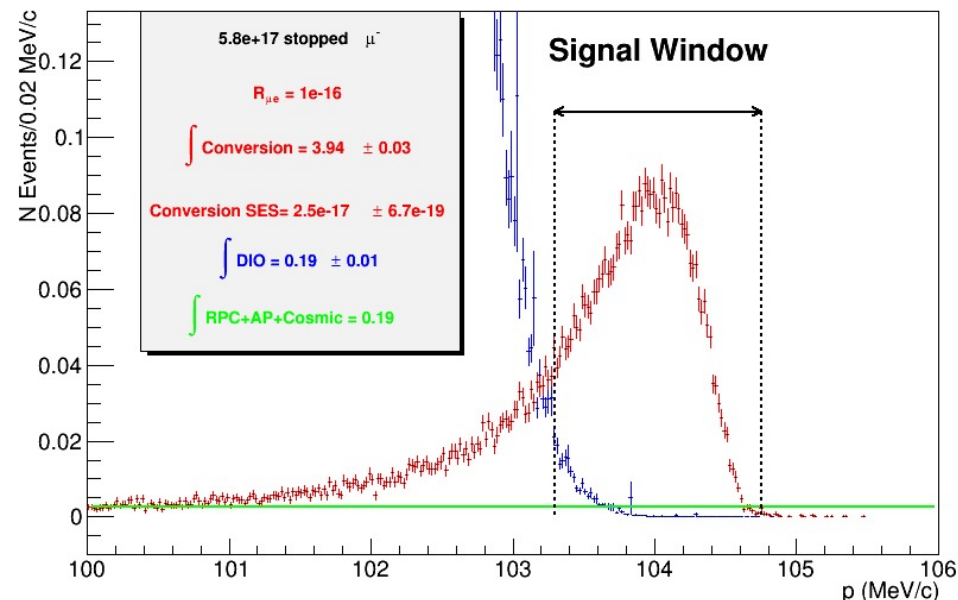
Rob Kutschke, SCD

FIFE Meeting


April 28, 2022

Mu2e in One Page

- Make muonic Al.
- Watch it decay:
 - Decay-in-orbit (DIO): 40%
 - Continuous E_e spectrum.
 - Muon capture on nucleus: 60%
 - Nuclear breakup: p, n, γ
 - Neutrino-less μ to e conversion
 - Mono-energetic $E_e \approx 105$ MeV
 - At endpoint of continuous spectrum.
- Measure E_e spectrum.
- Is there an excess at the endpoint?
- Quantitatively understand backgrounds



The Mu2e Detector



- Any observation is a clear signal of new physics
 - Standard model background $\sim \times 10^{40}$ less than our sensitivity
- Sensitivity $\sim 10^4$ times better than previous best experiment
- Sensitive to mass scales beyond direct production at LHC

Located on the Muon Campus

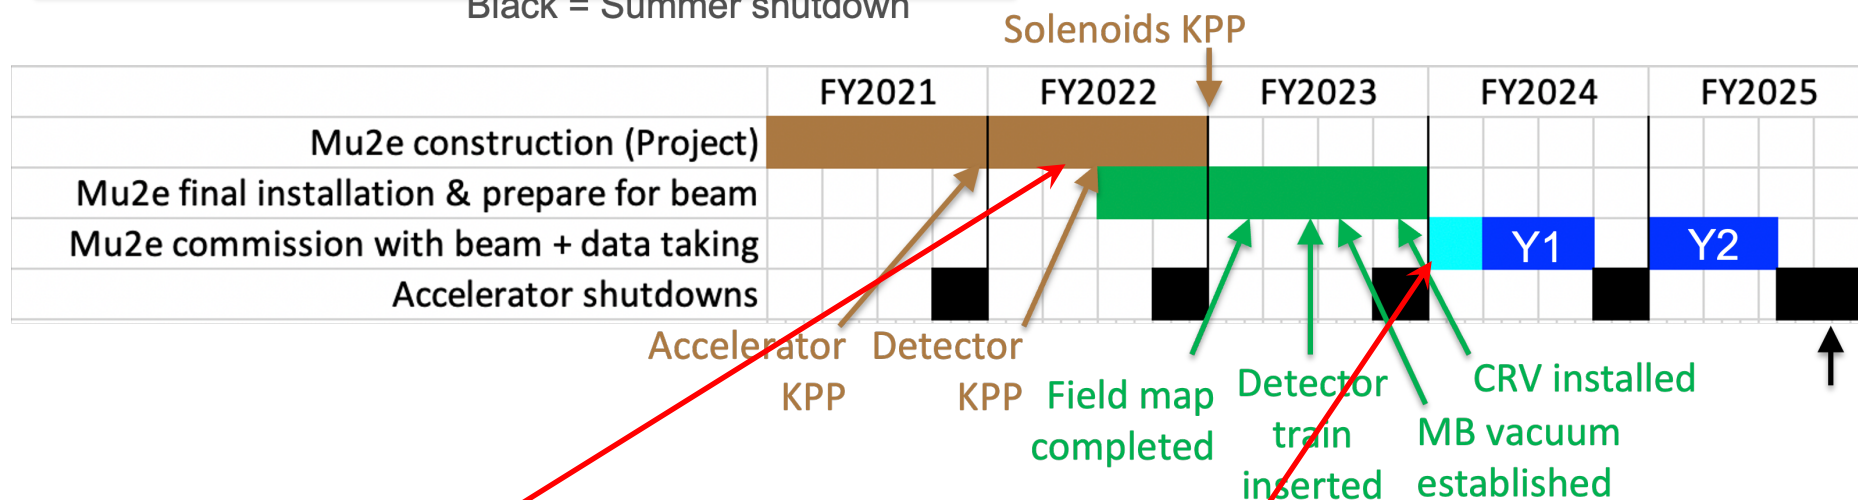
Mu2e
Hall



Pre-COVID Schedule (May 2020)

Key:

Brown = Project	Light blue = Commission with Beam
Green = Prep for beam	Dark blue = Physics Data Taking
Black = Summer shutdown	



- Routine cosmic ray data taking
- Q2 FY22
- First moderate volume data
 - Rough estimate: O(a few TB/day)

- Commissioning with beam
- Q1 FY24
- First high rate data
- Ramping up to 7 PB/year

Updated Schedule

- Re-baseline review:

June 22 - 23	DOE Review
May 10 - 12	Director's Review

- I am not allowed to discuss publicly the revised schedule until it is approved by DOE.

Computing and Software Status - 1

- Mature code base for
 - Simulations, including modeling of pileup
 - Trigger, reconstruction and physics analysis of events with perfect alignment and calibration
 - Studies of many types of backgrounds
- Several event displays including an REve based version.
- Code hosted in GitHub
 - Contributions via PR; CI using Jenkins triggered by GitHub comments
- Jenkins used for building production releases
- Nightly integration tests run on grid
- Ramping up code sanitizing
- Prototype DQM system
- Prototype system now in DAQ room; includes disk buffer.
- Reference documentation on wiki; new tutorials in development

Computing and Software Status - 2

- Support low-rate data taking from Vertical Slice Tests, test stands and sub-assembly testing
 - Intermittent. Peaks to 10's of GB/day.
 - Store on tape; FTS (old Fermilab FTS)
 - Analyzed interactively or by grid jobs (by users)
 - First conditions data
- Mock Data Challenge MDC2020
 - Simulation campaign to produce data sets with some of the diseases of real data: mis-alignment and mis-calibration
 - Will be used to learn how to calibrate and align
 - Our first project using POMS
- Infrastructure to support the above

What's Next

- Continue support for VST, test stands etc
- Continue MDC2020
 - Ramp up efforts on alignment and calibration
- Develop infrastructure to
 - Move raw data, logs etc from Mu2e Hall to long term storage
 - Process raw data as it arrives
 - Mirror curated information from online DBs to offline DBs
 - Run configuration; run status; slow control (EPICS).
 - Move more grid workflows to POMS.
 - Some will be complex and we may need ask for new features.
- When do we move from SAM to RUCIO/Metacat?
 - We need Data Dispatcher to use POMS
 - Last we heard was mid 2023. Is this still on track?

Moving Data from Hall to Long Term Storage

- Steady state
 - Strawman: 20 TB/day; 20k files/day
 - Peaks to 2x this (draining a backlog)
- This plus DQM, reco, calibration, simulation, analysis
 - O(50-100k) new SAM files per day; several operations per file.
- Prioritize which files are moved first
 - First offline jobs do nearline DQM
- We require an end-to-end checksum
 - Risk: limited disk bandwidth in counting room
- Manage free space on counting room disks and persistent dCache
- Risk: may need to split/join/reshape some files before moving to tape?
- Counting room disks are in-hand.

An Related FTS Question

- Part of [RITM1388520](#).
- The way we have used FTS requires two copy operations
 - Copy from online to FTS dropbox
 - Copy from dropbox to tapebacked dCache
- Is there a way to avoid the double copy?
 - What do other experiments do?

Other issues - 1

- dCache continues to be the primary culprit for job reliability
 - About once per month some of the nightly validation fails
 - In a recent campaign 5% to 10% of jobs failed
 - Automatically resubmitted by POMS
 - Net >99% complete OK without human intervention
- Jobsub management of python versions is fragile
 - setup jobsub
 - setup a Mu2e environment
 - Then jobsub often breaks due to python version issues
- We would appreciate improvements to fair use of tape drives
 - Requesting a handful of files from one tape can take days
 - Users are accessing pre-staged files to keep them in cache

Summary

- Well established base for Software and Computing
 - Critical contributions during the CD-1 ... CD-3
- We have already started some discussions about moving raw data to long term storage, data base mirroring etc
 - Ray Culbertson is leading a lot of this
- As our requirements firm-up we will be open new discussions
- I will remind our users to follow up on open tickets
- Once our new schedule is public (guess ~July this year) we can set milestones for having our infrastructure –in-place and tested before it is needed.