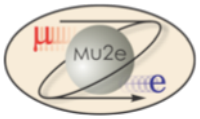

Mu2e Requirements

Rob Kutschke

November 28, 2012

SCD Projects Meeting #5

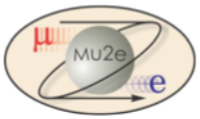

Outline

- Mu2e Overview
- Cost and Schedule
- Trigger and DAQ
- Offline Software
- Hardware Needs

The Mu2e Apparatus

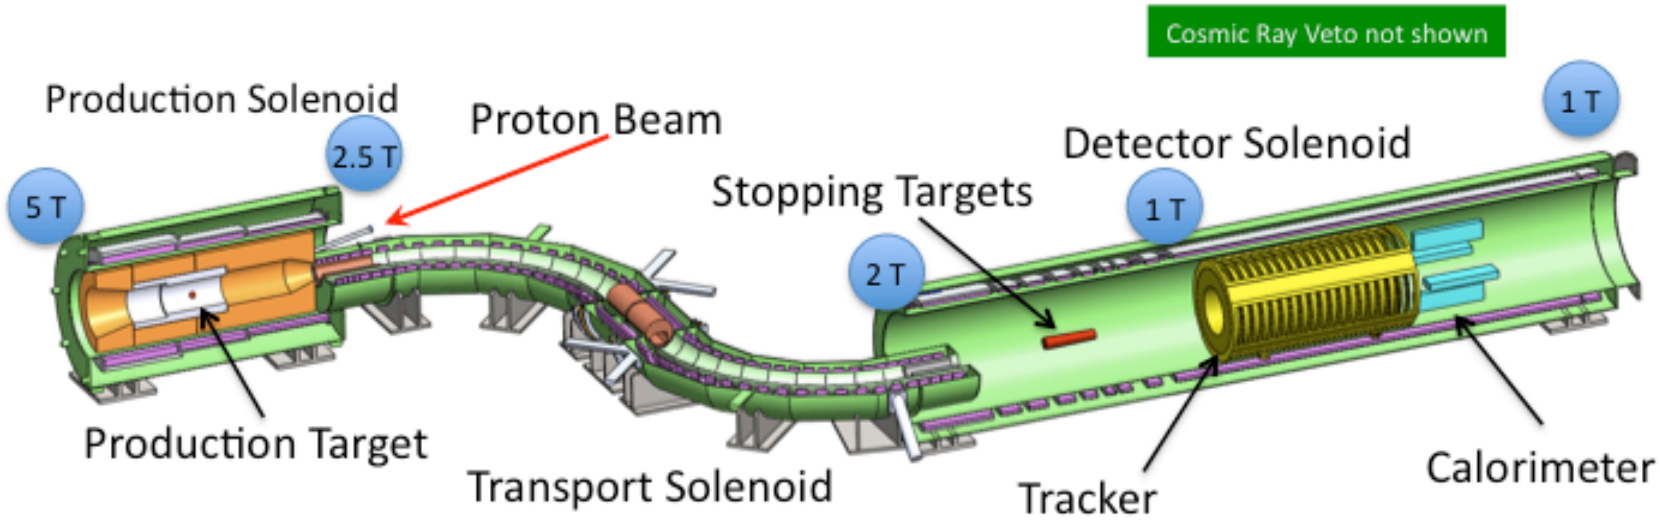

Comments on Previous Slide

- Proton micro bunch arrives every 1695 ns (590 kHz).
  - Slow spilled from Delivery Ring.
- Collect backscattered pions, which decay to muons
  - Muon beam is low momentum <p> ≈ 50 MeV
  - Lots of junk accompanies the muons
- S-bend solenoid
  - Separates charges to make μ^- beam
- Most muons range out in 17 thin foils stopping targets
  - These form muonic atoms, lifetime ≈ 864 ns
- Wait for prompt backgrounds to decay
- Measure the electron energy spectrum from the decay of the muonic atoms.

Muon To Electron Conversion

- Initial State
  - A muonic atom, at rest in the lab frame.
- Final state:
  - Electron recoiling against the intact nucleus
  - We do not detect the nucleus
  - Electron is mono energetic
  - Electron energy depends on nucleus
  - For Al:
    - E ≈ 104.97 MeV
    - Lifetime of muonic Al ~ 864 ns

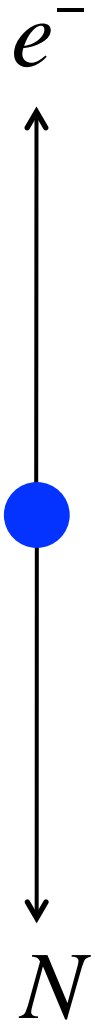

Cartoon of the Experiment - 1

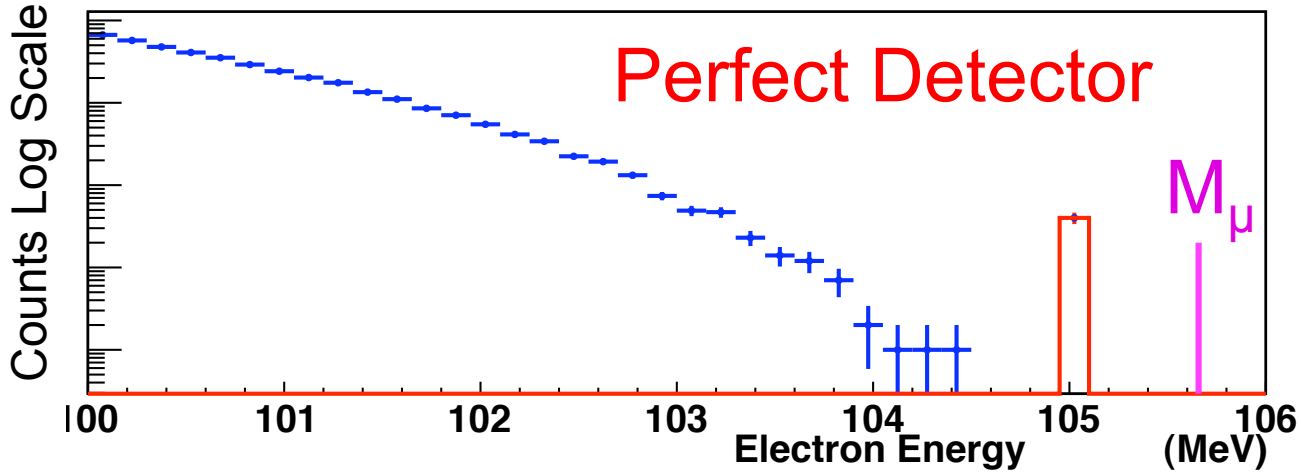

- Red box:
  - The mono-energetic signal for conversion
- Blue points:
  - Electrons from muonic Al decay in orbit
  - An irreducible background
  - Endpoint of this spectrum is at the conversion energy
  - Spectrum dropping fast: $dN/dE \approx (E_{max} - E)^5$

Cartoon of the Experiment - 2

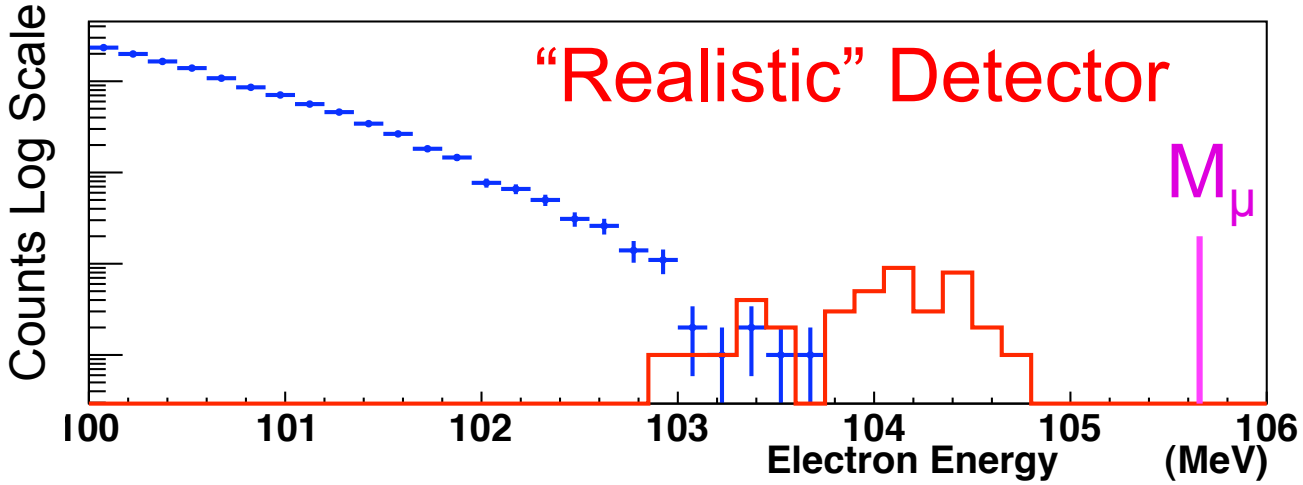

- Interactions in targets material and detector materials
  - Lower the energy of each electron
  - Smear the energy (mostly to lower E).
- Not shown: other backgrounds that vary slowly over this momentum range but << 1 event in signal region.

Cartoon of the Experiment - 3

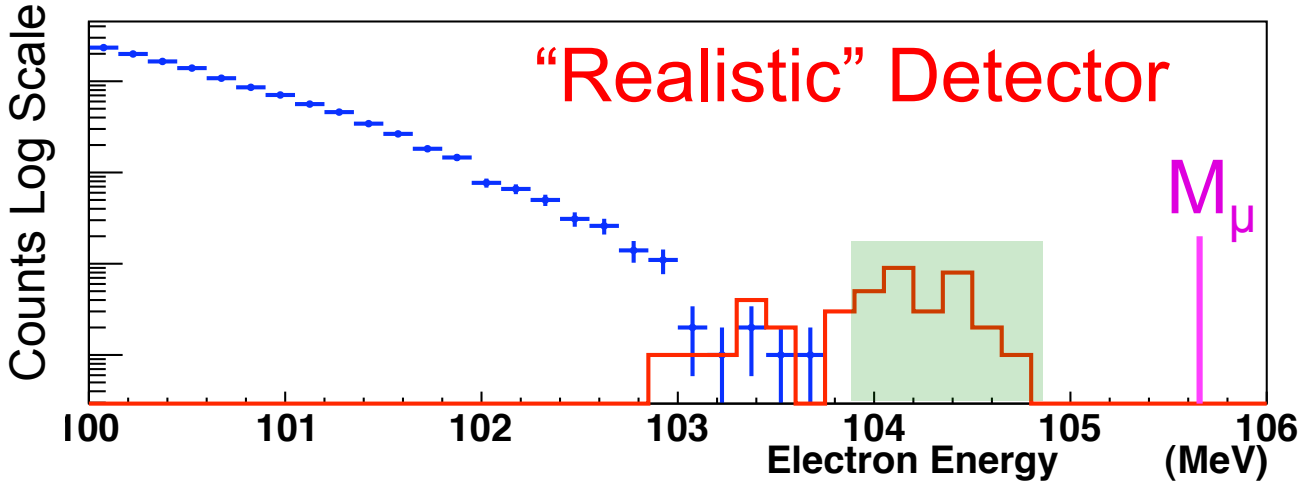

- Define a signal region
- Measure the electron Energy Spectrum
- Estimate background events in the signal region
- Count events in the signal region, subtract BG.
- Is there a significant excess above BG?

Micro Time Structure

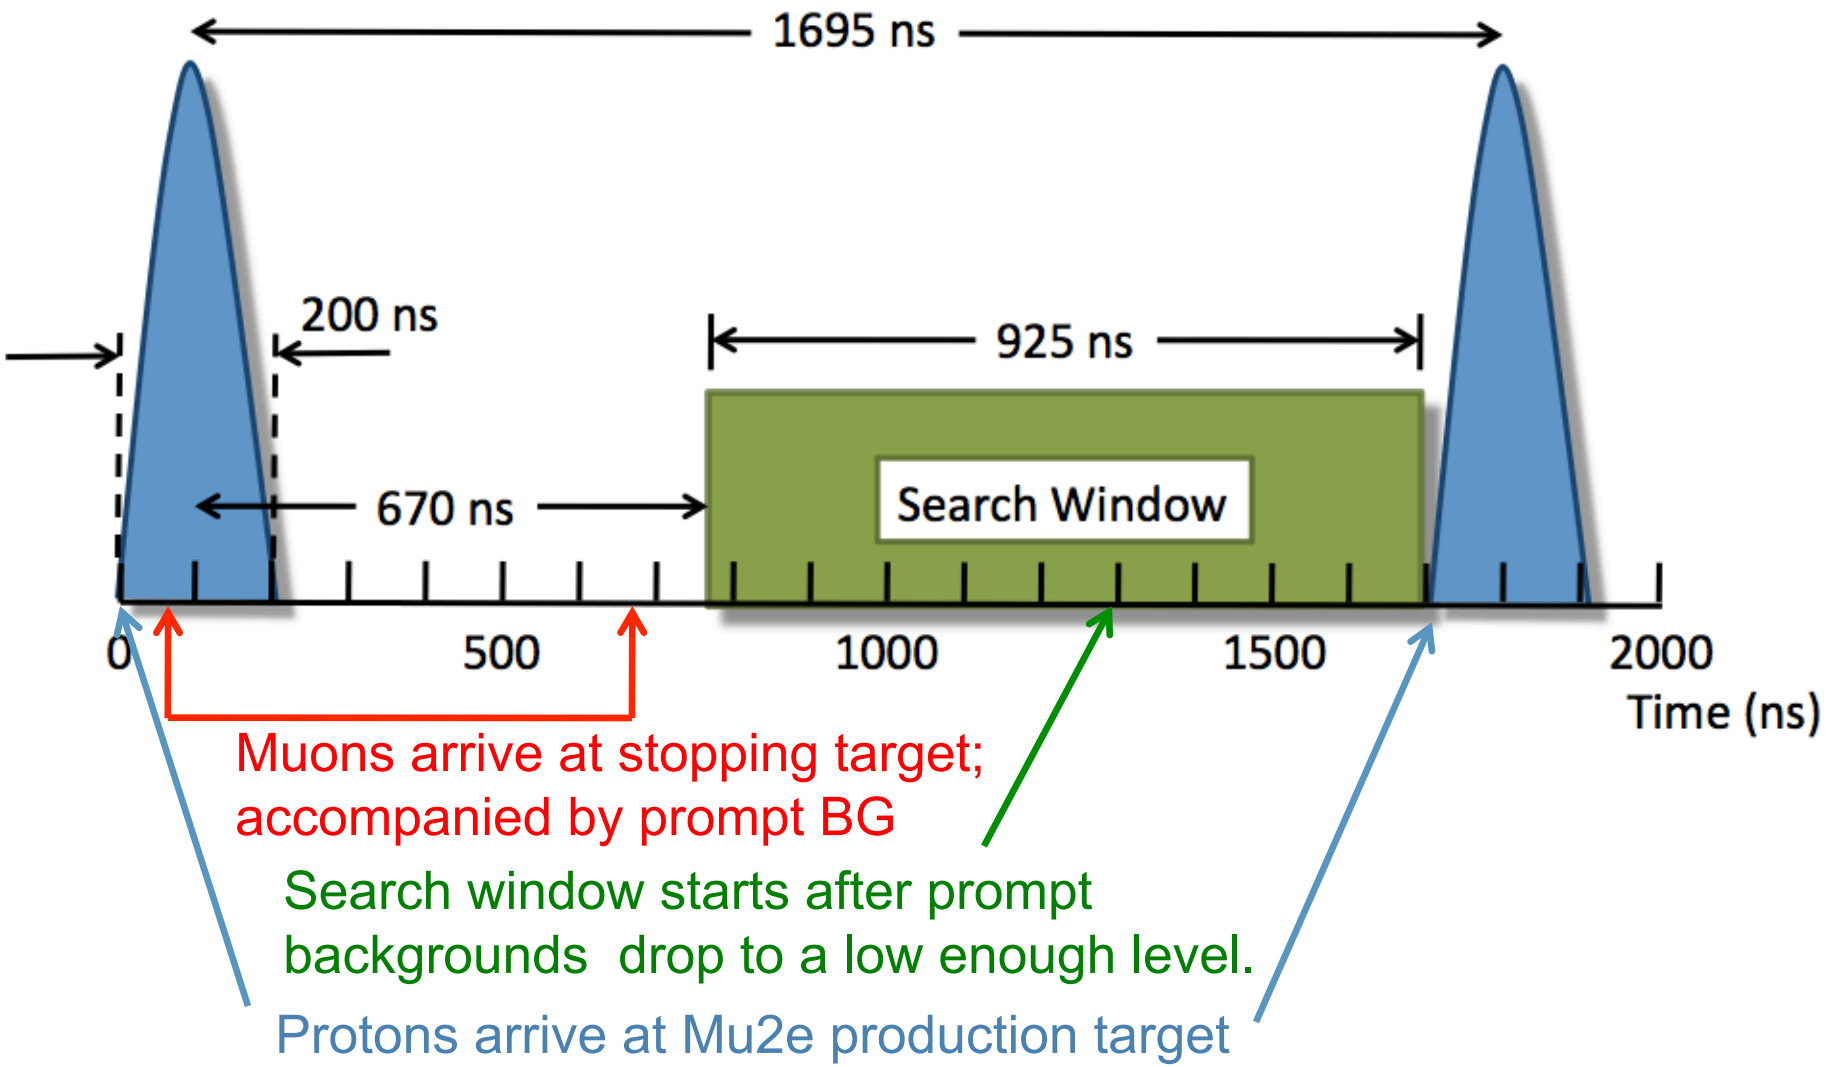

Macro Time Structure - 1

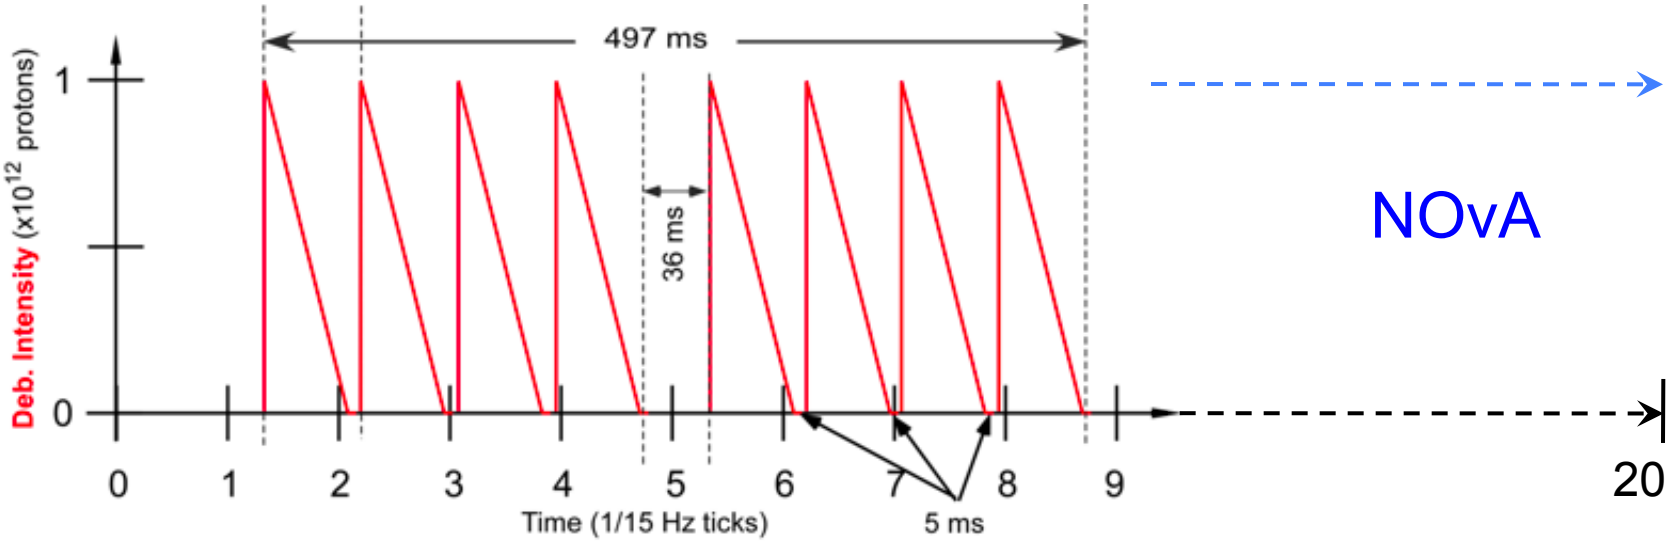

- NOvA era super cycle = 20 Booster cycles = 1.33 s
- Red curve:
  - Proton intensity in Delivery Ring: 8 slow spills per super cycle.
  - ~31k micro bunches per slow spill (1695 ns separation)
  - ~250k micro bunches per super cycle
- Mu2e protons are clear from Recycler before 8th tick.

Macro Time Structure - 2

- During the 12 Booster Cycles that NOvA beam is in MI and RR:
  - We will likely keep reading out the detector to compile background events.
  - May trigger on cosmic rays or may just read every cycle.

Cost and Schedule

Mu2e Schedule – as of June 2012

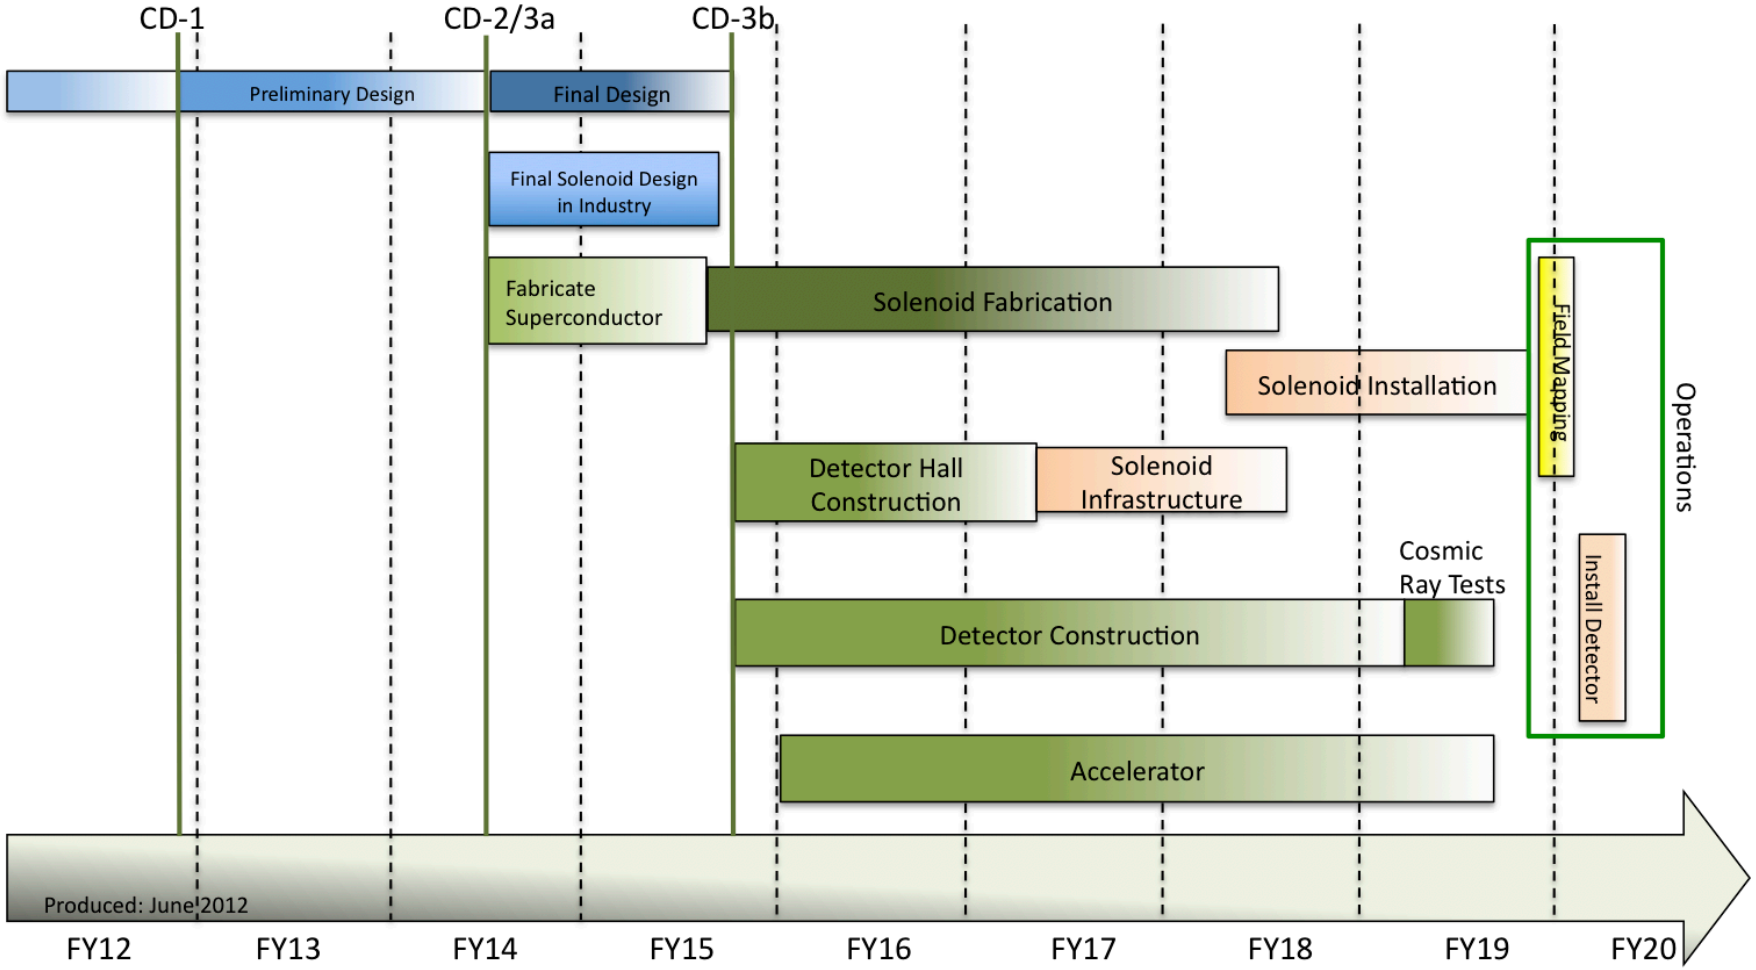

3 year run (to end of FY23); + final analysis phase 1 to 2 years?

Mu2e Total Project Cost – June 2012

	Total Cost (AY)
Project Management	$20M
Accelerator	$45M
Civil Construction	$25M
Solenoids	$98M
Muon Beamline	$14M
Tracker	$9M
Calorimeter	$6M
Cosmic Ray Veto	$6M
Trigger & DAQ	$6M
Total	$229M

June 2012

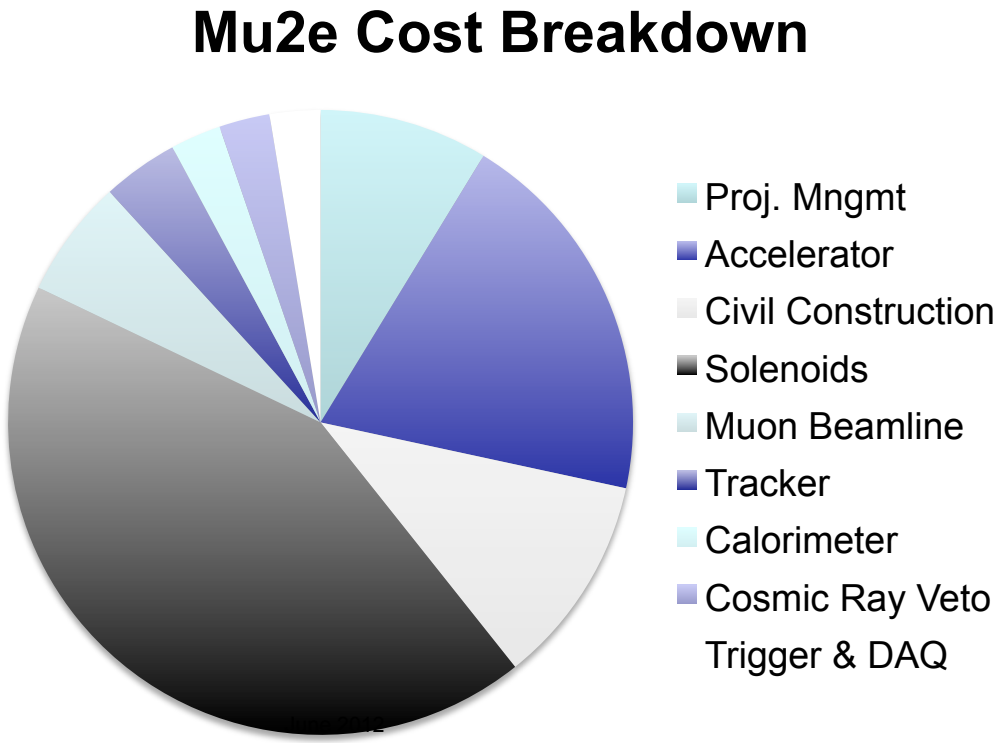

Offline Software is not on project

- All figures are escalated and include overheads and contingency

Mu2e Cost and Schedule

- Critical path and dominant cost is the solenoids.
- To match projected funding profiles, we may need to build some non-critical-path items early and let them sit on a shelf for a few years.
  - Possibly this could mean some elements of DAQ or trigger farm!

DAQ and Trigger

Streaming DAQ Mode

- Micro bunch cycle: 1695 ns
- Data rate off detector 30 GB/s
- 275 data sources
  - 2.5 Gbps optical link per source
- Software filter reduce rate by 1000.
- Raw Data to tape: ~(0.4 to 1.0 PB/year)
- Guess that reco data and MC events might increase the tape requirements by a factor of 2 or 3, giving O(1 to 3 PB/year) total tape.

Trigger Farm Size

- Existing Offline code on an Opteron 270 at 2 GHz.
  - 0.05 seconds/micro bunch on simulated all background microbunch
  - But code will be faster (x2?)
    - It is currently heavily instrumented
    - Algorithm is not stripped down for use in a trigger
- About 250K micro bunches per 1.33 second super cycle
- Implies server counts of:
  - ~12,500 Operton 270 (2.5 GFLOP)
  - 625 Quad Opteron 6128 (50 to 100 GFLOP) Reference Design
  - 156 Xeon E5 (200 to 300 GFLOP)
  - 52 Intel MIC (600-800 GFLOP)
- My guess: 25 to 100 FY19 servers.

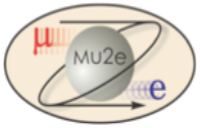

Offline Software

Offline Software

- By DOE rule Offline Software is in ops budget and is not part of the construction project.
- But …
  - We need fairly advanced offline software in order to prove the design and to evaluate design options.
  - The trigger farm will run a version of the offline software.
- Scientist labor is off project.
- CP labor is on project.
  - And we need CP labor.
  - Prior to CD-1 this was buried inside project mgt.
  - Not sure where it is now.
  - We will need integral ≈ 2 FTE DC

Offline Software Needs - 1

- Continued development of art.
  - I am concerned that existing art staff will migrate 100% to the many/multi core world and will leave art an orphan.
- Development of training materials for art – more later.
- Continued consulting on design of Mu2e code
  - At at present level or up to 2x higher.
- Continued support of:
  - art and its tool chain, ROOT, G4, CLHEP, ups/upd, SLF, g++, cdcvs(short term only), redmine, expwww or equivalent, docdb
- Release and distribution management
  - Now done by Lynn Garren

Offline Software Needs - 2

- Continued support of the work Krzysztof is doing:
  - He has a complete view of our simulation software and how it interacts with G4.
  - He can be plugged in anywhere.
- We will need improved build management with better support for test releases.
  - Do not want to use SRT
  - Will want some consulting and perhaps a little work.

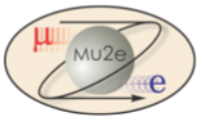

Training New People

- When you first encounter the Mu2e software, you need to master
  - The Mu2e software, proper
  - An intimidating beastiary of underlying tools …

The Underlying Tools

- C++ core language
- New features in C++11.
- C++ standard library
- art, our framework.
- FHICL, the run-time configuration language used by art
- CETLIB, a utility library
- The MessageFacility.
- ROOT
- CLHEP
- HEPPDT
- boost
- ups/upd for handling external products
- scons, our build system
- cvs
- The bash shell (both interactively and for scripting).
- Choices that we have made about how to use all of the above.
  - Some choices are carefully considered best practices while others are just choices.

How to get Through This

- CMS and BaBar experience is that the best solution is a self paced workbook
  - A series of:
    - An example
    - A story about the example
    - Suggested exercises
- The story should talk physics first and works down.
  - The story must introduce each tool as it is encountered and explain enough about that tool to get through the exercises.
  - Link into documentation about the tools.
- The workbook should give a complete tour of the detailed documentation suite.

My Vision

- We will write a Mu2e oriented workbook.
  - A good start already exists.
- We would like SCD to help us develop documentation of the underlying tools that we can link into from our workbook.
  - This keeps our workbook focused on Mu2e
- If we do this right, NOvA, g-2, LArSoft can develop their workbooks that also link into the documentation of the underlying tools.

Hardware Needs

Grid - 1

- Existing slot quotas:
  - 500 mu2e; 800 marsmu2e
  - 500 special quota to expire soon
- We can see saturating 1500 to 2000 slots.
  - We are just now starting to use opportunistic slots
- All work is simulation and reco of sim events.
- Most jobs are CPU heavy but IO light
  - Range from a few x 10 GB to a few x 100 GB per quota-day.
  - The few instances of 5 TB/day production were people not understanding what they were doing.
- Most disk files are out of date in a few months.
  - Higher fidelity simulations become available.

Grid - 2

- Why so many cycles needed?
  - Signal simulation is trivial.
  - But we need to simulate very large samples of many different backgrounds and overlay these on the signal events.
  - We need to repeat this for many variants on the detector design.
  - So far we have not discovered how to cut safely in the generator phase space.
  - We already factorize computations and reuse intermediate results where possible.

Disk

- /mu2e/data – 35 TB quota (24 TB used)
  - Would like it to double in FY13
  - Guess that we will need to be at 200 TB by FY17?
  - OK if this is split over 2 or more file servers.
- /grid/fermiapp/mu2e
  - We have avoided test releases so that we do not have problems with cross-stitched builds.
  - So we need a lot here: 120 GB
  - Need to add 60 GB from /grid/fermiapp/products/mu2e
  - Would like 250 GB ongoing.

Interactive Computing Needs

- We have enough GPCF interactive nodes for the near future.
- GPCF local batch
  - We have not been using it since it has been saturated every time we have looked.
- detsim
  - We heavily use this machine and would like to keep using it as long as we can.

Tape Needs

- For 1 year commissioning plus 3 year run
- O(0.4 to 1.0) PB/year raw data to tape
- Reco data:
  - Guess 0.1 times raw data
  - Guess we will reprocess a 2 times per year:
  - Therefore: (0.08 to 0.2) PB/year
- MC Data:
  - Guess 3 times the reco data each year?
  - (0.24 to 0.6) PB/year
- Total tape needs: 0.7 to 1.6 PB/year for 3 years

Tape Needs Before Running

- We have 10 volumes allocated. Barely using them.
- Guess that a few TB/year will be enough for test beam data.
- Guess that we may need a few 10's of TB/year for simulated events that we want to keep for future reference.
  - May not need this for a few years.

Looking Ahead

- We will want to use SAM
  - Not sure when.
- We will want support for a conditions database.
  - Not sure when.
- Electronic log book
  - Will want CRL or its replacement/evolution.
  - Will want it for commissioning.
  - Will want it the first time we do test beam work at FNAL. Not sure when this will be. Current test beam work is scheduled elsewhere.

Miscellaneous

- Complete the task to streamline the process of new user account creation.
- We would like the lab to support offsite users with windows boxes logging into GPCF.
  - The minimal solution might be to deploy and maintain a page that points to putty or cygwin and fills in any gaps in their documentation.

For more Information

- Mu2e home page:
  - http://mu2e.fnal.gov/
- Mu2e Conceptual Design Report (June 2012)
  - http://mu2e-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=1169;filename=CDR%20Final.pdf;version=14

Summary and Conclusions

- We need ongoing support for development of our Offline software and for the tools that it is built upon.
- Critical need for better documentation of the underlying tools.
- We would like to keep increasing disk and grid slots each year.
- Tape needs modest until data taking.
- Will need SAM, Databases, CRL: not sure when.

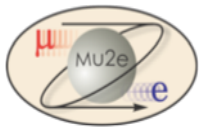

Backup Slides

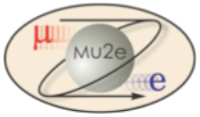

To Breach the C++ Barrier

- Need example driven, layered documentation
- Each layer should discuss the new ideas that are relevant for example at hand; then refer to the next layer for additional details.
- At the end of the chain you need the full documentation, both reference docs and conceptual ones.
- The problem is that you need to build it up from the bottom, which is tedious and boring.

Muonic Aluminium

- Muonic Al
  - Muon bound to an Al nucleus just like an electron would be
  - Lifetime 864 ns.
  - Free muon lifetime 2197 ns
- 2 Main Decay Modes
  - Decay-in-orbit (DIO): 40%
  - Capture on Nucleus: 60%

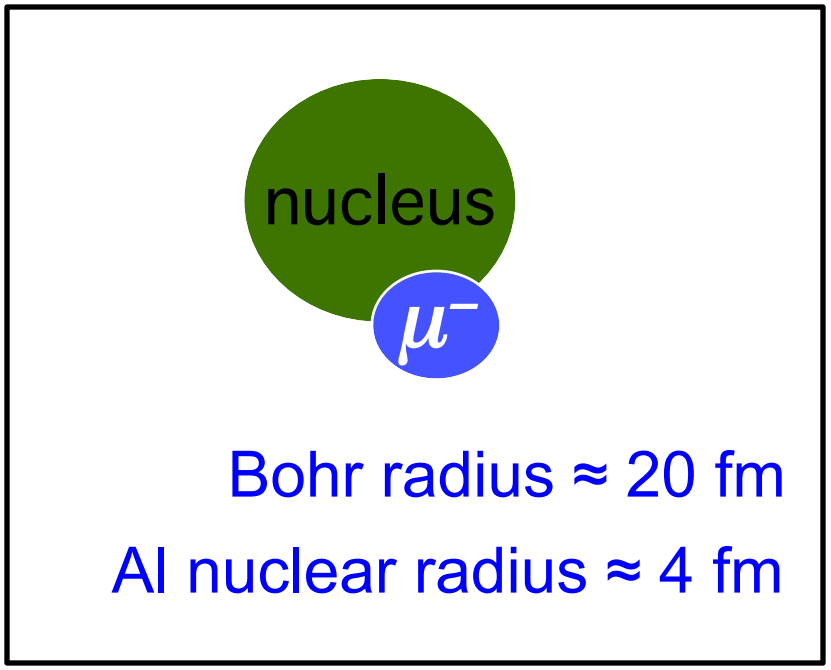

Cartoon of the Experiment

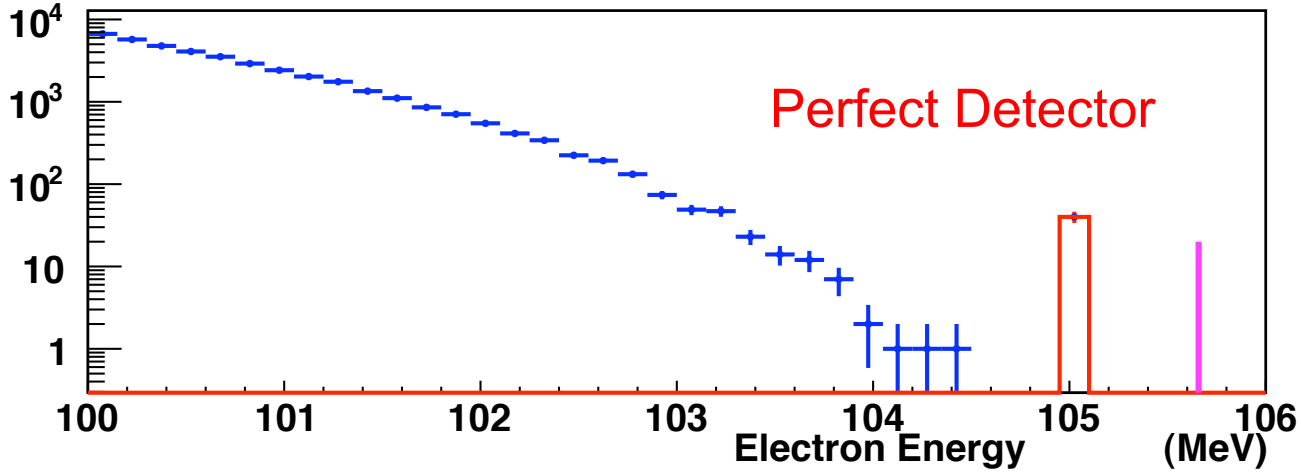

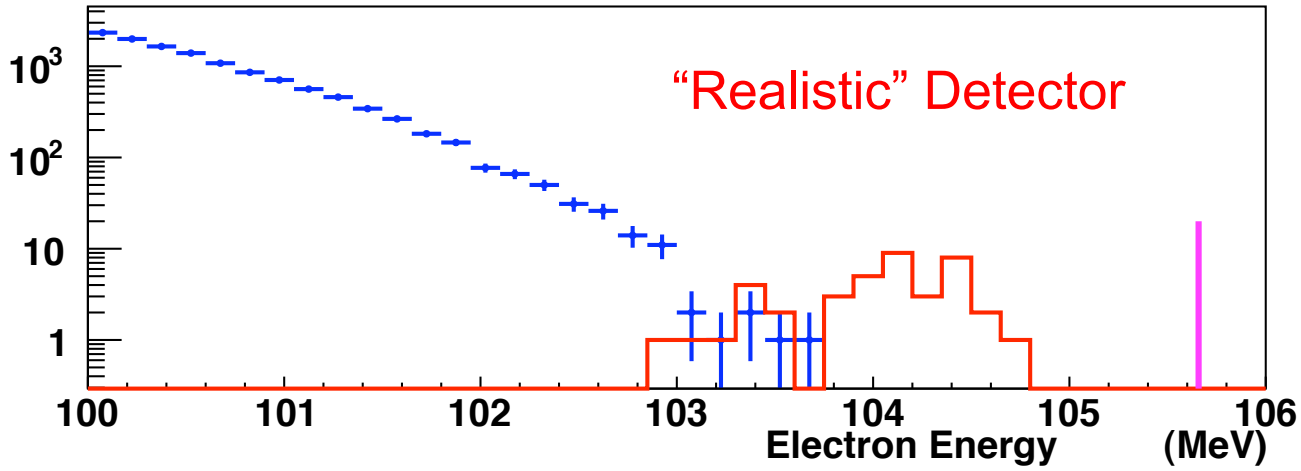

Decay of a Free Muon

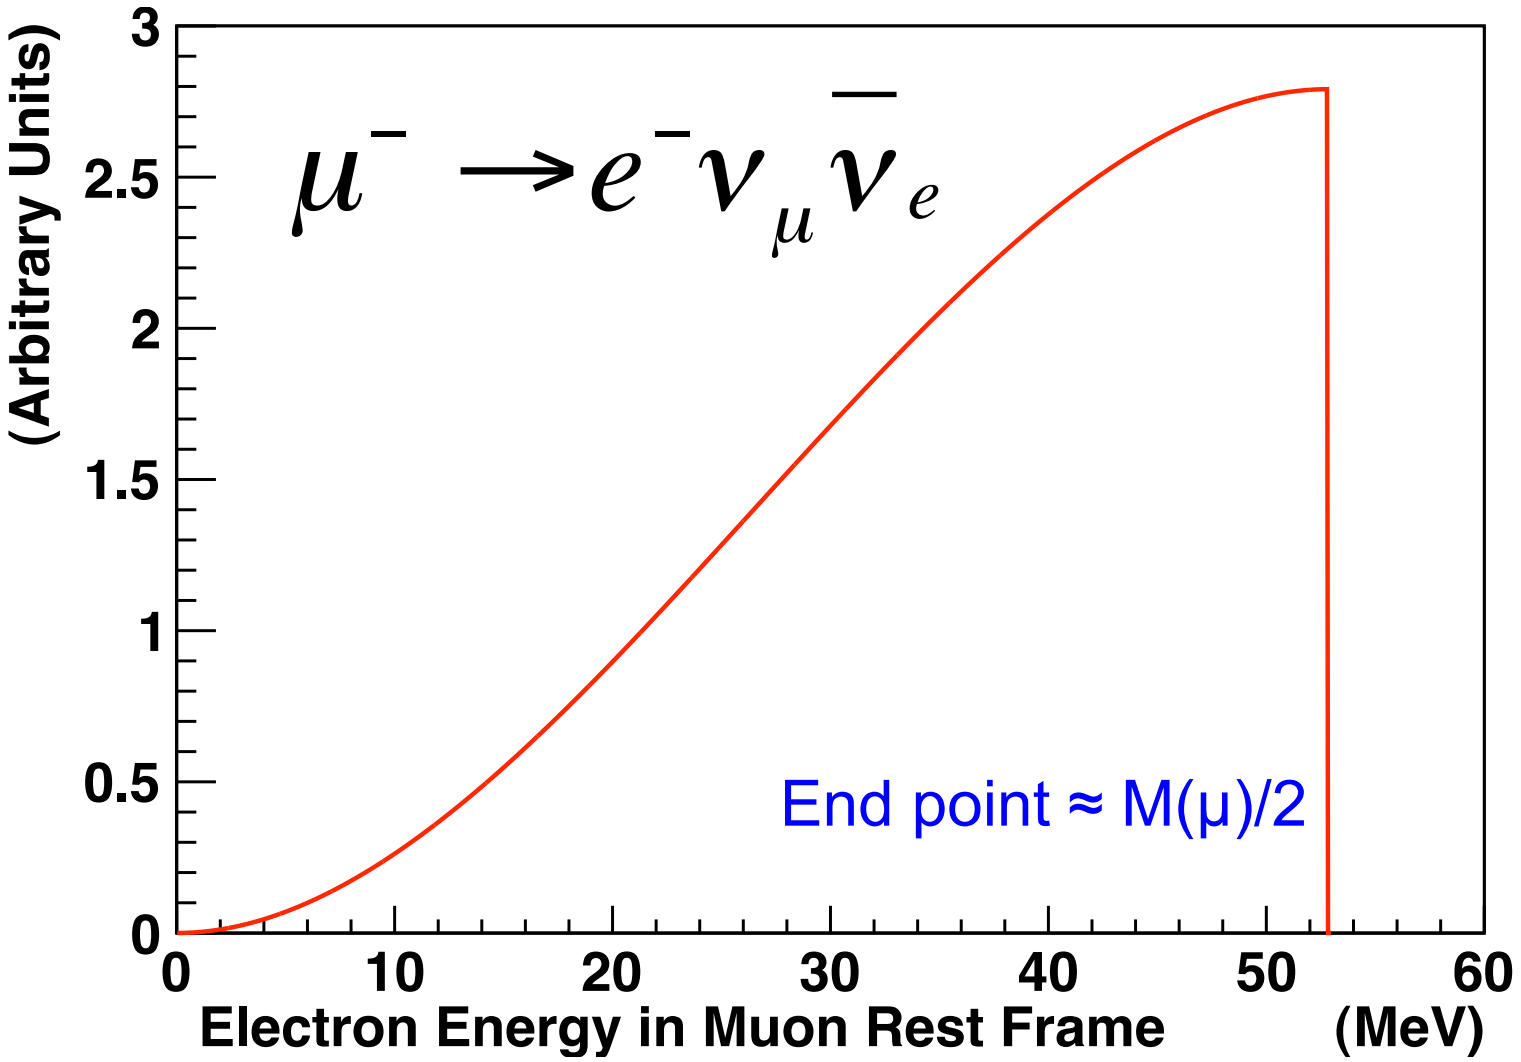

Energy Spectrum from Muonic Al

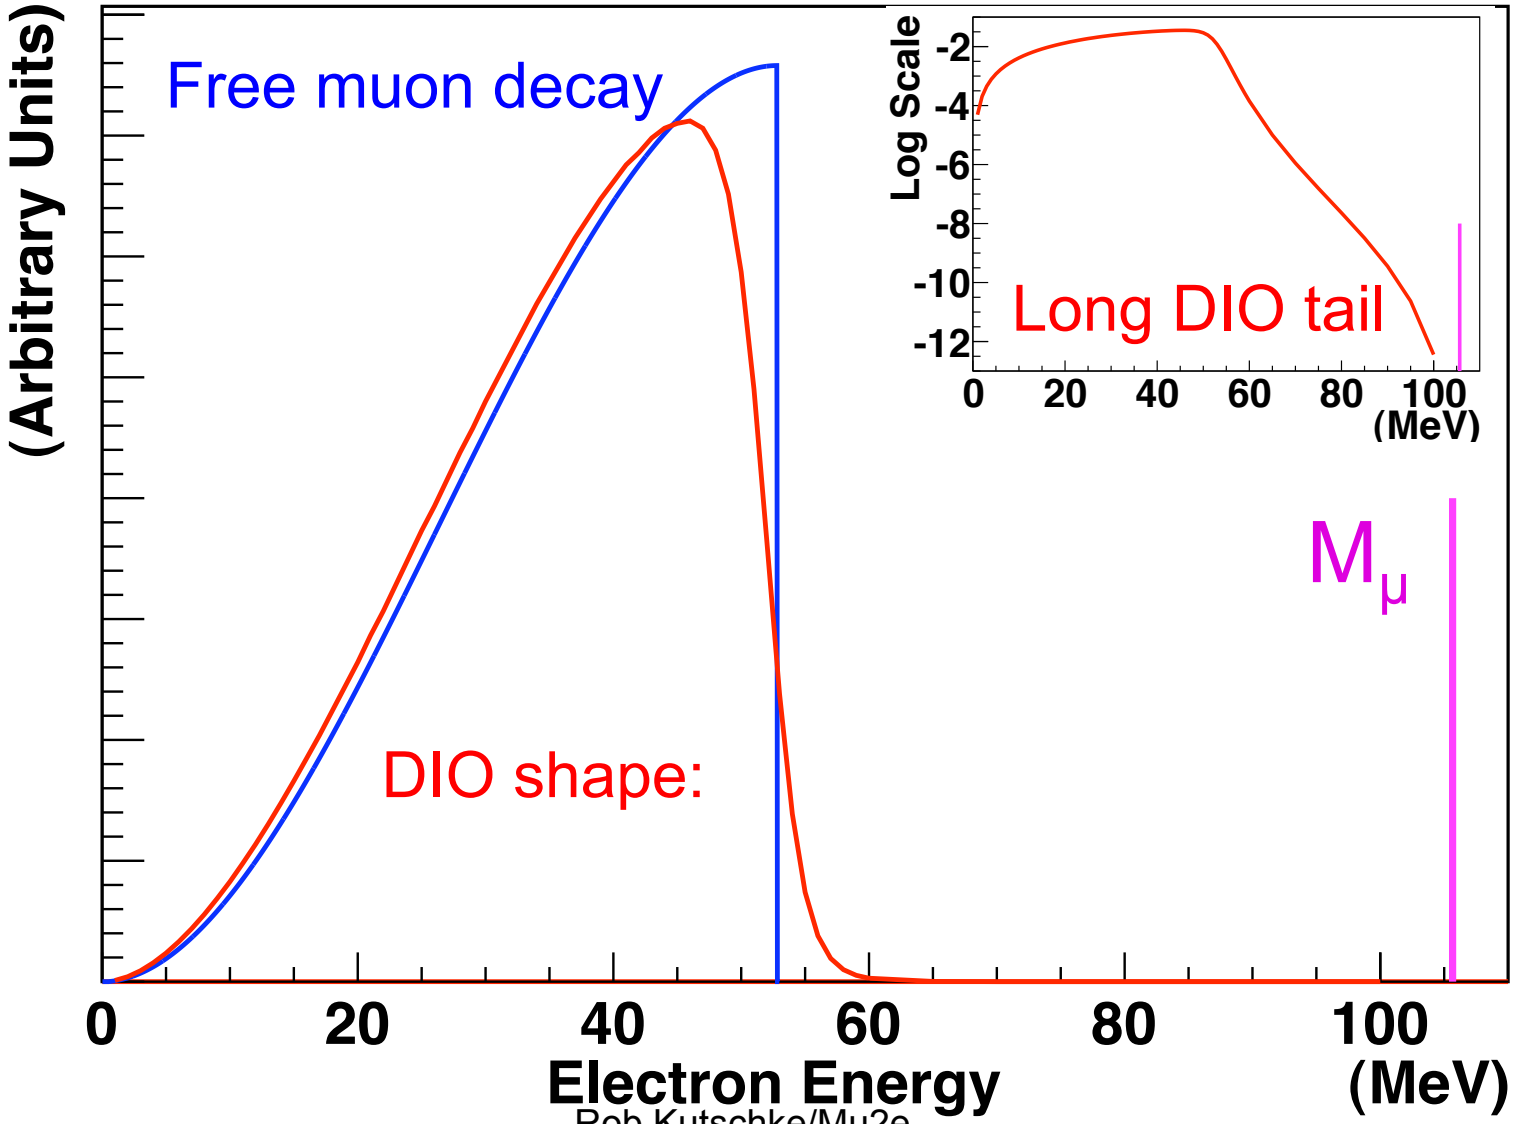

Cartoon of the Experiment

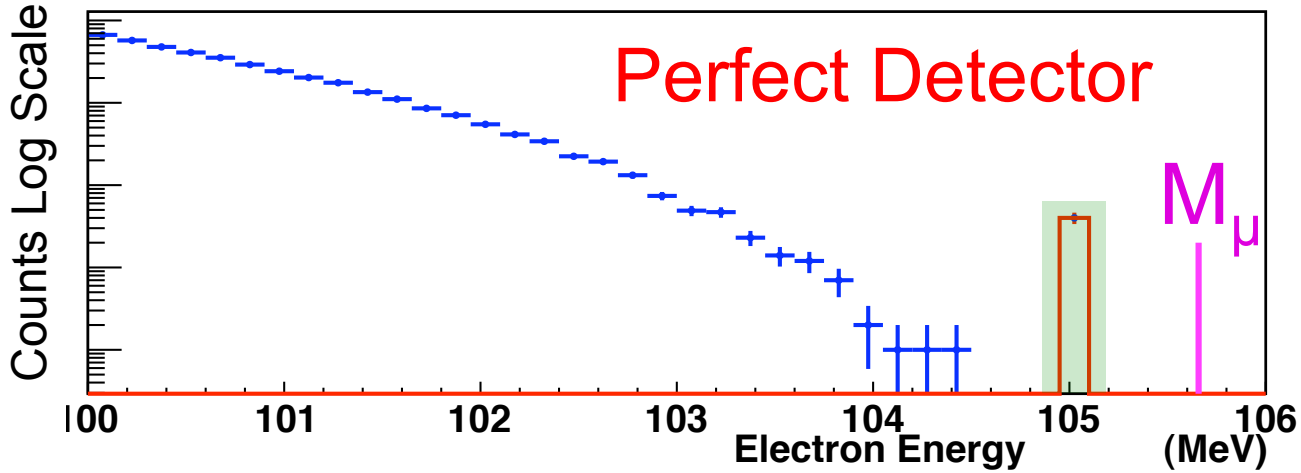

- Define a signal region
- Measure the electron Energy Spectrum
- Estimate background events in the signal region
- Count events in the signal region, subtract BG.

Software Tools

- Infrastructure Software
  - Framework proper; event-data model; persistency;run-time configuration; 3rd-party software (Geant4, ROOT, CLHEP, boost, …), build tools, release and distribution management, file catalog, database support for geometry, conditions-data, metadata, integration with grid workflow management …
  - Should be written by software professionals and designed in consultation with the physicist end-users.
- Physics Software
  - The event-data objects and relationships among them.
  - The geometry and conditions-data classes.
  - Algorithms: hit creation, reconstruction, monitoring, analyses etc
  - Will (mostly) be written by Mu2e people and designed in consultation with computing professionals.

Use Cases for Simulations

- A tool to:
  - Inform the design of the experiment
  - Once we have data: compute efficiencies, understand resolutions, backgrounds …
  - A tool to develop, debug and characterize reconstruction algorithms.
    - Design this in from the beginning!