The MuLan Experiment: Measuring the muon lifetime to 1 ppm*

Kevin R. Lynch[1,2,†]

(The MuLan Collaboration)

[1]*Department of Earth and Physical Sciences,*
York College, City University of New York,
94-20 Guy R. Brewer Blvd, Jamaica, NY 11451 USA
[2]*The Graduate Center, City University of New York,*
365 Fifth Avenue New York, NY 10016 USA

(Dated: April 18, 2016)

Abstract

The MuLan Collaboration has measured the lifetime of the positive muon to 1 ppm. Our result now drives the world average. Within the Standard Model framework, this permits a determination of the Fermi Constant to 0.5 ppm. I present our measurement method, our published results, and prospects for future improvements in the technique.

INTRODUCTION

The Standard Model of particle interactions is a triumph of modern physics. We find extremely impressive agreement between theoretical predictions and measurements in the realms of atomic physics, nuclear structure, high energy interactions, and astrophysics and cosmology. These theoretical predictions, of course, are based on precision measurements of a small number of fundamental input parameters. The electroweak sector of the model, in particular, rests on three very well measured values: the fine structure constant, $\alpha_{\rm em}$, the mass of the neutral weak gauge boson, $M_{\rm Z^0}$, and the Fermi constant, $G_{\rm F}$.

Striking improvements have occurred in the last decade in measuring these three fundamental inputs. The Gabrielse group at Harvard has measured the fine structure constant at the 0.37 ppb level [1], while the combination of the four LEP experiments has determined the $\rm Z^0$ mass to 23 ppm [2]. After extensive analysis of data collected in 2006 and 2007, the MuLan Collaboration published its final results, which give the first major improvement in our knowledge of the Fermi Constant in twenty years [3, 4].

The Fermi constant sets the strength of the weak $V-A$ interaction, and can be cleanly extracted at high precision from a measurement of the free muon lifetime [5]

$$\frac{1}{\tau_\mu} = \frac{G_{\rm F}^2 m_\mu^5}{192\pi^3}\left(1+\sum_i q_i\right) \ . \tag{1}$$

Here, we see the $V-A$ prediction factorized into a pure weak contact contribution (encoded in $G_{\rm F}$), plus a sum over non-weak corrections (the q_i). These include the massive phase space integrals (q_0), and QED and hadronic loops. This extracted value of the Fermi constant must then be connected with the weak interaction physics of interest - usually the Standard Model - through a loop expansion.

At the close of the twentieth century, extracting $G_{\rm F}$ from the muon lifetime was limited by theory - only q_0 and q_1 were known. The uncalculated higher order terms were esti-

mated to contribute 30 ppm, compared with a world averaged experimental uncertainty of 18 ppm. In 1999, van Ritbergen and Stuart [6] succeeded in calculating the second order QED corrections (q_2) in the limit of massless electrons, reducing the theory uncertainty to the sub-ppm level. On the heels of this result followed multiple proposals to improve the experimental uncertainties to comparable levels, culminating in the independent MuLan and FAST experiments at the Swiss Muon Source at the Paul Scherrer Institut (PSI) in Villigen Switzerland. Both efforts initially pursued 1 ppm measurements of the muon lifetime. In addition to extracting the Fermi constant, improved direct measurement of the free muon lifetime is critical for comparisons with bound muon lifetimes used in the extraction of nuclear physics parameters; see the talk on the MuCap experiment by B. Kiburg in these proceedings.

EXPERIMENTAL METHOD

Previous experiments have generally utilized low rate, "one at a time" methods: a beam of muons is focused on a stopping target, with average target occupancy of fewer than one muon. This approach does not scale to a 1 ppm measurement, where we need to record in excess of 10^{12} individual decays to reduce the statistical uncertainties. Instead, MuLan pursued a method with high stopping target occupancy, namely a pulsed source that permits us to perform many muon lifetime measurements simultaneously.

To this end, we developed a high-rate (7 MHz) beam tune in the πE3 beam line at PSI, and constructed a fast electrostatic kicker to chop the beam [7]. We collected polarized muons on a fixed stopping target for a 5 μs (two muon lifetime) accumulation period, and then deflected the beam away from the target during the next 22 μs (ten lifetime) measurement period.

The stopping target was surrounded by a large acceptance, point symmetric, high granularity detector. The detector had a truncated icosahedral, or "soccer ball", geometry; each hexagonal (pentagonal) face consisted of six (five) triangular dual layer plastic scintillators, connected through short light guides to fast photomultiplier tubes. The upstream and downstream pentagonal faces were not instrumented, to permit the entry of the beam transport corridor upstream, and the exit of beam electrons downstream. To reduce background contamination, we demanded a coincidence between the inner and outer tiles of each pair. For

the data collected in our 2006 and 2007 run periods, the PMTs were instrumented with custom, high-rate waveform digitizers (WFDs). The entire system was controlled and read out with a customized data acquisition system based on the MIDAS framework [8]. In all, we recorded 340 channels of PMT "physics" signals, along with a number of other diagnostic and monitoring channels. The result of the data acquisition was a set of lifetime spectra which, after application of a set of well controlled, data driven corrections, could be fit to a simple functional form where the ultimate goal was the measurement of τ_μ at the ppm level.

MAJOR SYSTEMATIC ERRORS

Systematics, not statistics, were the core concern for our 10^{12} event data sets. As such, the experiment was designed from the ground up to minimize or eliminate systematic errors. The main class of errors were those that systematically skewed count rate efficiency early-to-late in the measurement period. Any such errors could directly contribute to a shift between the fitted lifetime and the actual lifetime, usually without any diagnostic problems in the quality of fit parameters. There are a large number of effects that we identified, either from fundamental physics or from finite instrumental precision, which needed to be addressed; a few are recounted here.

Timing shifts

Rate dependencies (particularly in the PMTs) are known to cause small timing shifts of order a few picoseconds. We directly monitored this effect with a laser reference system. We illuminated a subset of the detector tiles, as well as an independent reference counter located well away from the detector, with a regular fast laser pulse. The laser pulses were injected at a very low rate, and independent of the state of the data acquisition cycle. They therefore uniformly populated the measurement period, leading only to a very small increase in the flat background. A systematic time shift on the detector would be measured as a change in the time difference with the reference counter. No measurable shifts were seen.

Additionally, we searched for timing shifts with a parallel "randoms" experiment. With an additional pair of tiles located away from the detector, we observed the decays of a radioactive source. This measurement should see a flat time spectrum across the observation

period; no deviation was observed.

Gain shifts

Any rate dependencies in the PMT gain or discriminator thresholds would manifest as a change in counting efficiency. With the WFDs, we could directly monitor the average amplitude response from the laser system, comparing the amplitude response with that of the reference counter. No measurable gain shifts were seen.

Beam-off backgrounds

The background term in the fit function arises primarily from muon arrivals during the beam off period and irreducible cosmic ray events. To minimize the number of these out-of-time arrivals, our beam kicker was designed to operate with a 25 kV potential difference between the plates, providing a beam extinction around 900. To eliminate any early-to-late variation in the background at the ppm level, the extinction (and hence the voltage difference) must be regulated at better than the 1 V level. This stability was achieved early on [7].

Pileup

In our experiment, pileup was the loss of counts or time shifting of counts due to the finite pulse time resolution of our detector. To combat this, our detector was a highly segmented design, with 170 dual layer tiles built from fast plastic scintillator to minimize event occupancy in each channel. Even with our high beam rate, the time structure and beam extinction ensured that the average measurement period occupancy of each channel was less than 0.1 hits per beam cycle. Combined with the good time resolution for pulse timing available with the WFDs, this reduced the pileup losses to about 10^{-4}. We had to understand these effects at roughly the 0.1% level to eliminate this as a dominant systematic.

In principle, we could have determined the effects of pileup processes on the lifetime spectrum, and modified our fitting function to account for them. The dominant contribution, an $\exp(-2t/\tau_\mu)$ term, has long been understood and applied in previous experiments.

Unfortunately, this dramatically reduces the statistical power for a given number of decays: the dominant pileup term alone reduces the statistical precision by a factor of two. Since pileup events were rare, however, we could directly reconstruct the pileup contributions from the data. This is possible because we could determine how various classes of events nearby in time in a single channel will lead to time shifts or event losses. We call our statistical reconstruction method the *shadow window reconstruction* procedure. While the details of its application depend on the details of the pulse reconstruction algorithm, we can easily describe the flavor of the pileup reconstruction.

Consider the dominant pileup term: what we call *normal pileup*. In this case, two decay events pass through a single channel within the resolution time of the pulse reconstruction. In this case, the reconstruction will "see" only one event - we have a small but finite deadtime. Since the probability of the individual events happening in the same time window are independent, it follows that this probability is the same as that of two events happening in the same time window, but in *different* measurement periods. The number of times the latter occurs matches the number of times events have been lost.

The correction procedure, then, was applied in the following way. First, we reconstructed every pulse we could find in the raw data. Then, we applied an *artificial deadtime* cut to remove nearby events, and put all the remaining events into a *pileup contaminated* histogram. We also built a second, *pileup correction* histogram by applying the approach discussed above: if we saw an event in observation period 1, we looked in the same window (the *shadow window*) of observation period 2. If we found an event in that second window, we also added an entry to the correction histogram. The sum of the contaminated and correction histograms should represent the true, uncontaminated lifetime spectrum. In addition to normal pileup, we considered additional pileup corrections from other sources, including (among others) three simultaneous events, timing jitter, and the accidental time coincidence of uncorrelated singles hits on an inner and outer tile pair.

The pileup correction procedure was repeated for many values of the artificial deadtime cut. The lifetime for each of all the corrected spectra will agree with each other if the procedure has accounted for all relevant pileup effects. For our final results, a small residual correlation between artificial deadtime and fitted lifetime of 0.008 ppm/ns remained after pileup correction. No conclusive source for this discrepancy was identified, but we showed that it was stable and linear with artificial deadtime; we extrapolated the correction to zero

deadtime to produce our final lifetime. We assigned a systematic uncertainty of 0.2 ppm for pileup corrections.

Muon polarization effects

Muon beams are naturally highly polarized: muons are born from the two body decay of spin-zero pions via the chiral weak interaction. Since the muon neutrino is only available in left-handed form, angular momentum conservation demands that all the muons have the same chirality. The large ensemble spin polarization can be either a blessing or a curse, depending on the application. In the case of MuLan, it carried the potential for disaster. If the ensemble maintains any of its polarization after stopping in the target, then 1. the component perpendicular to any local magnetic field will precess and decay with time, and 2. the parallel component will decay (longitudinal relaxation) with a different lifetime. Either effect could lead add an additional, unknown component to the spectrum for any individual detector tile.

We took a number of steps to both control and measure the size of these effects. The key observation is that detectors on opposite sides of the stopping target should have seen effects that cancel, up to (unavoidable) small differences in positioning, acceptance, and efficiency. The detector was designed to maximize the point symmetry, and opposing tiles were chosen to have closely matched efficiencies. To the extent possible, materials and construction of the elements surrounding the stopping target were chosen to minimize acceptance differences. Additionally, we performed systematic studies with stopping targets that maintain the residual polarization, such as silver, in order to measure the departures from design symmetry.

The physics targets were chosen to minimize the residual polarization, but in different ways. For the 2006 running period, we chose the ferromagnetic alloy Arnokrome-3 (AK3); the high internal fields cause many precession cycles between muon arrivals which "scrambled" the ensemble of spins. This picture was confirmed both with dedicated muon spin resonance studies and with analysis of point symmetric detector tile asymmetries. For 2007, we chose crystalline quartz in a strong applied field (of order 150 G). In quartz, a large fraction of the stopped muons are bound in muonium; the muon spin precesses a thousand times faster in muonium than does a free muon. The applied field then performs the same

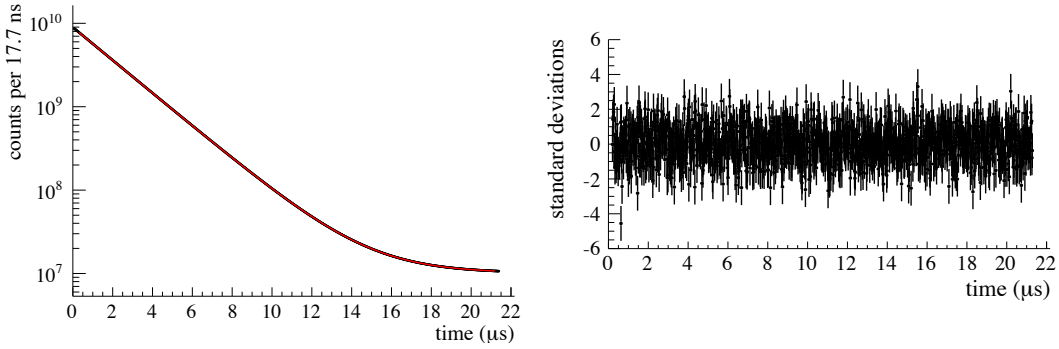

FIG. 1: A fit to the entire AK3 data set. The fits for the entire data sample began at 1 μs after the end of injection. We also display the residuals to the fit, showing no observable structure.

depolarization as the AK3 internal fields, for both the muonium and remaining free muons. Analysis of the efficacy of this target in reducing polarization effects showed that it gave us good measurement of and control over the polarization systematic.

THE STATUS OF MULAN

The experiment ended data collection with our 2007 run period. For both the 2006 and 2007 run periods, we analyzed in excess 10^{12} decay events. In addition to the dedicated systematic studies discussed above, we performed many consistency checks over subsets of the data, including various kicker, beam, and target conditions, discriminator threshold settings, artificial deadtimes, etc. Lifetime fits over various exclusive subsets were all consistent within statistics. Fit start and stop time scans revealed no structure in the residuals, suggesting that no time dependent effects of any significance were missed.

The final fit function for the AK3 (2006) data set was the same three parameter fit function: $f(t) = N\exp(-t/\tau)+B$. For the quartz (2007) data, fits were performed on individual tiles with a multiparameter function that accounted for measurable residual polarization effects that canceled in the sum of all data. All of the significant systematic errors for both run periods (both correlated and uncorrelated) are presented in Table I.

The results of this 2006 and 2007 run periods were published in 2010 [3], with the results

$$\tau_\mu(2006) = 2\ 196\ 979.9 \pm 2.5(stat) \pm 0.9(syst)\ \text{ps} \tag{2}$$

TABLE I: Sources of systematic uncertainties on the muon lifetime measurements in the R06/R07 running periods. The uncertainties listed in single-column format are common uncertainties and those listed in two-column format are uncorrelated uncertainties. The last two rows are the combined systematic uncertainties and the overall statistical uncertainties for the 2006 and 2007 datasets.

Uncertainty	R06 (ppm)	R07 (ppm)
Kicker stability	0.20	0.07
μSR distortions	0.10	0.20
Pulse pileup	0.20	
Gain variations	0.25	
Upstream stops	0.10	
Timing pick-off stability	0.12	
Master clock calibration	0.03	
Combined systematic uncertainty	0.42	0.42
Statistical uncertainty	1.14	1.68

and

$$\tau_\mu(2007) = 2\ 196\ 981.2 \pm 3.7(stat) \pm 0.9(syst)\ \text{ps}. \tag{3}$$

The combined result

$$\tau_\mu(\text{MuLan}) = 2\ 196\ 980.3 \pm 2.1(stat) \pm 0.7(syst)\ \text{ps} \tag{4}$$

is obtained from the weighted average of the two individual values with the appropriate accounting for the correlated uncertainties. A comparison of these results in the context of previous muon lifetime experiments is plotted in Figure 2.

Using our value of the muon lifetime, we can extract the Fermi constant

$$G_\text{F}(\text{MuLan}) = 1.166\,378\,7(6) \times 10^{-5}\,\text{GeV}^{-2}\ (0.5\ \text{ppm}). \tag{5}$$

This result represents a thirty-fold improvement over the 1999 PDG value obtained before the vRS theoretical work and the lifetime measurements pre-dating MuLan. The error in

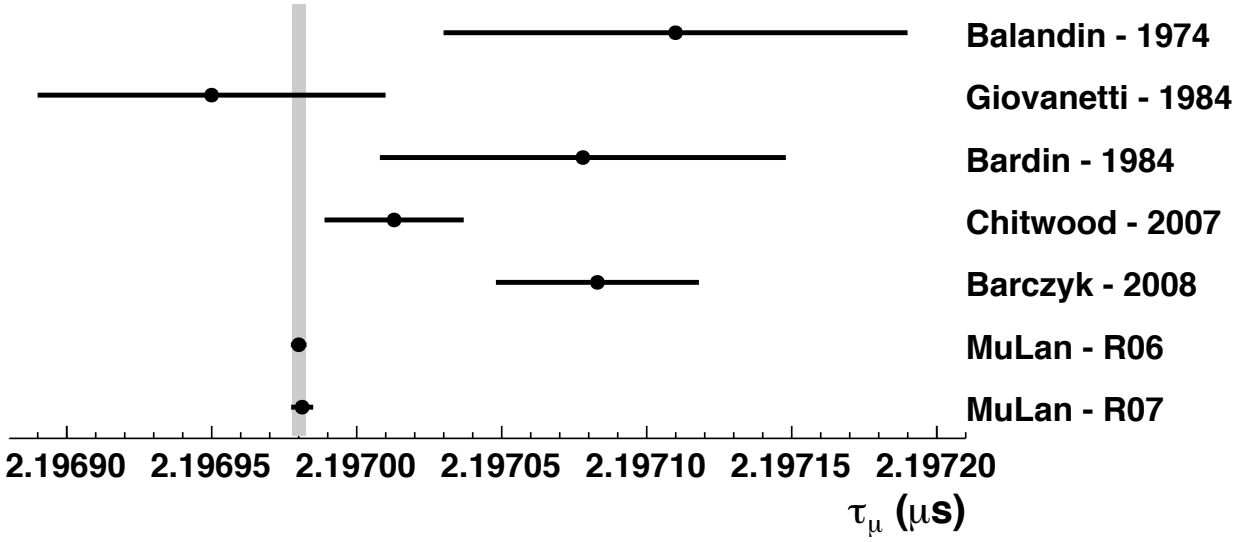

FIG. 2: A history of muon lifetime results. The narrow vertical bar shows the world averaged experimental muon lifetime, including our results.

G_F of 0.5 ppm is dominated by the 1.0 ppm uncertainty of the muon lifetime measurement with additional contributions of 0.08 ppm from the muon mass measurement and 0.14 ppm from the theoretical corrections.

FUTURE PROSPECTS

One goal of this workshop was to outline the prospects for new experiments at future neutrino sources. Although the MuLan experiment is finished and the detector has been dismantled, the techniques we used could be scaled up to provide continued improvements in the measurement of τ_μ and G_F should that prove valuable in the future. MuLan was statistics limited, but a new generation of lifetime experiment would need to take careful notice of the dominant systematics in our effort. Presumably a new experiment would run with much higher stopped muon rates, requiring greater detector coverage, segmentation, and uniformity to maximize the control of systematics given by the point symmetry of the detector. Special care will need to be taken to better understand pulse pileup corrections, and there will be need to better control or correct detector gain stability early-to-late in the measurement period. Finally, a more comprehensive understanding of muon polarization effects will be required. The extraction of the Fermi constant from the muon lifetime is

currently dominated by experimental uncertainty; with some modest improvements, we could likely bring the experimental contribution closer to parity with the theory contribution. A similar back-and-forth between theory and experiment in this measurement has a long history of productively enhancing our field.

---

* *Presented at NuFact15, 10-15 Aug 2015, Rio de Janeiro, Brazil [C15-08-10.2]*

† klynch@york.cuny.edu; Speaker

[1] D. Hanneke, S. Fogwell, and G. Gabrielse, Phys. Rev. Lett. **100**, 120801 (2008), 0801.1134.

[2] Phys. Rept. **427**, 257 (2006), hep-ex/0509008.

[3] D. M. Webber et al. (MuLan), Phys. Rev. Lett. **106**, 041803 (2011), [Phys. Rev. Lett.106,079901(2011)], 1010.0991.

[4] V. Tishchenko et al. (MuLan), Phys. Rev. **D87**, 052003 (2013), 1211.0960.

[5] T. van Ritbergen and R. G. Stuart, Nucl. Phys. **B564**, 343 (2000), hep-ph/9904240.

[6] T. van Ritbergen and R. G. Stuart, Phys. Rev. Lett. **82**, 488 (1999), hep-ph/9808283.

[7] M. J. Barnes and G. D. Wait, IEEE Trans. Plasma Sci. **32**, 1932 (2004).

[8] V. Tishchenko et al., Nucl. Instrum. Meth. **A592**, 114 (2008), 0802.1029.

[9] D. B. Chitwood et al. (MuLan), Phys. Rev. Lett. **99**, 032001 (2007), 0704.1981.

[10] A. Barczyk et al. (FAST), Phys. Lett. **B663**, 172 (2008), 0707.3904.