MuLan Progress Report 2006
and
Beam Request 2007

T.I. Banks[1], S. Battu[5], R.M. Carey[2†], D. Chitwood[3], S.M. Clayton[3], P.T. Debevec[3], S. Dhamija[5], W. Earle[2], A. Gafarov[2], F.E. Gray[1], G. Giovanetti[4], K. Giovanetti[4], T. Gorringe[5], Z. Hartwig[2], E. Hazen[2], D.W. Hertzog[3†], C. Hugon[2], P. Kammel[3], B. Kiburg[3], S. Knaack[3], J. Kunkle[3], B. Lauss[1], I. Logashenko[2], K.R. Lynch[2], R. McNabb[3], J.P. Miller[2], F. Mulhauser[3], C.J.G. Onderwater[6], Q. Peng[2], S. Rath[5], B.L. Roberts[2], W. Shiflet[4], A. Stavola[4], V. Tishchenko[5], D. Webber[3], P. Winter[3]

[1]BERKELEY – [2]BOSTON – [3]ILLINOIS –
[4]JAMES MADISON – [5]KENTUCKY – [6]KVI
[†]Co-spokesmen

January 17, 2007

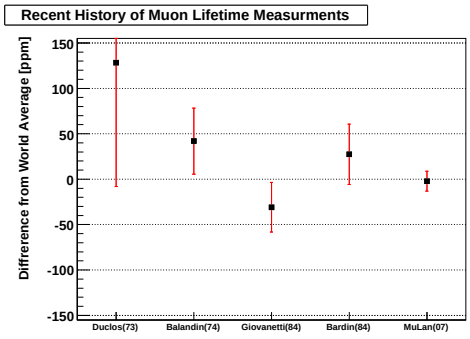

(a) The final result from our 2004 analysis yields the most precise value of the muon lifetime and Fermi constant to date.

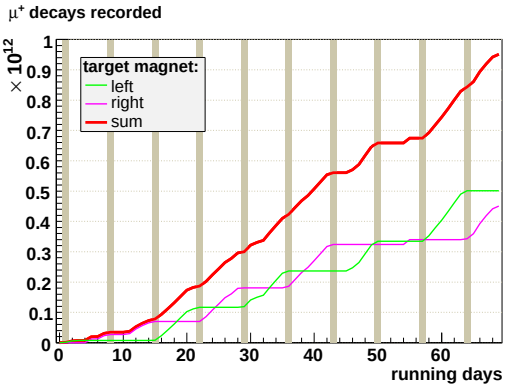

(b) We observed nearly 10^{12} decay events during our 2006 run.

Figure 1: Our 2006 efforts saw both the completion of our 2004 analysis and a successful data collection period.

Contents

1 Summary

The MuLan collaboration aims to measure the positive muon lifetime to a precision of one part per million (ppm), which would represent more than an order of magnitude increase in precision beyond the current world average, about 20 ppm. The muon lifetime is used to determine the Fermi coupling constant, G_{F}, which is the fundamental quantity governing the strength of any electroweak process. Recent theoretical work [1–3] enables a clean extraction

of $G_{\rm F}$ from τ_μ and paves the way for a new generation of high-precision measurements.[4–6] Additionally, by combining the MuLan result on the positive muon lifetime with its own measurement of the negative muon lifetime in high-pressure hydrogen gas, the MuCap Collaboration can determine the $\mu^- p$ capture rate and thus determine the induced pseudoscalar proton coupling constant, g_P (see the MuCap Annual Report for details).

In late July of 2006, after a month of installation work on the floor of the experimental hall, the MuLan experiment was craned into πE3 for a 10.5 week run. Apart from a week lost to kicker failure and another half week lost to PSI accelerator trouble, the run was a resounding success. We collected roughly 10^{12} muon decays, representing a statistical error of about 1.3 ppm. Our data collection efficiency improved substantially compared to 2005, both in average data throughput and tape-writing livetime as well as the global experimental "uptime." We also allocated one week to a program of systematic studies. Tests included the use of non-depolarizing targets, displaced-detector geometry, kicker stability tests, and measurements of beam properties. During normal data taking, a regular program of target orientation changes, laser runs (see below), and runs with adjusted sampling-time around pulses or measurement-period changes was carried out.

The analysis of the 2004 data set was completed in December 2006. The result, as yet unpublished (see cover Figure), is consistent with those of the 1970s and 1980s.[7–10] The error of roughly 11 ppm, dominantly statistical, is an improvement by roughly a factor of two over the previous world average. See Section 12 for details.

Our 2006 successes include:

- Completion of the 2004 data analysis (11 ppm precision), which is being prepared for publication. (Section 12)
- First full-scale data production run. Approximately 10^{12} muon decays were recorded, using a high internal-field "dephasing" target.
- The kicker worked at the design voltage of ± 12.5 kV during the entire run. We maintained an extinction factor exceeding 800 and a muon rate of approximately 7 MHz. (Section 4)
- The helium bag and wire chamber just upstream of the target were replaced by an all-vacuum muon corridor. The walls of the vacuum pipes were lined with the high-field AK-3 alloy, which also served as the main target material. (Section 3)
- The LED system, which had not worked reliably in 2005, proved invaluable in verifying the health of the scintillator tiles as well as the cabling, front-end electronics and DAQ. (Section 9). As a result, all 340 detector channels were operational and calibrated. (Section 7)
- A new pulsed nitrogen-laser system was installed with a common distribution to approximately 20 detectors. The calibration pulses, which propagate through the entire detector and electronics path, are being used to provide timing-stability data. (Section 9)
- The waveform digitizer (WFD) firmware was substantially upgraded. Problems noted in 2005 were fixed and new features, such as a readable fill count register, were implemented. Extensive quality assurance plots were available online. (Section 8)
- Lossless compression and multi-threading was implemented in the DAQ for 2006. The livetime, which was approximately 50% in 2005, rose to 85%. (Section 10)
- Considerable progress has been made on data analysis for the 2005 (Section 13) and

2006 (Section 14) data sets.

Principal challenges for 2007 include:

- We discovered that the muon rate varies significantly with changes in proton targeting at the E target station. Horizontal adjustments can change the rate by 30%. In the future, an agreement should be made between the users who share the target.
- It is unclear if the strong fringe fields of the MEG COBRA magnet in the neighboring πE5 experiment area will have any effect on the MuLan PMTs.[11] It was expected that the COBRA magnet would be energized near the end of our 2006 run, but as that did not occur, we did not have an opportunity to test mitigation strategies.
- We ran the kicker in two-cabinet mode throughout the 2006 run, which increases the transition time (on and off) of the kicker pulse from 45 ns to 60 ns. We need to reinstall a number of MOSFET cards, which were repaired at TRIUMF this past fall—and, do our best to prevent card failures in the future. The one extended kicker downtime was caused by an operator error. (Section 4)
- Although our online histograms proved extremely useful, they reflected a sample of only 10% of the data. Also, no detailed pulse fitting was performed. By the time of the next run, our offline production capability should be far more extensive than it is today.

The Collaboration remains strong. Scientists at every level helped to make our 2006 data taking run a success, including Ph.D., Master's and Bachelor's candidates:

- Dan Chitwood (University of Illinois) is writing his Ph.D. dissertation on the 2004 data set analysis.
- David Webber (University of Illinois) and Qinzeng Peng, (Boston University) are writing their Ph.D. dissertations on the 2005/2006 data set analyses.
- Sivaram Battu (University of Kentucky) completed a Master's program in Computer Science on the design and development of the MuLan data acquisition system.
- Sabyaschi Rath (University of Kentucky) completed a Master's thesis on studies of the LED calibration system.
- Josh Kunkle (University of Illinois) is writing his Bachelor's thesis on the effect of non-target stops in the muon corridor walls (Section 3).

We continue to inform the nuclear and particle physics communities of the progress of our efforts through seminar, colloquium, and conference presentations.[12, 13]

The Collaboration goals for 2007 include publication of our first physics result, from the 2004 data set, and major progress on the analysis of the 2005 and 2006 data. We request a 10 week run: two weeks for installation, beam tuning and systematic studies, and seven weeks of production running, plus one week of contingency to absorb possible kicker or accelerator failures. Between the data collected in 2006 and 2007, we should easily reach our statistical error goal of 1.0 ppm. The details are given in Section 15 of this report.

2 MuLan experimental concept

The MuLan experiment is simple in concept. During an accumulation period of 5 μs, a stream of approximately forty muons is brought to rest in a thin target. The muon beam is then switched off, and decays are recorded by a surrounding detector (the MuLan "ball") during

a measuring interval lasting approximately 10 muon lifetimes (22 μs). This cycle is repeated until more than 10^{12} decays are recorded. The time-structured muon beam is created by means of a high-frequency, high-voltage electrostatic kicker. During the measuring interval, the Michel positrons are recorded by a highly-segmented, symmetric detector, featuring 170 independent scintillator tile pairs. Each element is read out by a photomultiplier tube (PMT), whose signal is sampled at 450 MHz by a dedicated waveform digitizer channel. The time of arrival and energy deposited in each tile are derived from the WFD record. Decay time histograms, constructed from coincident hits, are then fit to extract the lifetime. The entire analysis is performed *blind*: the analysis team is not given the frequency of the master clock until the analysis is complete. Additionally, individual analyzers report their results with a secret, personal offset. In this way, we eliminate unconscious bias toward the world average while preparing our results.

The design of the experiment is driven by systematic error considerations. Primary concerns are related to multi-particle pileup, muon spin precession, muon decays outside the fiducial volume of the detector, time dependence of detector gains or electronic thresholds, and backgrounds. Pileup is minimized by the segmentation of the detector, by the relatively low pulse rate per element, and by the double hit resolution enabled by recording the entire waveform and fitting the pulse shape. Uncontrolled precession of the stopped, polarized muon ensemble will cause a rotation of the angular distribution of decay positrons, which is correlated to the direction of the muon spin. Any angular variation in detector response will then modify the ideal exponential time-spectrum. The very high internal magnetic field, approximately 0.5 T, of our target material, Arnokrome-3 (AK-3), results in a dephasing of the initial spin directions. Originally, we planned to minimize the residual polarization by using a depolarizing target material, such as sulfur, but after finding that the μSR signal in AK-3 was much smaller, we chose the latter for our first major production run. Finally, the detector features front/back matched, symmetric segments, where the time spectrum of the sum of corresponding elements is nearly immune to a change in spin direction. Non-fiducial muon stops are minimized by transporting the muons through vacuum all the way to the stopping target. The time-dependence of gains and thresholds can be directly monitored with WFD data.

The instrumentation naturally divides into subsystems: the muon beam (Illinois); targets (Berkeley); the detector ball and EMC (Illinois); the calibration system (JMU); the waveform digitizers and clock systems (Boston); and the data acquisition, logging and online analysis computers (Kentucky). Institutions with primary responsibilities for these items are indicated, but much of the work is shared by all collaborators.

3 Beamline

Our πE3 beamline has been described in detail in previous Annual Reports.[14–17] In 2006, we made no major changes to the beam tunes or configuration of elements. The extinction factor (EF) strayed little from a typical value of 830 and, apart from changes to the primary proton targeting, the muon flux held steady at 7 MHz. The choice of extinction factor and muon flux, which are inversely correlated, reflect a compromise between the competing demands of minimizing the statistical error and the systematic error contribution from the

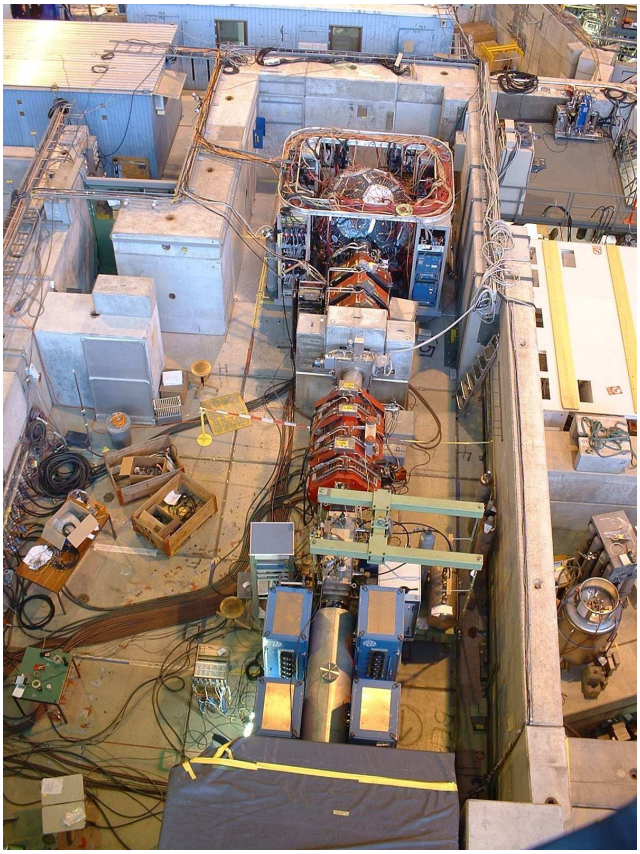

Figure 2: An aerial view of the MuLan experiment installed in the πE3 experimental area. The MuLan detector is visible at the top of the image; the MuLan kicker can be seen at the bottom.

flat background. In general, the properties of the beam are good: even at 7 MHz, the beam spot size is relatively small and no noticeable microstructures exist at any relevant time scale for our experiment.

One of the important findings in 2006 was the sensitivity of our rate on PSI proton targeting (horizontal steering on the E target). The E target is viewed on opposite sides by our πE3 line and the μE4 (LEM) beamline. During the run, the two experimental teams worked out a good compromise for steering the beam: while both experiments seek the highest rates, it was discovered that the MuLan rate could be increased by more than 25% with an acceptably small decrease to LEM. To optimize both efforts in the future, we should specify more precisely the desired targeting parameters. Special care should be taken following a shutdown for beam development or machine maintenance, after which the targeting tends to shift significantly. It is also important to note that the rate increase produced by optimal steering was accompanied by an improved extinction factor; that is, the increase was not just "more" beam, but "better" beam.

4 Electrostatic kicker

The electrostatic kicker [18] provides the pulsed beam for the experiment. Located just upstream of the electrostatic separator, it consists of two pairs of 0.75 m long by 20 cm wide electrode plates, mounted 12 cm apart. Normally, each plate is powered by one of four MOSFET stacks, in separate cabinets. For experimental operation, the kicker is typically operated at a constant frequency of approximately 37 kHz. For a 5 μs *accumulation* period, with the plates grounded, the beam can pass through to the target, undeflected. During the following 22 μs *measurement* period, with the plates at $\pm$12.5 kV, the muon beam is deflected and does not reach the center of the detector.

The kicker system worked very well during three months of operation by the MuCap collaboration prior to our experimental run; unfortunately we faced some difficulties, detailed below. However, after repairs, the kicker system proved very reliable and ran stably for the remainder of our run.

4.1 Failure of cards

After the very smooth and stable operation of the kicker during the MuCap beam time, we experienced unexpected trips of the system due to an overcurrent in one of the four cabinets driving the system. Despite the exchange of two broken MOSFET cards, the trips continued, always in the same cabinet. An extended search for the origin of the trouble, including a careful cleaning of crucial parts such as the high voltage vacuum feed-throughs, did not solve the problem. Eventually, suspicion turned to the electrostatic separator just downstream, which was sparking more strongly than during the MuCap run. The faster cycle time of the MuLan mode results in a higher current-draw than the MuCap "muon-on-demand" mode. Kicker current spikes caused by separator sparks were large enough to trip the kicker's overcurrent shutdown circuits. By carefully conditioning the separator on every ramp-up, we significantly reduced the frequency of kicker trips.

4.2 Vacuum breakdown

At the beginning of our production run, a loss of vacuum combined with operator error, resulted in considerable damage to several MOSFET cards. Since the number of spare cards remaining was insufficient to permit four cabinet operation, and repair of the cards could only be done by experts at TRIUMF, we decided to drive the system with just two cabinets. The increased capacitive load leads to a longer, but still acceptable, rise time for the high voltage. The work to inspect, modify and test functionality took us 3 days. After the repairs, the kicker was very stable once again.

In order to prevent future damage due to any vacuum breakdown, we installed a new interlock system which switches the kicker off whenever dangerous vacuum conditions are reached. TRIUMF has already repaired the damaged cards and returned them to us, along with detailed instructions for performing similar work in the future. We are planning to train several people within the collaboration to maintain, repair and calibrate the MOSFET cards on-site at PSI. In addition, we plan to build and calibrate additional replacement cards before the next run.

5 New vacuum muon corridor

In last year's report, we discussed two options to control or minimize errant muon stops or scatters from the material placed between the end of the beamline and the center of the ball. In the 2004 and 2005 configuration, muons exit the vacuum pipe through a thin window, pass through the EMC wire chamber, enter a helium-filled bag and then strike a large-diameter target. Detailed simulations, performed to investigate front/back asymmetry in the 2004 data, clearly show that some small fraction of muons (typically less than 10^{-3}) are scattered around the target, backscattered, or stop upstream of the target. Unlike target stops, where the large magnetic field and detector symmetry minimize the effects of muon polarization, non-target stops can produce changes in the measured lifetime which are difficult to quantify.

For 2006, we built an all-vacuum muon corridor to the target. The idea is to extend the vacuum pipe into and out of the MuLan ball and place a rotatable target in the pipe at the position of the ball center. This configuration required that the EMC wire chamber be placed *downstream* of the target, at the exit face of the ball rather than the entrance face. Muons can then only pass through to the EMC when the target is rotated out of their path. Special runs were taken each shift to measure the beam position, which was found to be very stable throughout the ten weeks of data taking.

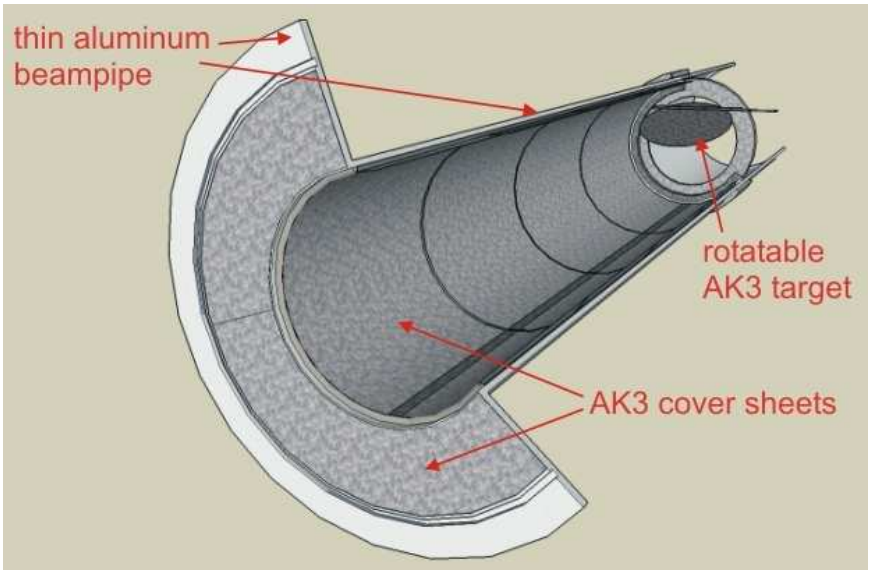

Figure 3: The "all-vacuum" muon corridor showing the rotatable target, viewed from upstream. The upstream portion of the beampipe is lined with AK-3 material.

Figure 3 shows the new vacuum pipe and rotating target. The target is made of the high-magnetic-field AK-3 alloy while the vacuum pipe is made from thin-walled aluminum. Although idealized beam dynamics simulations predict no stops outside the target, measurements demonstrate that roughly one in ten thousand muons scatter from upstream obstacles and then stop in the walls of the entrance corridor. To dephase the spin vectors of these errant muons, the entire vacuum pipe upstream of the target was lined with AK-3. Further,

a special set of plunging scintillator counters was used to measure the muon flux as a function of radius inside the vacuum pipe. Preliminary analysis of this data indicates that the errant muons have no significant effect on our lifetime measurement.

6 Target setups and systematics

The primary target for the 2006 run was a thin disk of Arnokrome-3 (AK-3), a ferromagnetic alloy composed of iron (60%), chromium (30%) and cobalt (10%). The disk, of radius 15 cm and thickness 0.5 mm, was located at the geometrical center of the MuLan ball and oriented with its face perpendicular to the beam axis; see Figure 3. The disk was mounted in the 20 cm diameter beam pipe on a vacuum feed-through that permitted the target to be rotated into the beam for production running, and out of the beam for beamline tuning. The internal magnetic field of roughly 0.5 T lies in the plane of the disk, perpendicular to the beam axis.

AK-3 was chosen in order to destroy the polarization of the muon ensemble in the target material. In the B-field of the AK-3 target, the muon spin precesses about the field axis with a precession frequency of roughly 100 MHz. Consequently, the original polarization of incoming muons, which is initially aligned along the beam axis, is dephased (i.e. scrambled) by the combination of the different muon arrival times and the fast muon spin precession - thus yielding an unpolarized ensemble of stopped muons. Note that AK-3 is well known from μSR studies as a polarization-destroying material. For example, in recent μSR measurements by Morenzoni *et al*,[19] using a DC muon beam and zero external field, they found no evidence of residual polarization at times $t \gtrsim 50$ ns after the arrival of the muon.

Due to the importance of any residual polarization of the muon ensemble - and its possible influence on the lifetime determination - the collaboration has conducted a number of different systematics measurements with different target configurations. These studies are summarized below:

1. By rotating the target section of the beampipe assembly through 180°, production data were collected with AK-3 field orientations of left-to-right and right-to-left. The change in AK-3 field orientation reverses the muon spin precession direction from initially rotating upwards to initially rotating downwards. Note that the top/bottom symmetry of the MuLan setup is slightly broken by the dipole magnet of the πE3 separator, leading to a small (few degree) vertical component in the spin direction of the muon beam.
2. Systematics data were collected with an elliptical AK-3 target that was oriented at a 45° angle relative to the beam axis. The 45° pitch enabled target configurations with a longitudinal magnetization component, i.e. a magnetization component directed either upstream or downstream along the beam axis. Such configurations thus permit the study of any effects of a longitudinal B-field component on the muon spin relaxation and the muon spin dephasing. At some level - because of the domain structure of the AK-3 material - such longitudinal fields are inescapable in the standard (left/right) orientation of the AK-3 production target.
3. Systematics data were also collected with targets of copper and aluminum. Both targets were ~ 19 cm disks, mounted in the vacuum pipe at the ball's geometrical center, and oriented at 90° to the beam axis. Both materials are well known from μSR

studies as polarization preserving materials, and therefore the initial muon spin will slowly rotate in the earth's magnetic field with a precession frequency ∼100 kHz. The small dephasing in the Al/Cu targets permits the study of the effects of the residual polarization and resulting counting-asymmetry on the lifetime determination.

4. To help study the effect of non-target stops, sets of measurements were made with both the AK-3 (polarization destroying) target and the Cu (polarization preserving) target moved either upstream or downstream of the geometrical center of the MuLan ball.

7 Detectors

7.1 EMC

The Entrance Muon Counter (EMC) is a wire chamber that records the time and position of beam muons. The EMC is filled with a gas mixture of 70% CF_4 and 30% isobutane. It has 25 μm aluminized mylar windows at the front and back. There are planes of 20 μm diameter tungsten wires; the EMC is mounted on a stand at the downstream end of the beam corridor, with wires oriented in the vertical and horizontal directions. There are one hundred wires (96 active) on each plane with 1 mm spacing. To reduce the number of channels, the wires are connected together in pairs; thus the detector has 2 mm resolution. The wire planes are separated from the outer windows and each other by 12.5 μm aluminized mylar sheets which are held at a positive high voltage. The EMC can function with instantaneous beam rates up to 10 MHz.

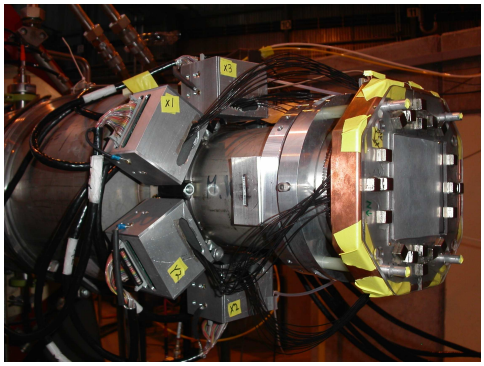

(a) A photograph of the EMC from 2005. The wire planes are housed in the octagonal structure at right.

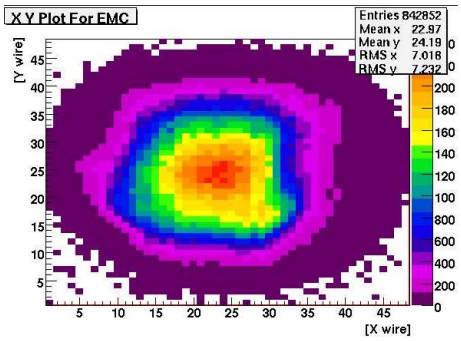

(b) An X-Y beam profile from the 2006 run, taken with the EMC.

Figure 4: The EMC is a wire chamber designed to provide an X-Y profile at high beam rates.

The chamber is read out by six amplifier/discriminator boards which produce logic pulses between 40 ns and 100 ns long, depending on the analog pulse's time-over-threshold. A custom FPGA-based readout module was used to digitize the time and position of hits, with a user selected prescaling factor during the beam on period. Prescaling is required during

the high rate accumulation period, but not during the measurement period, when the rate is about one thousand times lower. A beam profile from the EMC is shown at the right of Figure 4.

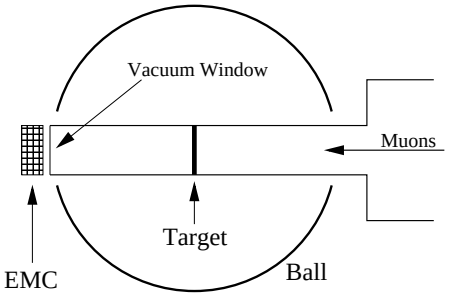

Figure 5: In 2006 the EMC was positioned at the back of the ball. The target was rotated out periodically to check the beam spot with the EMC.

As in previous years, the EMC was used to monitor the stability of the beam tune. However, since the beampipe was installed through the ball to bring muons to the center of the target in vacuum, it was no longer possible to make measurements with the EMC during data-taking. Instead, the EMC is placed at the back of the detector, as shown in Figure 5. Approximately three times a day, the target was rotated out of position, allowing beam profile measurements.

7.2 The scintillator tiles of the MuLan ball

The MuLan ball (visible at the top of Figure 2) was fully commissioned during the fall 2003 run and has been used successfully in the 2004, 2005 and 2006 data-taking runs. It consists of 170 scintillator tile pairs arranged in a truncated icosahedron. In 2006, all 340 tiles were fully functional and read out. The WFDs allow us to inspect the pulse size and shape from each detector *in situ*. This aids in calibration, monitoring, and debugging of the tiles.

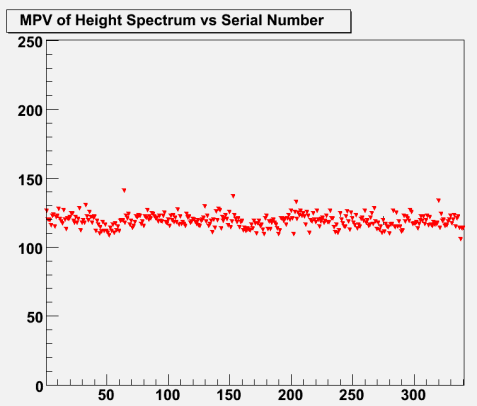

Figure 6: Snapshot of most probable value for Landau fits to the individual tile pairs after high-voltage adjustments. The detector tile number is on the horizontal axis; the WFD ADC value for the peak of the Landau is on the vertical.

We used online analysis of the WFD data to calibrate the PMT high voltages. Because essentially all decay positrons passing through scintillator tile pairs are minimum ionizing, the energy deposited per event is described by a Landau distribution. The pulse height distribution for each tile was therefore fit to a Landau function. The fit yielded a most

probable value (MPV) parameter for the distribution, which is proportional to the gain. Empirically, the MPV scales with the eighth power of the voltage; the PMT voltages were set so that the MPVs would all be approximately 120 ADC counts (full scale on the WFDs is 256). The high voltages were then iteratively adjusted. Figure 6 shows the fitted MPVs for all detectors at the start of the run. There was a slow rise in the gain of many of the PMTs over the course of the 2006 run which will need to be accounted for in the analysis; however, these changes take place on a long time-scale compared to the measurement period or muon lifetime and so should not pose a serious problem.

8 Electronics

The major MuLan electronics include timing control devices, a master clock and clock fan-outs, and the waveform digitizers used to read out the photomultiplier signals. None of these systems underwent major changes in 2006; incremental firmware updates dominated the evolution of these subsystems.

8.1 Clocks and timing control systems

There were three clocks driving the MuLan experiment in 2006 - two under our control, and a third which placed constraints on our control systems:

1. Fill and DAQ control: Constraints on data throughput require synchronizing the kicker cycle with the acquisition electronics, which take different types of start/stop control signals. To control a voltage stability systematic, the kicker must be operated in a continuously cycling mode, even when the DAQ is not collecting data. In the 2005 run period, these two tasks were performed by independent computer-controlled FPGA based Programmable Pulse Generators (PPGs) along with external synchronization logic. A major restructuring of the PPG firmware allowed us to combine both of these tasks into one device. This system allows the PPG to drive the kicker continuously in its 22 μs on, 5 μs off pattern, while simplifying synchronization with the DAQ electronics and the MIDAS control computer. The DAQ macro-cycles remain of order 100 ms.
2. Timebase signal: The main timebase for the electronics during the 2006 run period remained an Agilent E4400B frequency synthesizer operating near 451 MHz. We have previously confirmed that this device has absolute frequency precision and accuracy and both short and long term frequency stability at the 10^{-8} level, exceeding the experimental requirements. The absolute frequency of this clock is hidden from the analyzers at the 200 ppm level until the analysis is certified as complete.
3. Beam RF signal: The collaboration does not control this clock, but it can still have an impact on the systematic errors of the experiment. If we make poor choices in the frequencies we do control, beating against the beam RF can result in slowly varying signals in the reconstructed data stream.

The clock distribution network consists of a number of low skew fan-out units for the main WFD timebase, and a number of NIM and VME modules to distribute the control signals to the acquisition electronics. This equipment all performed as expected throughout the run period.

8.2 500 MHz waveform digitizers

8.2.1 Functional Description

As in 2005, the main acquisition electronics for the 2006 run period was a full 340-channel complement of 500 MHz waveform digitizers (WFDs). Having a digitized waveform for each event simplifies tile debugging and calibration, and allows us to minimize or directly monitor a number of important systematic errors, including dead time, pileup, and gain and timing shifts.

Each WFD channel consists of analog input buffers, a 600 MSps, 8 bit Maxim Flash Converter, half the resources of a Xilinx Virtex II FPGA, and a 128 Kibit $\times$ 40 FIFO memory. All channels on the board share additional power, control electronics, clock resources, and VME64x communication paths. The dedicated channel FPGA resources perform triggering and data formatting tasks.

FPGAs provide a great deal of power and implementation flexibility. An FPGA consists of a set of logic elements (logic gates, multipliers, multiplexers, look-up tables, etc.), clock management resources, and input/output resources. These resources communicate over a network of interconnect lines that are configured at power-on from an off-chip flash memory.

The configuration program (the *firmware*) is written using the industry standard VHDL hardware description programming language. In MuLan, we have used this flexibility to provide a number of behavioral packages, or *personalities*, to meet the digitization needs of different detectors.[1] During the 2006 run, the majority of channels used a "Fill and Offset" personality. This personality provides for zero suppression for the vast majority of uninteresting pedestal data with storage of fixed length blocks of pulse data when the analog signal rises over a programmable analog threshold. Each of these data blocks is tagged with a 16 bit fill cycle count within the DAQ macro-cycle and a 16 bit clock tick offset within the fill. These time stamps, along with the pulse data itself, allows us to measure the time and area of each pulse with high precision and accuracy.

The full complement of WFDs is split evenly among six VME crates, each of which is controlled and read out via a frontend control computer. Because a coincidence is required between inner and outer scintillator tiles, each pair is wired to the same WFD module, in order i) to minimize data lost in the event of a WFD failure, and ii) to guarantee identical clocks to paired channels. Similarly, control of systematics requires simultaneous collection of data from point-wise pairs on opposite sides of the target; thus, opposing pairs are wired to the same WFD module.

8.2.2 Progress and status

The 2005 run period was the first experiment-wide integration test of the WFD acquisition. While successful, it pointed out a number of subtle issues in both the WFD firmware and the DAQ control and readout. During the course of that run, approximately twelve WFD firmware revisions were released, solving most of the major issues. Unfortunately, the remaining bugs may have caused subtle, difficult to correct systematic issues in the 2005 data set.

[1] A number of these same WFDs are being utilized by the MuCap experiment to read out scintillators using a different personality from that of MuLan.

The time between the 2005 and 2006 runs was used to debug and correct the issues discovered during the run. During the setup period for the 2006 run, only three updates were installed to address issues discovered on-site. One of these early versions contained a timing bug that only affected data integrity on a small number of boards when operated at 500 MHz; a satisfactory workaround which eliminated the problem on all boards was to operate the experiment near 450 MHz. While this bug was fixed before high rate data taking commenced, we chose not to adjust the already blinded clock frequency. The final firmware version used for the bulk of the run operated without any known problems that affected the data integrity, although it had an intermittent start-up issue; the effects of this bug were easily identified by the frontend and corrective action could be taken immediately. There are no known data integrity issues with any of the 2006 firmware revisions.

There was one desirable feature missing in 2006, which we call *reflexive triggering*. In the 2006 run, the WFDs were operated in a mode where each channel was independently triggered. However, understanding singles distributions, small amplitude pulses, and pileup distributions would be easier if both channels in a pair could be recorded when either crosses threshold. We intended this feature to be available in 2006, but bugs in the implementation prevented us from using it. This feature, and a number of minor cosmetic issues, has been successfully addressed since the 2006 run. We believe the firmware is already prepared for 2007 production running.

9 Calibration and testing systems

In 2006, the reliability of the LED testing system was improved, and we installed a nitrogen laser system for timing calibration and timing stability tests.

9.1 LED testing system

The LED system is used for testing many aspects of the detectors, electronics, and DAQ, and for monitoring timing features. An LED is mounted on each of the photomultiplier tubes of the MuLan ball. Fast light output (a few nanoseconds) is made possible by a high-current driver circuit mounted with the LED in the detector housing.[20] The electronics to power and trigger these boards are contained in 12 driver boxes, which receive triggers from either a VME module controlled by one of the data acquisition computers, or from a pulser. The computer-controlled VME module fires a selected set of channels at programmable rates, which allows for a multitude of tests. For instance, individual channels can be pulsed to test maps of the detector configuration, or random time spectra can be generated to look for anomalies in the data processing.

Improvements to the LED driver circuits greatly increased their reliability. During the installation phase, they helped provide diagnostic tests of problematic channels, running asynchronously or as part of the standard data acquisition. Indeed, just before the MuLan ball was craned in, the LED system provided a source of pulses for a *mock run*, in which most aspects of the standard data acquisition and readout were tested. An example of accumulated LED signals as measured by the WFDs is shown in Figure 7. The 2-D histogram shows the PMT response to LED pulses with the development in time along the x-axis, the pulse height

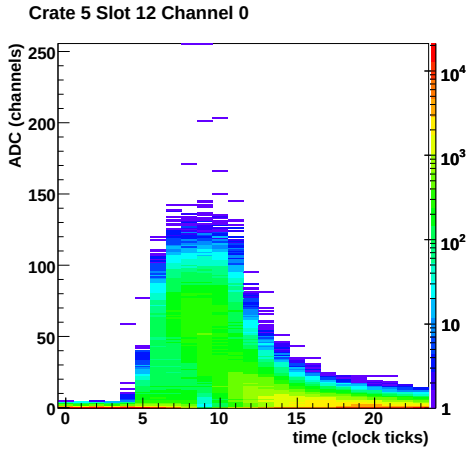

Figure 7: Accumulated LED pulse height spectra collected during the 2006 run.

along the y-axis and a color-coded intensity.

9.2 Laser timing calibration system

In the 2006 run, a laser system was added to allow systematic error studies on the timing accuracy and stability of the counters and WFDs. The laser data can be used to check whether the recorded time of pulses are pulled by other nearby pulses, which could cause early-to-late shifts in the recorded times of events. The nitrogen laser light is split to a reference PMT (remote from the detector sled, to shield it from potential detector generated interference), a reference photodiode, and 24 detector tiles on the ball (10 tile pairs and 4 individual tiles). During data-taking, the laser was run approximately 10% of the time. It was pulsed at a constant rate, unrelated to the kicker period, to spread the laser pulses evenly over the kicker cycle. This results in a small increase in the constant background but permits these data to be used in the lifetime analysis without fear of systematic effects.

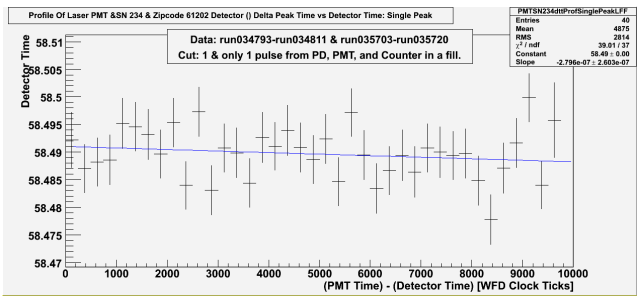

Figure 8: The average time difference for laser pulses between a single tile on the detector and a reference PMT shows no trend versus time in the measurement period (the slope is consistent with zero at the 1σ level).

Since the reference PMT is unaffected by the kicker cycle, we can check for a change in the time offset between the laser in the PMTs on the ball and in the reference PMT. If this time offset is stable within the measurement period, there is no early-to-late shift in detector timing. Figure 8 shows a preliminary plot of the average time offset between a single detector and the reference PMT versus time in the measurement period. Although a more detailed

analysis will be performed, there is no indication of a systematic shift with time-in-fill. We will also use this system to determine the absolute timing precision of the entire acquisition chain.

10 Data acquisition

The MuLan data acquisition system (DAQ) provides:

- control and readout for the 340 channels of 500 MHz WFDs and the EMC,
- control for the kicker,
- control and readback for the beamline magnets and PMT high voltages,
- control for the LED and laser systems,
- readback for separator status, and
- readback of various environmental conditions including temperatures, magnetic fields, etc.

The DAQ is based on the MIDAS system [21] and has a distributed, modular design with multiple frontends (FEs) and a single backend (BE) for event building and data storage.

During 2006, our DAQ efforts were focused on resolving some difficulties and limitations that we encountered during the 2005 run, caused primarily by high data rates from the new WFDs [17]. Improvements to the data acquisition for 2006 include:

- DAQ layout: All components of the DAQ were thoroughly investigated for possible bottlenecks. The main problem concerned data throughput on the BE processor: there was a severe restriction in logging data to physical storage over the local computer network. To avoid transferring composed events over the network, both the event builder and the event logger were run on the BE. Distributing various tasks among DAQ processors and subnetworks produced additional improvement. The final DAQ configuration in 2006 is shown in Figure 9.
- Data stream compression: In the original design [17], we planned to compress the digitized pulses on-line. After careful study of various compression algorithms, both custom and off-the-shelf, we decided in favor of the ZLIB library [22] which provides both high speed and good compression ratios. However, even the fastest compression algorithms we considered were not able to compress a single high rate (up to 30 Mb/s) data stream on the backend. Therefore the task of data compression was moved upstream, to the six dedicated WFD frontends, each of them compressing the data from its VME crate.
- Multi-threaded frontend: Imposing the additional task of compressing the data stream on each FE processor is a significant bottleneck at very high data rates. To increase the throughput of the WFD FEs, we developed a multi-threaded version of their programs which executes two tasks in parallel (the data readout from WFD modules and data compression) on dual processor machines.

Additional minor improvements included: i) improving the stability of the DAQ by fixing bugs in various DAQ components and optimizing various parameters (run stop/start transient time, watchdog timeouts, buffer sizes, etc.); ii) upgrading the BE PC; iii) upgrading the storage systems with an additional LTO3 tape robot; and iv) upgrading the connection to the PSI network to gigabit speed.

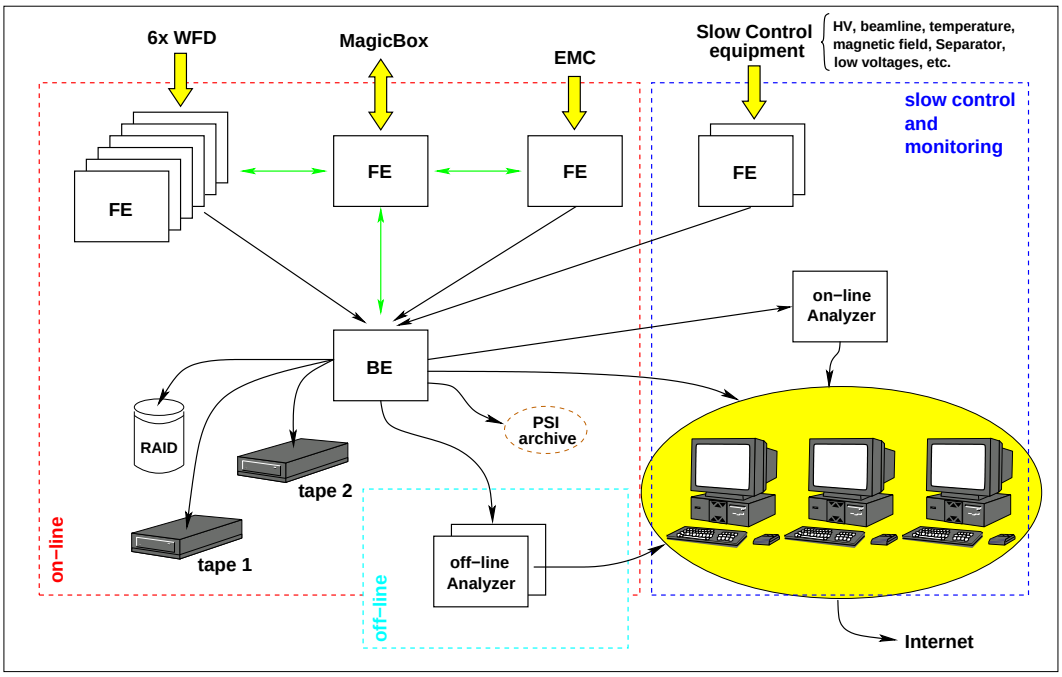

Figure 9: Schematic of data acquisition system (DAQ) showing the frontends (FE's), backend (BE), analyzers, slow control and setup for data monitoring and control. To optimize performance, the DAQ hardware was distributed between on-line, off-line, slow control/monitoring subnetworks.

The performance of the 2006 DAQ, as studied in a test setup, is illustrated by Figure 10. Almost dead-time free operation was possible at muon decay rates (per front end) up to 1 MHz. At very high rates (above 1.4 MHz) the operation with no data compression (red triangles) was problematic due to buffer overflows, whereas by compressing the pulses on-line, it was possible to handle decay rates up to 1.8 MHz. In real experimental conditions, the highest number of detected muons per day was about 2.5×10^{10}. At this rate, the dead time of the DAQ was about$\sim$ 20%, which is a significant improvement relative to 2005.

11 Real-time data monitoring

Since the real-time data monitoring system developed in 2005 was very successful, it did not undergo any significant changes in 2006. The major upgrades concern the user interface, which was thoroughly rewritten to increase functionality and ease-of-use:

- A user-friendly framework was developed based on the PHP Web programming language.[23] It serves as an interface between a user's ROOT macros [24] and a Web server. By following a simple set of rules, any user can now easily add his own ROOT analysis macros into the system. The output of the ROOT macros (both graphical and textual) is displayed in an HTML page via HTTP, which is very useful for remote monitoring of data taking. For local monitoring, the regular graphical interface of ROOT can still be used instead, which makes the system more universal. There is an option either to

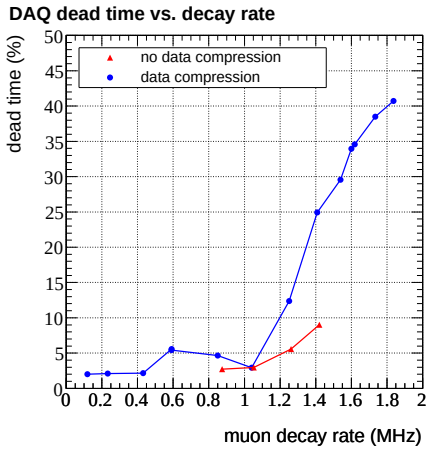

Figure 10: Dead time of the DAQ as function of muon decay rate measured in test setup where pulses from electrons were simulated by LEDs. Two modes of operation were studied, with data compression (blue circles) and with no data compression (red triangles). The DAQ is able to acquire up to 1.0×10^6 muon decays per second with negligible dead time. At higher rates the throughput of the DAQ decreases due to CPU overloading by the task of data compression, whereas the operation with no data compression becomes problematic due to buffer overflows.

retrieve the histograms directly from the on-line MIDAS analyzer or to read the data from run-based ROOT data files stored on disk.

- The possibility of logging the slow control data into output ASCII data files was added into slow control frontends, which can be controlled via the MIDAS on-line database. The new feature allows us to monitor the status of slow control equipment on-line, without the need to analyze MIDAS data files.
- A standalone application was written in the Java programming language for local monitoring of data from slow-control frontends. The application allows us to export the ASCII data files into the on-line MySQL database [25] of the experiment, which is useful for off-line data analysis.

An example of the new user interface of the data monitoring system is shown in Figure 11. A user macro can be executed with a mouse click on the item listed in the left frame. The right frame contains the output of the macro. In this example, the macro `WFD_hits_ball` serves to monitor the angular distibution of hits and the counting rates of all 340 MuLan detectors.

Although there were crashes when the rates were changed suddenly, the DAQ generally ran smoothly, with little or no operator intervention required. During periods of smooth data production, off-site collaborators ran *remote shifts*, easing the burden on those at PSI. Just giving on-site collaborators a few hours for a leisurely dinner proved a great boost to their morale. In case of emergencies, which arose from time-to-time, experts could quickly be summoned to the barrack by phone.

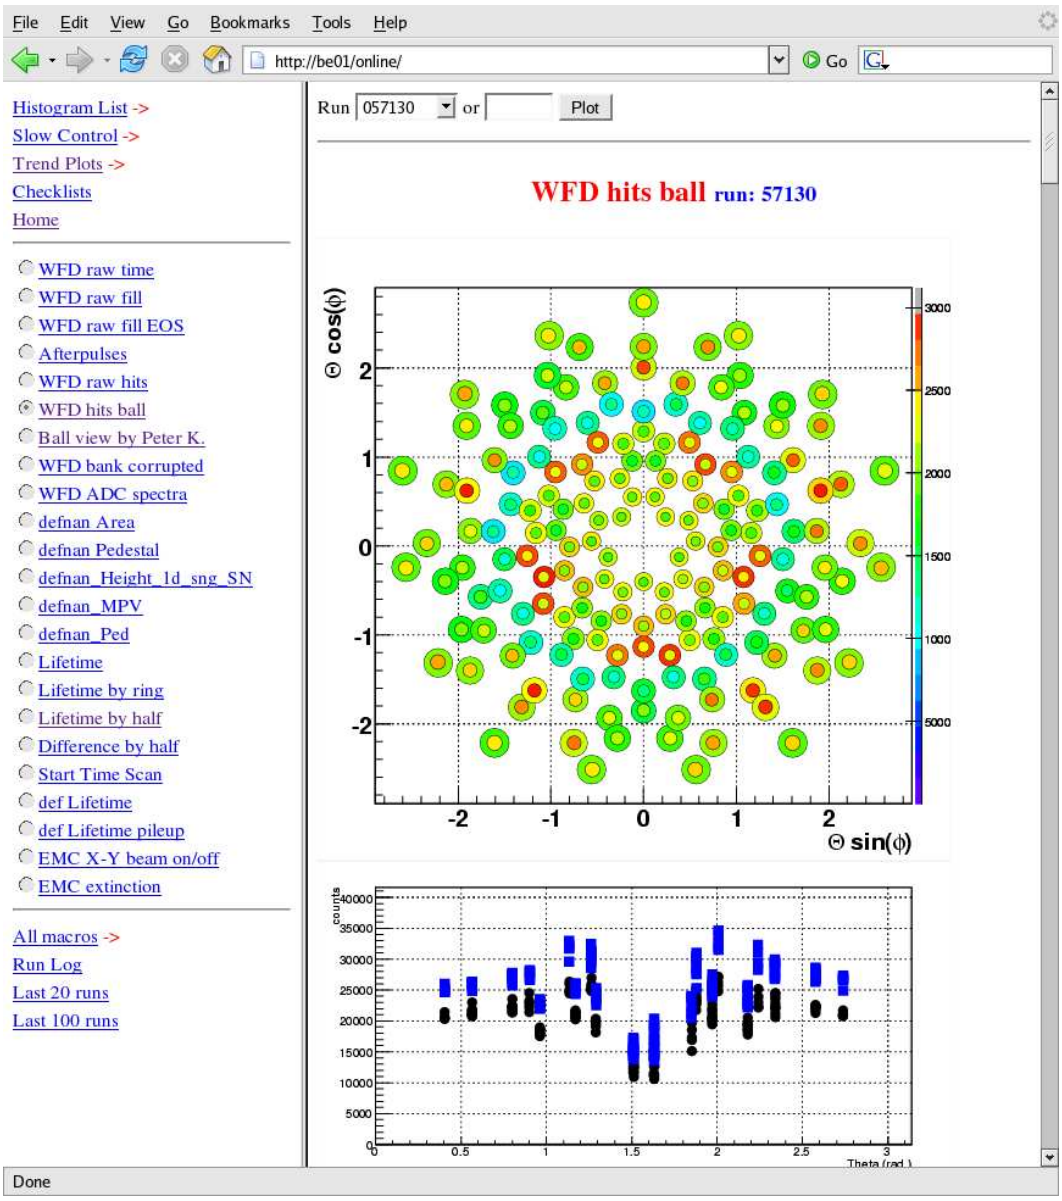

Figure 11: Real-time data monitoring for the experiment is possible through a standard web browser.

12 Analysis of the 2004 data

The analysis of the μ^+ lifetime data (see Figure 12) collected at the end of 2004 has recently been completed. At the Users' Meeting in February we will present the most precise value of the muon lifetime to date, plus an updated value of the Fermi constant based on the new world average. Although we do not present a final number in this document, it is clear from the figure on the front cover that our result is in excellent agreement with the previous world average.

The 2004 data set was handicapped by frequent failures of the kicker and the use of

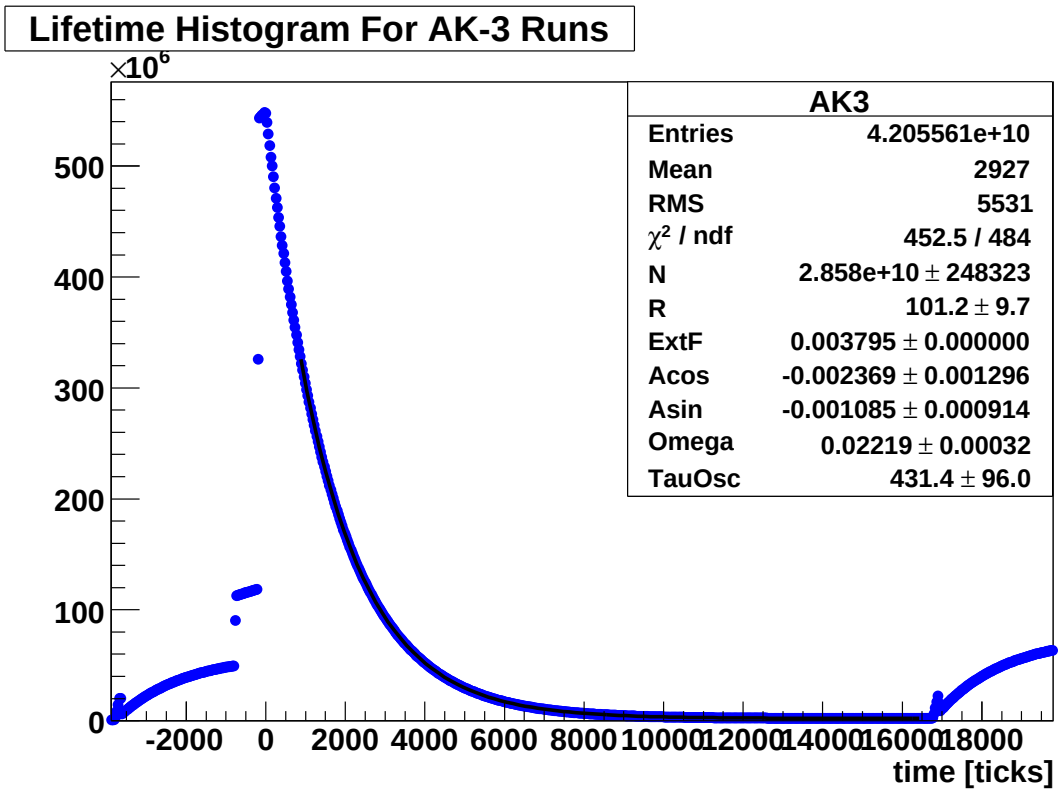

Figure 12: Full statistics lifetime histogram for the AK3 data set with fit. The error on the fit parameter R (a blinded measure of the lifetime in units of ppm) shows the statistical precision of this data set to be 9.7 ppm.

multihit TDCs instead of WFDs. This resulted in substantial time spent investigating systematic uncertainties associated with a larger background, higher pileup rates, and problems introduced by the electronics, all effects that will be greatly diminished in the later data sets. Despite these obstacles, we have achieved a (statistically dominated) measurement error of roughly 11 ppm, a precision roughly two times better than the current world average.

We have addressed all the principal systematic errors listed in our last report:

- TDC timing shifts,
- Pileup and dead time corrections,
- Flatness of the background,
- Gain changes in PMTs,
- Changing of the ground on signal lines,
- Rate dependent timing shifts,
- Residual polarization of the muons, and
- Detector asymmetry from non-target muons stops.

Full details on the systematic studies can be found in Appendix A, including a preliminary systematic error table, Table 1. A final systematic error table will be presented at the Users' Meeting.

13 Analysis of the 2005 data

In 2005, for the first time, PMT signals from the entire ball were read out with waveform digitizers (WFDs). Like any new system, the WFDs had their share of bugs, on which we reported last year.[17] However, the lessons learned from the 2005 data set proved to be invaluable for making strategic choices for the 2006 run. For example, it was clear from the 2005 data that to resolve overlapping pulses, we would have to lower the phototube gain, which had been set to a level appropriate to the discriminators used in 2004. In 2006, we reduced the number of ADC samples per trigger from 32 to 24. By discarding these extra, trailing ADC samples, we reduced the data rate without significant loss of information about pileup, pedestal, or background pulses. The 2005 WFD data also provided the full-width, half-maximum (FWHM) of the decay electron pulses, which was used in the pileup analysis of the 2004 dataset. WFD data from 2005 confirmed that the front/back lifetime asymmetry seen in the 2004 analysis was not an electronics artifact but a property of the target geometry. Finally, the high rates associated with WFD readout provided valuable experience for the data acquisition team. Upgrades to the frontend programs and computer layout increased the livetime of the DAQ from 50% in 2005 to 85% in 2006, and despite an increased beam rate, extended the mean time between operator intervention from a few hours to several days.

A few issues make us wary of including the 2005 dataset in our final physics analysis. Foremost is the target geometry, which changed between 2005 and 2006. The systematic uncertainty on non-target stops in the 2004 and 2005 data is a few ppm. This uncertainty, discovered after the 2005 run, led us to redesign the muon corridor and target for 2006. In addition, during the 2005 run, several changes were made to the WFD firmware to correct bugs; all of them introduced, changed, or removed subtle behavioral problems:

- We discovered that having the full complement of boards installed in a VME crate caused silent readout errors leading to uncorrectable data corruption. As a result, we could only read out 70% of the channels in 2005.
- There was a "hole" in the WFD trigger, where the WFD was blind to a second trigger pulse between 32 and 36 clock ticks after the primary trigger. This causes a large distortion of the channel autocorrelation and pileup spectra, and a subtle distortion of the lifetime spectra.
- Because paired inner-outer channels were run in different clock domains within the WFD, there was a 4 clock-tick uncertainty in the relative times of the paired channels which dramatically complicates the pileup correction and timing stability studies.

By themselves, these issues are not fatal; a full data analysis could still be performed and the lifetime extracted. Given our manpower resources and the limited statistical power of the 2005 data set (the 2005 data comprises less than 10% of that collected in 2006), we feel that the significant additional effort to understand the systematic challenges is not justified. We learned much from the 2005 run that made the 2006 run a success; we have since turned all our analysis efforts to the latter.

14 Analysis of the 2006 data

14.1 Pulse reconstruction

The shift from TDC data to waveform data gives us access to a vast amount of information that was not available to us in our 2004 analysis. Most obvious is the much higher timing resolution on individual events. The most important impact, however, will be on reduction of systematic uncertainties:

- We can separate pulses that are much closer together in time than with a discriminator/TDC combination, giving us much better pileup reconstruction and rejection.
- We can directly monitor pulse shape structure and PMT gains throughout the run, and more importantly, search for early-to-late changes in the fill that might affect the lifetime histograms.
- Along with the laser system, we can directly monitor the timing stability within the fill caused by potential auto and cross-correlated effects (afterpulsing, pulse-to-pulse timing pulls, etc).

To realize these goals, we apply a modified version of the template fitting method pioneered by the Brookhaven E821 $g-2$ experiment. We start from the empirical observation that PMT pulse shapes are nearly independent of pulse amplitude (see Figure 13). We can

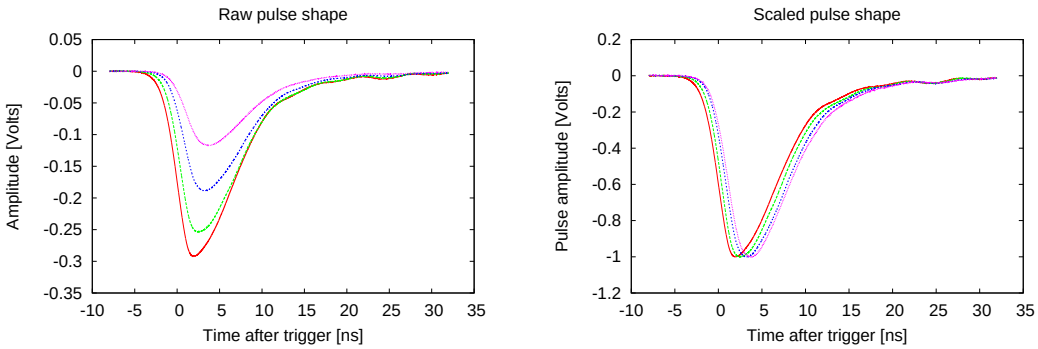

Figure 13: Empirical studies show that PMT pulse shapes are largely independent of pulse height, allowing us to use a template pulse fitting method to extract peak timing from digitized data.

use this observation to build a physically meaningful *average pulse template* with a well defined pulse shape and peak time. We can then use that template to extract individual decay event timing and pulse shape data, including pedestal, height, area, rise and fall times, etc. Both of these procedures must necessarily start with a pulse finding algorithm.

14.1.1 Pulse finding

Before characterizing a pulse event, we have to find pulse candidates. The zero suppressed data stream is scanned, and time-adjacent blocks are merged to form *pulse islands*. Each island is then searched for local maxima that exceed a software trigger threshold, and an initial guess is made as to the pedestal. The local maxima are labeled as pulse candidates

if the rise, fall, and separation from neighboring maxima are large enough. Additional heuristics are applied to handle off-scale pulses and to address islands with evidence of multiple candidates.

14.1.2 Constructing an average pulse template

The template method is based on the statistical reconstruction of an average pulse shape. To this end, we first need to ensure that the true pulse peak is distributed uniformly between clock boundaries; this we have done by careful choice of DAQ and digitization clock frequencies (see Section 8.1), and the use of an analog trigger which is independent of all clocks. Next, we define a proxy measure of peak position within the clock tick, which we call the *pseudotime*, ψ: label the ADC samples a_i, with $a_{\max}$ the maximum sample, then

$$\psi = \frac{2}{\pi} \tan^{-1} \left(\frac{a_{\max} - a_{\max-1}}{a_{\max} - a_{\max+1}} \right) \ . \tag{1}$$

The function varies continuously on $[0, 1)$, minimizes the impact of pedestal and amplitude variations, and is in 1-to-1 correspondence with the true sub-tick peak-offset within the bin.

We then run over the data, channel by channel, looking for islands where the pulse-finding algorithm finds only one pulse candidate. For those candidates, we build the pseudotime histogram. We invert this histogram to build a mapping function that takes the non-flat pseudotime into the flat sub-tick offset: $\tau : \psi \to t$. Using this mapping τ, we can shift each candidate by the calculated offset so that the true peaks of all pulses line up and add the samples to a finely binned (10 ps) histogram which extends from about 10 ns before the pulse peak, to 40 ns beyond. After processing several million pulses, the average value in the first few bins, a pedestal of sorts, is subtracted from every histogram bin. Because of the analog trigger, the average true amplitude of the pulses does not depend on the pseudotime. Thus, the pedestal-subtracted histogram (usually scaled to unit amplitude) reflects the true average pulse shape. If the template is insufficiently smooth, we repeat the procedure, but this time using the preliminary template of the previous iteration to refine our estimate of the pedestal and amplitude.

We test the stability of the pulse shape, and our understanding of the detector and software, by varying algorithm parameters and looking for template deviations, as well as testing pulse shape stability across the fill and across the run. It is important to note that laser pulses, which have different shapes from those produced by decay electrons, must be treated separately.

14.1.3 Pulse fitting

Once the average pulse shape has been constructed, we can extract fit times and amplitudes for all events in the data set. First, we apply the pulse finding algorithm, without pulse isolation constraints. In principle, for each pulse candidate on the island, we calculate an initial guess for the pedestal, pulse height and time (from the pseudotime), and then perform a χ^2 minimization against the pulse template. In those cases where the χ^2 value is poor, we have generally missed a pulse that is close in time to one of the identified pulses. In this

case, we add an additional pulse candidate at the point of maximal deviation, and refit; if the χ^2 improves substantially, we accept the additional pulse, otherwise, we reject it.

This procedure excels at identifying and quantifying all types of pulses, dramatically reducing the effective dead time and pileup from nearby events. In the vast majority of cases, however, islands have only a single pulse. For these, there is a time efficient shortcut which we call the *fast fit*. For well separated pulses, the pseudotime mapping gives an accurate value for the true pulse time. In this case, the fitting problem can be reduced to a solution of simultaneous equations, with the pedestal and pulse amplitude as the unknowns, without performing a full χ^2 minimization.

After pulse finding, we perform the fast fit island-by-island. If the χ^2 is sufficiently good, indicating that we have a well separated set of pulses, we do not perform the full fit. In the remaining cases (and a small fraction of all islands, for systematic studies), we resort to the full fit. In all cases, a set of fitted pulse parameters is stored in ROOT Trees for further analysis and histogram construction.

14.2 Lifetime histogram construction

Fitting pulses to the 2006 WFD data results in a list of times, pedestals, and amplitudes for each pulse. Next, pulses from pairs of inner-outer detectors must be joined into coincidences. The data stream of coincidences is then binned to make a lifetime histogram. Events missed or over-counted due to pileup can be reconstructed from the datastream using a *shadow window* technique, just as in 2004: hits in a fixed width window at a fixed offset from a trigger hit are added or subtracted (depending on the type or source of the pileup) to correct the original lifetime histogram. A typical deadtime analysis uses several shadow windows to check the consistency of the deadtime correction. The WFD data has the additional advantage that one can look at the noise under the hardware trigger by characterizing islands from systematic runs where a single trigger produces a long data block. The pileup corrected lifetime histogram can then be fit using a simple three parameter exponential plus background function to extract the lifetime.

14.3 Production Infrastructure

Two algorithmically similar analysis codes are being developed in parallel at separate institutions as cross checks on each other. Once both sets of code are complete and in agreement, the fastest, most computationally efficient functions from each will be combined into a production code that will run at the National Center for Supercomputing Applications (NCSA) at the University of Illinois, on an Intel Xeon cluster of over 1000 nodes. An initial production pass will extract physics events by interpreting the raw ADC data stream as individual pulses. The output of the production code will be a list of hits in ROOT tree format, with hits characterized by pulse amplitude, time, and pedestal, as well as by bookkeeping parameters such as quality of fit and pulse template used. Lifetime and diagnostic histograms for each run, using cuts on physics and bookkeeping parameters, will then be built in a pass through the reconstructed data. The parameters of rare or interesting pulses can also be saved for each run. These smaller histogram files will then be moved from the NCSA cluster to local interactive machines for fitting and further analysis.

14.4 Preliminary Schedule

The 2006 dataset comprises roughly 70 TB on 175 LTO3 tapes, and will take approximately one month to transfer to the mass storage system at NCSA. The transfer will be complete in late January, 2007. On a similar timescale, the first draft of two independent pulse reconstruction codes will be ready for comparison. Once the production code is ready, it will be tested on approximately 25% of the 2006 dataset. Based on our tests, we will submit a proposal in July 2007 to NCSA for the computation resources necessary to process the full 2006 dataset.

15 Beam request

One attractive feature of the MuLan setup is the ability to measure the muon lifetime in different targets. Different categories of target materials include: i) polarization preserving materials such as aluminum, ii) polarization destroying materials such as AK-3, in which the muon spin orientation is quickly scrambled by its fast precession in the high internal B-field of the magnetized material, and iii) polarization destroying targets such as quartz, in which the muon spin orientation is quickly scrambled by muonium formation and its fast precession in a moderate (10–100 G) external B-field. The constancy of the measured muon lifetime across the different target categories is a powerful cross-check on key systematic uncertainties in the muon lifetime determination.

In our 2006 production run, we collected $\sim 10^{12}$ decay electrons from muon stops in magnetized AK-3. In our 2007 production run we propose to collect roughly 10^{12} decays from muon stops split between i) quartz, a muonium formation target, and ii) aluminum, a polarization-preserving target. Again we stress that together, the consistency of lifetime measurements in different targets - AK-3, quartz and aluminum - will yield an extremely convincing result.

Assuming similar conditions to our successful 2006 AK3 run - recording approximately $\sim 2 \times 10^{10}$ μ^+ decays per day - we will reach our statistics goal of 10^{12} decays in roughly seven weeks of running. In addition, based on experience from 2006, we would require i) one week for installation, debugging and calibration, ii) one week for various systematics studies, and iii) one week of contingency to absorb kicker and accelerator failures. Based on our timetable to be ready for the run, and the availability of the Collaboration members during the summer period, a suitable schedule for Mulan running would be ten or eleven contiguous weeks from the beginning of June to the middle of August. Based on informal discussions, this schedule is acceptable to the other main users of the πE3 beamline (i.e. MuCap and ALC).

References

[1] T. van Ritbergen and R. G. Stuart, Phys. Lett. **B437**, 201 (1998), `hep-ph/9802341`.

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

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

[4] R. M. Carey et al. (1999), PSI Proposal R-99-07, URL http://www.npl.uiuc.edu/exp/mulan/.

[5] F. R. Cavallo et al. (1999), PSI Proposal R-99-06, URL http://fast.web.cern.ch/fast/.

[6] S. N. Nakamura et al., Hyperfine Interact. **138**, 445 (2001).

[7] J. Duclos, A. Magnon, and J. Picard, Phys. Lett. **B47**, 491 (1973).

[8] M. P. Balandin, V. M. Grebenyuk, V. G. Zinov, A. D. Konin, and A. N. Ponomarev (1974), CERN-TRANS-74-9.

[9] G. Bardin et al., Phys. Lett. **B137**, 135 (1984).

[10] K. L. Giovanetti et al., Phys. Rev. **D29**, 343 (1984).

[11] The MEG Collaboration, *The MEG experiment at PSI*, URL http://meg.web.psi.ch.

[12] F. Mulhauser (MuLan), Nucl. Phys. Proc. Suppl. **155**, 358 (2006).

[13] K. R. Lynch (MuLan), AIP Conf. Proc. **870**, 333 (2006), hep-ex/0608045.

[14] D. Hertzog, R. Carey, et al. (MuLan) (2003), PSI Experiment R-99-07, URL http://www.npl.uiuc.edu/exp/mulan/documents/ProgressReports/ProgressReportDec2002.pdf.

[15] R. M. Carey et al. (MuLan) (2004), PSI Experiment R-99-07, URL http://www.npl.uiuc.edu/exp/mulan/documents/ProgressReports/MuLanProgressReportJan2004.pdf.

[16] T. I. Banks et al. (MuLan) (2005), PSI Experiment R-99-07, URL http://www.npl.uiuc.edu/exp/mulan/documents/ProgressReports/ProgressReport2004.pdf.

[17] T. I. Banks et al. (MuLan) (2006), PSI Experiment R-99-07, URL http://www.npl.uiuc.edu/exp/mulan/documents/ProgressReports/report2005-hyperlinked.pdf.

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

[19] E. Morenzoni et al., *Private communication* (2006).

[20] G. Przybylski, H. Steiner, F. Bieser, and J. Wolf (KamLAND) (2001), KamLAND Note: Calibration-011107.

[21] S. Ritt et al., *MIDAS data acquisition system*, URL http://midas.psi.ch.

[22] *ZLIB compression library*, URL http://www.zlib.net.

[23] A. Gutmans and Z. Suraski, *PHP: Hypertext preprocessor*, URL http://www.php.net/.

[24] R. Brun and F. Rademakers, Nucl. Instrum. Meth. **A389**, 81 (1997), URL http://root.cern.ch.

[25] MySQL AB, *MySQL database*, URL `http://www.mysql.com/`.

[26] J. F. Ziegler, J. P. Biersack, and U. Littmark, *The Stopping and Range of Ions in Solids* (Pergamon Press, New York, NY, 1985), URL `http://www.srim.org`.

A Analysis of the 2004 data set

In the sections which follow, we present more details on the analysis of the 2004 data set. At the outset, it should be noted that like many high precision measurements, the 2004 analysis was performed *blind*. The analyzer was given the master clock frequency only to a precision of 200 ppm. The exact value is known by a few members of the collaboration who are not directly involved in the analysis. In the same spirit, the analyzer uses his own private offset, ω_0, in presentations, so that collaboration members who know the true clock frequency are not biased in their reactions.

A histogram showing the sum of all data for the 2004 result, along with a fit (Figure 12), appears in the main part of this Report. Because we chose not to include the result from the sulfur target, the statistical uncertainty on the lifetime, 9.7 ppm, is somewhat larger than estimated in our previous report. Note that the reduced χ^2 is less than one, satisfying a necessary (but not sufficient) goodness-of-fit condition. As explained below, unfittable structures, introduced into the data by the electronics, prohibit fit start times earlier than 1.05 μs. Any fit start time from 1.05 μs on gives consistent results, as seen in Figure 14.

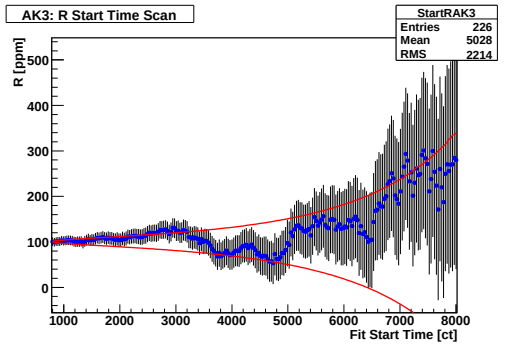

(a) The fitted lifetime in ppm with the red bands indicating the allowed 1σ deviation from the initial fit time.

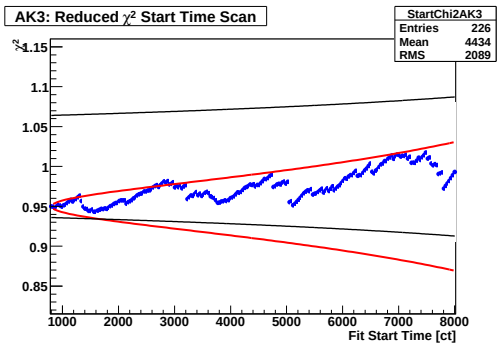

(b) The reduced χ^2 with the red region showing the 1σ allowed deviation about one.

Figure 14: Variation of fit parameters as a function of fit start time.

The 2004 total data set included many different running conditions, allowing for consistency checks over a wide range of experimental settings. Below is a list of parameters that were varied during the run:

- Extinction factors varied significantly because of kicker malfunctions.
- Two different stopping targets were used: Arnokrome-3 (AK-3) with a large internal magnetic field, and sulfur, which is known to be a good depolarizer.
- The magnetic field in the target region was rotated on a regular basis.
- The magnetic field near the EMC was also rotated at regular intervals.
- The discriminator thresholds were set at two different values, 80 mV and 200 mV.
- The measurement period was set at 22 μs and 27 μs.

The data subsets created by these variations all yielded statistically consistent lifetimes.

Much of the analysis effort focused on the various systematic errors. These are described below and summarized in Table 1.

Category	Error	Post-Fitting Correction
Pileup	2.0	
Fit parametrization of CAEN oscillation	0.5	
Long time constant non-linearity	0.9	
Duplicate words	1.0	1.0
Afterpulse effect on CAEN buffer	0.7	
Extinction changes	3.5	
Timing shift	0.8	
Gain and threshold changes	1.4	
Residual polarization	2.0	
Total systematic	5.2	1.0
Statistical	9.7	
Combined statistical and systematic	11.0	1.0

Table 1: Preliminary systematic error table for the 2004 analysis; all entries are in units of ppm.

A.1 TDC effects

The use of the CAEN v767 multihit TDC led to some of the largest systematic problems, which were explained in last year's report.[17] Here we summarize our studies of some of these issues:

- While the TDC is capable of recording both the leading and trailing edges of a discriminated signal, the small event buffers and slow readout required us to save the leading edge only. This reduced the dead-time, but resulted in some loss of information about pulse pileup.
- The small time bins of the TDC, produced with an on-board interpolation circuit, had a differential non-linearity of 0.5%. To avoid problems associated with this time variation, the coarse time was used in determining coincidences and as the bin width for our histograms.
- Every hit in a CAEN module results in a small oscillating shift in the time given to subsequent hits. This effect, which persists for a few microseconds after the initial pulse, affected the data in two ways. First, because the time of the kicker was recorded in each TDC module, there was an oscillation term found on all channels. We learned through extensive lab tests that the oscillation could be parametrized in a simple functional form, and its effects could largely be removed by the fit. The residual error which arises from using the functional parametrization of the TDC's non-linearity following start pulse is 0.5 ppm. The other issue, that one positron can affect the recorded time of another one in the same channel, is discussed below in Section A.5.
- In addition to the oscillation term observed after each hit, the possibility of a long time constant component was investigated in the lab tests. The results were consistent with no statistically significant non-linearity associated with the kicker signal. The study suggested an error on the fitted lifetime of 0.9 ppm.

- The front end buffers of the TDC are only two hits deep. Despite a relatively slow readout speed, our overall rates were so low that the resulting loss of data was negligible. However, this was complicated by the occasional correlated afterpulse from either cable reflections or ion feedback within the photomultiplier tubes. Queuing losses in the second stage buffers were negligible because of their larger size and the low data rates. Based on data collected, we concluded the error on the lifetime from queuing losses was less than 0.7 ppm.
- Although we set up the modules to record leading edges only, bursts of spurious transitions with the trailing edge bit set sometimes appeared in the data. Since they are associated with earlier transitions, at definite times many clocks periods earlier, it was not possible to remove the affected fills. Fortunately, these trailing edge bursts have no measurable effect on real data and can be safely ignored.
- Very rarely, duplicate events were discovered in the data stream. These events had a time spectrum that differed from the normal time spectrum and resulted in a potential pulling of the lifetime. But because these events affected less than 6×10^{-6} of the total events, the impact on the lifetime was estimated to be small: a 1 ppm shift with an uncertainty of 1 ppm.
- Several other potential problems were investigated, but they had no effect on the result and are not reported here.

A.2 Pileup and dead time

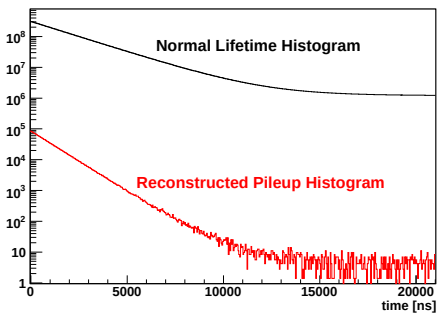

(a) Lifetime and reconstructed pileup histograms.

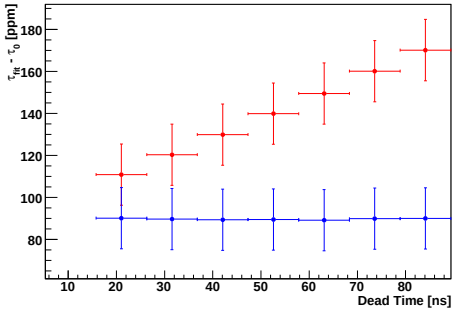

(b) Lifetime versus artificial dead time. Red(blue) points have not(have) been corrected for pileup.

Figure 15: Corrections for pileup and TDC dead time.

When one or more positrons pass through a detector element during the dead time from a previous event, the event(s) are not included in the lifetime histograms. If unaccounted for, this *pulse pileup* results in a large systematic pulling on the fitted value for τ_μ. To correct for pileup, we statistically reconstruct the distribution of missing events from the data itself and, before fitting, add them back into the lifetime histogram. To construct the pileup distribution, we first trigger on an inner/outer coincidence and then look in a fixed window at a later (or earlier) time for the existence of another hit. If another hit is

found, we make an entry in the pileup histogram. Figure 15(a) shows the reconstructed pileup distribution relative to the lifetime histogram. As expected, the pileup distribution has a characteristic lifetime of $\tau_\mu/2$. The amplitude of the pileup correction is about 10^{-3} of the singles amplitude and is entirely below the background level. To ensure that the correction is performed correctly, we artificially varied the dead time to amplify the effect of pileup. The magnitude of the pileup effect and stability of the correction can be seen in Figure 15(b). If left uncorrected, the pileup would lead to a 76 ppm error on the lifetime. Based on simulations and the varying of pileup reconstruction parameters, we conclude that the pileup correction adds a systematic uncertainty of 2 ppm to the 2004 data set.

There are other effects that can be corrected using a statistical reconstruction. One example is what we call *positive pileup*. This occurs when the inner and outer tile elements of a single detector form a coincidence from two events that are unrelated. Just like normal pileup, this has a $\tau_\mu/2$ distribution, but instead of a missed event, we count an event that should not be included. In reconstructing these events we first trigger on a single inner, which is not in coincidence with its outer, and then record the times of single "outer" events found in a delayed window. Other types of systematic errors that fall into this category are afterpulses, inner/outer time shifts that cause the loss of a hit, and the shifting of the recorded time due to pileup. All of these effects are an order of magnitude smaller than normal pileup and therefore make a negligible contribution to the systematic error.

A.3 Flat background

During the measurement periods of 2005 and 2006, the kicker typically reduced the beam rate by a factor of approximately one thousand, but in 2004 these extinction factors were more typically in the range of 100-600 . The resulting background is included in the fit function and, provided it is constant, will not cause a systematic error. However, if the kicker voltage is not stable, then the nominally flat background will change with time, causing a systematic shift in the lifetime measurement. At the beginning of the 2004 run the kicker voltage was measured to be stable to 0.25 V during the measurement period. The change in extinction factor versus kicker voltage was measured during the run. From these results, we estimate that the maximum possible change in background rate over the measurement period is about 0.8 ppm of the total rate. Monte Carlo simulations indicate that the maximum resulting systematic error in the lifetime measurement is 3.5 ppm. We estimate that mechanical distortions of the kicker, which arise from the electrostatic forces on the plates, produce lifetime changes which are more than an order of magnitude smaller.

A.4 Gain and threshold shifts

Gain and DC offset shifts during the measurement period would cause an effective shift in the discriminator threshold. This, in turn, would result in a changing efficiency for the detection of decay positrons during the fill. Simulations indicate that a 1 ppm shift in detector efficiency during the measurement period could result in up to a 0.25 ppm systematic error in the lifetime.

Given the very low rates in the MuLan detector at all times (typically, only a few percent of the tile pairs were occupied in a given fill), we don't expect to see any changes in gain over

the course of the measurement period. Likewise, there is no reason to expect changes in the electronic threshold. Unfortunately there were only a limited number of channels from the 2004 data set which include WFD data[2] that could serve to monitor the gain and pedestal. Therefore we chose to use data from 2006 for our studies.

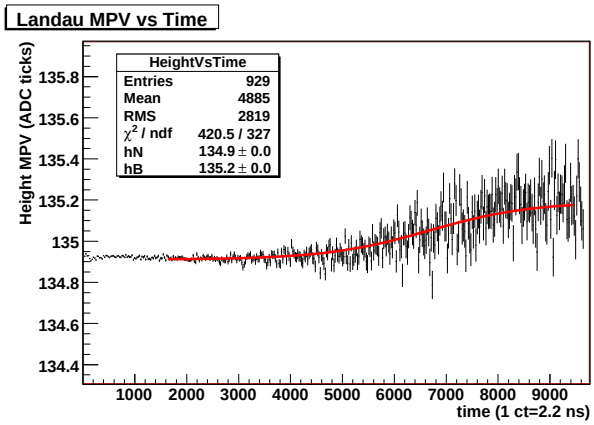

Figure 16: The position of the Landau peak in a fit to pulse heights allows us to track the gain stability as a function of time in fill.

We used the the position of the Landau peak as means of characterizing the gain. Over the course of a single fill that peak shifts by more than a percent; see Figure 16. This is dominated not by true gain shifts, but by the difference in average pulse height between signal and background events. At early times, the data is dominated by decays from the target while at late times, the non-target backgrounds dominate; the different energy and acceptance distributions for the two populations result, on average, in slightly different pulse parameters.

Once the transition from signal to background was accounted for, there were no signs of gain shifts. Pedestal shifts were also negligible. Based on this analysis we were left with an uncertainty of 2.6 ppm for the low discriminator threshold data and 1.0 ppm for the high discriminator threshold data. The combined effect is a 1.8 ppm uncertainty on the lifetime due to potential gain and threshold changes during the measurement period.

A.5 Timing shifts

If the difference between the time of a muon decay and the recorded time of the detected decay positrons shifts during the measurement period, due to effects in the PMTs or readout electronics, then there will be a systematic error in the lifetime measurement. Early-to-late timing shifts could occur if the recorded time of a hit is affected by recent hits in the same detector. In particular, simulations indicate that a 3 ps shift in the timing could cause a 1 ppm systematic error in the lifetime measurement. We estimated the timing shift by comparing the time between coincident hits on an inner/outer pair when one of the channels was recently hit to the time between the two when there was no other recent hit in those same channels. Figure 17 shows the surprisingly large and long-lived shift, which was later reproduced in bench tests. Luckily, the relatively low rate in any one channel implies a very small effect on the average timing, and a systematic error less than 0.8 ppm.

[2]This small data sample was collected using an older WFD from the Brookhaven $g-2$ experiment.

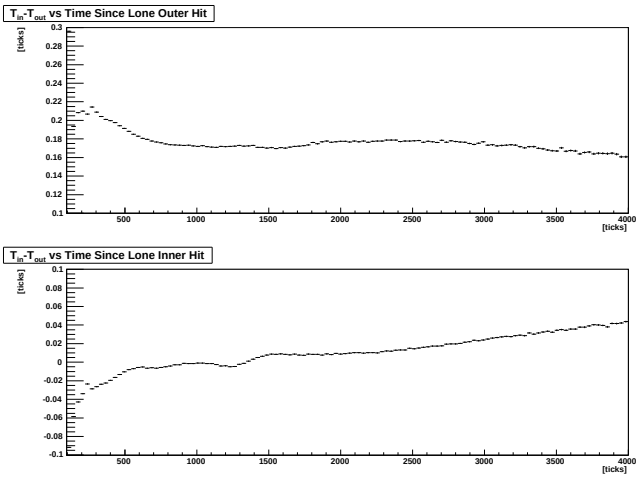

Figure 17: Behaviour of the time between inner and outer hits as a function of time in fill showed a statistically significant trend in the 2004 data set.

A.6 Residual polarization and stopping distributions

One of the most vexing problems in the analysis was the observation of a significant forward/backward asymmetry in the measured value of the lifetime, for both the AK-3 and sulfur targets, but particularly the latter. As described earlier in Section 2, the spins of muons in both targets dephase during the accumulation period and the effect of the residual polarization, if any, is minimized by the forward/backward symmetry of the detector. However if a muon stops outside the target, in a region of small magnetic field, the precession period of the spin is large relative to the measurement period and no significant dephasing occurs. Moreover, the further the muon stops from the target, the less the detector symmetry will help minimize the effect of the strong correlation between the direction of the average muon spin and the direction of the decay positron. As an extreme example, consider the lifetime measured from muon stops (spins pointing upstream) far upstream in the entrance corridor. The effects of the slow precession and possible depolarization are exacerbated by the fact that most downstream decays are recorded but most upstream decays escape up the beampipe and are not seen.

At first, we attempted to trace the observed asymmetry to entrance corridor stops. Indeed, anticipating trouble with muon stops both in the corridor and in the EMC region, we had performed a few dedicated studies using a plastic plate placed upstream of the target region, with and without applied magnetic fields. We also measured stopping distributions in air using the MuCap detector, from which we deduced the expected stopping distribution in the helium bag. We learned that muon stops at several points along the muon corridor, where the magnetic field is roughly 10 Gauss, could produce large lifetime asymmetries over the detector. However, both the stopping studies and the simulations indicated that the number of anticipated stops was insufficient to produce the observed asymmetry. Moreover, it was unclear how entrance corridor stops could produce noticeably different asymmetries for the two targets.

It was perhaps just that discrepancy which inspired us to consider an alternative source of problematic stops - those which arise from muons which are scattered around the target and come to rest in the PVC walls of the inner surface of the MuLan ball. As in the entrance corridor, the magnetic field on the inside of the ball is small. The residual polarization in

PVC is about 1/3 and the depolarization lifetime of roughly $5\,\mu$s is dangerously close to that of the muon lifetime. Moreover, whatever trouble is created by these downstream stops would be larger for the sulfur target, which is significantly smaller in radius than the AK-3 target ($R_{\mathrm{Sulfur}} = 13.4\,\mathrm{cm}$, while $R_{\mathrm{AK-3}} = 27.6\,\mathrm{cm}$). Our SRIM [26] Monte Carlo indicates that roughly 1/200 of our incoming beam scatters around the sulfur target. For AK-3, the fraction of downstream stops is reduced by roughly a factor of 5.

Using the SRIM Monte Carlo, we estimated the number of muons which

1. stop in the entrance corridor,
2. are backscattered into the helium bag or the upstream half of the detector,
3. end up in the target holder, or
4. scatter around the target and end up downstream.

Starting with this "cocktail" of stops, the GEANT-based simulation predicted well the $\cos\theta$ asymmetry observed for both targets. However, because of our large solid-angle coverage, the effect on the lifetime measured from the summed signal of all the tiles was not particularly large. In AK-3, the lifetime was pulled much less than 1 ppm; for sulfur, approximately 5 ppm. Perhaps fortuitously, the corrections result in near perfect agreement for the two data sets. However, uneasy with the large correction required by sulfur, we quote only the AK-3 data in our final result. Additionally, since there is some uncertainty in modeling the stopping positions, magnetic fields, and relaxations, we place a 2 ppm uncertainty on our correction.