MuLan Progress Report 2005
and
Beam Request 2006

T.I. Banks[1], E. Bartel[5], S. Battu[6], R.M. Carey[2](Co-spokesman), D. Chitwood[3], C. Church[5], S.M. Clayton[3], P.T. Debevec[3], S. Dhamija[6], M.L. Dantuono[5], W. Earle[2], R. Esmaili[5], A. Gafarov[2], F.E. Gray[1], K. Giovanetti[5], T. Gorringe[6], Z. Hartwig[2], D.W. Hertzog[3](Co-spokesman), J. Jungmann[7], 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], M. Miller[5], F. Mulhauser[3,8], C.J.G. Onderwater[7], M. Ojha[6], C.S. Özben[4], Q. Peng[2], S. Rath[6], B.L. Roberts[2], V. Tishchenko[6], W. Vreeland[2], D. Webber[3], P. Winter[3]

BERKELEY[1] – BOSTON[2] – ILLINOIS[3] – ITU[4] – JAMES MADISON[5] – KENTUCKY[6] – KVI[7] – PSI[8]

January 18, 2006

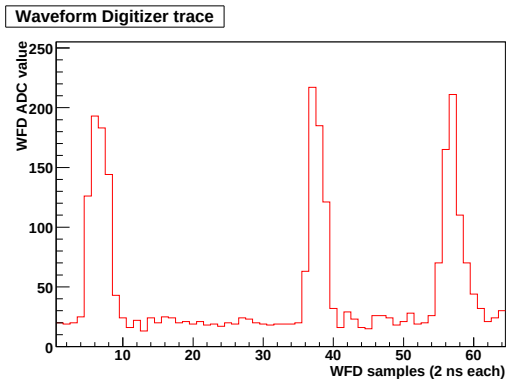

(a) Three pulses close together in one channel are easily distinguishable using waveform digitizers.

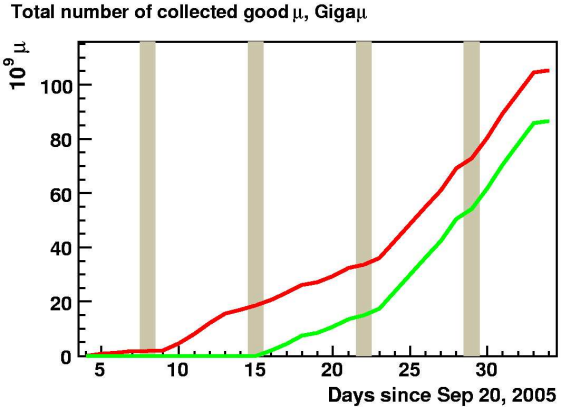

(b) Online monitoring tools permit real time tracking of observed muon decays as a function of time.

Figure 1: The data sample of 8×10^{10} muons collected in 2005 should yield a muon lifetime measurement with a statistical uncertainty of 4 ppm.

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] enables a clean extraction of G_F from τ_μ and paves the way for a new high-precision measurement. Indeed, MuLan is only one of three current efforts to measure the muon lifetime.[2, 3]

In 2005, the MuLan experiment ran twice in πE3. This included a two-week beam study program in June and a six-week final commissioning and physics run in the fall. These runs were successful and the experiment is ready for production running in 2006 and 2007. The final two technical hurdles, outlined at the Users' Meeting last year, were the failures of the MOSFET switches in the kicker system and the missing waveform digitizers (WFDs). The MOSFET cards were repaired by our TRIUMF colleagues and the addition of over-current protection as well as a careful program of calibration and maintenance have rendered the system far more robust (section 4). In the fall run, using tunes established in a summer beam test, the extinction factor was typically better than our systematic requirement of one thousand. All 340 WFD channels were produced in July and were delivered for use in the fall engineering/physics run (section 6). Apart from a few modest firmware and power issues, they worked extremely well. The pulse and timing information provided by the WFDs has already proved to be a spectacular success, again meeting the demands outlined in the original proposal.

We continue to make good progress on the analysis of the fall 2004 data and expect a result to be ready in the early summer of 2006 (section 9). We have identified a potential systematic problem related to the few per thousand muons that stop just upstream of our target (between the Entrance Muon Counter (EMC) and the target, inside the 40 cm long helium corridor). We confirmed the stopping distribution by using the MuCap detector (and a day of their running time, graciously offered by the MuCap collaboration) to measure the stopping distribution along the corridor. Although we estimate the effect of these errant muons to be of order 10 ppm, they remain a subject of investigation. Before production running in 2006, either the helium bag will be replaced with a vacuum tube or we will impose a magnetic field in that region, perhaps with a saddle coil wound around the outside of the helium bag corridor (section 3).

Our 2005 successes include:

- Dedicated extinction factor (EF) studies were carried out before the fall run (section 3). The twin goals of an EF > 1000 and a rate of approximately 7 MHz were achieved.
- The kicker worked flawlessly, over a several month-long period, at the design plate voltage of ± 12.5 kV (section 4). This included use in "Muon on Demand" mode for the MuCap experiment.
- The EMC was equipped with a custom Field Programmable Gate Array (FPGA) readout (section 5).
- An real-time, online monitoring system was implemented to track the health of the system and help track down any problems quickly (section 8).
- A data set of more than 8×10^{10} positron decays was obtained in approximately two weeks, and analysis have begun (section 10).

Disappointments include:

- Because of power/cooling problems, we had to run with 10 percent of the WFDs removed from the array of six VME crates (section 6). We are investigating this anomaly and can correct it by adding a seventh crate in 2006, if required.
- Early in the run, firmware problems on the new WFDs caused minor flaws in the data stream.
- The data acquisition live time was a modest fifty percent (section 7). We expect it will be much greater in 2006 when fewer data words will be written per decay positron and online data compression will be implemented.

The Collaboration remains strong. Scientists at every level helped to make our 2005 data taking run a success, including Ph.D. and Master's candidates. Dan Chitwood of the University of Illinois is writing his Ph.D. dissertation on the 2004 analysis. David Webber, also of Illinois, Qinzeng Peng, of Boston University, and Sabyasachi Rath of the University of Kentucky will tackle the 2005 and 2006 analyses. The design, development, and commissioning of the WFD-based data acquisition was the Master's research project for University of Kentucky graduate student Sivaram Battu, who is expected to graduate in summer 2006.

The Collaboration goals for 2006 include completion of the analysis of the 2004 data, further development of lossless online compression software to enable a high-rate, high-statistics run, development of the improved "muon corridor" to ensure that no errant muons stop in the low field region just upstream of the target, updating the WFD firmware to enable shorter data blocks and to eliminate the few remaining errors. We request a 10.5 week run: one half week for installation, two weeks for beam tuning and systematic studies and eight weeks of production running. This last should be sufficient time to collect roughly half the data required by our proposal. The details are given in section 11 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 kicker. During the measuring interval, the Michel positrons are recorded by a multi-segmented, symmetric detector, featuring 170 independent scintillator tile pairs. Each element is read out by a photomultiplier tube (PMT), whose signal is sampled at 500 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 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 peak rate per element, and by the double hit resolution enabled by recording the pulse height in both tile elements for each event. 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 decay 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 found that the μSR signal in AK-3 was much smaller. 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 limiting the amount of material upstream of the stopping target; a 9.3 cm, 76 μm mylar vacuum window and helium bag spanning the volume between the EMC and target accomplish this goal. A small permanent magnet array helps to dephase the small fraction of incoming muons - roughly one per thousand - which are known to stop in the foils of the EMC.

The instrumentation naturally divides into subsystems: the muon beam (Illinois); magnets and 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 Beam studies

Our πE3 beam line configuration has been described in detail in previous Annual Reports.[4, 5] In 2005, we made one major change to the physical layout. We operated in a "kicker first" mode, meaning the kicker was placed upstream of the electrostatic separator, which yields a higher extinction factor (EF). In general, the properties of the beam are adequate for the experiment: the beam spot size is relatively small (see figure 6), the flux is approximately 7 MHz, the EF exceeds 1000, and there are no noticeable beam microstructures at any relevant time scale.

We made a comprehensive study of the basic tune in June. First, running the kicker in DC mode, we made a thorough sweep of each slit element in the beamline, registering the muon flux and extinction factor for each setting. Together with similar sweeps of critical quadrupole pairs in the line, we arrived at a working beam tune, one that was verified and then later used in the fall data-taking period. We made a few important and surprising discoveries. In particular, with several asymmetric upstream slit settings, the extinction factor could be dramatically improved while allowing the main muon beam to pass. The final settings are close to our theoretical models but differ in a manner that suggests unknowns in the source distribution or in the description of the magnets in the line. We believe we can increase the rate slightly in 2006 and, of course, the rate rises linearly with the PSI proton beam current.

As described briefly in last year's report, we used the MuCap spectrometer to determine the location of muon stops in air just downstream of our EMC detector. The low-momentum tail revealed in the measurement implies that a small fraction (a few per 1000) of the muons which make it to the entrance of the MuLan ball, might not make it to our stopping target. In the MuLan setup, these straggling muons would stop in helium in a region with a magnetic field of no more than a gauss. The large precession period of those stopped muons, many tens of microseconds, induces an effective acceptance change on the detector tiles that is not canceled out by the point-like symmetry of the MuLan ball.

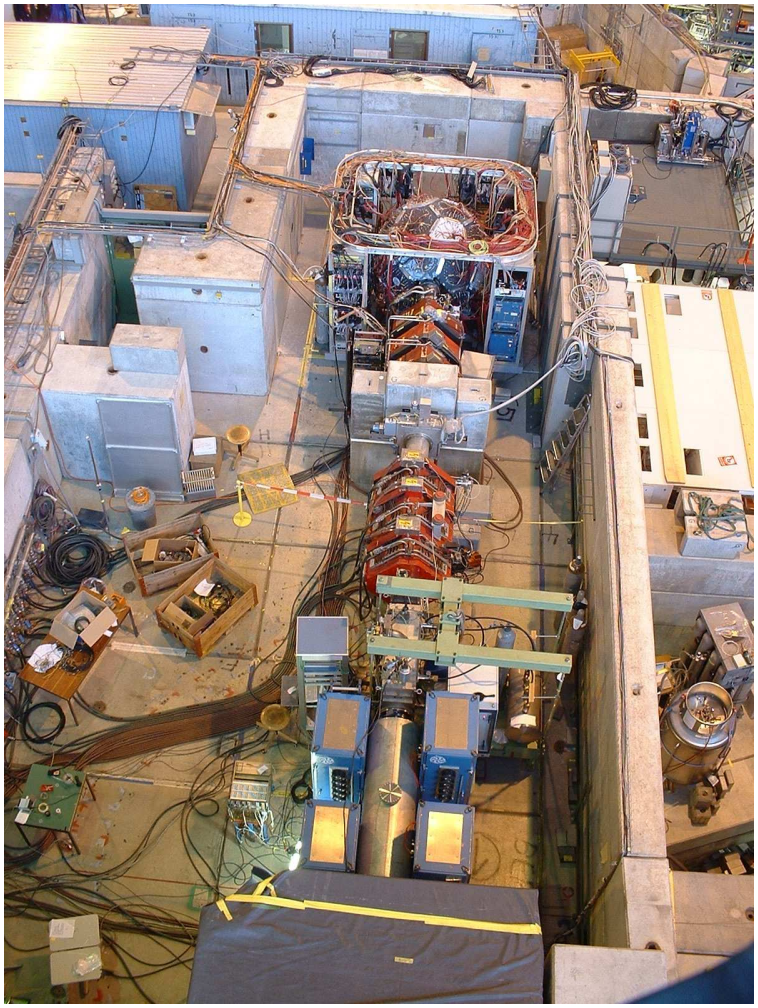

Figure 2: An aerial view of the MuLan ball and πE3 beamline.

The 2004 test was cut short when the PSI accelerator had a problem just at the end of the MuCap run. This year, we used a day of MuCap running in December 2005 to reinvestigate the low-momentum tails and the anomalous momentum shift of those muons that make it through the kicker. With improved reconstruction software and more controlled conditions, we measured the momentum stopping range in air and mapped out the long, low-momentum tail that had been observed previously. Figure 3 is a composite plot giving the kicked (red) and unkicked (black) stopping distributions in air (approximately ten times helium in density). The horizontal axis is in millimeters and the beam "enters" the plot from the left, arriving in air at the dotted vertical line. A gaussian fit is overlaid on the unkicked distribution and a clear low-side (momentum) tail is evident. These tail muons would lie in a region of relatively low field.

It is important to note that, for reasons as yet unknown, the kicked distribution peaks at a lower momentum than the unkicked distribution. Given the kick supplied by the 1.5 m long plates with a potential of 25 kV across the (reduced) gap of 12 cm, our models indicate that no muons should pass. The measured muons - the $^1/_{1000}$ that do pass - are shifted downward in momentum. We do not know their polarization direction. We used both SRIM [6] and GEANT [7] to predict the stopping fraction in our helium corridor based on the series of plots in figure 3 and similar plots made at other nominal momenta.

There are two options for controlling the muons that might stop in the muon corridor between the EMC location and the AK-3 target. The first involves building a saddle magnet such as the one used to surround the TPC in the MuCap experiment. These Berkeley magnets have undergone several iterations and one can

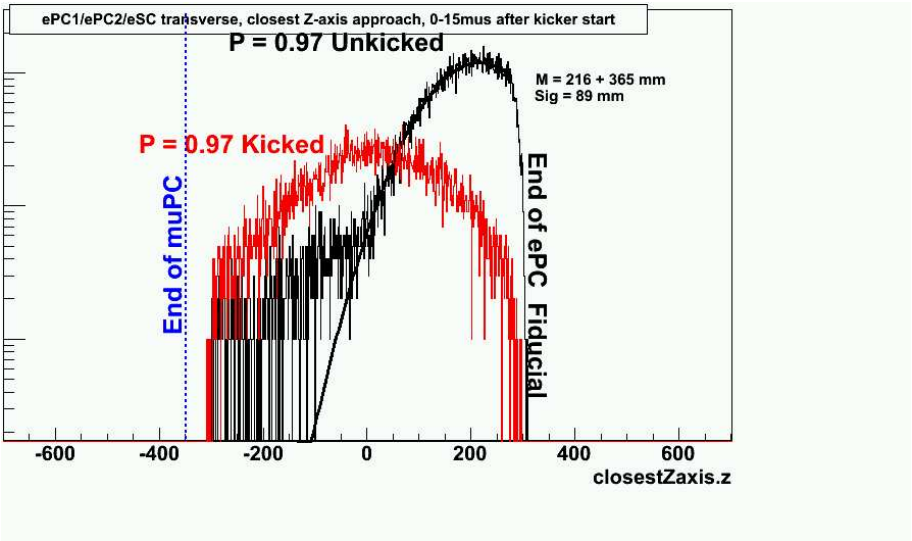

Figure 3: A 2005 air-scan run using the MuCap spectrometer to image the stops in air downstream of the μPC wire chamber. At a momentum of 28 MeV/c, the peak of the unkicked stopping distribution occurs 580 mm downstream of the wire chamber. The width of the Gaussian peak suggests a $\Delta P/P$ of 2.6 percent. The kicked beam appears to have a central momentum of 26.4 MeV/c, a 5.7 percent downward shift.

be made for MuLan that would use the existing cooling and power infrastructure available for the MuCap magnet. The transverse magnetic field it would produce along the muon corridor, approximately 40-50 G, should be enough to precess and dephase the muons during the accumulation period but remain small enough to permit the PMTs to function without additional shielding. The magnet would be placed symmetrically in the MuLan ball and a new AK-3 target would be placed at the midplane. As before, a helium volume would occupy the upstream region between the EMC and the target.

An alternate plan involves the elimination of all matter between the entrance window and the target by means of a long thin vacuum tube, which would also extend from the target to the back of the MuLan ball. The EMC would be placed in a symmetric position downstream of the center of the ball; that is, at the exit pentagon face rather than the entrance pentagon face. With a new AK-3 target inside, no muons would pass through to the EMC. However, the target will be designed to rotate out of the way of the beam for beam tuning and monitoring runs. The advantage of this plan is that all the muons would travel to the target in vacuum. The disadvantage would be the loss of fill-by-fill monitoring of the event flux and sneaker muons (those coming during the kicker cycle). The Collaboration will investigate both plans and prepare a hardware solution by the time of the summer run.

4 Electrostatic kicker

The electrostatic kicker provides the pulsed muon beam which is essential to our experiment.[8] Located just upstream of the electrostatic separator, the kicker consists of two pairs of 0.75 m long by 20 cm wide electrode plates, mounted 12 cm apart. When opposing plates in a pair are biased at $\pm$12.5 kV the beam is deflected into a collimator. When the plates are shorted to ground, the beam passes undisturbed. The kicker is typically operated at a cycle frequency of approximately 37 kHz. There is a 5 μs accumulation period with the plates grounded and a 22 μs measurement period with the plates biased. The kicker will be used in a "muon-on-demand" mode for the MuCap experiment and possibly for some experiments mounted by the μSR community. After some repair and upgrade work performed in early 2005 (detailed below), the system worked well during the entire experimental run of both the MuLan and MuCap experiments.

4.1 MOSFET card failure and solutions

Fast-switching, high voltage MOSFET cards provide a quick means of charging and discharging the kicker plates. MOSFET card failures during the 2004 experimental run were traced to three different problems:

- High voltage breakdown between the Fischer feed-through connector and the deflector plates, caused by outgassing of the cable insulation.
- Breakdown within the Fischer connector itself, on the output of the modulator: by mistake, sealed environmental plugs were ordered and the sealed plugs were not correctly wired.
- Timing problems of the MOSFET cards. Timing offsets of more than a few nanoseconds between adjacent cards can quickly destroy the MOSFETs themselves.

The original TRIUMF kickers did not have an overcurrent interlock and hence, in the event of one of the above mentioned breakdowns occurring, a MOSFET stack would likely attempt to turn-off high current (up to 54 A) and therefore avalanche. Tests at TRIUMF showed that, although the MOSFETs are avalanche rated, they tend to oscillate and self-destruct during avalanching. As foreseen in our preceding progress report [5] we have implemented the following electronics upgrades:

- TRIUMF engineers designed, constructed and tested fiber optic based overcurrent detection cards. An overcurrent card is connected on the DC end of each stack of MOSFET stack (Figure 4). A TRIUMF TTL splitter box has been designed, built and tested. This splitter box receives fiber optic feedback from each of the eight overcurrent cards and, under approximately 95 percent of operating conditions, prevents the stacks of MOSFETs turning off if a current of more than approximately 10 A occurs for a period of more than approximately 160 ns.
- In the course of recalibrating the MOSFET cards, several timing mismatches were discovered. The problems were traced to damaged fiber optic interconnections, which were replaced
- In order to make the MOSFET cards more insensitive to timing mismatches, we have incorporated a resistance of approximately 3.7 Ω on the gate of every MOSFET. This gate resistance slows down the turn-on voltage collapse across a MOSFET, from approximately 3 ns to 6 ns, without significantly (less than 1 ns) increasing the deflector plate voltage rise time. The gate resistance reduces the maximum drain-source voltage (V_{ds}) across a MOSFET, during turn-on, which can result from a timing mismatch. With a smaller V_{ds}, the MOSFET is less likely to avalanche. After adding the new gate resistance, all the MOSFET card timing offsets were remeasured.

4.2 Preparing the kicker for the fall run

As already described earlier in this report (section 3), we made three important changes to the kicker system during and just after the beam test run of June 2005. At that point, while the MOSFETs cards were still at TRIUMF for repair, the kicker was used only in a DC mode. The default position of the kicker and separator were reversed so that the former is now upstream. Although the order of the two should not, in the ideal case, make a significant difference, we learned that it was far easier to achieve good extinction factors with the kicker upstream. In addition, given the fact that nearly all the muons passing through kicker are located within 6 cm of the horizontal midplane, we increased the kicking field by reducing the gap between the deflector plates from 15 cm to 12 cm. We narrowed the opening of the collimator accordingly. An end view of the lower kicker plate and a glimpse of the new over-current protection are shown in figure 4.

Many tests were performed offline to characterize the timing of the kicker. The most important ones, later performed on a regular basis, were measurements of the drain-source voltage and timing for each of the 136 MOSFET cards. Before the kicker was used in the fall, all the timing offsets had been carefully matched. We also cleaned the inside of the kicker cabinets at least once a week throughout the run, removing metallic dust (generated mostly by the Hallendienst workshop) which was brought inside the cabinets by the cooling system fans.

During the beginning of the fall 2005 run period, various additional measurements were performed on the kicker. Perhaps the most important concerned the flat top circuit, intended to minimize any voltage

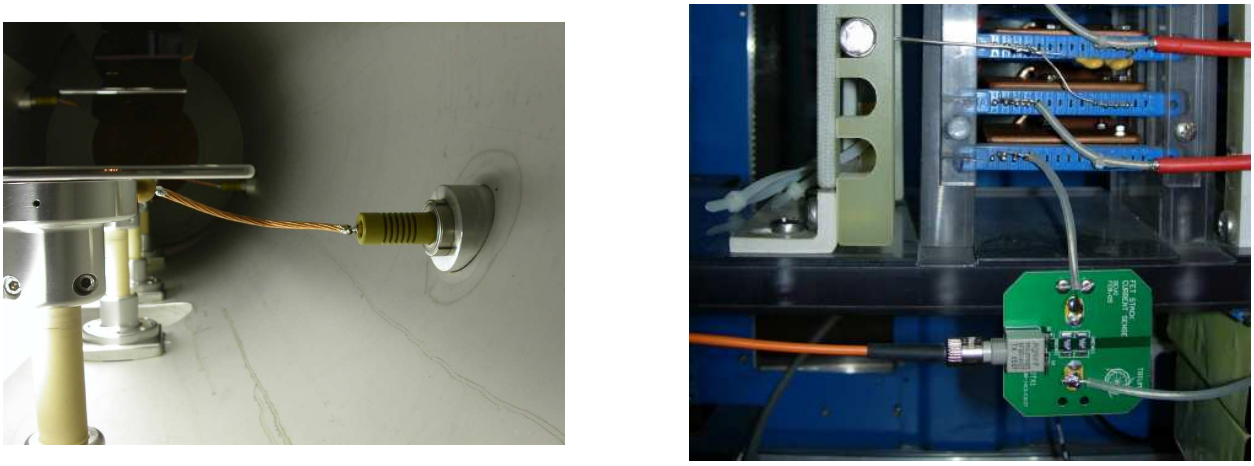

Figure 4: Left: Inside the beampipe: copper connector between the deflector plate and the feed-through. Right: Over-current protection circuit connected and installed in a MOSFET stack.

sag during the measurement period, which had not been tested *in situ* since its installment in 2004.[1] The 12.5 kV flat top sags less than 300 mV, a fractional decrease of 2.4×10^{-5}; see figure 5. Measurements of the extinction factor versus kicker voltage were made with various beamline configurations. Additional measurements, such as timing delays between various control signals and kicker transitions were made using the capacitive pick-ups mounted on the inside of the kicker tank.

Careful inspection of the MOSFET cards, maintenance and timing calibration, as well as vigilant monitoring of the extinction factor helped make the TRIUMF kicker an extremely reliable device in 2005. As noted above, it ran for nearly four months, without incident. The new fault-protection circuit worked perfectly, when, during the MuCap run, one of the plates was inadvertently shorted to ground through a loose collimator. (The collimators will be better secured for the 2006 run.) Results on the extinction factor are given elsewhere in this report.

[1] A description of the flat top circuit appears in our 2004 report[5].

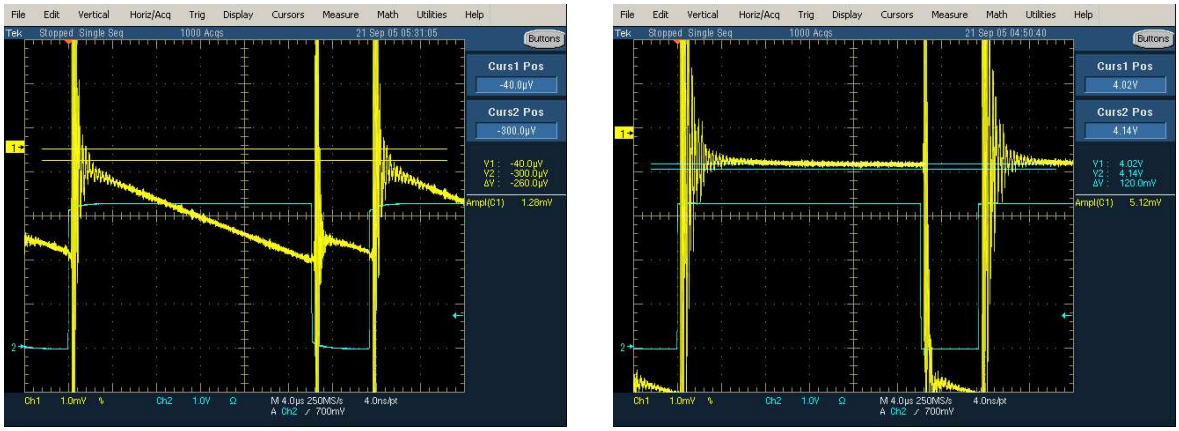

Figure 5: Effect of the flat top circuit: The left photo shows a non-working system, he right photo a good one. Both measurements were done at 12.5 kV. The flatness is good to approximatively 100-300 mV, depending on the MOSFET stack.

5 Detectors

5.1 EMC

The Entrance Muon Counter (EMC), a wire chamber positioned at the entrance to the MuLan ball, records the time and position of muons entering the detector. The beam profile provided by the EMC is used to focus the beam. Timing information on the incoming muons allows us to study and possibly reject fills in which a muon entered the detector during the measurement period. The EMC can function with at instantaneous beam rates up to 10 MHz.

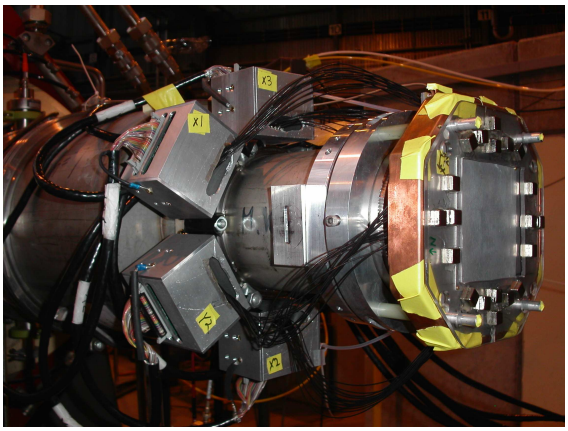

(a) The EMC mounted on the beampipe. The magnets used to create a 50 G field in the EMC are visible on the front face.

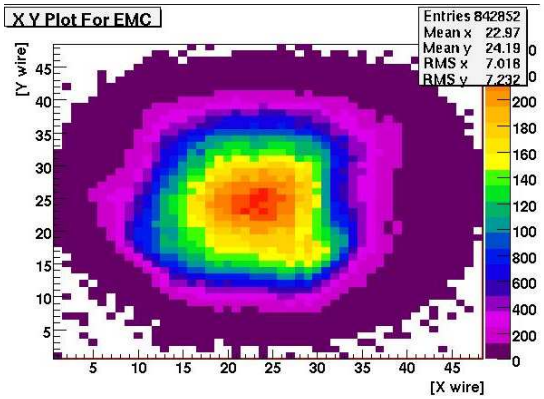

(b) An X-Y profile of the EMC from the 2005 run.

Figure 6: The EMC is a wire chamber designed to provide an X-Y profile at high beam rates.

The EMC is filled with a 70:30 gas mixture of CF_4 and isobutane. It has 25 μm aluminized mylar windows at the front and back. There are two planes of 20 μm diameter tungsten wires, one in the vertical and one in the horizontal direction. 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 sit at a positive high voltage.

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. In 2005 the readout scheme for the EMC data was modified. An FPGA 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 one thousand times lower.

A beam profile from the EMC is shown at the right of Figure 6.

5.2 The detector tiles of the MuLan ball

The MuLan ball (figure 2) was fully commissioned during the fall 2003 run and has been used successfully in the 2004 and 2005 data-taking runs. Calibration, monitoring and simple debugging of the individual tiles became much easier with the arrival of the WFDs in 2005, when we could readily inspect the pulse response and gain of each element. Because essentially all decay positrons passing through a scintillator tile pair are minimum ionizing, the energy deposited per event is well described by a landau distribution. In order to set the high voltage correctly, a small sample of data was obtained at a series of different voltage settings for each detector. In each data set, for each detector, a histogram of the energy deposited was fit to a landau distribution. The fit yielded a most probable value (MPV) parameter for the distribution, which

is proportional to the gain. Empiracally, the MPV scales with the eighth power of the voltage; the PMT voltages were set so that the MPVs would all be approximately 350. Operationally, the voltages were then iteratively adjusted, as necessary, until all the fitted MPVs were within 10 percent of 350. Figure 7 shows the fitted MPVs for all detectors, with all but ten at satisfactory values.

As we slowly increased the beam rate during the 2005 data taking, we noticed a gentle rise in the average MPV for each element. As the MPV does not drop immediately when the rate is lowered, this effect is most likely related to heating in the enclosed tile cans. We are investigating this effect and plan to make corrections, perhaps involving simple venting or cooling. The slow gain change is not a concern for the timescale of measurement.

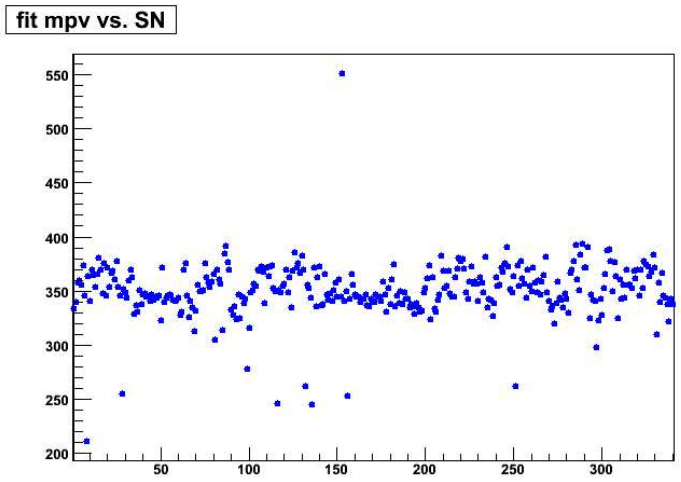

Figure 7: Snapshot of most-probable value for landau fits to the individual tile pairs after high-voltage adjustments. Some detectors will require special adjustments or repair.

6 Electronics

The MuLan electronics include timing control devices, a master clock and clock fanouts, and the waveform digitizers used to read out the photomultiplier signals.

6.1 Clocks and timing control systems

There were four main clocks in the MuLan experiment in 2005:

1. Fill control: In the AC mode of the experiment, the kicker cycle (or fill) must be time synchronized with the various acquisition electronics, all of which take different types of start/stop control signals. We maintain this synchronization with an FPGA based Programmable Pulse Generator (PPG). To control kicker based systematics, the kicker operates in a continuously cycling mode, but we gate the synchronized control signals to the electronics by means of a second, DAQ control clock. The repetition rate of the fill clock is on the order of 30 μs.

2. DAQ control: This clock gates the control signals to the DAQ electronics and synchronizes the back end MIDAS controller with the hardware. These tasks are performed by a PPG that is identical to the fill controller. These DAQ macro cycles are on the order of 100 ms.

3. Timebase signal: The main timebase for the 2005 run period was provided by an Agilent E4400 frequency synthesizer generating a sinusoid at approximately 500 MHz. In previously reported tests, we have confirmed the absolute frequency and short and long term frequency stability of this unit at

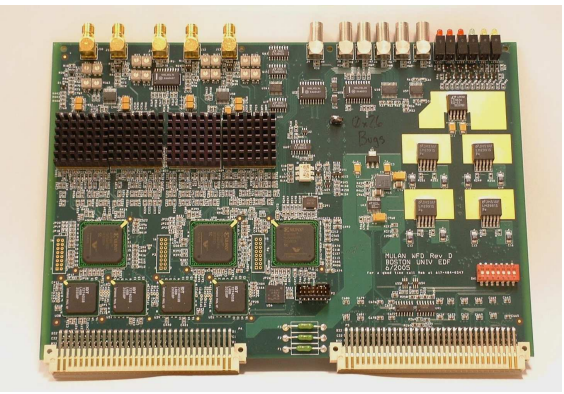

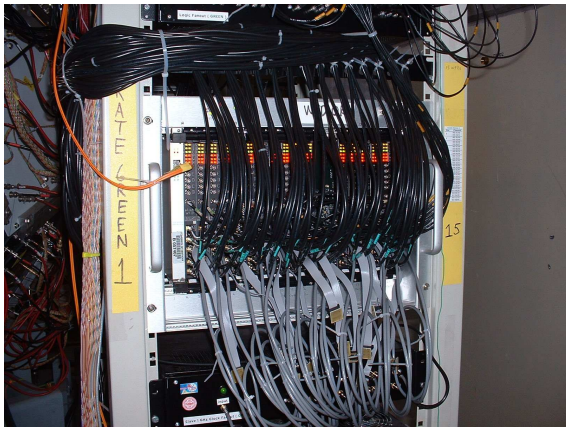

Figure 8: Pictures of the MuLan WFD: on the left, the front-side of the WFD VME module; on the right, one of the six VME crates after installation and cabling.

the 10^{-8} level; this stability exceeds the requirements of the experiment. Prior to the running period, we reconfirmed that the absolute frequency remained within this tolerance.

4. Beam RF signal: While the collaboration does not control this clock, it must be considered in the control of systematic errors. We must avoid "slow components" in the data stream that might be generated by harmonics of our clocks beating against the beam RF. Careful choice of the other clock frequencies prevents this from being an issue.

The main hardware components of the timing system were unchanged from the 2004 running period. In addition to the Agilent E4400 and the two PPG units, the clock system consists of a number of high precision, low skew fanout units for the main WFD timebase, a divider module to down-convert the main timebase for the MHTDC, and a number of NIM fanouts to distribute control signals to the acquisition electronics. This equipment performed as expected; as in the past, we expect no significant systematic errors attributable to these subsystems.

6.2 500 MHz waveform digitizers

6.2.1 Functional description

In 2005, we ran the experiment with a full complement (340 channels) of four channel, 500 MHz waveform digitizers (WFDs); see figure 8. While TDCs were adequate for the 2004 analysis, simple debugging, calibration, pileup, backgrounds, gain and timing stability are far more easily monitored with the full waveform. Each channel consists of analog input buffers, a Maxim MAX106 600 MHz 8 bit flash converter, half the resources of a Xilinx Virtex II FPGA, and a 128 Kbit × 40 FIFO memory. The FPGAs perform triggering and data formatting tasks. The four FIFO memories are connected over a shared bus to a third Virtex II FPGA that links the onboard memories to the VME64 bus, providing a high speed, 64 bit wide data bus. In addition to these high speed data paths, each module contains a number of lower speed buses for control and setup functions; these include a control bus shared between the three FPGAs, a front panel NIM control subsystem, and front panel status LEDs.

FPGAs provide a great deal of power and 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 can be configured at power-on.

In MuLan we have used the configurability of the FPGA to provide behavioral packages called "personalities". We have provided three of these runtime selectable personalities in the initial firmware release:

1. Oscilloscope Mode: When the front panel RUN input is pulled active, the data stream is copied directly

from FPGA input to FPGA output. This is a classical ADC mode, where the data stream contains no time stamps.

2. Fill and Offset Mode: When the front panel RUN input is pulled active, an internal timer is reset, and starts counting clock ticks. Whenever an input rises above a configurable analog or digital threshold, a block of data is copied to the FPGA output. Each block consists of a configurable number of data samples, followed by the time stamp associated with the first sample. In this mode, the time stamp consists of a 16 bit fill count in the data segment, and a 16 bit clock tick offset within the fill. This is the mode used in the experiment.

3. Large Time-word Mode: Operationally, this mode is identical to the Fill and Offset Mode. The difference lies in the format of the time stamp - here a simple 32 bit clock tick counter. This mode is being used by the MuCap experiment.

The inherent flexibility provided by the FPGA allows the construction of additional or modified personalities as needed.

As deployed in the experiment, the WFDs are divided into roughly equal groups, each fed by approximately $1/6$ of the detector. This required six VME crates with fourteen or fifteen modules each. Because a coincidence is required between inner and outer scintillator tiles, each pair is wired to the same WFD module, in order to minimize data lost in the event of a WFD failure. 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.

6.2.2 Progress and status

By the spring of 2005, it was clear that our second prototype worked. The late spring and early summer of 2005 were devoted to tests of the hardware and incremental refinements of the firmware and, in July, we proceeded with the production and assembly of the one hundred circuit boards. The finished boards were tested then shipped to PSI in groups of roughly thirty. By early September, 92 working boards had arrived, 86 of which were installed on the MuLan ball. A waveform from the decay positrons is shown on the cover of this document.

The arrival of the waveform digitizers and the accompanying huge increase in data volume was a tremendous challenge to the the data acquisition and the on-line analysis team. Perhaps the most vexing problem concerned apparently random and intermittent failures in some of the VME crates. Data registration would fail in isolated channels or whole cards, resulting in unreadable data for several runs, and then seemingly correct itself, only to fail again a few minutes later. This failure mode was ultimately traced to an apparent limit in the number of boards one can safely install in a given crate. This problem was not seen in bench tests, where as many as sixteen boards were successfully installed in the test crate. For reasons that are still under investigation, the insertion of subsequent WFDs results in a rapid increase in the frequency of this failure mode. The addition of a seventh VME crate to the system should provide a simple remedy.

The firmware used by the FPGAs on the board underwent almost a dozen revisions over the course of the fall run. One particular problem was our inability to run the WFD FIFOs in concurrent mode, where the FIFOs are written to and read from at the same time. Instead, to ensure the integrity of the data, we were forced to serialize reads and writes, thus severely limiting the livetime. The corruption of the data in concurrent operation was traced to a timing error in the readout, and was corrected by the middle of the data taking run.

The firmware used to configure the FPGAs is now adequate but needs further refinements. There are still hints of stale flash ADC samples in some channels and extra 32 bit words. We see no sign of time word errors. Perhaps the most critical upgrade concerns the block length. During the fall 2005 run, the minimum pulse length was 32 samples, or 64 ns of data. The minimum length could have been changed to 24 samples, or 48 ns but the optimal length, 16 samples or 32 ns, which will allow us to handle the much higher rates that we expect in 2006, doesn't work yet.

6.3 LED calibration and testing

An LED is mounted on each of the PMTs on the MuLan ball. Fast light output (few nanoseconds) is made possible by a triggering circuit mounted with the LED and embedded in the detector housing. The individual LED pulsers are controlled by driver circuits which provide both the trigger pulse and power for the light flash. The trigger pulses are generated in a VME-based programmable pulser, and routed along with the power to each LED through a driver box.

Twelve driver boxes were constructed and installed during the summer of 2004. Each box controls up to 32 LED pulsers. Six of the boxes were fully operational and test data was recorded during the fall 2004 run. In 2005, the remaining six boxes were commissioned to the system and approximately 85 percent of the detectors are now responding as expected to LED light flashes.

A few problems were encountered during the process of installation and testing. Individual channels have failed due to occasional component failure on the LED boards or driver circuits. There are also a few channels where electronic noise seems to be capable of generating false triggers. These problems will be addressed before the upcoming run.

In general the system is capable of pulsing individual detectors with light pulses comparable to the signal produced by decay positrons, at rates of several kilohertz, in an arbitrary computer controlled pattern. Before the summer run, the system will be used to study the detector system response, both as a preparation for upcoming data recording and as a method for assessing systematic gain and timing errors.

7 Data acquisition

During 2005 our DAQ efforts were focused on the development, commissioning and running of the data acquisition for the fall MuLan run. The acquisition provided readout for two new components of the MuLan experimental setup: the 500 MHz WFDs instrumenting the 340 scintillator tiles of the MuLan detector ball, and a custom built FPGA instrumenting the 96 wires of the EMC. In addition, the DAQ provided the control and readback for the electrostatic kicker, beamline magnets, high voltage system, LED calibration system and various current, field and temperature probes.

The data acquisition setup is based upon a local gigabit network of eight frontend processors, one backend processor, two redundant disk arrays (RAIDs) and a LTO3 tape robot. We developed the DAQ with a modular, distributed design philosophy using the MIDAS software library. The DAQ was designed to accumulate data in sequential chunks or data segments with fixed lengths of typically 80 ms. Shorter data segments produce more manageable events sizes while longer data segments reduce acquisition dead times.

Utilizing the on-board memories in the WFDs and the EMC electronics, each data segment represents a 80 ms deadtime-free history of the activity in the MuLan scintillator ball and the EMC. Between each data segment, the frontends were responsible for completing the readout of the data fragments from the detector hardware and sending the fragments over the gigabit network to the backend processor. One frontend – denoted the master – was responsible for synchronizing the sequence of data segments across the various frontend processors. A schematic of the system is shown in figure 9.

Six frontend processors were responsible for the data readout from the six VME crates housing the 340 channels of waveform digitizers. At the start of the 2005 run the readout was operated with sequential writing (to the WFD memory from the WFD FADCs) and reading (from the WFD memory onto the VME bus) of digitizer data. Near the end of the 2005 run the readout was modified for faster, concurrent read/write operation. Data rates per WFD frontend processor were typically of order 10^5 pulses, about 4 MB, per second.

The EMC frontend provided readout of all chamber hits during the measurement period and prescaled chamber hits during the fill period. As noted earlier, the prescaling of fill period hits was accomplished with a custom built FPGA. An additional frontend processor was dedicated to the control and the readback for the LED calibration system and the slow control components.

The backend processor was responsible for the assembly of the data fragments from the different frontends and the transfer of the resulting events to our RAIDs.[2] The backend also managed the user control of running conditions via the MIDAS web interface and the MIDAS online database. Since our data storage on local disk

[2]Note that our RAIDs also doubled as our online data analysis cluster.

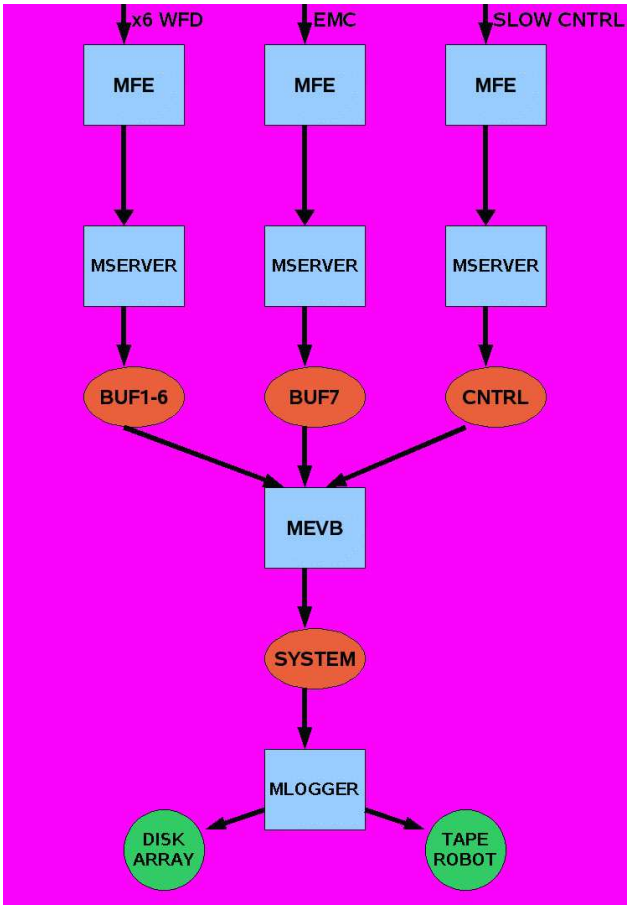

Figure 9: Schematic of data acquisition system showing the frontends (MFE's), event builder (MEVB), and data logger (MLOGGER).

arrays was limited, we asynchronously migrated the accumulated data to permanent storage on our LTO3 tape robot system and backup storage on the PSI tape robot archive. The high data rate (approximately 25 MB/s) from the frontend processors, made the backend throughput the rate-limiting bottleneck for the data acquisition in our 2005 run.

During the fall 2005 run we accumulated data at a time-averaged decay-positron rate of about 0.5 MHz and recorded a total of about 10^{11} decay positrons with digitized pulse shapes. This represented a major improvement over our 2004 run which yield approximately 10^{10} decay positrons with digitized leading-edge times. However, some difficulties and limitations were encountered with our 2005 DAQ system. One problem noted above was the occasional misformatting of waveform digitizer data. Another problem was significant deadtimes associated with the data throughput on the backend processor. These problems are currently being investigated using waveform digitizer and acquisition setups at BU and UK, and will be remedied for our 2006 running. Related work on lossless compression of digitized pulses is also underway.

8 Real-time data monitoring

The fall 2005 run featured a new data monitoring system. The MIDAS Analyzer framework is used to process from ten to one hundred percent of the raw data in real time and produce histograms representing the most important features of the data. At the end of each run, another program analyzes the histograms and extracts and stores in a database the integral parameters which describe data quality. Examples of such parameters include the number of dead channels, the extinction factor, the data collection rate, etc. A

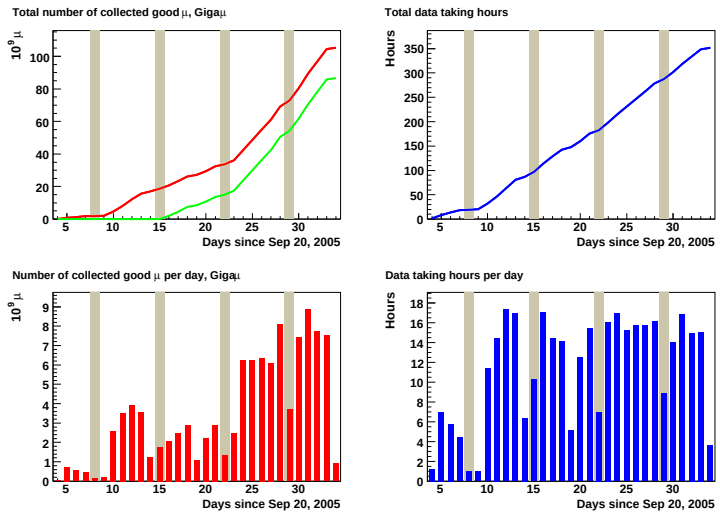

Figure 10: Overview of 2005 data collection. We estimate at least 8×10^{10} good muon decay positrons were collected.

custom web-interface provides access, for both operators and experts, to the runs themselves, as well as the beamline data-monitoring databases. Histograms for a particular run can be selected, or the behavior of any parameter stored in the database can be graphed over a period of time or range of runs. The online data monitoring kept data-quality high during production running and helped us to spot and correct occasional problems. A sample of plots is shown in figure 10.

9 Analysis of the 2004 data

The analysis of the μ^+ lifetime data (see figure 11) collected at the end of 2004 continued throughout 2005; by the summer of 2006, we hope to present the most precise value of the muon lifetime to date. The 2004 data set was handicapped by frequent failures of the kicker and the use of multihit TDCs instead of WFDs. This resulted in substantial time spent investigating systematic uncertainties associated with a larger background, higher pileup rates, and effects introduced from the electronics. Despite these obstacles, we expect a fractional measurement error of 8-10 ppm, a precision two times better than the current world average.

The systematic issues include:

- 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
- Detector asymmetry from non-target muons stops

Details on all of these topics can be found in appendix A of the report.

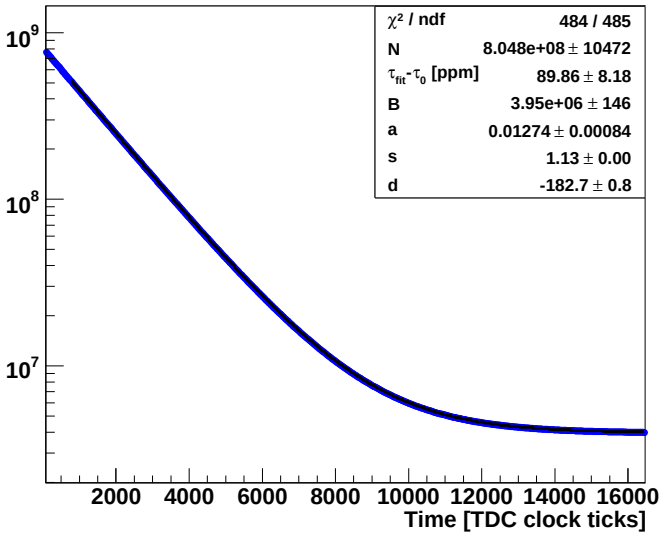

Figure 11: Full statistics lifetime histogram with fit. The fit parameter $\tau_{\rm fit} - \tau_0$ shows the statistical precision of this data set to be 8.2 ppm.

10 The 2005 data set and analysis plans

We estimate that the 2005 dataset contains at least 8×10^{10} decay positrons (see Figure 10), allowing us to extract a result to better than 4 ppm statistical uncertainty. Because the new WFDs record samples of every waveform for every event (enabling detailed study of gain, pileup, and backgrounds), the data set is large. We recorded 16.5 TB on 32 LTO3 tapes. The data is backed up to the PSI tape archive, and the original tapes have been returned to Illinois for analysis. We note that the final data set for a 1 ppm result, uncompressed, will be at least an order of magnitude larger.

The analysis will be carried out in three steps. The first step, which was recently completed, is to build an average pulse shape for each the 340 individual detector elements, using a small part of the data; see figure 12. These average pulse shapes are then used in the second step, a full pass through the data using a custom pulse finding algorithm, to determine pulse areas (proportional to the deposited energy) and times. The full pass will also sanitize the data by rejecting pulses from phototube breakdown (rare, but evident in the data for several of the PMTs) and fills in which muons enter the detector during the measurement period. In the third step, a pass through the processed data will be used to explore systematic effects and to extract the lifetime. To prevent bias, the exact WFD digitization clock frequency is hidden from the analyzers at the 200 ppm level, just as in the 2004 analysis.

Results from the 2004 data analysis allowed us to fix some parameters during the run and explore others. For example, the 2004 data analysis gave us enough confidence in the AK-3 target that it was used exclusively in 2005. Additionally, the kicker was run exclusively with a time structure of 5 μs beam on, 22 μs beam off. In the summer beam studies, our best beam tune achieved a maximum rate of approximately 3 MHz, with good extinction. We began the fall data taking using this same tune, but with continued studies, managed to raise the flux to 7 MHz for the last week of the run, without compromising the extinction factor. We also investigated several systematic effects outlined below.

- Pileup: With the 2 ns sampling of the waveform digitizers, we should be able to resolve pulses whose centers are separated by as little as 3 ns. We are investigating the use of an area (deposited energy) cut to resolve pulse pairs whose centers are even closer.

- Gain Changes: To examine the low-energy pulses from the phototubes, the WFD digitization threshold was set just above the noise. This threshold setting allows the recording of an abundance of pulses

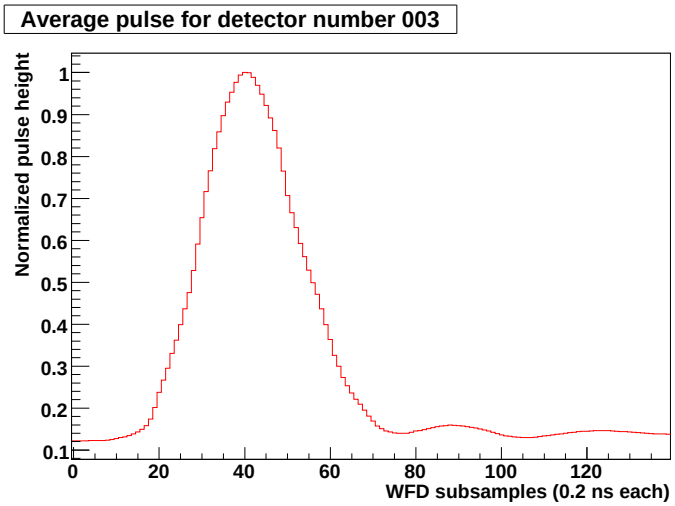

Figure 12: The average pulse for one detector, shown at ten times the WFD timing resolution. The first step in the analysis involves determining average pulse shapes for each detector.

to cross-check tube gain versus measurement time. At 5 MHz DC beam rate and above, a slow gain change in the tubes was observed. The gain change took place over several hours, and will not affect the determination of the muon lifetime. More analysis and bench tests will be done to investigate the drift.

- Background stability: Various measurements were made of the background time spectrum which, on the time scale of the muon lifetime, should be as flat as possible. Dedicated runs were taken to check the kicker extinction versus kicker high voltage. Direct measurements of kicker voltage versus time were made using capacitive pickups near the plates. Randoms counters and pulsers test the stability of the readout electronics. Overall, kicker flat top measurements, combined with measurements of the beam extinction versus kicker voltage, indicate that the error due to non-flat beam backgrounds will be below 1 ppm.

- Polarization effects: Since the surface muons collected by the πE3 beamline are highly polarized, slow muon precession could distort the extracted lifetime. A magnet around the EMC and the high internal magnetic field of the AK-3 target ensure that the muons precess at a high rate in any high-Z materials. During the run, the magnetic fields in the EMC and the target were flipped at regular intervals. However, as in the 2004 data set, we are still vulnerable to stops in a low magnetic field in the helium bag between the EMC and the target.

11 Beam request 2006

The MuLan experiment is ready for physics production running. The time required to reach our statistical error goal is determined by the beam rate, the kicker period (27 μs), the muon lifetime itself (approximately 2.2 μs), the fit start time (roughly 500 ns), the solid angle acceptance of the detector (75 percent for MuLan) and the data acquisition livetime. Assuming that we operate sixteen hours per day at a DC beam rate of 8 MHz, we expect to record about 4×10^{11} muon decays in an eight week production run during 2006. This is almost one half the statistics of the original proposal. We would probably need an additional two weeks for systematic studies and one half week for installation and assembly.

We have discussed the running schedule with the other main users of the πE3 line, MuCap and ALC, and have mutually agreed on a schedule that is convenient and appropriate for all parties. This schedule places the continuation of the MuCap run at the beginning of the running cycle, followed by a 4.5 week

intermediate run of ALC, a 10.5 week run of MuLan, and a final lengthy run for ALC. The exchange time is included and is based on our real experience in dismantling and setting up the experiments. The suggested dates are given in the following table:

User	Dates	Duration (weeks)
MuCap	April 5 - June 19	10.5
ALC	June 21 - July 24	4.5
MuLan	**July 26 - October 9**	**10.5**
ALC	October 11 - shutdown	9.5

References

[1] T. van Ritbergen and R. G. Stuart, Nucl. Phys. B, **564**, 343–390. (2000)

[2] F.R. Cavallo *et al.*, "Precision measurement of the μ^+ lifetime G_μ with the FAST detector", PSI proposal R-99-06 (1999).

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

[4] MuLan Progress Report 2003, MuLan Collaboration, R. Carey *et al.*, `http://www.npl.uiuc.edu/exp/mulan/documents/ProgressReports/MuLanProgressReportJan2004.pdf`

[5] MuLan Progress Report 2004, MuLan Collaboration, T. Banks *et al.*, `http://www.npl.uiuc.edu/exp/mulan/documents/ProgressReports/ProgressReport2004.pdf`

[6] J.F. Ziegler, J.P. Biersack and U. Littmark, Stopping and Range of Ions in Matter, Pergamon Press, New York, 2003. `http://www.srim.org/`.

[7] Geant4 Collaboration (S. Agostinelli *et al.*), Nucl. Inst. Meth. A, **506**, 250–303 (2003).

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

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 is done "blind". In this case, the analyzer is 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, τ_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, along with a fit (figure 11), appears in the main part of this document. The statistical uncertainty on the lifetime, a lower bound on the ultimate precision, is 8.2 ppm. Note that the reduced χ^2 is nearly 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 900 ns. Any fit start time from 900 ns on gives consistent results, as seen in figure 13.

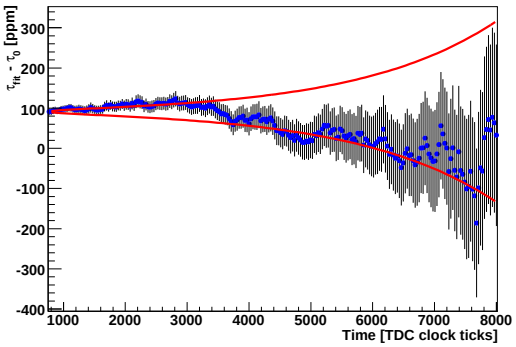

(a) The fitted lifetime in ppm with the red bands indicating the allowed 1σ deviation from the initial fit time.

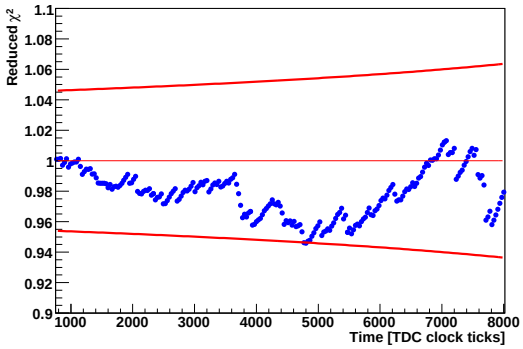

(b) The reduced χ^2 with the red region showing the 1σ allowed deviation about one.

Figure 13: Variation of fit parameters as a function of fit start time.

The 2004 total data set included many different running conditions. 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 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 fitted lifetimes of these subsets are all statistically consistent.

A.1 Systematic studies

Much of the analysis effort has been spent focusing on the various systematic errors.

A.1.1 TDC effects

The use of the CAEN v767 multihit TDC led to some of the largest systematic issues. Most issues related to this TDC were known before the run and had already been investigated to some degree.

- The TDC is capable of recording both the leading and trailing edges from a discriminated signal, but the small input buffer, along with the slow readout, forced us to use the leading-edge-only mode, just to reduce the number of missed events. The use of only the leading edge results in a loss of information for pileup reconstruction, but it effectively reduced the overall dead time by a factor of two.
- The TDC operated with an input clock of 23.7 MHz and uses a delay locked loop (DLL) to interpolate the frequency by a factor of 32. At best the DLL had a differential non-linearity of 5 percent. This small shift in time is addressed by using bins that correspond to the 23.7 MHz input clock for the lifetime histograms. An additional problem occurs if the input clock signal is interrupted or deformed, causing the DLL to stop functioning. To avoid this problem, great care was taken to keep the clock signal quality high.
- The most troubling systematic effect was identified during post-run analysis: every hit in a CAEN module results in a small oscillating shift of the time given to subsequent hits. The oscillation continues for a few microseconds after the initial event (see figure 14). This had two effects. First, because the time of the kicker was recorded in each TDC module, there was an oscillation term found on all signals. Second, it raised the possibility of an early to late effect by having one positron affect the recorded time of a later one. Errors associated with this latter problem were investigated in several months of dedicated lab tests. We found that by including the lab results in the fit, we could alleviate the effects of the oscillation component. Moreover, the effect of one positron on another is negligible because of the low rate of hits per detector and the dephasing of the electrical oscillation.

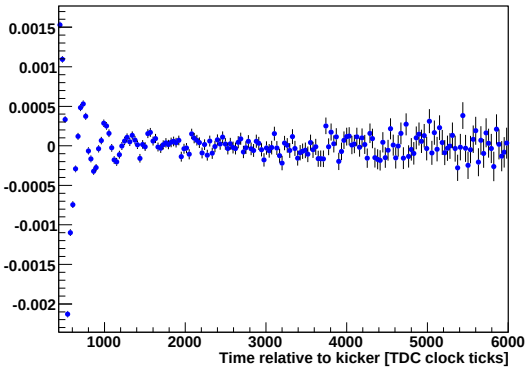

(a) Residuals from a simple fit with no TDC noise term.

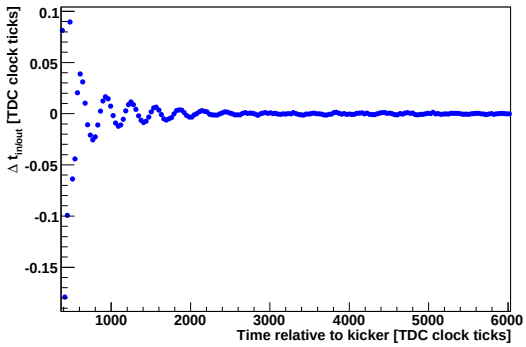

(b) Lab results showing the time between a split signal sent to two different TDC modules.

Figure 14: Oscillations from electronics seen in the data.

- There were several other potential sources of errors from the TDC that are not fully addressed here, but have been addressed either in the analysis or during the setup before the run. These include spurious trailing edges, recording of rollovers, and the detection of missing clock ticks.

A.1.2 Pileup and 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. These missing events are called pileup, and if

unaccounted for, result in a large systematic pulling on the fitted value for τ_μ. To correct for pileup, we artificially reproduce the distribution of missing events 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 at least a 20 ppm shift in the lifetime. Based on simulations and the varying of pileup reconstruction parameters we conclude that the pileup correction adds a systematic uncertainty of 1 ppm to the 2004 data set.

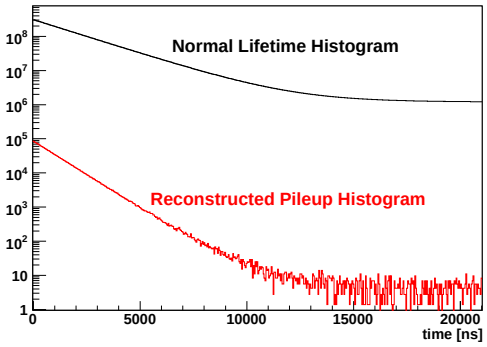

(a) Lifetime and Reconstructed Pileup Curves.

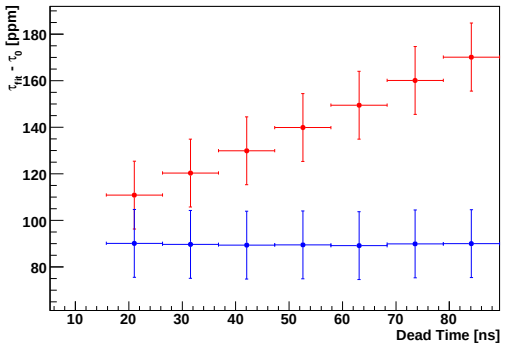

(b) Lifetime versus artificial dead time. Red(blue) points have not(have) been corrected for pileup.

Figure 15: Corrections for pileup and TDC dead time.

There are other effects that can be corrected using a statistical reconstruction, as is done for pileup. 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.1.3 Flat background

During the measurement periods of 2005, 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. Measurements of the change in extinction factor versus kicker voltage were performed during the run. These can be combined to estimate the maximum possible change in background rate over the measurement period at about 0.8 ppm of the total rate. Toy Monte Carlo simulations indicate that the maximum resulting systematic error in the lifetime

measurement is 3.4 ppm. Since this is an upper limit based on the lowest kicker voltage and extinction factor from the 2004 run, we expect to reduce this error with further analysis of the 2004 data.

A.1.4 Gain and pedestal 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. Preliminary simulations indicate that a 1ppm shift in detector efficiency during the measurement period could result in up to a 0.25 ppm systematic error in the lifetime. In order to measure the change in detector efficiency with time, it is useful to study the spectrum of pulse heights produced by the detectors. There are a limited number of channels from the 2004 data set which include WFD data that will allow offline study of the changes in gain and pedestal with time. However, the full ball was read out via WFDs in the 2005 run. Analyzing the 2005 WFD data will allow us to calculate a systematic error due to gain and pedestal shifts from the 2004 run.

A.1.5 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. Preliminary simulations indicate that a 3 ps shift in the timing could cause a 1 ppm systematic error in the lifetime measurement. An estimate of the timing shift will be obtained 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 hit in those channels recently. Given the results from bench test measurements made several years ago, we do not anticipate any problems.

A.1.6 Residual polarization and stopping distributions

Most recently, we have been focusing on the error presented by the residual polarization of muons which stop in the helium corridor between the EMC and the target. As described earlier (section 2), the spins of muons in the AK-3 target 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

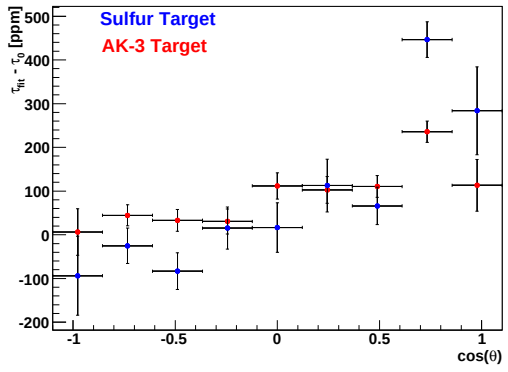

(a) Sulfur and AK-3 data sets. The symmetry of the detector results in a large reduction of this systematic error when the tiles are summed together.

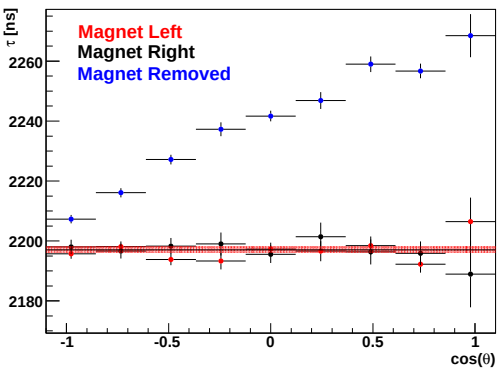

(b) Special study with a plastic plate at the EMC. The current world average is shown for reference but the data itself contains a hidden offset.

Figure 16: Lifetime versus cos(θ).

target region in a low 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. With conventional fitting techniques, it is difficult identify the precession frequency for muons stops in the corridor, but the effect is apparent when plotting the fitted lifetime as a function of position on the detector.

Anticipating trouble with muon stops both in the corridor and in the EMC region, we performed a few dedicated studies with a plastic plate placed upstream of the target region. We made tests with the plate placed just after the EMC and half-way between the EMC and target. In the first, we applied several magnetic fields to the EMC region. With no magnetic field the variation in lifetime versus theta was several percent; but applying a magnetic field significantly reduces the polarization effect of these muons stops. By rotating the field at regular intervals we were able to average out the effect at this position. Using the data when the plate was midway between the EMC and target we can infer the error caused by muons stopping in the corridor. Based on this study, along with stopping distributions measured using the MuCap detector, we estimate the residual polarization effect from muons stopping at this location to have less than a 10 ppm effect on the lifetime.

A.2 Current status

The analysis of the 2004 data set is quickly coming to a close. We have determined that we can produce a fractional statistical uncertainty of 8–10 ppm (see figure 16). Most systematic issues are well understood with the remainder under active investigation. We hope to finish the systematic studies by May 2006, unblind the analysis, and release the final value from this data set by the summer of 2006.