

Cosmic Ray Muon Capture and Decay With Modern Digital DAQ

Wojtek Skulski, SkuTek Instrumentation

Summary

The goal of this project is to detect, record, and analyze muons decaying within a detector, using a modern Digital Data Acquisition (DDAQ) device developed and manufactured by our company. We are presenting a *modern digital signal processing alternative* to conventional analog NIM electronics.

When the protons in cosmic rays interact with Earth's atmosphere, unstable particles are produced in the upper atmosphere. These particles decay into μ -mesons we seek to detect (*Experiments in Modern Physics*, Adrian C. Meilissinos & Jim Napolitano, p. 399). Solomon Shulman (SkuTek Fall 2025 intern) developed an experiment shown in Figure 1, using our Vireo digitizer. Solomon developed a data acquisition script in Python and collected event data over 21 hours. This writeup presents the analysis steps applied to the event data file in order to measure the muon half life.

Experimental Setup

We used a 15" x 15" x 2" plastic scintillator with a PMT, shown in Figure 1. The scintillator was manufactured by Thermo Fisher. We acquired it on Ebay. Thanks to a large detector area (15" x 15" = 1,452 cm²), we expect about 1,452 μ -mesons per minute hitting the detector. Based on our previous studies, we expect that one out of 6000 such mesons will stop and then decay within the detector volume. We thus expect 24 punch through mesons per second, and 15 stopped mesons per hour. We can thus demonstrate the physics after about one day of running. In addition to the cosmic rays, the detector also responded to ambient radioactivity of potassium-40 and possibly radon.

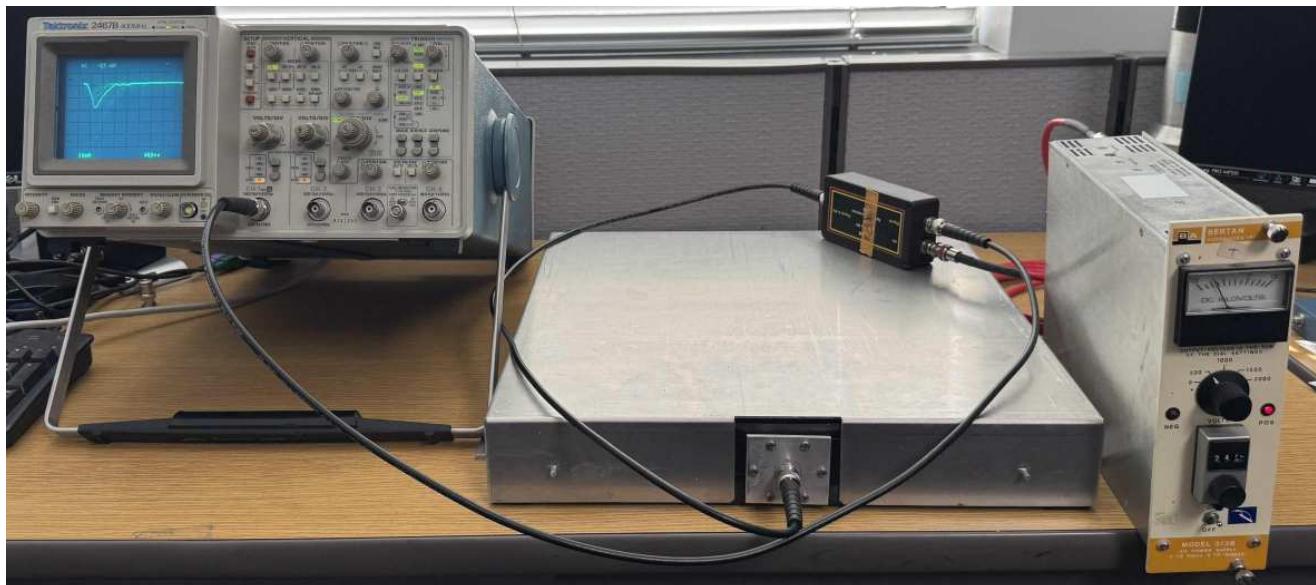


Figure 1. The 15"x15"x2" plastic scintillator + PMT in a metal box (middle) powered via HV splitter box (next to the HV supply) is delivering negative pulses to an analog scope.

We verified the detector operation using an analog scope shown in Figure 1. We then connected the detector to our 2-channel Vireo digitizing the signals with 14 bits at 100 MHz. The digitizer could record the waveforms to disk in ASCII format for subsequent offline processing.



Figure 2. The 15" x 15" x 2" plastic scintillator + PMT in a metal box (middle) powered via HV splitter box (far right) is delivering negative pulses to the Vireo digitizer. The triggering and signal processing was performed online in real time by Vireo, forming a complete and modern Nuclear Physics experimental system. No external electronics was used other than the photomultiplier HV supply.

How Is Our Approach Different From Analog NIM Experiments

Cosmic ray capture and decay is a classic experiment commonly used to train students. Several comprehensive descriptions of such student labs are collected and summarized in Appendix A. These apparatuses used traditional analog NIM electronics such as discriminators and Time to Amplitude Converters (TAC), some of which were decades old. The manufacturers of those units went out of business long time ago. (Neither Phillips Scientific nor LeCroy are in Nuclear Physics business anymore.) This is unfortunate. On one hand, it is imperative that students learn how to use traditional Nuclear Physics electronics which is as valid today as it was decades ago. On the other hand, there is some irony in teaching modern physics with old electronics reclaimed from decaying electronic pools. We are presenting a modern alternative to such a traditional approach.

Analog NIM electronics is enabling the students to measure the muon decay time distributions, but not the amplitudes of the signals. Amplitude measurement would require using CAMAC or VME converters, crate controllers, and event recording software. None of the listed experiments went in this direction. Also, random coincidences did not receive adequate attention in most measurements mentioned in Appendix A, even though randoms are an inevitable admixture to any muon experiment. Our approach provides a consistent approach to both the signal amplitudes and to the random coincidences. Both were recorded in our event files and analyzed offline.

The subject of ambient radioactivity was barely tackled in traditional experiments. This contamination can be diminished or eliminated with increased thresholds, but the methods of adjusting the thresholds were seldom elaborated in sufficient detail, due to lacking amplitude information. We are presenting a consistent approach to ambient radioactivity when the detector can deliver spectroscopic information.

All in all, the conventional muon experiments described in Appendix A do not come close to what we could accomplish thanks to having the *event-by-event time-amplitude correlation* in our data. This may be seen as either the advantage or disadvantage of our approach. On one hand, we are showing the richness of information which is hidden under the surface of this allegedly simple experiment. On the other hand, will the students appreciate the beauty of the modern digital data acquisition and intricacies of the analysis? Is it likely that students will see the detailed information as being the undue burden caused by the modern, data-rich approach? We are hoping that students will appreciate both the contemporary electronics as well as digital signal processing over the old methods used in teaching modern physics at many educational institutions.

Data Acquisition Software Script

Event triggering and waveform recording was performed by our Vireo digitizer, running the FPGA firmware developed at SkuTek and a data acquisition script written in Python by our intern Mr. Solomon Shulman. The Vireo and the script formed a complete DAQ. The script generated four files:

1. *Waveform event file*

When a trigger was detected on the channel connected to the detector, the firmware examined the number of subsequent triggers within the *pileup detection time window* of 1500 samples (15 μ s). In case this number was greater than one, the candidate muon decay waveform was recorded to the event file. The digitizer also recorded the single-trigger events prescaled by 5,000 to provide an unbiased sample of single-trigger events.

2. *Histogram file*

Every time an event was triggered, the height of the differentiated leading pulse was accumulated in a histogram within the Vireo FPGA firmware. The histogram was periodically saved to disk every 100 events. (In future experiments we will increase this number to 500.)

3. *The settings files*

The script generated a file logging all the settings used for general data and troubleshooting. The pre-run file was generated directly before the main program loop. The post-run settings file was recorded to disk at the end of the experiment. Both files should be identical to indicate that the settings were not corrupted during the run.

Display

Throughout a run, the user could check the status, looking at the text periodically printed to the terminal once every PRESCALE_DISPLAY triggers. The script reported the number of single-trigger and multiple-trigger events. Additionally, the firmware histogram was periodically written to disk in ASCII. It could be quickly displayed to verify progress. Newer versions of the script will use either Igor or gnuplot ASCII format for this task. (Gnuplot format will be added soon.)

Running the Experiment

To run the experiment, the user must *ssh* into the digitizer. Ensure that the script is downloaded and modified to your specifications, and run it with *sudo*.

Example: *sudo python3 muon_decay_detection.py*.

The script will now be running and collecting events on every trigger.

The Event File

We chose Igor for our offline processing because of its rich graphics and analysis-friendly language syntax. We accumulated several examples of data analyses over the years, as well as a stash of graphics format macros. Other users will undoubtedly use their own data analysis software, like Python, Matlab, Mathematica, Physics Analysis Workstation (PAW), or ROOT. The event file format will need to be adjusted accordingly, what should not pose much of a problem because the source code of the DAQ script is available to the Vireo users. In our case we proceeded with what we already had.

The ASCII event file was structured as follows. I added line numbers to the file, which are not present in the actual file. The numbers highlighted in yellow are parameters of the Igor function calls.

```

1.    IGOR
2.    X // Format    = "IGOR WAVES"
3.    X // Datetime = "UTC Time: 2025-11-13 13:55:2"
4.    X InitProcessing(1000) // reserve tables in-memory
5.    WAVES/o/D ADC          // "ADC" is the wave name, repeated below
6.    BEGIN
7.    -641
8.    -642
9.    ... more samples
10.   -643
11.   -643
12.   END
13.   X ProcessOneEvent(ADC, 1, 2) // wave name, evt num, how many triggers
14.   Repeat the sequence WAVES ... ProcessOneEvent for every event.

```

The file starts with the declaration `IGOR`, which tells Igor to sequentially process the file. Each line of the file proceeds with Igor syntax, which could be as well typed into Igor from its command line. Thus, the entire file is one big Igor macro. This ingenious trick is built into Igor. It lets the user drag-and-drop the file into Igor to start it being processed.

The letter “X” in the 1st column means “execute”. It can be followed by any Igor command. You can display figures, add labels, etc., with every such command. Since the “//” means “comment” under Igor, the lines starting with “X //” are comments added for documentation. The first “real” command is `X InitProcessing(1000)`, where “1000” denotes the expected number of events. You can specify more than expected, but not less. This line will call a function `InitProcessing` written ahead of time and provided to Igor in text format in the Igor Procedure File (extension `.ipf`). In our case this function will reserve in-memory tables for pulse times of occurrence, areas, or pulse heights. These tables are known as “ntuples” under other data processing packages, such as ROOT. The tables are then

filled by another function `ProcessOneEvent`, which will calculate the required quantities and add them to the in-memory tables, one set of values for each event at a time.

The numbers stored at locations n across all tables are correlated, because they describe the same event. This is the foundation of the correlation analysis of the event-by-event files.

The event file proceeds for the requested number of events, until the end. You can add another function call at the end, which is `ShowResults()`, if you prepared such a function ahead of time. Both the `InitProcessing` and `ShowResults` can be also executed “by hand” from the Igor command line.

All in all, the Igor approach is analogous to Python, albeit with a different syntax of the Igor language. The main Igor asset is that it is “ready to go” with comprehensive graphics and numerics, all under one roof. Realistically, using Igor is more efficient than a Do-It-Yourself development of all these features would be. The main downside is that Igor is not free, although its academic pricing is quite affordable.

The Outline of the Offline Analysis

The offline processing consists of calling the `InitProcessing` once in order to reserve the in-memory “ntuples”. This function can be either called “automagically” as shown above, or typed by hand under the Igor command prompt, prior to drag-and-dropping the event file. The function `ProcessOneEvent` is then called for every stored waveform. Most waveforms are known to contain two pulses detected by the FPGA firmware. The processing function is then searching for them using a very simple leading edge software discriminator. After the pulse is found, its height and time are both calculated and stored in the tables. The baseline is subtracted from the waveform prior to these calculations. The baseline is determined from the first 128 samples of each waveform.

The most important quantities are of course the times of occurrence of the pulses. They are found with a simple weighted average calculation (a.k.a. the center of gravity), and stored. The pulses are also fitted with gaussian functions to provide a consistency check for the center of gravity algorithm. We found that both methods provided consistent results.

In case the waveform contained two pulses (the majority of the recorded events), we also calculated the time difference dT between the two pulses, and stored it in the table. This quantity is expected to show the exponential distribution of the leading-to-trailing time differences, demonstrating the muon decay.

The Results of a Short Test Run

We started with recording a short test run over 40 minutes, and running the event file through the Igor analysis. Initially we did not pay a particular attention to fine tuning the analysis settings. The “somewhat blind” analysis yielded an encouraging result shown in Figure 3 on the next page. It proved that the approach worked through all stages of the experiment. We could then proceed to collecting more data.

Joyful Life of a Cosmic Ray μ -Meson

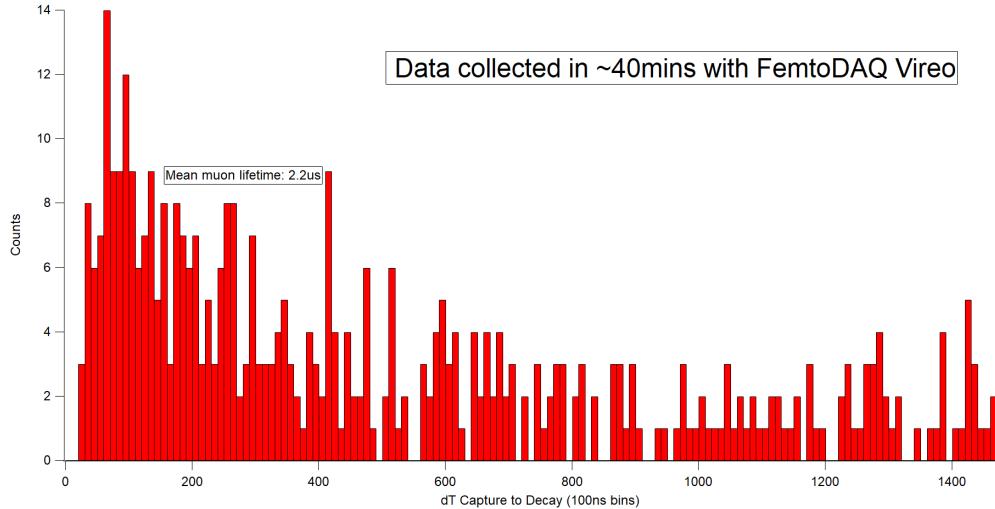


Figure 3. The time interval distribution between the leading and the trailing pulses collected during a short test run.

Analysis of the Long Run

Having demonstrated that the experiment was working, we recorded ~ 18 thousand events in 21 hours, storing 168 megabytes on disk. Then we started fine tuning the analysis and examining the results .

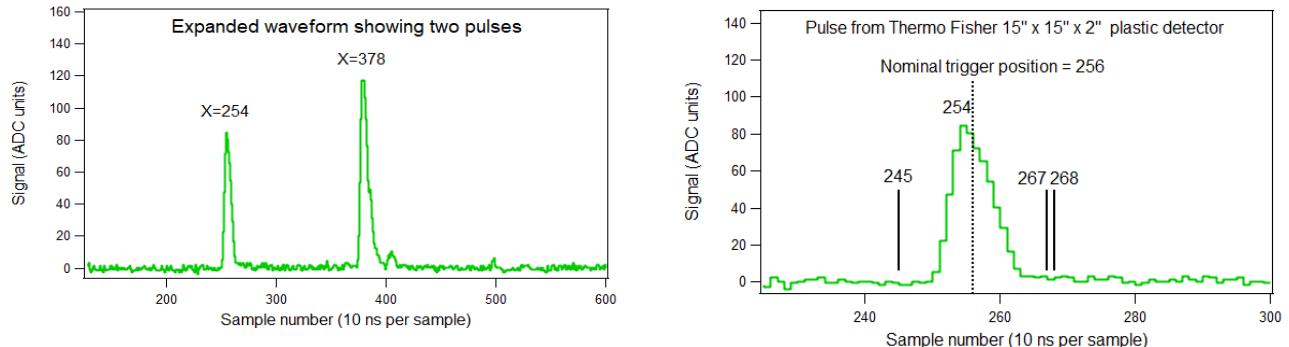


Figure 4. Left: two pulses in a single waveform. Right: the expanded leading pulse shows how we defined the time boundaries for analysis. The leading pulse was expected between samples 245 till 267, while the search for the trailing pulse started at sample 268.

We started with examining a few waveforms and defining the time boundaries where the pulses were searched for. This is shown in Figure 4. The time boundaries were decided looking at just a single event. We need to verify that our decision is robust versus the so-called “time walk” of the pulses caused by their amplitude variations. The time walk is inherent in the leading edge trigger implemented in our FPGA. We need to know how large was the time walk in this data set. This is shown in Figure 5.

Joyful Life of a Cosmic Ray μ -Meson

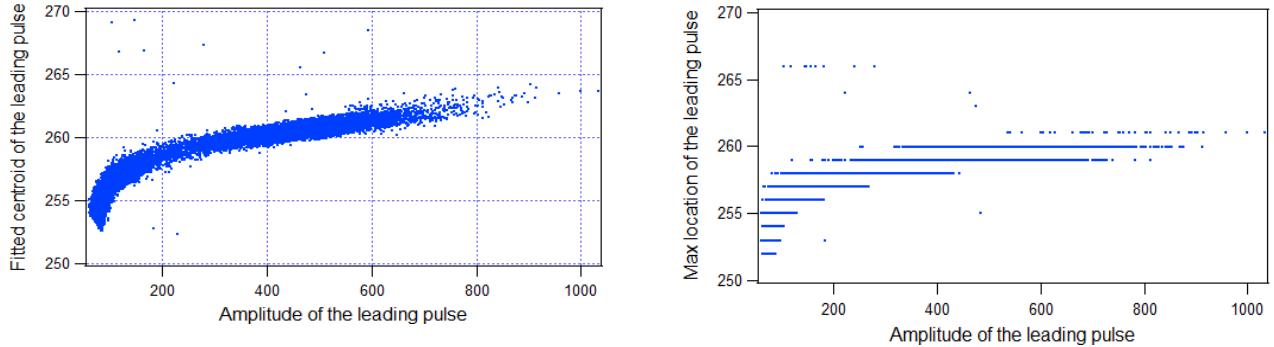


Figure 5. The “time walk” is defined as the change of the pulse location caused by its amplitude. Left: The fitted location as a function of the pulse amplitude. Right: The maximum of the pulse as a function of the amplitude. Maximum is limited to integers. Both plots demonstrate the same effect.

We draw attention to two facts.

- 1) The *time walk* is known to be present in all analog leading edge discriminators. It was remedied in the analog domain by the so-called Constant Fraction Discriminators (CFD), which work beautifully in the analog electronics. The digital version of the CFD is plagued with numerical instabilities, which are often overlooked by the NP community. The discussion of these instabilities is beyond the scope of this report. Please inquire the author for details. Fortunately, in this analysis we do not have to dive into the digital CFD because we have two other methods at our disposal: fitting the pulse with a gaussian profile, and calculating the weighted average (a.k.a. the center of gravity of the pulse). We settled on the fitted positions in this analysis. The COG would yield the same results.
- 2) The time walk ranges from 255 to 265 in the left panel of Figure 5. It means that the time walk extends over ten ADC clock ticks, which is one hundred nanoseconds. This is far from being insignificant! A decent analysis must take the time walk into account. In this analysis we always calculated the time differences dT between the leading and the trailing pulses, whose locations were found with the same algorithm. In other words, we did not assume that the location of the leading pulse was equal to the nominal trigger.

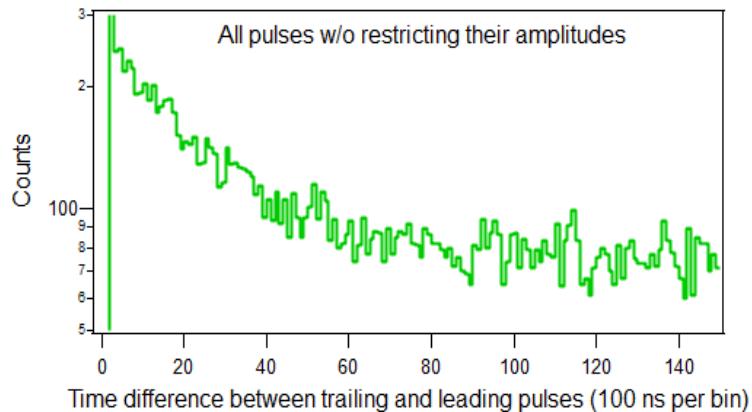


Figure 6. The time interval distribution dT between the leading and the trailing pulses collected during a long “production” run. Note the constant random background ~ 80 counts per 100 ns.

The First Look at the Time Interval Distribution

Having defined our pulse search boundaries and other details, we then looked at the time intervals between the leading and the trailing pulses, without restricting their amplitudes. We saw a pronounced exponential distribution up to about 8 μ s, see Figure 6. Now we could look deeper under the hood.

Leading Pulse Amplitudes

Next, we looked at the amplitudes of leading pulses, shown in the left panel of Figure 7 as a function of time (i.e., the event number). We clearly saw two bands of the amplitudes. We also saw that this feature did not depend on time. Neither band could be attributed to transient events such as turning on or off the air conditioning in the lab. This seemingly trivial observation is in fact of fundamental importance for the analysis. Should we have seen any transient behavior, we would need to question the validity of the data and we would be forced to debug the experiment.

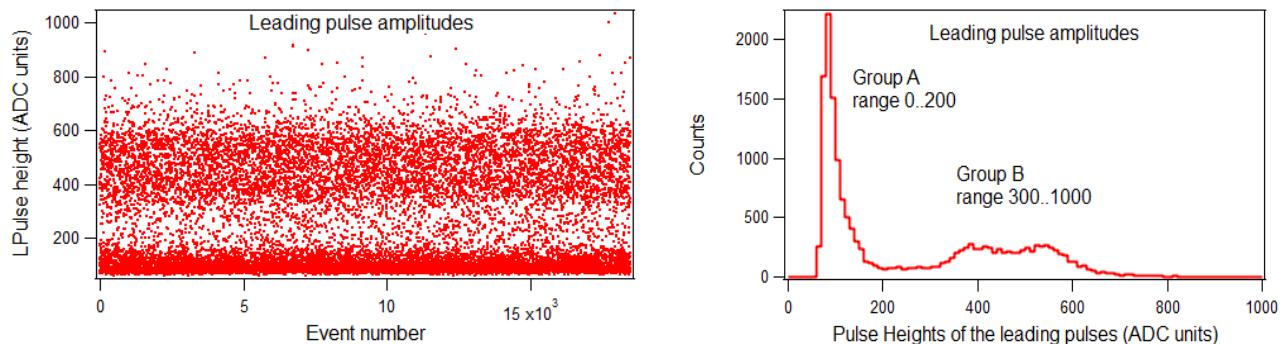


Figure 7. Left: the time list of the amplitudes of all leading pulses. Right: the histogram generated by projecting the time list onto the Y axis.

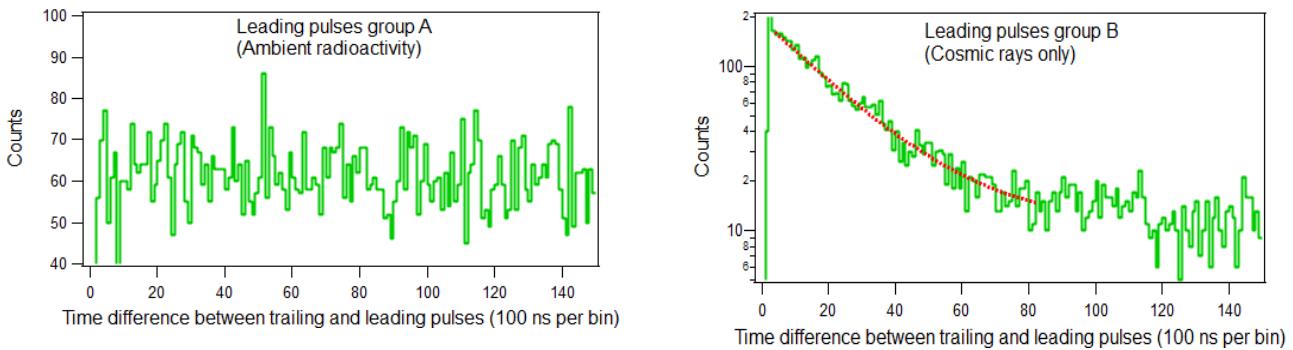


Figure 8. The time interval distribution dT between the leading and the trailing pulses. We selected the leading pulses belonging to group A on the left, and to group B on the right. The dT histogram on the left does not indicate any muon decay. The red line in the right panel is the exponential fit.

The distribution of the leading pulse amplitudes was obtained by projecting the time list on the Y axis. It is shown in the right panel of Figure 7. We saw two groups A and B of events, the low-amplitude group A and the high-amplitude group B. At this point we did not know the nature of these events, so we wanted to investigate how they affected the measured time intervals. We constructed the time

interval distributions separately for these two groups. (These two histograms were projected from the same interval list, with an additional “if” statement in the projection procedure.)

In Figure 8 we plotted the time interval distributions separately for groups A and B. We are now seeing that the low energy pulses A do not lead to the “capture and decay” pattern. The time differences between the *leading* pulses belonging to the group A, and the subsequent *trailing* pulses of any amplitude, are uniformly random, without even a trace of the “capture and decay” pattern. This is the opposite to the group B, where we see a “capture and decay”, which we have already seen in Figure 6.

The “physics behind” these results seems pretty clear. The low energy pulses (group A) were caused by either noise or ambient radioactivity, presumably due to ^{40}K (potassium-40), which is known to be present in the environment. The high energy pulses (group B) were due to cosmic ray muons either punching through, or being stopped in the detector. Both the stopped and the punch-through events were present, as demonstrated by the exponential decay component up to $8 \mu\text{s}$, and the uniform random distribution beyond this time interval. The random events must be also present under the exponential decay, but we cannot clearly separate them from under the exponential decay curve.

Reducing the Random Coincidences

Comparing Figure 6 with the right panel of Figure 8, we see that the random component was decreased from about 80 events per bin to about 10 events per bin for dT larger than $8 \mu\text{s}$. These extra random events were moved to the left panel of Figure 8. The right panel of Figure 8 appears much cleaner than Figure 6, after we excluded group A from the dT histogram. This cleanup will improve the muon life time calculation.

Trailing Pulse Amplitudes

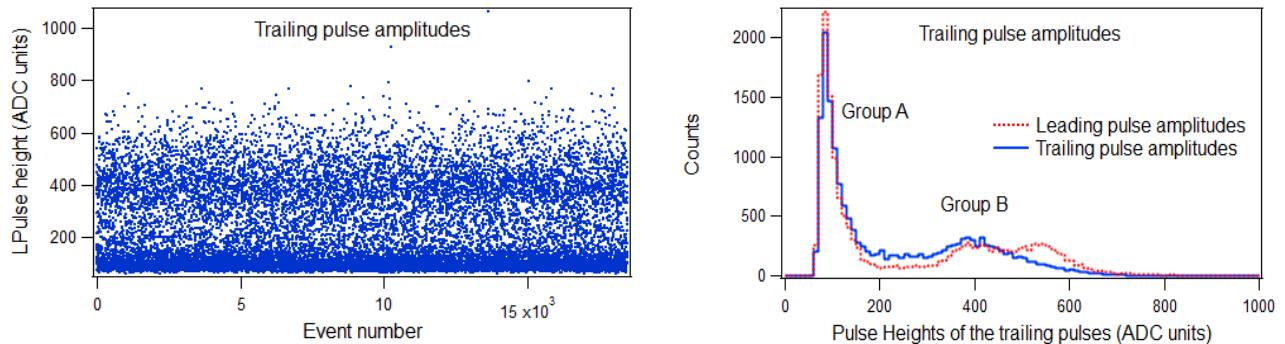


Figure 9. Left: the time list of the amplitudes of all trailing pulses. Right: the histograms generated by projecting the time lists onto the Y axis. Red histogram: leading pulse amplitudes from the earlier Figure. Blue histogram: trailing pulse amplitudes.

Next, we looked at the amplitudes of the trailing pulses, shown in the left panel of Figure 9 as a function of time (i.e., the event number). The figure looks almost identical to Figure 7 which showed the leading pulse amplitudes. We expect that the trailing pulses will show the group A (radioactivity) and a new group of pulses (call it C) due to electron emission from muon decay. We will try to find this new group of events.

Cleaning up the Trailing Pulse Amplitudes in Search for Electrons

The trailing pulse amplitudes in Figure 9 consist of the low energy group A and of the high energy group, which presumably is a mix of the group B (random punch-through muons) and the new group C (electrons from the muon decay). In a series of steps we will remove the “unwanted events” in order to bring up the group C. We will remove accidental coincidences, using the in-memory list of time intervals dT and pulse amplitudes of both the leading and the trailing pulses. We will thus exploit the time-amplitude correlations which are present in our in-memory event lists (ntuples).

Under Igor, the phrase “remove the event” is most conveniently implemented by setting the ntuple entry to “not a number” (NaN), which effectively removes the entry from further processing. If we were using another analysis software then we would use a series of nested “if statements” in the code.

1) Remove the events initiated by the leading group A (noise and radioactivity). We have shown in the left panel of Figure 8 that these events do not involve muon capture and decay.

2) Using the remaining events, create two trailing pulse histograms to separate the “true” from the “random” coincidences which were apparent in the right panel of Figure 8.

2.1 Events with dT between 20 ns and 7.5 μ s. These events are a mix of the muon decay below 7.5 μ s (exponent in Figure 8), and random background, whose admixture must be present.

2.2 Events with dT between 7.5 μ s and 15 μ s. These are mostly random events, as shown in the right panel of Figure 8. We have thus divided the full 15 μ s range into two equal parts.

NB: There is still some residual muon decay admixture beyond 7.5 μ s. In the future we will widen the range to 10 μ s. Adding a few microseconds will better separate the two components.

3) *Crucial step:* subtract the histogram 2.2 from 2.1 in order to correct for random coincidences.

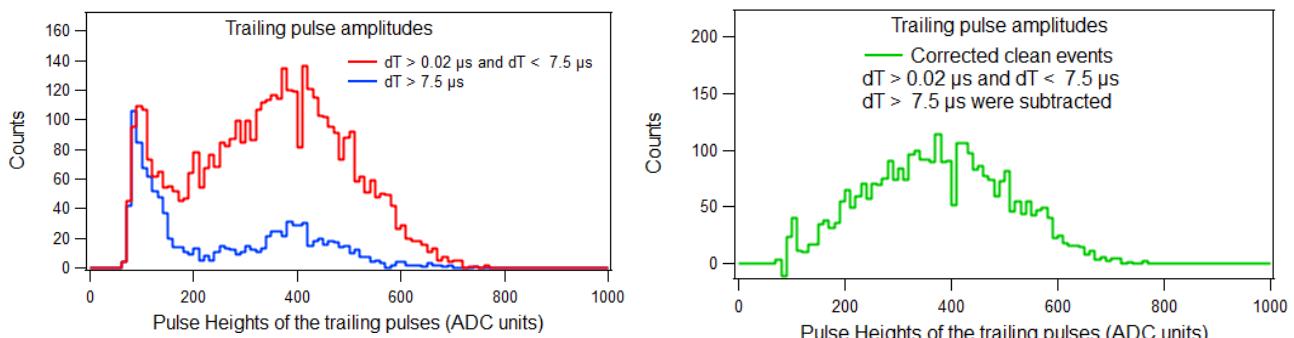


Figure 10. Left: The trailing pulse amplitude distributions for 1st and 2nd halves of the full dT range. Both categories of events were initiated by cosmic rays. The leading “radioactive pulse” events were eliminated before creating these histograms. Right: Subtracting the 2nd half from the 1st half removed the group A component, as evidenced by a disappearance of the residual low energy peak.

In the left panel of Figure 10 we can still discern the peak around PH=400 in the blue histogram. Most likely these are the residual muon decays which are still present at $dT \geq 7.5 \mu$ s. Subtracting the 2nd half from the 1st half histograms has removed all the group A pulses. In principle we could have split

the full 15 μ s interval some other way (like 10 μ s and 5 μ s, and scale the 2nd histogram by two), but we preferred to use the half / half split for simplicity. In the future we will setup wider dT time windows.

Did We Find The Electrons?

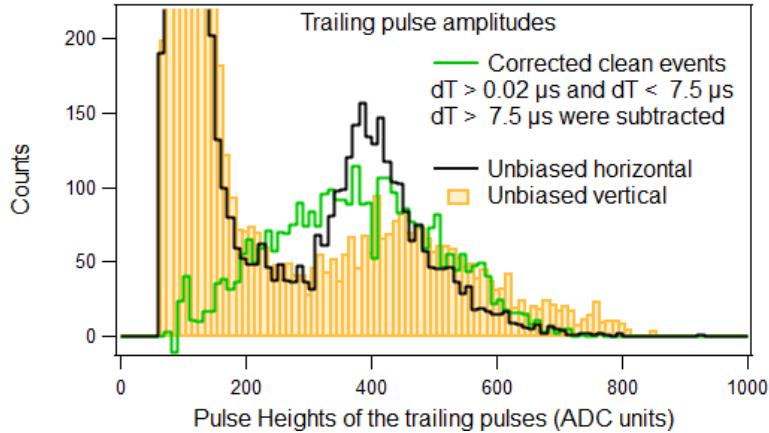


Figure 11. **Green**: cleaned trailing PH histogram with random coincidences removed in the previous step. **Black**: the unbiased PH histogram with the detector in horizontal position. **Orange**: the unbiased PH histogram with the detector in vertical position.

In the previous section we followed a classic method of subtracting the random background. We now need to look critically at the final result (green histogram in Figure 10) and ask “did we indeed find the electron emission from the muon decay”? In Figure 11 we compare the green “*electron histogram*” with the “*unbiased histograms*” in order to discern the difference. The procedure to construct the latter is explained in the next section. But first let us note the differences and the similarities.

1. The green histogram lacks any indication of the low energy group A, which are the dominant components in the black and orange histograms.
2. The high energy peak around PH = 400 is much narrower in the black unbiased histogram measured with the detector laying horizontally flat on the table.
3. The high energy component is shifted to higher energies in the orange unbiased histogram measured with the detector positioned vertically on its side.

So how can we understand this figure? Both the muon and the decay electron were relativistic, what implies they were both minimum ionizing. I am putting forward the hypothesis, that both the muon (black) and the electron (green) deposited about the same amount of energy in the 2"-thick plastic, which was too thin to stop either particle. The energy deposit shifted to higher energies (orange) when the detector was standing vertically on its side. The remaining difference was that the electrons were emitted isotropically, while the muons were mostly coming from around the zenith, resulting in a narrower distribution of their paths in the black histogram.

These observations may qualitatively explain the differences between the histograms.

Inclusive (Unbiased) Pulses

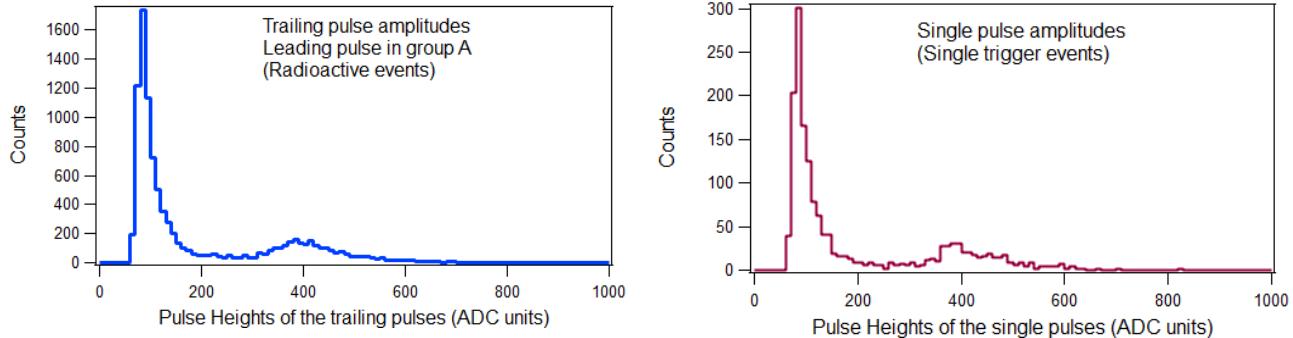


Figure 12. Left: The “leading pulse is radioactive” provided an unbiased sample of inclusive pulses. Right: The events with only a single pulse in the waveform, selected online by the firmware. Apart from the number of events, both distributions look identical.

In the previous section we compared an unbiased histogram of the muon pulse amplitudes with the supposed pulses produced by the electrons. In this section we will explain the unbiased histogram. The unbiased event samples were delivered in two different ways.

- 1) The explicit sample was selected online by the FPGA (events with only a single trigger within the coincidence window). These were the most common events because only one out of ~ 6000 muons stopped within the detector. These events were prescaled by 5000 before the recording, what has reduced their statistics.
- 2) Within the double-trigger events, the low amplitude “radio” *leading* pulses provided an unbiased collection of *trailing* pulses, equivalent to (1). We illustrated this in the left panel of Figure 8.

We compared these two pulse height distributions in Figure 12. The distribution (1) turned out more abundant. We juxtaposed it in Figure 11 against the supposed electron events.

Further insight was provided by comparing the horizontal and the vertical detector orientations.

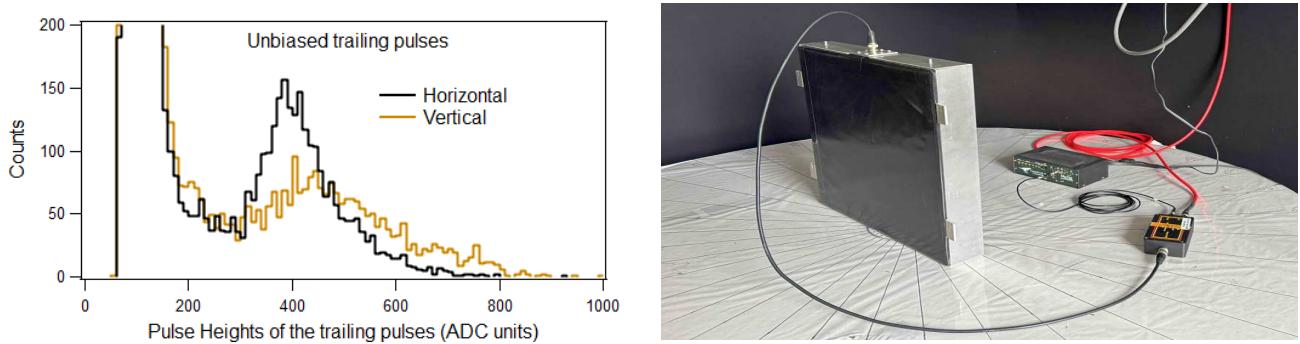


Figure 13. The unbiased distributions of the trailing pulses for two orientations of the detector: horizontal (black) and vertical (orange), shown at the right. The latter histogram was shifted towards higher values and became wider because there was more plastic for muons to traverse when the detector was vertical.

The Decay of the Trapped μ -Meson

We looked at the μ -meson decay fitted with an exponential function in Figure 8. The actual fitting proceeded as follows (copied from Igor output). We need to pay attention to the histogram binning into 100 ns wide bins in Figure 8.

• **CurveFit/M=2/W=0 exp, dT_hist_high_energy_P0[4,82]/D**

```
Curve fit with data subrange: dT_hist_high_energy_P0[4,82]
fit_dT_hist_high_energy_P0= W_coef[0]+W_coef[1]*exp(-W_coef[2]*x)
W_coef={10.882,181.99,0.046849}
Coefficient values ± one standard deviation
y0      = 10.882 ± 1.97
A       = 181.99 ± 3.87
invTau = 0.046849 ± 0.00231
```

This formula needs to be compared with the usual notation $N(t) = N_0 \cdot \exp(-\lambda t)$, where the decay constant λ is related to the half-life: $\lambda = \ln(2) / t_{\text{half}}$. Solving for t_{half} we obtain:

$t_{\text{half}} = \ln(2) / \lambda$, where $\lambda = W_{\text{coef}}[2] = 0.046849$ from the above fit.

$t_{\text{half}} = 0.693 / 0.046849 = 14.79$ in units $\Delta t = 100$ ns due to histogram binning.

The upper/lower limis are calculated with the standard deviation provided above, from $(0.046849 - 0.00231) = 0.044539$ to $(0.046849 + 0.00231) = 0.049159$, that is

$t_{\text{half}} = 0.693 / 0.049159 = 14.097$ to $0.693 / 0.044539 = 15.559$ (units of $\Delta t = 100$ ns).

So finally, we report the half life $t_{\text{half}} = 1.479 \mu\text{s} \pm 0.073 \mu\text{s}$. Mean life = $2.135 \mu\text{s}$.

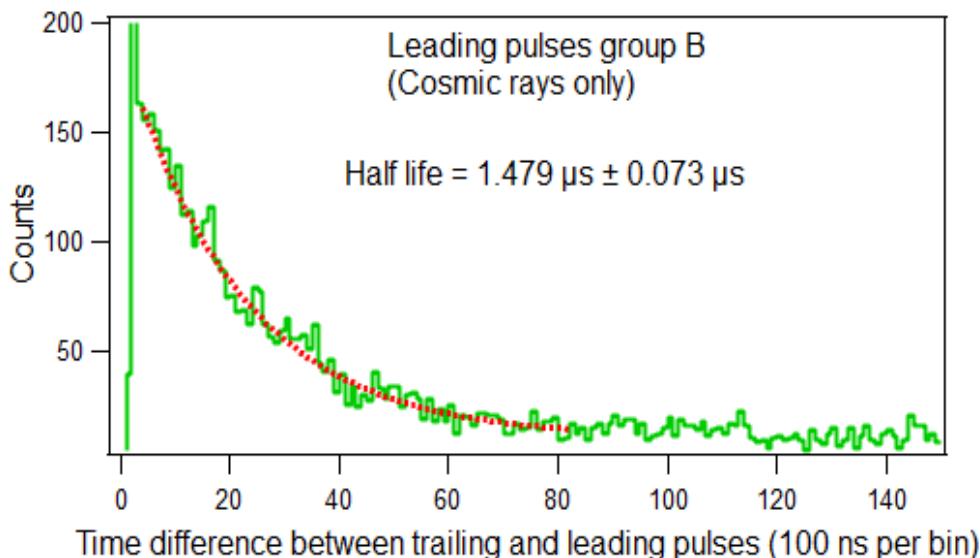


Figure 14. The time interval distribution dT between the leading and the trailing pulses, the same as Figure 8, but plotted in linear scale. We selected the leading pulses belonging to group B. The red line is the exponential fit.

Joyful Half Life of a Cosmic Ray μ -Meson (As Seen By AI)

We need to compare our result with the prior body of knowledge, maintained by Google.

Query to Google: *half life of mu mesons*

The AI Overview answered: The **half-life** of a mu-meson at rest is approximately **2.2 μ s**.

However, a little bit down that same page I found the following:

The **half-life** of a muon at rest is approximately **1.5 μ s**, though its **mean lifetime is 2.2 μ s**.

It appears that AI has some problems with choosing the right answer. Let me try another query.

Query to Google: *mu meson half life formula*

The AI Overview responded with the statement:

t_{half} is the **half-life** of the muon, which is approximately **2.2 μ s** in its rest frame.

So now we can choose which AI answer better fits our experimental result. The screen dump of AI knowledge is reproduced in Appendix B, where we can see both answers on the same screen.

Appendix A: References To Measuring Muon Capture and Decay

The following survey is suggesting that most advanced labs are teaching this subject using pretty ancient equipment from Phillips Scientific or LeCroy. Neither company is manufacturing the Nuclear Physics equipment anymore. Modernizing these labs would not be a bad idea. Students would not only gain physics insight, but they would also be exposed to modern digital data acquisition and digital signal processing.

Ref 1. Experiments in Modern Physics by Adrian Melissinos and James Napolitano (2003), p.399.

The detector consisted of a 35-cm high liquid scintillator tank and a phototube. The electronics was sketched on page 407. It consisted of a discriminator providing both the START and STOP signals to the Time to Amplitude Converter (TAC), converting the time interval to a voltage pulse, digitized with custom electronic card. Both the TAC and the Multichannel Analyzer (MCA) were integrated on this custom card. The emphasis was on collecting the muon decay histogram and fitting it with an exponent.

Ref 2. Carla Bonifazi et al, Muon Lifetime Laboratory: Measurement of the muon lifetime, 2018.

https://www.ictp-saifr.org/wp-content/uploads/2018/08/Lab_MuonLifetime_2018.pdf 8 pages.

The detector used three scintillator bars coupled to PMTs, and an aluminium plate to slow down the muons and increase the number of stopped muons. The signals were processed with a 4-channel QuarkNet data acquisition card. Pulse discrimination was performed with the on-card discriminators, while coincidence logic was processed with custom firmware running in a Complex Programmable Logic Device (CPLD), which can be viewed as a simplified FPGA.

Joyful Life of a Cosmic Ray μ -Meson

Crude pulse amplitude information was provided with a Time Over Threshold algorithm, but it was very crude. The board was operated over the serial RS232 interface using text commands.

Ref 3. UC Davis: Tony Tyson, Determination of the Muon Lifetime, 2021 (20 pages).

<https://122.physics.ucdavis.edu/course/cosmology/sites/default/files/files/Muon%20Lifetime/Muon%20Lifetime%202021.pdf>

The detector consisted of an unspecified scintillator coupled to a PMT, followed by NIM discriminators, NIM logic modules, and the TAC followed by an MCA. The emphasis was on students solving a bunch of problems on their own, collecting the muon decay histogram, and fitting it with an exponent. The manual provides comprehensive photographs of the detectors and the NIM electronics. The block diagram on page 10 specifies all the modules. The somewhat mysterious remark on that page (*Set the sampling rate to 1.25MS and the number of samples to 20*) describes that the rectangular TAC signal was sampled 20x to improve the resolution. It did not imply using a waveform digitizer in their setup.

Ref 4. Boston University, Michael Carrigan and Zach Collins, Measurement of the Muon Lifetime

https://physics.bu.edu/adlab_and_elab/wp-content/uploads/2019/05/Muon-Lifetime-Michael-Carrigan.pdf 2019, 3 pages.

The detector consisted a water-based Cherenkov detector with 16 PMTs (!!). The electronics consisted of NIM discriminators and NIM logic modules, a TAC, and an MCA. The resulting time exponent was distorted “due to imperfections in our stop trigger coincidence”, which were not systematically identified in the paper.

Ref 5. https://phylab.org/wp-content/uploads/2016/04/MounLifetimeNIMManual_v2.26.pdf

The detector consisted of three scintillation paddles operated in coincidence realized with NIM discriminators and NIM logic modules. The capture-to-decay time was measured with a TAC. The approach emphasized “do not touch the thresholds” which were preset ahead of the lab. The emphasis was on collecting the muon decay histogram and fitting it with an exponent.

Ref 6. University of Wisconsin, Madison. MEASUREMENT OF THE MUON LIFETIME (4/2/09)

<https://www.physics.wisc.edu/courses/home/spring2020/407/experiments/muon/muon.pdf> 7 pages.

The detector consisted of two plastic paddles, a tank with 10 gallons of liquid scintillator, and PMTs. The electronics consisted of NIM modules: LeCroy 621 quad discriminator, and LeCroy 365 4-fold logic unit, and Ortec 566 TAC. The Pulse Height Analyzer (PHA) was not specified, except for the information that it provided 256 channels and 8 volt full scale (page 5). The student’s experimental activities seemed pretty limited. The procedure consisted of recording the time distribution and fitting the exponent using software named SigmaPlot.

Ref. 7. Laboratory study of the cosmic-ray muon lifetime, T.Ward, M.Baker, J.Breeden, et. al. 1985

https://www.sfu.ca/phys/431/references/muons/ward_muon_AmJPhys1985.pdf

Detectors: Scintillator, water, and lead glass. Electronics: Analog NIM modules and TAC.

Ref. 8. University of Florida, Cosmic Ray Muons and the Muon Lifetime, March/14/2025.

https://www.phys.ufl.edu/courses/phy4803L/group_I/muon/muon.pdf 19 pages.

Joyful Life of a Cosmic Ray μ -Meson

The apparatus consisted of four plastic scintillators. The experiment addressed both the angular distribution from zenith and the muon lifetime. Most of the writeup concerns the angular distribution. The lifetime part starts on page 16. Two pairs of paddles were used as two independent detectors, consisting of the upper “capture detector” and the lower “veto”. The analog processing chain was barely described, other than mentioning “*Phillips model 730 five-channel discriminator*” on page 8. It implies that the analog chain consisted of NIM modules. The DAQ device was “*National Instruments USB- 6341 multifunction data acquisition module — together with a LabVIEW Muon program*” (top of page 9). This DAQ was in fact a set of counters operated at 100 MHz (basically, Time-to-Digital converters counting a 10 ns clock).

The experimental procedure is described in great detail. It consists of recording the number of coincidences for various detector configurations. The amplitude information was not mentioned in the writeup, because the electronics was not capable of such measurements.

Ref 9. MIT Department of Physics: 8.13-14 | Fall 2016 | Undergraduate

<https://ocw.mit.edu/courses/8-13-14-experimental-physics-i-ii-junior-lab-fall-2016-spring-2017/pages/experiments/the-speed-and-mean-life-of-cosmic-ray-muons/>

https://ocw.mit.edu/courses/8-13-14-experimental-physics-i-ii-junior-lab-fall-2016-spring-2017/resources/mit8_13-14f16-s17exp14/

Similar approach to the ones mentioned above. Also uses NIM modules and a TAC.

Ref 10. University of Washington, Muon lifetime measurement

Physics 433, Fall 2025 Advanced Laboratory: Nuclear and Particle Physics Experiments
https://courses.washington.edu/phys433/muon_life/muon_lifetime.php

Experiment instructions 2015: Muon Lifetime Experiment, H. Lubatti and D. Pengra, Oct/2015.
https://courses.washington.edu/phys433/muon_life/Muon_Lifetime_F15.pdf 10 pages.

The apparatus consisted of three plastic paddles and an absorber between the paddles #2 and #3. The PMTs were coupled to the paddles. A digital oscilloscope was used to examine the signals. This provided the opportunity to adjust the threshold, using a ^{60}Co source. Other than the scope, the “discriminators” were mentioned on page 6. The LeCroy 365AL logic unit and a TAC were mentioned on page 7, without providing a connection diagram. It means that the electronics was composed of NIM modules. The unspecified “pulse height analyzer” is mentioned on page 8. The analysis was mentioned on pages 9 and 10, using LabVIEW Norland Interface, Mathematica, Maple or KaleidaGraph.

Ref 11. The University of Edinburgh, Muon Lifetime Measurement

<https://www2.ph.ed.ac.uk/~muheim/teaching/projects/muon-lifetime.pdf> 7 pages

Equipment: scintillation detectors with PMTs, plastic and $\text{NaI}(\text{Ti})$. This is the 1st time I saw the $\text{NaI}(\text{Ti})$ in this role. The description is somewhat bleak: *you will use electronics to make a coincidence selecting the later pulse of the emitted electron/positron. This signal allows to apply a gate to the ADC card which resides in a PC where the data taking is steered. You will need to calibrate the energy scale of the ADC card using radioactive sources. A more extensive guide is to be written, but you can also find descriptions of several very similar apparatus on the web.*

Ref 12. Syracuse University of Washington, Muon lifetime measurement, 12 pages, 2012.

Joyful Life of a Cosmic Ray μ -Meson

<https://hep.syr.edu/wp-content/uploads/2021/04/MUON-LIFETIME-SU.QuarkNet.2012.pdf>

SU QuarkNet Workshop 2012 — Lab Activity 1, MEASUREMENT OF THE MUON LIFETIME

The detector consisted of a liquid scintillator tank and two PMTs. Electronics consisted of one multi-channel discriminator (LeCroy 321B), one dual 3 fold logic unit (LeCroy 162); one level adapter (to convert from NIM logic level to TTL level), one scaler (ORTEC 770) to monitor the counting rate independent of the computer, one interface board used to measure time (Multi-Function I/O Card NI 6023E), and LabVIEW software to interface with the card and record timing signals on the disk. So it was again a classic NIM apparatus, but without using a TAC+MCA. Their functions were provided by the LabView board serving as a Time To Digital Converter (TDC).

Ref. 13. Simple technique for determining the mean lifetime of the cosmic ray μ meson

A. Owens and A. E. Macgregor, Citation: Am. J. Phys. 46, 859 (1978); doi: 10.1119/1.11407

Detector: 15x15 cm plastic cylinder + 2" phototube. Electronics: NIM modules and TAC.

Ref. 14. Chunyang Ding, Cosmic Ray Muons TeXtbook, University of Washington Interlake High School 2015. https://phys.washington.edu/sites/phys/files/documents/wiki/ding_muon_textbook.pdf

A comprehensive 81 page survey of muon physics and experiments. There is very little in this thesis concerning actual measurements. This text is amazingly broad given the fact that it was written by a high school student.

Appendix B. Artificial Intelligence Opinions About Muon Half-Life.

Google mu meson half life formula

AI Mode All Images Videos Shopping Short videos News More Tools

AI Overview

The formula for a muon's half-life is based on exponential decay, given by the formula $N(t) = N_0 \left(\frac{1}{2}\right)^{t/t_{half}}$. Here, $N(t)$ is the number of muons remaining at time t , N_0 is the initial number of muons, and t_{half} is the half-life of the muon, which is approximately $2.2\mu s$ in its rest frame. This formula can also be expressed using the decay constant, λ , as $N(t) = N_0 e^{-\lambda t}$, where $\lambda = \frac{\ln(2)}{t_{half}}$.

Half-life formula

- **Exponential decay form:** $N(t) = N_0 \left(\frac{1}{2}\right)^{t/t_{half}}$
 - $N(t)$: The number of muons remaining after time t .
 - N_0 : The initial number of muons.
 - t : The elapsed time.
 - t_{half} : The half-life of the muon (approximately $2.2\mu s$ at rest).
- **Using decay constant:** $N(t) = N_0 e^{-\lambda t}$
 - The relationship between the decay constant and half-life is $\lambda = \frac{\ln(2)}{t_{half}}$.

Key values for muons

- **Mean lifetime (τ):** Approximately $2.2\mu s$.
- **Half-life ($t_{1/2}$):** Approximately $1.56\mu s$ to $1.523\mu s$ (calculated as $\frac{\ln(2)}{2.2\mu s}$). It is important to use the correct value for the specific calculation. For most general purposes, the mean lifetime of $2.2\mu s$ is used in the exponential decay formula, and the half-life is the time it takes for half the muons to decay.

Figure B1. A collection of AI opinions concerning the muon half-life. In order to cover all bets, AI is proposing both the 2.2 and 1.56 microseconds for this quantity. Not a bad strategy for an intelligent beast which consumed several billion dollars thus far and it is supposed to rule the world.