PHYS 4430 Advanced Lab
Muon Lifetime and Investigation of Cosmic Rays Lab

Version 1.1

Fall 2017

1 Introduction

In this experiment you will detect charged particles with a plastic scintillator and photomultiplier tube (PMT). There are a variety of different things you can study. The writeup will suggest a setup to measure the muon lifetime.

IMPORTANT WARNING : IN THIS EXPERIMENT YOU WILL BE WORKING WITH PMTS WHICH WILL BE DESTROYED IF YOU TURN ON THE HIGH VOLTAGE WHILE TOO MUCH LIGHT IS GOING INTO THE TUBE, OR TURN THE VOLTAGE TOO HIGH, OR APPLY THE WRONG POLARITY VOLTAGE. SEE INSTRUCTIONS TO AVOID THIS, SEE BELOW FOR SAFE TURN ON PROCEDURE.

2 Muons and Cosmic Rays

Muons are fundamental particles, like electrons. They interact through the same forces as electrons and have the same charge (and are also spin 1/2), but are approximately 200 times heavier ($m_\mu c^2$= 105 MeV). Due to their heavier mass, they are unstable, and decay with a mean lifetime of approximately 2.197 μsec[3], to an electron and a pair of invisible neutrinos.

Primary cosmic rays are mostly very high energy protons in space, often from outside the solar system. When they reach the upper levels of the Earth's atmosphere they interact, producing showers of short-lived secondary particles, especially pions, which rapidly decay to other particles, especially muons (a mixture of both positively and negatively charged muons). Many of the relativistic muons live long enough reach us on the surface of the Earth. Muons can travel through many meters of material before stopping. Once they are at rest they eventually decay to electrons. (Negative muons can also sometimes capture on protons in nuclei.)

The rate of muons passing through a horizontal detector at sea level is approximately 1 muon per cm^2 per minute. (In Boulder it is about twice this due to the higher elevation.) The muons at the surface of the Earth have a mean energy of approximately 4 GeV. The kinetic energy spectrum is fairly flat below a few GeV, so for our purposes we can assume that about half of muons are distributed uniformly in energy between 0 and 4 GeV.

Muons lose energy and slow down as they travel through material due to collisions with atomic electrons and elastic scattering from nuclei. The rate of energy loss per unit distance

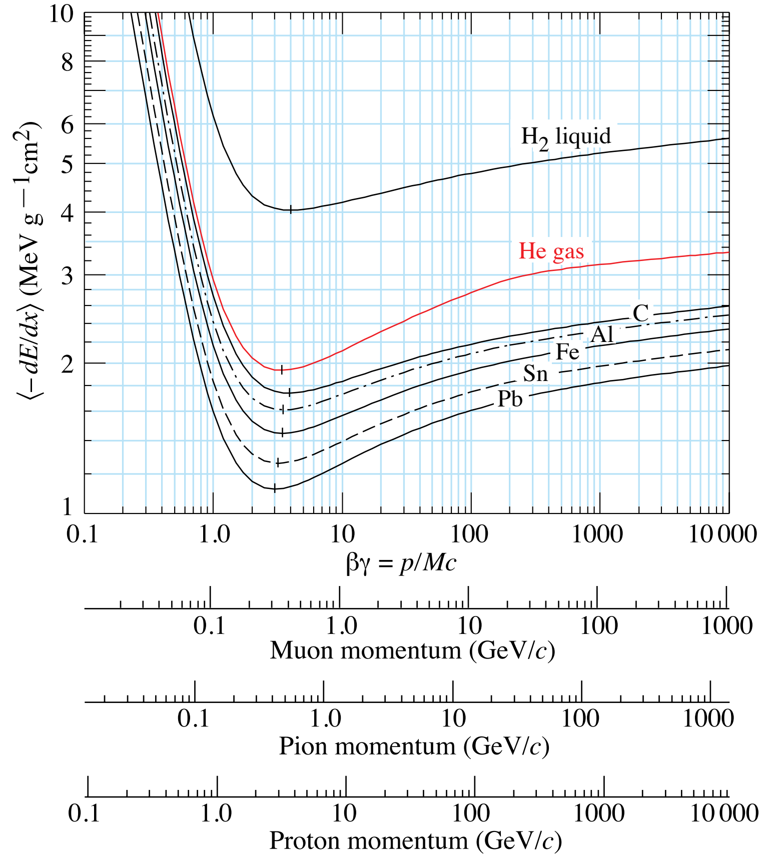

Figure 1: Energy loss per unit distance travelled for various charged particle types. From [3].

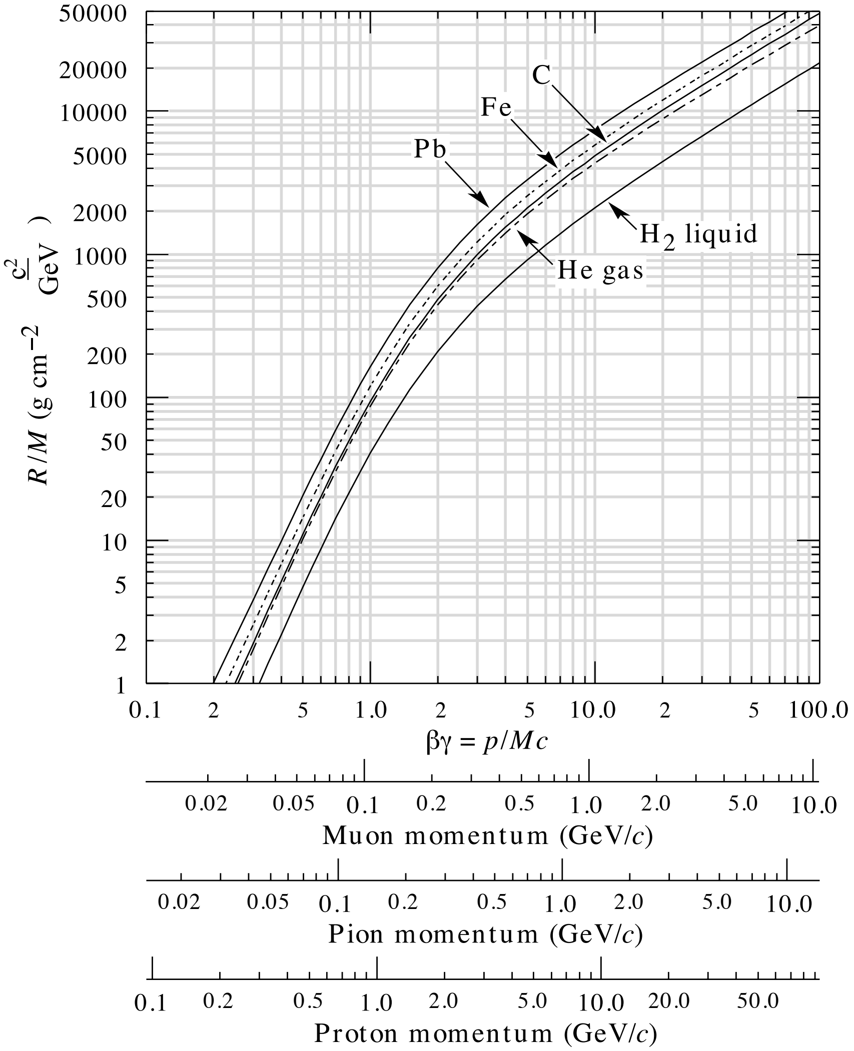

Figure 2: Range/M (where M =particle mass in GeV/c^2) for charged particles traveling through material. From [3].

for various charged particles is shown in Figure 1, while Figure 2 shows this integrated up to indicate the range, or how far a charged particle will travel on average in a given material before coming to rest.

A brief aside about units.... Due to the tiny masses involved, units of electron volts are typically used for energy (or really MeV or GeV), instead of Joules. You should recall from relativity that the energy of a relativistic particle is given by

$$E^2 = p^2c^2 + m^2c^4, \tag{1}$$

where E is the total energy, p is the momentum, and m is the rest mass of the relativistic particle. Because of this relationship, particle physicists will often use the even more bizarre units of GeV/c^2 for mass, and GeV/c for momentum. In these units, all of quantities already have the correct units to be plugged right into in Equation 1. No need to multiply anything by factors of c since it is already taken care of by the units! [1]

Question 1:

a. Given that the average cosmic ray muon has a kinetic energy of 4 GeV, how much energy do you expect an average muon to lose while traveling through 10 cm of plastic scintillator (which we can treat as a hunk of carbon with a density of 1.05 g/cm^3)?

b. Based on this, would you expect most of the muons that enter the detector will stop inside it?

c. According to Figure 2, below what momentum (in GeV/c) do muons to stop in 10 cm of scintillator? What kinetic energy does this correspond to? Hints: What R/M values does 10 cm correspond to on the y-axis? How can you convert a range in cm into g/cm^2? Recall that muons have an mc^2 of 0.105 GeV.

d. Given what you know about the energy distribution of these muons on the Earth, roughly what is the rate at which you would expect muons to stop and decay in the detector? If you want to collect roughly 5,000 decays, how long should you take data for?

A muon will decay to an electron with the same charge sign as the muon, and an unobserved neutral neutrino and anti-neutrino. These electrons typically have a kinetic energy between 30 to 50 MeV. Electrons lose energy much more rapidly than muons. Since they are much lighter than muons, in additional to collisional energy loses, they also radiate electromagnetic energy due to their deflection by the electric fields of nuclei (Bremsstrahlung). So the electrons resulting from muon decays in the scintillator are likely to deposit most of their kinetic energy inside the scintillator.

3 Equipment: Scintillator and Phototubes

Both scintillators and photomultiplier tubes are discussed in detail in the reference text [1]. Here we shall just give a brief introduction. As charged particles travel through material,

---

[1] Sometimes people can get lazy and drop the c's in the units, and will talk refer to a mass in GeV, etc. But really the c's are implied in the units.

they can excite the electrons in it. In scintillator, visible light (usually) is given off as the electrons de-excite. The amount of light produced is proportional to the amount of energy deposited in the scintillator material. As discussed in [1](2009 edition on reserve in the library, earlier editions on the lab shelves), some particles will pass through the scintillator depositing a fraction of their energy in the process, while others can deposit all their energy and stop. Scintillators can be solids, gases, or liquids (quinine is a natural scintillator so tonic water will glow blue under UV light). The plastic scintillator used in this experiment has the advantages that it is relatively inexpensive, so that large pieces can be used, and the response is very fast - the conversion process takes place in a fraction of a nanosecond. This makes it useful for experiments involving timing. Its primary drawback is that it has fairly poor energy resolution (the amount of light produced by a given amount of energy deposited will vary). In contrast, if you do the gamma ray spectroscopy experiment, you will use a NaI scintillator crystal which has much better energy resolution, but is quite slow and much smaller in size.

When a high energy particle deposits energy in the scintillator it produces a pulse of light, part of which hits a photomultiplier, causing the PMT to produce a current pulse. This experiment will be concerned with using these current pulses to learn as much as you can about the original high energy particle events. See the references for a more extensive discussion of PMTs. A photomultiplier is an instrument for detecting visible light (a good one can easily detect single photons), and works in the following manner. Photons from the pulse of light strike a photocathode inside the PMT. Due to the photoelectric effect, if the photon energy is high enough, a photon will kick and electron out of the photocathode. The PMT uses a high voltage power supply to create an electric field outside the photocathode, so that the emitted electrons feel a field, which accelerates them into the first dynode. Because they have gained a hundred volts or so of energy, each electron which hits this dynode knocks off two or three more electrons. These electrons are then accelerated into the second dynode, and so on down the dynode chain. A typical PMT might have 10 or a dozen dynodes. Each dynode must be biased at a higher voltage than the previous one. In this manner, a single photon will produce an electron but this electron will be tremendously multiplied (i.e. amplified). (There is a PMT that has been broken open in the lab, and you can see the interior structure.)

The amount of amplification depends on the type of PMT and the applied voltage, but is something like 3 for each dynode, so for an 11 dynode chain is around 3^{11} , around 10^6 or so, resulting in a substantial pulse of current at the output. As you might guess from the previous explanation, the amount of amplification rises very rapidly with the applied voltage, and the largest gain (where it is best to work) is quite close to the voltage which will spontaneously rip electrons out of the photocathode or the dynodes, which is the reason for the following warning. **Applying too large a voltage or the wrong polarity of voltage to a PMT will generally destroy it**. The polarity (i.e. positive or negative voltage) and max voltage should be given on the tube. If you have any doubt, check with the instructor before applying voltage. The PMT produces current pulses but we normally observe signal voltages, so what you observe depends on the termination you put on the PMT output. If it goes directly into a scope, there is a 1 Megohm termination from the input impedance of the scope. This determines the current to voltage conversion. The 1 Megohm, in combination with the capacitance in the PMT base and the capacitance of the cable determines the length of the pulse. For fast timing work, 50 ohm termination is always used, in which case the PMT pulses are only a few nanoseconds long.

Another issue with PMTs is that they rapidly deteriorate if you allow them to produce too much current. The common way people burn out PMTs is by turing on the high voltage while shining a lot of light into them. You should never turn on the high voltage while the tube is exposed to room lights, even indirectly. In general it is a good idea to always expose a PMT to as little light as possible, even when there is no bias voltage applied to it, since large amounts of light tend to make the tube noisier later on. All of these issues are traced to the fact that the photocathode and dynode surfaces are made of materials (like Cs coated GaAs) that are optimized for electron emission, but which are fairly chemically reactive. They are all in sealed glass vacuum tubes, so handle them with care. Often one accidentally leaves tiny light leaks which could burn out the tube. To avoid this, turn the high voltage on using the following procedure:

SAFE PMT TURN ON PROCEDURE, READ THIS CAREFULLY: The PMT has an SHV high voltage input and one or more BNC outputs. Sometimes there are outputs for both the Anode and the first Dynode. If this is the case, you want to connect to the output signal from the Anode. Send the output from the PMT base directly into the oscilloscope, with no termination resistance. (Make sure the output is not also connected into the preamp, which acts as a 50 ohm terminator. Check that the scope input is set on DC.) Make sure that the scope channel is set to have **1 MΩ** resistance first. The 1 megohm input resistance of the scope will convert the PMT current to voltage, with a conversion of 0.1 V per 10^{-7} amps out of the tube.

Be sure that you are using the **CORRECT POLARITY** for the high voltage. PMTs often require a negative high voltage. Check with the instructor if you are unsure. Put the scope on the most sensitive scale and watch the output voltage as you slowly turn up the high voltage bias. If there are no light leaks, the output current should stay in the nanoamp range, giving a few mV or less on the scope. If you have a light leak the voltage you see will start to rise as the bias voltage is increased because the amplification increases. Also, if the light present is mostly coming from the fluorescent lights you will start to see 120 Hz modulation on the voltage. If you see any such modulation it tells you immediately that you have a light leak. In any case, you should never turn up the high voltage so that the current from the PMT exceeds 10^{-7} amps (0.1 V DC on scope). If you see the scope voltage start to rise, stop turning up the high voltage. Locate and eliminate the light leaks.

Question 2: What is the rate of pulses that you see? Sketch the shape. What is the duration and height of the pulses?

4 Experiment

4.1 Overview

The electronics that you need are enclosed in a NIM electronics crate, which also supplies ± 6V, ± 12V, and ± 24V into the lower back of each module. Make sure each module is firmly seated in the crate. The power switch for the crate is on the lower left. Before doing anything, use a DMM to probe the test points on the front left side to check that all of the

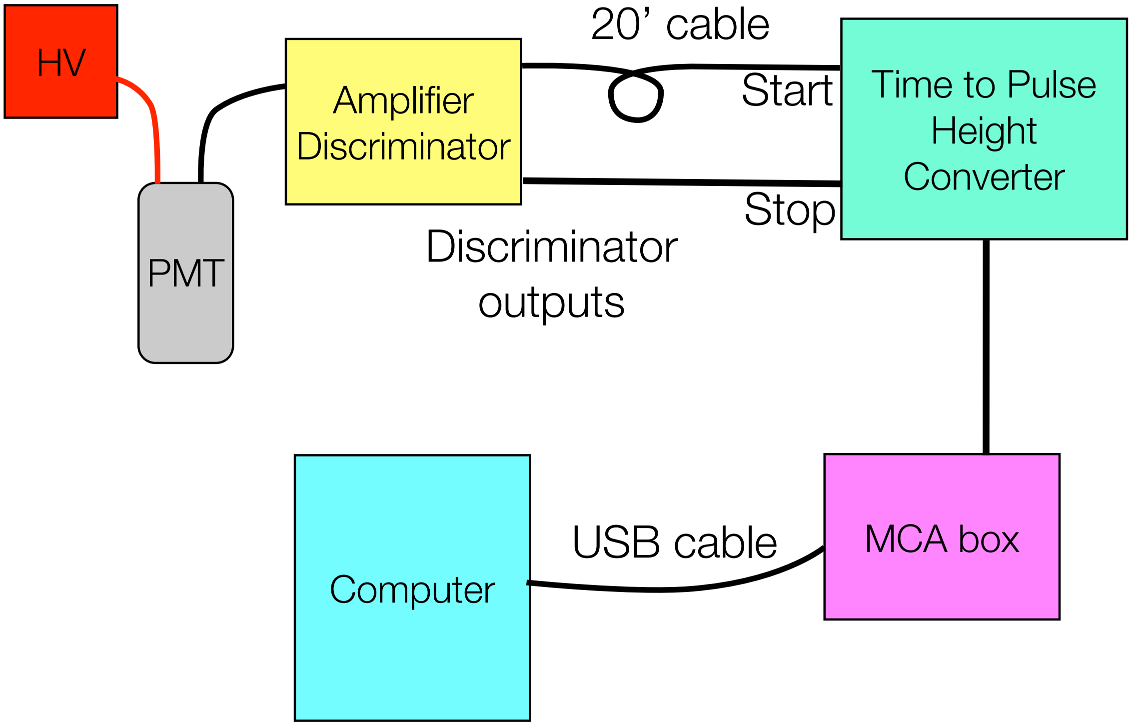

Figure 3: Basic setup for the lab.

voltages are being properly supplied. If they are not, speak to your instructor.

NIM electronics modules in particle physics typically use one of two different digital logic conventions:

TTL This uses voltage to indicate true and false. False is a low voltage signal (from 0 to +0.8 V), and true is a higher voltage signal (in the range of +2.2 to +5 V).

NIM This is a *current-based* logic signal. False is 0 A, and true is -16 mA, which for a 50 ohm termination will be -0.8 V. (Yes, it is negative.)

You can't mix these logic signals. Most of the units that you will use in this lab use NIM logic. If you are unsure, look up the manual for that unit and check the input and output specifications. There are links to many of the module manuals on the "Hints" section of the website.

4.2 Discriminator Amplifier

Connect the PMT output to the Ortec 9302 Amplifier discriminator input (it is on the back of the unit). You likely only need to lower gain setting. There are several outputs on the back. In addition to the amplified pulse, there are also two discriminator logic signal outputs. They use NIM logic. (See above.) You can adjust the discriminator level with the screw adjust on the front. The discriminator output pulses are of very short duration.

Question 3: Using PMT pulses as input, look at the input signal and discriminator output and determine the threshold of the discriminator.

You should adjust the discriminator level so that the rate of discriminator pulses above threshold is ~50 Hz or less. There is a visual scaler module that might make it easier to count discriminator output pulses to check this. (The scaler module can be set to accept positive or negative input signals.)

4.3 Time to Amplitude Converter (TAC)

We are using the Ortec 437 Time to Pulse Height converter. This module takes a start pulse and a stop pulse and converts the time difference between them into a voltage and outputs an up-down pulse of that amplitude. The longer the time difference, the larger the signal amplitude. The time range in question is controlled by knobs and a multiplier switch. If no stop pulse is received in the full time range then no pulse is output and it resets itself to await the next start signal. We are looking for the time between a muon pulse and a delayed pulse from a decay electron. So we want the discriminator pulses to be the start and the stop, but we want to delay the start pulse (using a long cable) so that the start is from a muon, and the stop signal is from the subsequent electron pulse.

Question 4: Using the two discriminator output signals, determine the time shift produced by the long cable. Sketch the output. Is the shift long enough so that that the"start" comes after the comes after the "stop"? What is the delay? Since we are delaying the start, this is a small correction that will need to be applied to the decay times.

Question 5: Use the discriminator outputs (from PMT pulses or small negative pulses from a function generator) as the start and stop, and sketch the output from the TAC. Verify that as the time between the start and stop pulses increases, the amplitude increases. Given the muon lifetime, what settings should you use for the full timing range?

4.4 Multi-Channel Analyzer

The output of the TAC should be connected to a multi-channel analyzer. This analyzes a stream of voltage pulses, and sorts them into bins of pulse height. (In our case that pulse height represents the time between the muon stopping and decaying to an electron.) ORTEC / AMETEKs easy MCA (Multi-Channel Analyzer) connects to a computer in the lab, which uses the software MAESTRO-32 MCA Emulator.

Connect the 1 ohm output of the TAC to the input. Use the Maestro software and verify that you can collect data. You might need to go to Display→Detector to select the detector if no window is visible. Be sure to do an Acquire→Clear before you do an Acquire→Start to collect a new dataset. To save your data, go to File→Export and save it to a ASCII SPE file. This gives a list with the number of counts recorded in each channel. There are thousands of channels.

4.5 Timing Calibration

The MCA records the number of peaks in each bin of pulse height. So you need to convert the bin number to a time between start/stop on the TAC. To do this, you should use the delay generator. The Ortec 416A can take a negative or positive input. For a NIM logic signal use negative. Take two coincident pulses (first verify on the scope that they really are coincident!) and then compare the un-delayed one to the delayed marker output of the gate generator. Note that there are three different delay ranges multipliers possible (0.1, 1, or 10 μsec, controlled by a 3-way switch) and then the delay can be tuned with the dial on the front. (The dial goes up to 11!) Verify with a scope that you understand how to read the delay settings and that you see the expected delay.

Question 6: What is the smallest possible delay that you can obtain? Collect data with the MCA for a variety of different delay settings. Fit this to a line to obtain a conversion between the recorded bin and the time delay. Does the response look linear? Recall that the muon lifetime is roughly 2 microseconds, so your calibration points should span a few lifetimes.

5 Data collection and Analysis

You are now ready to collect muon data. Since the rate of stopping muons is low, you'll likely need at least 24 hours of data. Export your data into a file and analyze it in Mathematica. Convert the bin number to a time using the calibration that you found earlier. If there are any empty bins at the start, cut those points out. (There is some threshold below which pulses are not recorded.) Since there are thousands of bins, it will likely also be helpful to group together 10 bins to have more points per time bin. In Mathematica, for example, the `Total[mylist[[i;;i+9]]]` command will sum together elements i to i+9 of a list. You might find it useful to insert this inside of a `Table[function,{i,1,max,10}]` command to sum the bins.

Question 7: Each bin represents a count of a number of events. Thinking back to PHYS 2150, what is the approximate statistical uncertainty on the number of counts recorded, assuming that the counts obey a Poisson distribution? (You might want to consult Ch 11 Section 2 of Taylor's book [4] to remind yourself about some important properties of the Poisson Disitrubtion.)

Fit your data and determine the muon lifetime in scintillator.

Question 8: What is the uncertainty on your lifetime? How does your observed value compare to the accepted value in vacuum?

Question 9:

a. As mentioned above, in addition to decaying, negative muons can also capture on protons in nuclei, resulting in a neutron and muon neutrino (which is unobserved). Why can't this occur for positive muons?

b. What effect, if any, would this capture process on the observed muon lifetime? Hint: Do the lifetimes or the rates add when there are two competing processes?

c. Assuming that the muons in the lab are 45% negatively charged muons and 55% positively charged, by comparing your observed value to the vacuum value, can you estimate the rate of μ^- capture on nuclei in the plastic?

6 Suggestions for Further Study

One possible improvement is to use the constant fraction discriminator. The discriminator that we have been using fires once the PMT pulse gets above a given value. But as you many have noticed earlier when looked at the PMT pulses on the scope, larger amplitude pulses are wider in time than the smaller ones. Thus a large amplitude pulse will cross the discriminator threshold earlier than a small one, and thus will open the TAC timing gate a bit earlier. The constant fraction discriminator rather than firing at a given voltage, fires at a given fraction of the pulse height. So using this discriminator on the amplified signal output will reduce this jitter and potentially will give a better lifetime measurement.

This lab has assume that the signals are coming from muons. But perhaps you'd like to know more about these muons or see some evidence that they aren't just noise in the scintillator. There maybe a smaller piece of plastic scintillator (a paddle) which also has a PMT attached, and can be easily moved around if you are careful. You can learn a great deal by comparing the signals from the two scintillators. The simplest experiment is to just look at the signals from the two detectors on two channels of an oscilloscope. If you trigger on only one channel, how often do you see a signal from the other? (There are NIM logic modules that have and AND gate feature.)Does the rate of coincidences change significantly if you place lead between the two pieces of scintillator? Does it depend on where you put the second detector, or how it is oriented? Are the muons coming in from all directions, or only from above? Can you measure the angular distribution?

References

[1] Experiments in Modern Physics, A. Melissinos and J. Napolitano, Academic Press, 2nd ed., 2003.

[2] “A Simplified Muon Lifetime Experiment for the Instructional Laboratory”, R.E. Hall, D. A. Lind, R. A. Ristinen, American Journal of Physics, **38** 1196 (1970).

[3] K.A. Olive et al. (Particle Data Group), Chin. Phys. C, **38**, 090001 (2014).

[4] An Introduction to Error Analysis, J. R. Taylor, 2nd Edition, (1996).