Cosmic Ray Counting

Although we (almost) never think about it, we are bombarded by high-energy ionizing radiation every second of our lives. These high energy particles originate in outer space in the "solar wind". Very high in the atmosphere, the dominant particle species are protons or alpha particles, of extremely high energy (10^8 to 10^{20} eV). When these "primary" cosmic ray particles meet an atom in the upper atmosphere, the collision produces a "shower" of energetic hadrons, which then decay into energetic leptons. The most common of these "secondary" cosmic ray reactions are

$$\begin{aligned} \pi^+ &\rightarrow \mu^+ + \nu_\mu \\ \pi^- &\rightarrow \mu^- + \bar{\nu}_\mu \ , \end{aligned}$$

that is, the decay of pions into muons and neutrinos/antineutrinos.

A muon itself is unstable, and will decay (at rest) into an electron or positron (and a couple of neutrinos) in a little over 2 microseconds. However, because of relativistic time dilation, the muons in high-speed flight from the upper atmosphere survive for much longer (according to our reference frame). Indeed, the lengthening of the muon lifetime in flight is a vivid example of relativistic kinematics in action. The discovery of the cosmic ray muon was one of the first indications that there could be particles other than the ones that make up ordinary matter: protons, neutrons, and electrons.

In this experiment, you will attempt to measure the flux of cosmic rays at about sea level (since Seattle is nearly there), and compare your value to a widely accepted one known since the 1940's. Cosmic rays come out of the sky in all directions, which means there are a couple of ways to talk about the flux. If we restrict the angle to straight down (zenith = 0), then the accepted intensity of high energy (hard) cosmic ray flux is

$$I_v = 0.82 \times 10^{-2} \text{cm}^{-2}\text{s}^{-1}\text{str}^{-1},$$

where the subscript v indicates that this is the vertically directed flux. This says that the number of cosmic ray particles passing through a vertically-facing surface is 0.82×10^{-2} particles per square centimeter per second per steradian. The steradian is a unit of solid angle in the way that a radian is a measure of planar angle. The solid angle intercepted by a sphere is 4π steradians, by the upper half of a hemisphere 2π steradians, and so forth. The steradian measure is necessary because the cosmic rays can come from all angles, not just straight down out of the sky.

If we integrate the intensity in the downward direction (i.e., the vertical component of the velocity) across a horizontal surface over the solid angle from horizon to horizon in all directions, we get the downward flux:

$$J = \int_{\text{hemisphere}} I(\theta) \cos\theta d\Omega = 1.27 \times 10^{-2}\, \text{cm}^{-2}\text{s}^{-1},$$

where the $\cos\theta$ term accounts for the fact that the flux across the surface from particles traveling at an angle to the surface normal will not be as large as the flux from particles traveling straight down.

If $I(\theta)$ were a constant, then the integral would give π times the intensity I_v, but it is notably less. This is because the flux intensity varies according to zenith angle. An approximate empirical relationship is

$$I(\theta) \approx I_v \cos^2\theta \ .$$

The angular variation in the flux largely comes from the fact that muons arriving from angles toward the horizon must pass through a greater thickness of atmosphere than those that come from directly overhead.

The muon telescope

Although cosmic rays come in a variety of particle types with a variety of energies, research has shown that most (about 80%) of the detectable high energy particles (the so called "hard" cosmic rays) to reach the surface of the earth are muons, both positive and negative. Thus, for the remainder of this write-up, I will use "muons" as a shorthand for "cosmic rays", since that is what we will mostly look at.

To measure the muon flux, we will use a "telescope," which in this context means an array of scintillator "paddle" detectors arranged with one scintillator overlapping another, as shown in Fig. 1. A muon traversing the array will, with some probability, cause a simultaneous pulse to occur in more than one detector, since with its high energy, it will usually not be stopped by a single paddle. Thus, we can use coincidence gating to count muon fly-bys as a way to distinguish these events from background radiation and photomultiplier noise.

Because the expected muon count rate is rather small, it is helpful to use a coincidence level of three or more in order to really suppress accidental coincidences. One of the aspects of this experiment you will explore is the rate of "accidentals" versus true coincidences. In addition, by comparing the rate of coincidences using 2 versus 3 versus 4 paddles, you will be able to estimate the counting efficiency each paddle, which you can then use to estimate a correction to your overall data set.

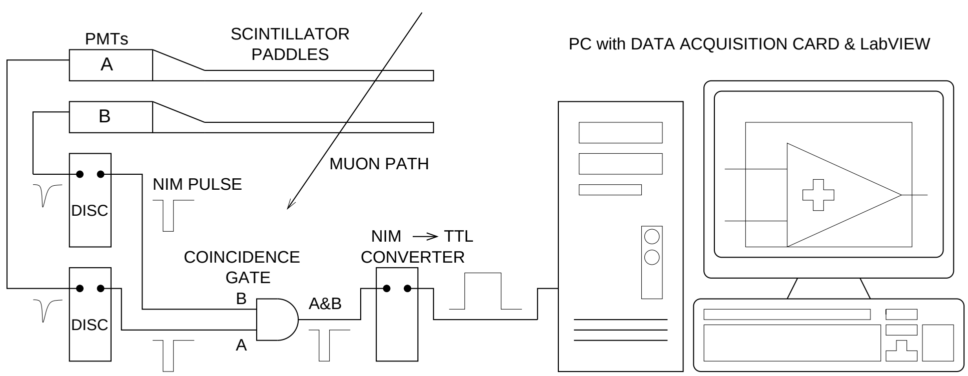

Figure 1: Schematic of the muon telescope apparatus. Although there are only two scintillator paddles shown in the figure, the actual setup uses between three and four paddles. The AND gate is a 4-fold NIM logic unit. Data are collected with the LabVIEW Interval Counter software.

Some details about the scintillator paddles: The scintillator material is Bicron BC-408 plastic scintillator cut into a sheet that is 1 cm thick, with an area of 12 inches by 6 inches. One end of the sheet is glued to a Lucite light pipe, which is then attached to the face of a 2 inch diameter PMT with optical grease. The PMTs are Electron Tubes model 9266KB. The paddle assembly is wrapped in aluminum foil, and then wrapped with heavy black paper and tape. A picture of an unwrapped paddle is shown in Fig. 2

Each paddle assembly rests on a shelf in a large wooden box. The shelves can be moved so that the vertical extent of the paddle array can be varied. The detectors are arranged so that the PMTs

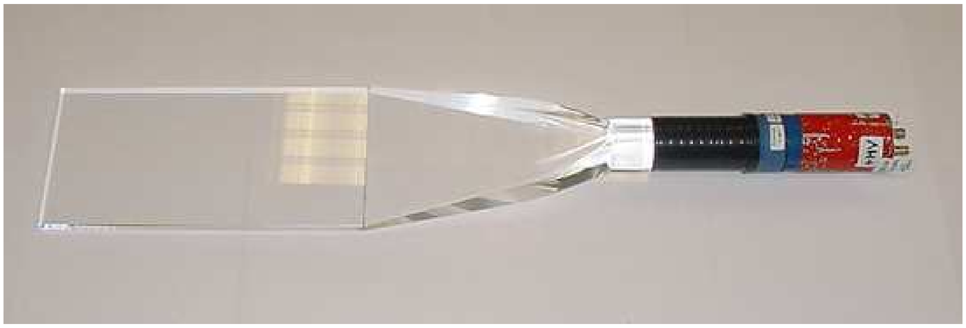

Figure 2: Scintillator paddle detector used in the muon telescope apparatus, before being wrapped with aluminum foil and black paper..

alternate which side of the paddle they are on—this minimizes the effect of a muon causing a coincidence by passing through multiple PMTs, rather than multiple scintillator paddles!

Setting up the telescope

The paddle array box will already be set up at your station with the paddles in place. You will need to connect the cables and set up the electronics.

First, note the tags attached to each PMT base indicating a voltage. This voltage has been chosen by the lab technician to produce a pulse of sufficient height when a ^{22}Na source is brought nearby. The variations in voltage from one paddle to the next are due to differences in the PMT manufacturing. (It is hard to make PMTs all exactly alike.)

If not already plugged in, obtain SHV cables from the cable rack. Since you have two or three high voltage supplies at your station, you must match, as well as you can, pairs of detectors to be plugged into each high voltage supply. For example, if your labels say 780V, 930V, 950V, and 960V, you might plug the lower value 780- and 930-volt units into one supply, and the higher value 950- and 960-volt units into the other.

Obtain 50Ω BNC cables from the cable rack. ***Important: the cables must all be the same length.*** (Remember, you are setting up coincidence counting.) Attach each cable appropriately, and double check to make sure the right HV cable goes to the right HV supply.

By using the Tektronix MSO2024 digital oscilloscope, you can make use of the scope's memory capabilities to use the cosmic rays themselves to set the threshold/HV combination and and check the relative timing of each detector.

Start by plugging in the signal cables from each detector to each channel of the 4 input scope. Make sure each input is terminated with 50Ω. Turn on the scope and let it boot up. Set the inputs to "1x" input scaling, and then set the channel sensitivity to 50–100 mV/Div (not critical, but make all channels the same). Set the trigger level on the first channel a few 10s of mV below zero and use a negative slope.

Slowly bring up the HV supply on the triggered-input paddle, following the usual protocol, while watching the scope. You are looking to make sure there are no light leaks or other problems. If you think you do see a problem, ask the instructor or TA for help.

If you are using a supply with two paddles connected to it, check the signal on both paddles, and then if all looks good, set the high voltage to the geometric mean of the two settings recommended for your PMT pair. (The "geometric mean" is the square root of the two values multiplied together.) You will probably see a lot of background and/or noise pulses, along with some muon pulses too.

Repeat the above process so that you can see signals from all four paddles on the scope display. It is helpful to order the traces on the display from top to bottom in the same sequence as the paddle arrangement itself. Once you have all paddles working, it is time to fine tune the HV settings and set the thresholds.

The first step is to study the pulses made by muons versus background radiation pulses. With the trigger set to the top paddle (A), adjust the trigger level from just below zero towards a lower value. You should see the size of the displayed pulses increase for that channel. (This happens because the trigger level will block pulses that are smaller, so the display appears to favor the larger pulses. But it is important to realize that this does not mean the signal itself is growing in amplitude.)

As you make the trigger level lower (selecting larger-amplitude negative pulses), you will begin to notice many *simultaneous events*—pulses that show up in more than one channel. These are the signals from muons as they go through multiple paddles.

Keep lowering the trigger level until what you see are mostly simultaneous pulses in all four channels. When you see such an event, STOP data collection on the scope and study the pulse heights in all channels.

Now your goal is to set the high voltage on each supply so that the pulse heights for the collection are well within the threshold ranges of the discriminators which typically have minimum of 30 mV and a maximum of 1.0 V. Start with the pulse that appears to be the smallest. If it is below 100 mV raise the supply voltage for that detector by 50 volts of so. Then let the scope run and repeat the above process for the remaining pulses. You may find it helpful to lower the voltage for one set of detectors, as well, in order to bring all with in a better agreement. In any case, do not exceed 1200 volts on any one supply. It is not important that the pulses be identical in amplitude, but you should be able to make the detectors respond more uniformly than before.

After setting the HV values, use the scope trigger level to estimate the discriminator threshold to use for each detector as follows. As before, start with the top paddle (A) and set the scope to trigger in this channel. Adjust the trigger level so that about half of the time pulses that meet the trigger are also simultaneous in the other channels. Then repeat this process by triggering on the other channels and looking for simultaneous events. (You may notice a useful feature of the MSO2024 scope: it remembers the trigger level last used for each channel.)

Now, feed the signal from each PMT into a discriminator channel. Turn the discriminator thresholds down and set the width for each channel to 40 ns. After setting the width, set the discriminator threshold on each channel to the values you recorded above minus a few (5–10) millivolts. Use the test point on the discriminator to dial in these numbers.

After setting the widths and the thresholds, rout the discriminator outputs into a channel in the 4-fold logic unit. Use a NIM counter/timer unit to first measure "singles" rates for each paddle: set the coincidence to 1-fold, and use the pins to turn off all but one input at a time. Look for any paddle that seems to be either much higher or much lower in count rates than the others. If you have such a paddle, try adjusting the discriminator to make it match the others reasonably well.

Then measure 2-fold coincidence rates between the top two paddles (A and B), the middle two paddles (B and C) and the bottom two paddles (C and D), over a reasonable counting period (e.g., $\sim$ 30 seconds, or a time that gives about 100 counts or more). You should see roughly the same count rates in each pair (with the appropriate pins set in the logic unit). If one pair seems excessively low, check the singles rates for each paddle in the pair and adjust the discriminator level and/or voltage to the PMT for that paddle. Ask for help if you don't think it is working properly.

Characterizing the muon telescope

If all is well, you should be able to collect data in order to address two questions:

- What number of muons per unit time pass through the scintillators of the telescope? In other words, what is the *acceptance* of the telescope?
- What percentage of these muons does the telescope actually count? In other words, what is the *efficiency* of the telescope?

You will address these questions below.

Acceptance of the telescope

The acceptance of the telescope depends on the surface area of the scintillator paddles and the solid angle subtended by surfaces of the paddle array. Then one must include the angular dependence of the muon flux to estimate the total flux through the telescope. Clearly a larger paddle area will see more muons, and it is not hard to believe that the counting rate would be directly proportional to this area, all else being equal.

Start by making some measurements of the apparatus. Use a ruler or tape measure to find the separation of the scintillator paddles. Also measure the area of the scintillator region of one of the paddles (they are all the same), and check their alignment in the box by eye. You want to insure that they overlap as much as possible. Record your measurements with a sketch in your notebook.

To calculate the solid angle seen by the array is a bit complicated. For a muon to make a coincidence between the top and bottom paddle, it must pass through both. Certainly a muon coming straight down will do this. But a muon coming in at an angle will also do this, if the angle relative to the vertical is is not too big. The biggest angle would be along path from one edge of the top detector to the opposite edge of the bottom detector. If the top and bottom paddle are far apart, this biggest angle would be smaller, and so the the solid angle will be smaller than if they are close together. However, calculating the solid angle seen by the whole array quantity is tricky for a couple of reasons. First, not every element of the detector sees the same solid angle. A small element of area dA in the center of the bottom paddle would register a count if the muon went through it and *any* spot on the top paddle, covering the whole surface A of that paddle. If you imagine a point on the bottom paddle, and rays intercepting this point and all points on the top paddle, you will get the solid angle for that portion of the array. But another element of the bottom paddle dA' located, say, near the edge, would see a different solid angle by the same construction. Moreover, the directions of rays passing through the top paddle and dA versus those passing through the top paddle and dA', with respect to straight up (zenith angle $\theta = 0$) would also, in general, be different. It is known that the muon flux varies with zenith angle proportional to (about) $\cos^2\theta$. Thus, a good calculation of the expected flux of the telescope requires integrating the flux as a function of θ over the solid angle $\Omega(x, y)$ at each point (x, y) on the surface of the paddle, and then integrating this function over that surface A. At best, this is tedious; at worst, it requires a computer.

So, we will estimate the flux by making a couple of simplifying assumptions:

1. Assume that each small element of the bottom of the paddle array sees the *same solid angle*, which is equal to the solid angle seen by the element at the center of the bottom of the paddle array.
2. Assume that the rectangular (6"×12") paddles can be modeled as circular paddles having the same area.

The purpose of this second assumption is that it is easy to calculate what the flux per unit area would be for this cylindrical geometry. The solid angle would be that subtended by a cone. If the cone has zenith angle θ between the axis and the rim, then it is easy to show that

$$\int_{\text{cone}} d\Omega = \int_0^{2\pi} \int_0^{\theta} \sin\theta' d\theta' d\phi = 2\pi(1-\cos\theta)\ . \tag{1}$$

Likewise, to find the estimated flux, we need to first integrate $I(\theta) \approx I_v \cos^2\theta$ over the solid angle. Hence,

$$J(\theta) \approx 2\pi \int_0^{\theta} \left(I_v \cos^2\theta'\right) \cos\theta' \sin\theta' d\theta' = \left(\frac{\pi}{2}\right) I_v (1-\cos^4\theta)\ . \tag{2}$$

The estimated acceptance from these assumptions would be found by multiplying $J(\theta)$ by the area A of the bottom paddle.

Exercise 1 *(A) Verify the calculation that gives Eq. (2). (B) Verify that $J(\theta = \pi/2)$ gives the expected result of 0.0127 particles per cm^2 per second. (C) Calculate the expected flux through a single paddle from this result. This will give you an expected "singles" rate for one paddle (and an upper bound on the expected count rate from your array).*

From your paddle dimensions, calculate the effective angle θ seen by cone whose area is the same as the area of one paddle and whose height is the same as the separation between the top and bottom paddle. Use this value, and the results above to calculate an estimate of the expected counting rate of hard muons.

Comment (in your notebook) on whether you think this estimate would be higher or lower than a more careful calculation. (Assume 100% efficiency of your detector paddles).

Efficiency of the paddle detectors

The goal here is to compare 2-fold coincidences with 3-fold and 4-fold coincidences in order to estimate the paddle efficiency for counting muons.

Here is a definition of *efficiency*: The efficiency of a detector is the *detected* number of events divided by the *actual* number of events. In other words, if a detector counts 50 particles during a time when there were 100 particle hitting it that it *could* count, then the efficiency is 50% or 0.5.

Another way to say this is that the efficiency is the *probability* of registering a count, given that there is a valid count to register. If we have two detectors arranged to count in a coincident manner, a particle may be detected by both detectors, by the first but not the second, by the second but not the first, or by neither. Only in the first case will a count be recorded, since coincidence is required. This implies that the probability of registering a coincident count in both paddle A and B, $P(A\&B)$ is the probability of registering a count in A times the probability of registering a count in B, or $P(A)P(B)$.

You will want to use a two-channel (or more) NIM counter and two sections of a logic unit. This will allow a direct comparison of one paddle arrangement with another. One channel will be set up to count, say, 2-fold coincidences from the "outside" paddles A and D and the other will be set up to count, say, 3- or 4-fold coincidences from the outside paddles plus one or more paddles from in-between. This will give a direct comparison of the effect of the middle paddle(s) on the count rate, since any muon traveling from A to D must also pass through B and C.

To make sure you have enough counts, estimate the counts needed to determine the efficiency to better than 5% uncertainty based on Poisson statistics. (How many counts is this?)

Use the NIM counter/timer set up as described to perform the following tests:

1. 2-fold coincidence with (A and D) versus 3-fold coincidence with (A and B and D)
2. 2-fold coincidence with (A and D) versus 3-fold coincidence with (A and C and D)
3. 3-fold coincidence with (A and B and D) versus 4-fold coincidence with (A and B and C and D)
4. 3-fold coincidence with (A and C and D) versus 4-fold coincidence with (A and B and C and D)

From your results of the above measurements, calculate the efficiency (and the uncertainty on that efficiency) for the middle paddles, B and C. Use the different combinations from the above measurements to check consistency in your determinations.

Then ask for help in order to swap the middle (B and C) paddles with the outside (A and D) paddles of the array, and repeat the tests to find the efficiencies of these paddles. The paddles are somewhat fragile. You will need to turn of the HV outputs and disconnect all cables. Then you may slide the shelves out of the box and rearrange the order. ***Do NOT pull the paddles themselves off the shelves.***

After you have completed the second round of tests, reassemble the paddle array as originally set.

In your notebook describe the method and reasoning you used to find the efficiency of your paddles.

Accidental coincidences

the efficiency analysis assumes that only real muons get detected by our array. But, as you will probably already have noticed, each paddle individually records a large number of counts per unit time. Only a small fraction of these are high-energy muons; most can be ascribed to PMT noise and low-level background sources. But with the low count rate for muons in this setup, the question of accidental coincidences becomes important.

As derived in the tutorial on counting statistics, the rate of accidental coincidences can be easily estimated, given the assumption that the coincidences are truly random. For two detectors, the rate of 2-fold coincidences is

$$R_2 = 2T r_A r_B \ , \tag{3}$$

where T is the common pulse width, and r_A and r_B are the rates of detectors A and B, respectively. For 3-fold coincidences the rate is

$$R_3 = 3T^2 r_A r_B r_C \ . \tag{4}$$

To test these relationships, we need to force any real coincidences out of the experiment. We can do this by adding in sufficient delay to the discriminator signals so that array is no longer in good synchronization.

Choose two paddles in your array (say B and C). Between the output of the discriminator and the input to the logic unit, add sufficient cable delay to move the one paddle out of time with the other (for 40 ns pulse widths, 100 ns is probably sufficient.)

Use the pins on the logic unit to set up counting of the singles rates for each of the three detectors, and then count doubles (2-fold coincidence) rates (R_2) for the paddle pair used.

Calculate the the theoretical accidental rate predictions for 2- and 3-fold coincidences. Compare your calculation for the 2-fold result to what you measured. Since these are statistical processes, your numbers will not match exactly, but you should be able to define a reasonable uncertainty based on Poisson statistics.

From the above, is it worth checking the 3-fold rate? To try, you will need another length of cable to get the random triples rate R_3. You should see that the rate of random coincidences, especially for the triple, is much lower than when they are in time. This is a good way to show that you really are counting cosmic rays by the coincidence method!

Also, based on the formulas (3) and (4), what would you expect to be the rate R_4 for 4-fold accidentals? How long would you need to wait to get, say 10 counts?

The long count

Reset correct timing in your paddle array in order to make a good coincidence count with 3 paddles. (You may use 4, if you find that their efficiency is very good.)

You could use the simple NIM counter timer for this part, but it is tedious to collect statistics with it. Instead, use the Random Interval Counter program to do this work for you.

You will first need to use a NIM gate generator to convert the fast negative NIM pulses to "slow NIM" or TTL pulses. Connect an output of the logic unit to the START input on one of the gate generator channels, Then feed the TTL output to a scope to set the width to about 0.5 microseconds, and connect this output to the input of the computer interface box supplied for this purpose.

Start data collection with the Interval Counting program and let it run for a while. After about 5 minutes, stop data collection and look at the histograms for the following settings:

(a) Count distributions (green histogram) for interval times equal to 2 times, 10 times and 20 times the mean interval time between counts. (For example if the mean interval between counts is 0.4 seconds, 10 times this would be a 4 second interval, and 20 time would be an 8 second interval.)

(b) Interval distributions (red histogram) for count intervals of 1, 2, 5, 10, and 50 counts, using the "scaling" summing method on the interval distribution calculation.

After collecting these histograms, restart the data collection, ***but do not clear the previous data.*** You want to see how the histograms build up over time. Continue taking data for at least a half-hour if possible (or until the end of the lab period if not).

While the data are being collected, you can explore the settings on the program. What is the difference between "Scaling" and "Interleaving"? How can you explain the evolution of the histograms as you change the interval number? What is the effect of changing the bin number on the interval distribution? You can change the viewing parameters (interval time for green, count number for red, and binning for either distribution) while data are being taken.

If you click on the "Show features" button, you will reveal additional controls that will allow you to place a theoretical distribution over your histograms, based on different theories.

Ideally, you should have a final data set corresponding to 20 minutes or more of counting data.

When you are done taking data, and have time, you may also explore a feature that allows you to simulate data according to different interval distribution functions. It is interesting to study the interval distribution evolution for the "uniform" function.

In your notebook:

- Discuss the relationships among the various histograms that you see. In particular, note how the mean and variance change, as the data sets get larger, and as you vary how you look at the same data set. Also how do the means and variances of the two different types of distributions (interval vs. count) compare with each other?

- Calculate the flux measured by your telescope, and compare it to the ranges expected from your earlier calculation. Don't forget to calculate an uncertainty in the number. (It will probably be somewhat less.) From your efficiency measurements, determine the efficiency of the telescope as a whole. Use this to correct your data to account for the estimated "missed" muons. Also, consider how much error you get by ignoring the rate of accidental coincidences?

Extended Analysis

With a computer language, such as Python and its powerful numerical libraries, it is relatively straightforward to improve on the gross simplification we used to estimate the theoretical flux we should get through the telescope. You may work with your instructor to investigate this problem.

Prepared by D. B. Pengra

`counting_telescope_F21.tex -- Updated 23 November 2021`