PHYSICS 180F

Supplementary Material

Rene A. Ong & William E. Slater
UCLA, March 2017

Version 1.3

Introduction

This document provides supplementary material for the Physics 180F labs. Most of the material relates to analysis aspects – e.g. binning, background subtraction, fitting, etc. The material is somewhat anecdotal in nature and is not comprehensive, but it should be of help in doing the labs. This document assumes that students have read the Physics 180F Lab Manual.

The figures shown in this document come from a set of complicated ROOT macros which were difficult to support and, in any case, provided an approach that was too cookbook in nature. Thus, students should consider these figures as useful ways to approach the analysis, but should not necessarily try to duplicate everything shown here.

A. μ^+ Experiment

The point of the μ^+ experiment is to detect its spin precession in a magnetic field due to the torque, $\vec{\mu} \times \vec{B}$. The precession is rendered observable as a result of angular momentum conservation, neutrino helicity, and the shape of the density of final states factor, dN/dE, in the decay $\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$. The experimental data consist of TDC values for the time between detecting a muon that stops in a magnetized target, and the decay of the same muon. Analysis consists of plotting a histogram of these TDC values separately for "UP" and "DOWN" decay events, referring to the vertical hemisphere into which the positron enters from the target. The muon spin (momentum) is predominantly UP (DOWN) initially. Analysis consists of fitting the adjustable parameters of an appropriate theoretical curve to the difference of the UP and DOWN TDC plots.

You should not wait until the whole experiment is finished before you develop and run analysis software but rather work on it early on and run it periodically throughout the quarter in order to check data quality. This is the best way to monitor the progress of your experiment and to learn how to extract results. You can analyze data either on the muplus machine or on one of the cluster machines (p180f-a or p180fb). Since the muplus internal disk on which data are recorded is not seen by either p180f-a or p180f-b, if you use a cluster computer you have to ftp the data file from muplus to the cluster. ROOT is a somewhat intimidating edifice so you are provided with several example macros that can be used to start the analysis of your data. You will need to add a significant amount to these examples. The program statements are typical of C and C++. You can find a Wikipedia reference for just about any such statement. The favored learning path is to start with something that works. By convention, the file extension '.C' is used to distinguish a program (often called a macro) that runs under the ROOT interpreter called CINT.

You will want to produce a program that will plot the UP(U) and DOWN(D) distributions and fit them to a determine lifetime values. The program also needs to plot U minus D and fit that to the appropriate oscillating form (see below). At the end of this section, Figure 1a shows the runtime for nearly 1000 hours of data taken in 2012. Figure 1b(c) histograms the TDC values for U(D) events. Figure 1d shows the difference of the U and D data, scaled by exp(T/2197ns). In other words, in this plot the muon lifetime has been removed. It is a challenge to see the parity-violating signal (Figure 1d) and it takes at least several weeks of data to get a good distribution and fit. The raw

U-D distribution can be fit to the theoretical expression: $\exp(-\Gamma t)\,(A+B\cos(\omega t))$, where Γ = (1/2197ns) is the muon decay width. A way to see the precession more clearly is to multiply the difference distribution (both content and uncertainty) bin-by-bin by a factor of $\exp(+\Gamma t)$; this is what has been done in Figure 1d. The smooth fit in Figure 1d is of the form $(A + B\cos(\omega t))$.

Before going further, you should recognize that there are at least three important backgrounds in the μ^+ experiment. To remove them or to correct for them can only be done as part of the computer analysis. Data management is part of this exercise also. The backgrounds are:

(1) Fake **M** signals may be generated by non-stopping muons because the C4 veto counter is not 100% efficient or, more likely, an oblique muon may hit C1, C2, and C3 and exit through the gap between C3 and C4. If a second oblique muon enters through the same gap and hits C4 during the 20-microsecond G2 gate, an accidental event will be recorded. Their time distribution is random, i.e. flat. These events can be removed statistically from the desired μ^+ distribution by fitting the data to the form:

$$f(t) = A + B\exp(-\Gamma t)$$

Accidentals determine the A coefficient; the decreasing exponential is the muon decay signal.

(2) Counters are not sensitive to the sign of a particle's charge. Background from μ^- decay and capture exhibits a shorter effective "mean life" due to the capture channel, which is not possible for μ^+ . The μ^- capture rate depends strongly on the atomic number of the material in which the μ^- stops.

(3) A μ^+ will often stop in part of the apparatus other than the target. Other than a possible stray background field, there is no magnetic field outside the target. Therefore the spin of any μ^+ that stops outside the target will not precess, and such events can not affect either the frequency or the amplitude of the precession signal. But these non-target events do affect the *denominator* of the parity-violating asymmetry (U-D)/(U+D), i.e., what you are measuring in this experiment. So they must be counted and corrected for. This background is removed by running with the target swung out of position for long enough to measure it. There are a few complications to take into account. You cannot afford to run with target out to get as many incident muons as you do with target in place so you must correct for running time. The computer time recorded for each event, makes this easy. The support structure is made of Fe or Al and a muon can even stop in C3.

There is no way to know specifically where non-target events originate and they will contaminate both U and D decays. As in normal data, these target-out events are a mixture of μ^+ and μ^- . While this is irrelevant for the μ^+ which have the same mean life regardless of the stopping material, any μ^- background will tend to have a longer effective mean life for the target-out data than for Cu target-in data since its capture probability for Cu is greater.

After you take enough target-in events to demonstrate beyond doubt that you see the muon precession and have good enough statistical precision for the magnetic moment measurement, you will need to run the *muplus* program for two or three days with the target out. To subtract target-out data from U+D one must scale the target-out data to match the run-time of the target-in data. The computer time of the first event of the input data file establishes the lower bound of a histogram as the starting point of your experiment. From this you can figure out the total run time; for example, if you use bins of 3600 seconds then each bin will be one hour in duration.

Some distributions of data from a previous quarter are shown at the end of this section. Figure 2a(b) shows the first 200 TDC bins (4 microseconds total) of the target-in, magnet-on distribution for U(D) decays. Both decrease as the sum of two exponentials, one fixed at 2197 ns (μ^+) and the other at 160 ns (μ^- capture in Cu). The data points oscillate about the fitted curve but in opposite directions. A target-out run produced Figure 3(a) and Figure 3(b) shows an expanded view of the first 200 TDC bins (4 microseconds total). The fit is a single exponential for muon decays plus a constant for accidentals. This form fits the data quite well though the decay constant is not quite the same as the accepted μ^+ value. The small difference is presumably due to the background from μ^- captures in the relatively low-Z material of C3 and the aluminum frames where the capture rate is comparable to the decay rate. It is fortuitous that an exponential is an adequate representation for the target-out data. The motivation to fit this background plot to anything is to smooth out the statistical fluctuations. With the runtime correction, the exponential fit is then adequate to calculate the actual background to be subtracted from each bin of the target-in plot.

In the original version of this experiment done by Professor Harold Ticho in 1972, the subtraction was accomplished by assuming a pure a μ^+ time distribution for the target-out events. The rate of these background μ^+ events was determined from a scalar count of the triggers. The count was

corrected for accidentals by also counting chance coincidences of U and D positron candidates with μ^+ stops. You will find it instructive to read that old write-up for the details.

The difference signal, U-D, reveals the spin precession in better detail. Figure 4a is a U-D histogram scaled by exp(Γt) . After scaling, a simple μ^+ decay distribution would appear as a horizontal straight line. The plot is fitted to (A + B cos(ωt)) . In principle it is important to use the correct time-zero, synchronous with **M**, so that the phase of the fitted function is not artificially changed by cutting the data.

The frequency parameter ω allows a measurement of the muon magnetic moment since you can calculate the magnetic field independently. Refer to the Ticho write-up for an example calculation; the target winding and dimensions are the same as then. If you are concerned, you can count the turns when the target is out. The precession amplitude is given by the B parameter. The fit to the data is good. Figure 4a, by itself however, is not adequate to measure the polarization which is proportional to the up-down asymmetry: (U-D)/(U+D). The denominator must be corrected for background from both target-out events and μ^- capture events, as discussed above. Figure 4b is corrected by subtracting target-out events; this shows up in the fit because the A parameter is smaller by a factor nearly 6. Within errors, the other two parameters (B and ω) are indistinguishable when you compare Figure 4a and Figure 4b. The μ^- capture contamination in U+D is apparent in the first two bins in either Figure 4c or Figure 4d, which both show the U+D distribution. The effect of non-target contamination (nearly 40% of U+D) is evident in the plot of Figure 4d (corrected) compared to Figure 4c (not corrected). While you might feel that the fit in Figure 4c cannot be valid because all the points lie below the fit, bear in mind that the data include unseen points beyond TDC=200.

The total (U + D) decay rate for muons that stop in the Cu target is given by the single parameter p0 in the fit to Figure 4(d). In addition, the first bin was excluded from the fit to remove contamination from μ^- capture. The result for the asymmetry measurement is:

$$\alpha = \frac{B \text{ (from Fig 4b)}}{p0 \text{ (from Fig 4d)}} = 0.044 \pm 0.012$$

This fits does not agree very well with the original Clark & Hersil result (G.W. Clark and J. Hersil, Phys. Rev. **108**, 1538 (1957)). Figure 4(d) suggests a possible reason for the discrepancy, namely the poor fit to points above 40 bins that systematically lie below the fitted line. Clearly you will find an even lower result for α if p0 is not corrected for target-out contribution.

Historically, the relative amplitude of the parity violating signal was crucial to the development of the electro-weak V-A theory of Glashow, Weinberg and Salam. In a clean experiment it is referred to as a maximal parity violation. Unfortunately, the effect is substantially washed out in our 180F experiment because the initial π^+ does is not generally at rest when it decays and so can boost an up-going m+ in the inertial frame of the π^+ to down-going in the laboratory frame. In addition the μ^+ is somewhat depolarized by scattering during its journey through the atmosphere.

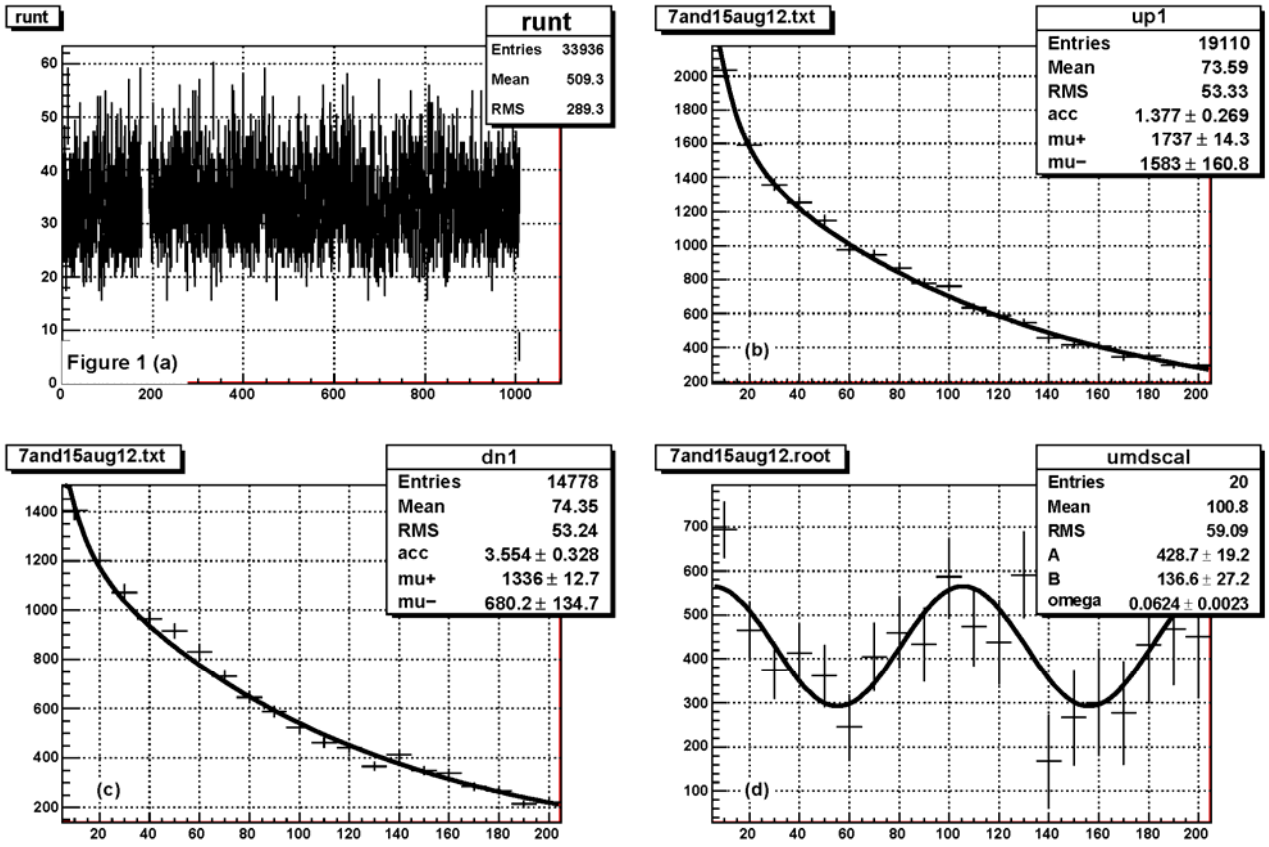

Figure 1: (a) Histogram of run times (hours) for the μ^- experiment; (b) TDC distribution for U events; (c) TDR distribution for D events; (d) U-D distribution, scaled by exp(T/2197ns).

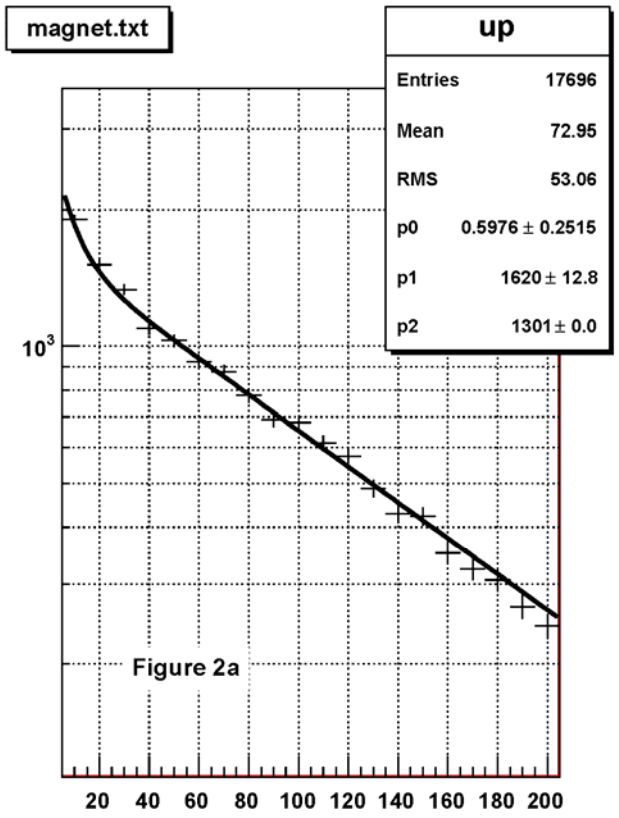

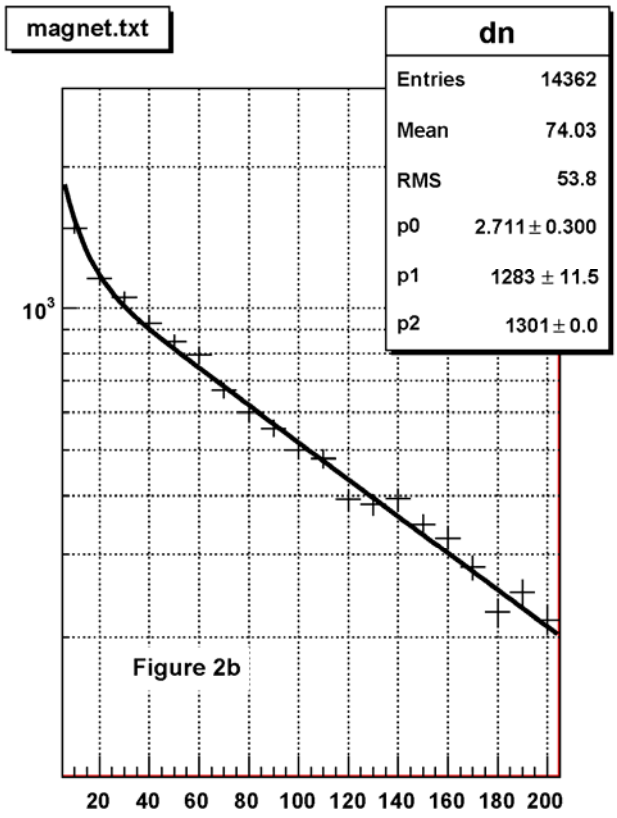

Figure 2: (a) TDC distribution (first 200 bins, 4 microseconds total) for the target-in, magnet-on configuration for U decays; (b) same distribution for D decays. The fits are the sum of two exponentials.

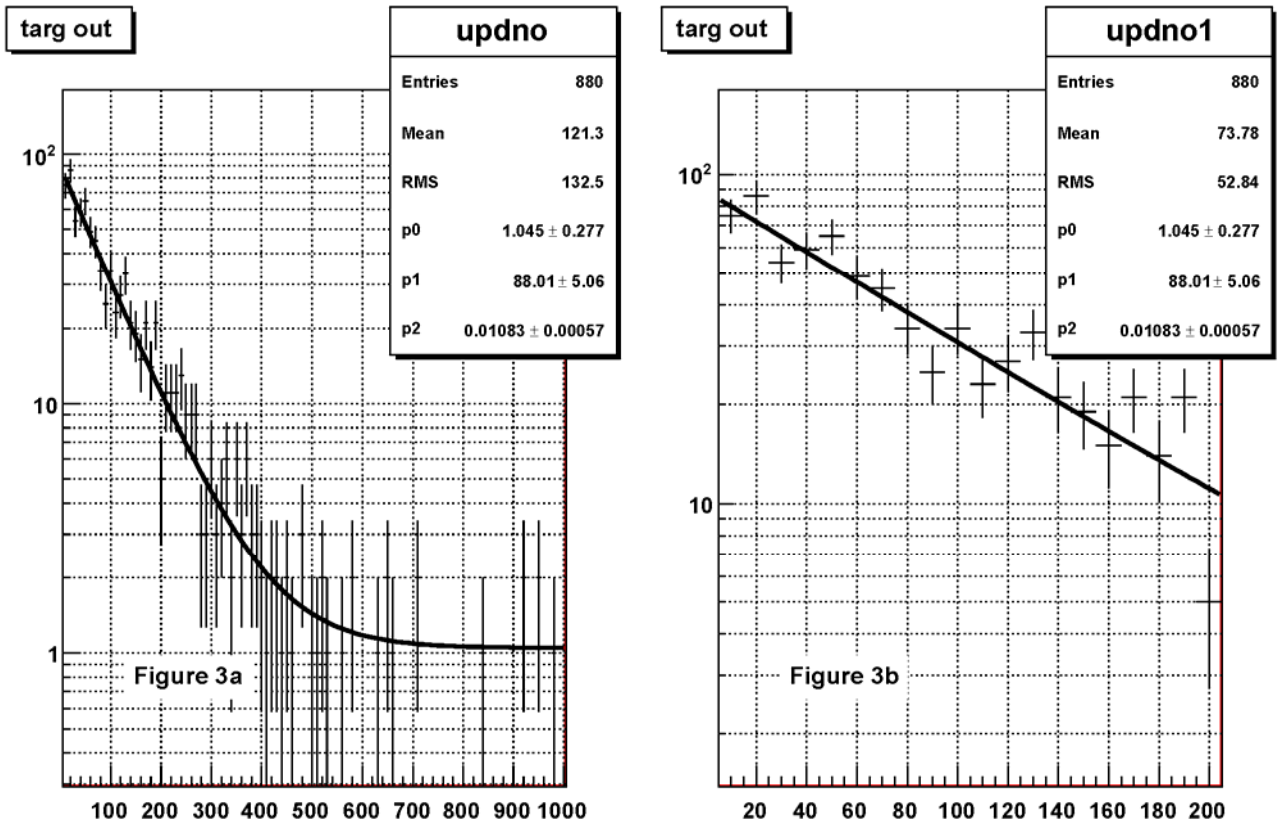

Figure 3: (a) TDC distribution for the target-out configuration; (b) same distribution, showing only the first 200 bins (4 microseconds). The fits are a single exponential plus a constant term.

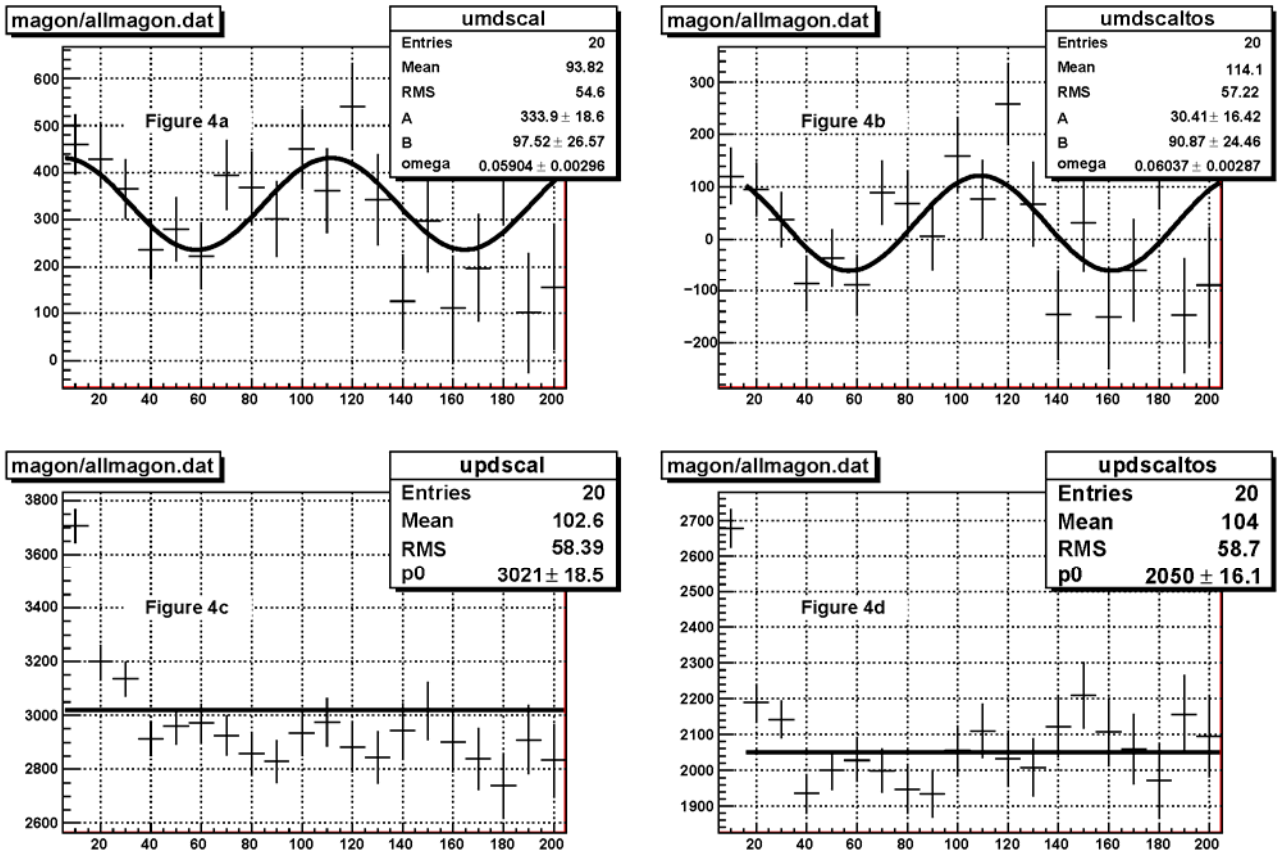

Figure 4: (a) U-D distribution, scaled by exp(T/2197ns); (b) same (a), except corrected by subtracting target-out events; (c) U+D distribution, scaled by exp(T/2197ns) with no correction for the non-target contamination; (d) same as (c), except corrected for non-target contamination.

B. μ^- Experiment

At the quark level, the capture process is: $\mu^- + u \rightarrow \nu_\mu + d$ while the competing decay process is $\mu^- \rightarrow \nu_\mu + e^- + \bar{\nu}_e$ which occurs at the same rate as $\mu^+ \rightarrow \bar{\nu}_\mu + e^+ + \nu_e$. Negative muons disappear more rapidly than positive muons because there are two channels instead of one. In addition, capture is exothermic because the muon mass is about 105.7 MeV while the muon neutrino is nearly massless. The light quark (*u* and *d*) masses are nearly equal, so almost the entire 105.7 MeV is released in the capture. In the real world, free quarks are never observed so what you have in this experiment is capture by nuclei; for example, if the target is the most abundant iron isotope Fe^{56}, then the most likely final state is Mn^{55} + n + ν_μ. The two atoms, Fe and Mn, are massive enough not to end up with much kinetic energy so the neutron and neutrino share nearly all the final state energy. To do the calculation more carefully (e.g., by accounting for binding energy and the Fermi density of final states) is not so relevant to this experiment; the main point is that the observable result is the low-energy neutron that you try to detect in a large block of scintillator.

Neutrons from muon capture in the target enter the neutron detector where they can be detected if they interact in any of three ways:

(1) quasi-elastic scattering on hydrogen (i.e. the proton) in the plastic,
(2) radiative capture on a proton to make a deuteron,
(3) absorption leading to low energy gamma or neutron emission.

Processes (1) and (2) both give prompt signals, i.e., they emit a charged secondary whose ionization can be detected within a fraction of a nanosecond of the capture. Proton recoils from reaction (1) yield at most 5 MeV, in the case a neutron from the capture process hits a nearly unbound proton. Process (2) is slightly exothermic and can give a signal even for close-to-zero kinetic energy incoming neutrons: n + p → d + γ. At most, the 2.3 MeV binding energy of the deuteron is converted to kinetic energy and shared by the outgoing deuteron and photon. The light produced in the neutron detector scintillator from any of these processes, however, is very much smaller than the background light from muons that directly hit the scintillator. These signals can deposit as much as 40 MeV (2 MeV/cm * 20 cm) due to dE/dX ionization. Alternatively, a muon (either positive or negative) may stop outside the target, then decay ($\mu \rightarrow e + \nu + \bar{\nu}$) and send into the scintillator counter an e^+ or e^- that in turn loses up to 50 MeV. The signals you want (from the

low energy neutron) are at least an order of magnitude smaller than these considerably more frequent backgrounds. This fact of life means that *the veto counter C4 must be highly efficient* and the detection efficiency of the neutron system must be carefully controlled. Setup of the counters C4, C5, C6, C7, and C8 therefore requires more care and thought than you can expect to achieve with rudimentary plateau runs.

Process (3) is subject to a further complication because it may start a "cascade" of neutron interactions. In an extreme case where the "counter" material is instead the right isotope of uranium instead of plastic scintillator, you get a self-sustaining chain reaction, which, if uncontrolled, makes a nuclear explosion! Many of the initial neutron interactions in the counter are not detected but their secondary neutrons still interact there. A sequence of connected events is started, only the last of which is detected. The distribution of these signals is a long tail clearly visible in the time distributions for all the targets. It is well described after about 10 microseconds simply by a linear fit, $a + bt$, or by a constant plus decreasing exponential $A+B\exp(-Ct)$. The exponential is what you would expect if there is effectively a constant probability (in time) for a neutron interaction. If you expand the exponential to first order, it looks like $a + bt$. You can thus describe the shape of the time dependence in the time region of interest, but you need to know how to extrapolate to t=0. In addition, accidentals due to muons that happen to enter the side of the neutron scintillator result in coincidences among PMT's 5, 6, 7, and 8 that occur randomly within the 16 microsecond TDC maximum count following a muon stop. Thus you should expect a flat background in the TDC time distributions from the neutron counters. The bottom line of all of is that there will be several backgrounds that will show up in the neutron detector lifetime distributions – if you can estimate the level and shape of each background you can subtract it out correctly; in the worst case, examining the lifetime data itself may permit a reasonable background subtraction.

In light of these obstacles, it is worth discussing some of the ways in which to ensure more robust data and to also discuss ways to fit that data that you do take. The Lab Manual discusses the basic strategy for doing the counter plateaus. This will work fine for the counters C1-C3.

a) Neutron Counter Plateaus

For the counters C5-C8, there are a couple of strategies that you can try. The simplest is to carry out a standard plateau for a given counter (using the coincidence method) and then examine the plateau shape and increase the HV of the counter gradually to make it more efficient as a neutron

detector, but at the same time monitor the singles rates to ensure that the counter has not entered the noise regime.

What exactly is the PMT noise that we are talking about? First, there can be extra photoelectrons created at the photocathode from spurious photons (light) in the scintillator counter. This light could come from small light leaks as well as extra charged tracks passing through the scintillator. However, the photoelectric effect is not the only source of electrons emitted from the photocathode and not necessarily the dominant source. There is thermal emission that depends strongly on the cathode temperature and *cold field emission* (CFE) due to electron tunneling, a quantum mechanical effect that depends strongly on the electric field just outside the cathode surface. There are even sub-effects in the CFE category that depend on the initial state of the electrons in the cathode material. This topic has been around since the dawn of quantum mechanics and is still in a state of incomplete theoretical understanding. For the 180F lab, CFE is a source of PMT noise that you should know exists but it is also beyond the scope to discuss in great detail – you can find historical articles on the subject by Googling “electron work function”. The work function is the difference of potential energy of a free electron outside the cathode and one inside in the Fermi sea.

A second, and generally better, way to determine the HV values for C5-C8 is to make use of a two-step process. First, plateau each of the four neutron counters C5-C8 against the 4-fold coincidence of C1, C2, C3, and C4. When you do this, bypass the Phillips 776 amplifiers for counters C5-C8. The goal is to set the HV’s high enough to achieve good efficiency for through-going muons but low enough to avoid excessive CFE. In the experiment, you will detect a coincidence between one of the neutron counter PMTs and the 4-fold (C1, C2, C3, and C4). Since noise pulses from different PMTs are unrelated in time and thus very unlikely to be simultaneous, the discriminator widths for C5, C6, X7, and C8 should be as small as is consistent with the extended geometry of the scintillator block. Once you finish this first step, the second step to achieve enough gain to detect neutron interactions is to put back the Phillips 776 amplifier stage in each of the C5-C8 signals.

A third possible way to set the high voltage values for C5-C8 is to use an additional charge-to-digital converter (QDC) to measure the integrated charge of the signals on each of C5-C8. This way requires a fair bit of extra work but can yield good results. To do this, the PMT signals (bypassing the Phillips 776 amplifier) are fed into four channels (5-8) of a Phillips 7166 QDC that is

gated by the coincidence of C1-C4. The relative delays between the gate and the C5-C8 PMT signals need to be appropriately adjusted by using the oscilloscope. The DAQ code, muminus.c, also needs to be modified by adding a statement "#define QDCflag" right after the "include" statements and the muminus code must be rebuilt (i.e. by running "make"). QDCflag is a conditional compilation flat to add code that reads the 7166 module and write its data to the output stream. If you have any difficulties rebuilding the DAQ code, discuss with the instructor or TA. When you run this version of the DAQ code, your output is longer by 16 bytes, the four 4-byte QDC data. Plots of the charge data for C5-C8 should show a clear peak for the through-going muons and a noise peak at low charge values. By taking different runs with different HV values, you should be able easily see how the separation between the signal (through-going muons) and the noise changes with HV. It should thus be straightforward to set the HV so that you are as efficient as possible, maintaining a good separation between the signal and noise.

b. Anti-Counter (C4) Plateau

The HV for the anti-counter C4 also requires some care. Use the logic (discriminator widths and thresholds) as shown in Figure 13 of the Lab Manual. Plateau C4 against the 3-fold coincidence of C2, C3 and C5. To get C5 into coincidence, you will need to temporarily bypass its 96 ns delay cable. Connect the output of the (C2, C3, and C5) coincidence into one scaler and connect the output of the (C2, C3, C4, and C5) coincidence into a second scaler. Then try to adjust the HV for C4 so to reduce its inefficiency to below 0.5% (i.e. so that it is 99.5% efficient) for through-going muons, while monitoring the C4 singles rate to ensure that you are not introducing more accident and noise events than you can handle.

c. Ringing and Positive Ion Feedback

Very large light pulses incident on a PMT photocathode often lead to "ringing" in the analog output due to high frequency instabilities of the resistor voltage-divider circuit that distributes HV to the dynodes. In addition, large "after pulses" can be made by positive ions (produced in the residual gas inside the PMT) that drift back and smash into the photocathode. When ringing or after pulses occur, C4 can fail to detect and veto a muon which is thereupon falsely identified as stopping. If that muon goes all the way through the neutron scintillator there can be a large photon yield in the neutron counter that results in multiple discriminator pulses. These generally die out after a few tens of nanoseconds and are not of much concern for the Al target, where the expected lifetime far exceeds 100 ns. For Pb and Fe targets, this is a potential problem, and it can be reduced by requiring coincidences among two or more of C5-C8 and by making C4 as efficient as possible.

On the other hand, the rate of through-going muons is so large compared to the rate of detectable capture events that the performance required of C4 is quite severe.

Data Runs and Analysis

The remaining discussion of the mu-minus section summarizes the results from a set of canned analysis programs that existed several years ago. While we present some of the results of these programs and a description of how they were obtained, we encourage students to derive their own sets of plots – by, e.g., developing their own ROOT macros. We also provide some hints on general topics such as backgrounds, analysis strategies, and so forth.

In this lab you will need to measure the capture rates in Fe, Al, and Pb. The best target to do first is iron since the rate to be measured is far greater than the muon decay rate but not so large as to be potentially limited by the Phillips 7186D TDC resolution. After that a target-out run should be done, though the effects of this run for Fe will be less severe than for Al.

You will save time and confusion if you organize your data runs carefully. It is hard to know from the outset how much time to spend on each target. It is likely to be different for the three targets and plots may need different scales. You can easily end up with disjoint sets of data that you will want to run thru analysis programs more than once.

Backgrounds and Distortions.

A pmt time distribution can be distorted by afterpulsing from positive ion feedback. These pulses accumulate near 800 ns in the five-inch tubes. Thru-going muons outnumber stopping muons by an order of magnitude or so. If the C4 veto fails occasionally despite your care with setup, the event triggers as a mu-stop. If enough light from one of these muons is seen by two or more neutron-pmts, there is a good chance to generate a coincidence of feedback pulses.

A small amount of background from μ^+ decays is to be expected. This can be accounted for in fitting the time distribution by including a term with μ^+ decay, i.e. fixed decay time but variable amplitude. Beyond 16 microseconds (4000 counts of the TDC and 8 μ^+ mean lifetimes) this contribution is effectively dead but there are generally still some events in a tail that extends beyond the TDC upper limit. The physics of slow neutrons suggests that there are many

secondary interactions that give off more low energy neutrons before a detectable interaction occurs. In the three examples below, a flat distribution seems almost adequate.

A measurable and subtractable background due to muons that are captured in material outside the target should be estimated and removed. Therefore take some data with the target removed. Any target-in distribution can be corrected by subtracting the corresponding target-out distribution scaled for runtime. It is often not practical to get smooth background distributions and bin-to-bin fluctuations can reduce the statistical significance of the corrected target-in data. It is tempting to smooth the background but that can build in a prejudiced idea of what the target-out data should look like. It may be better to accept these fluctuations; this is your judgment. The target-out analysis depends on target type since the μ^- fit is done to only for t<tcut**,** which should be chosen appropriately for each target. The run time can be found from the computer time that is recorded for each event.

Other sources of background can affect your experiment. When a muon hits C1, C2, and C3 (above the target slab) and misses C4 (below the target), there is a false trigger. In the great majority of these there will be no TDC stop signal and the DAQ (muminus.c) will ignore them (and hence no event will be recorded). Occasionally an unrelated muon will enter the neutron detector obliquely, having missed C1, C2, C3, and C4. These accidental events contribute to what can be called the “tail”; in contrast to the neutron cascade events.

Sample Histograms

Fe Target:

Figures 12 and 13 show TDC distributions for the Fe target. The subsequent no-target runs are shown in Figures 14 and 15. The TDC distributions for the eight counters (C1-C8) are plotted in Figure 12b to Figure 12h on a log scale. Figure 12a is the accumulated events as a function of time in hours. You can read the runtime from this plot to be 336 hours. Figure 13a shows the TDC distribution on a linear scale. The units on the horizontal axis are in 4 nsec, so the x-axis goes from 0 to 16 μsec. Events with TDC > 6 μsec (value of 1500 on the x-axis) are used to determine the flat background or tail distribution, as indicated by the bold red horizontal line. Events with 0< t < tcut (the signal part of the distribution) are fitted to a function that is a sum of three terms: tail plus μ^+ decay A*exp(-t/2.197) plus B*exp(-Ct) for μ^- capture. The tcut parameter is chosen appropriately for each target. For Fe, based on Figure 13a it was set to a value of 2.8 μs (value of

700 on the x-axis). Figure 13(b) shows just the signal region, with the three fitted functions: tail (red), μ^+ (blue) and μ^- (green). The fitted values are given in the legend where A = ampl+, B = ampl- and C = rate-. The fitted time constant (i.e.,1/rate-) for Figure 13b is 54.67+/- 2.8 or (219+/- 11) ns. Figure 13b shows an additional broken line that is the scaled and smoothed contribution from the target-out run that of course you don't know until you do that run.

Figure 14 shows the same distributions as in Figure 12, except that the target was removed. Here the runtime was 193 hours. Figure 15 shows the target-out data scaled by runtime; this is the histogram of the fitted background values to be subtracted from the raw Fe distribution of Figure 13. The result of the subtraction is shown in Figure 16a and the fit for the signal region is shown in Figure 16b. The final result on the μ^- time constant is (202 +/- 11) ns, which is somewhat smaller than the value obtained in Figure 13b where the target-out data had not been considered.

Al Target:

The results for the Al target are shown in Figures 17-19. The target-out subtraction was most important for Al because the μ^- disappearance rate for this target is the closest one to the target-out effective rate.

Pb Target:

The results for the Pb target are shown in Figures 20-22.

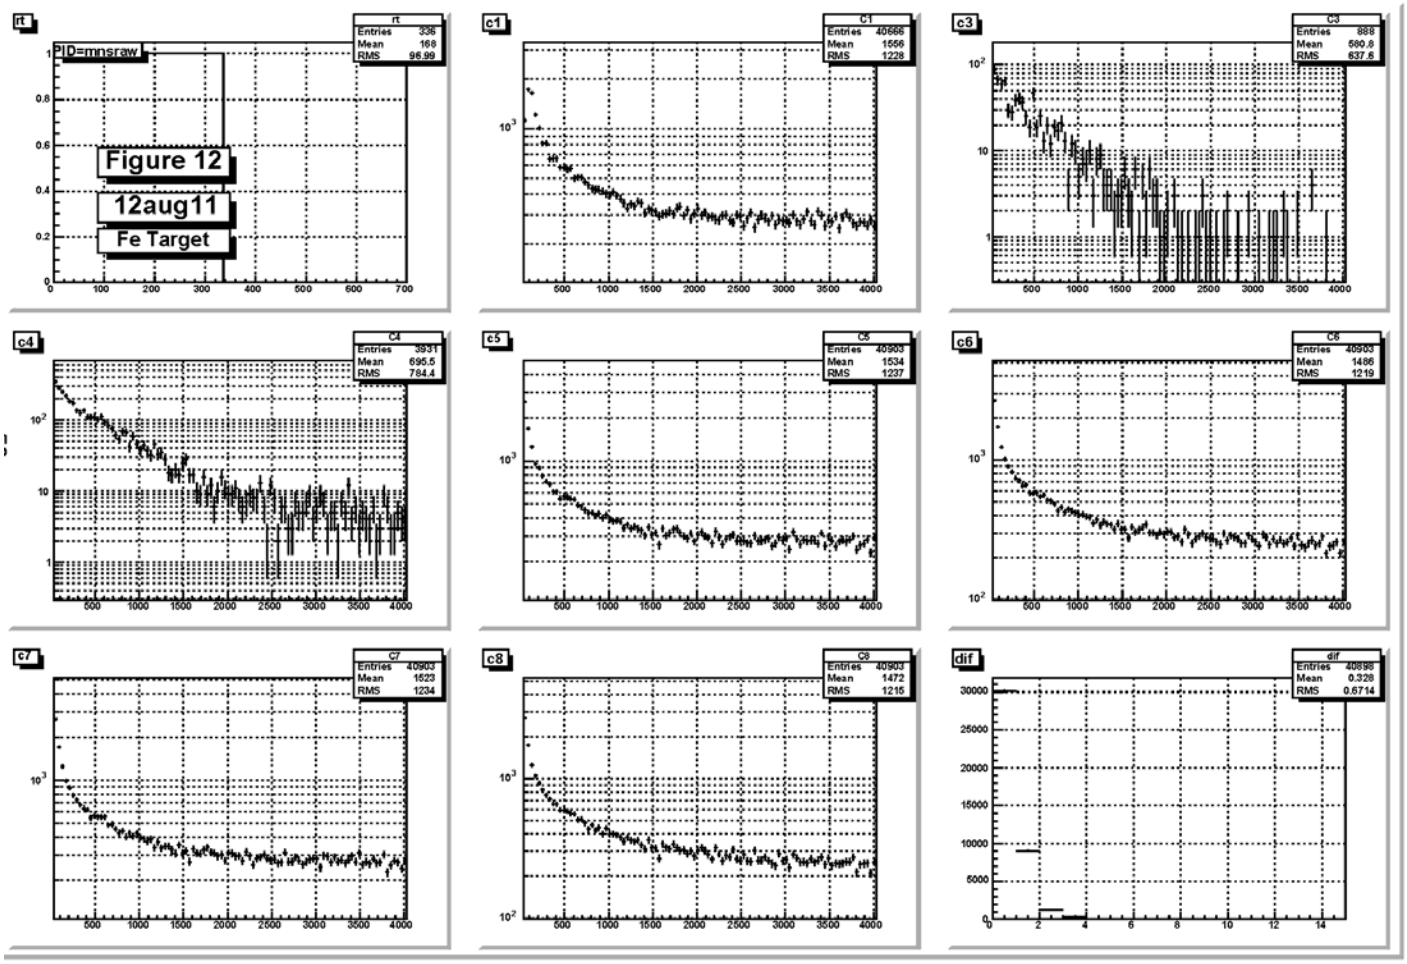

Figure 12: Histograms of raw data for μ^- experiment with the Fe target. a) shows the runtime in hours and b) - h) show the TDC distributions for counters C1-C8, respectively.

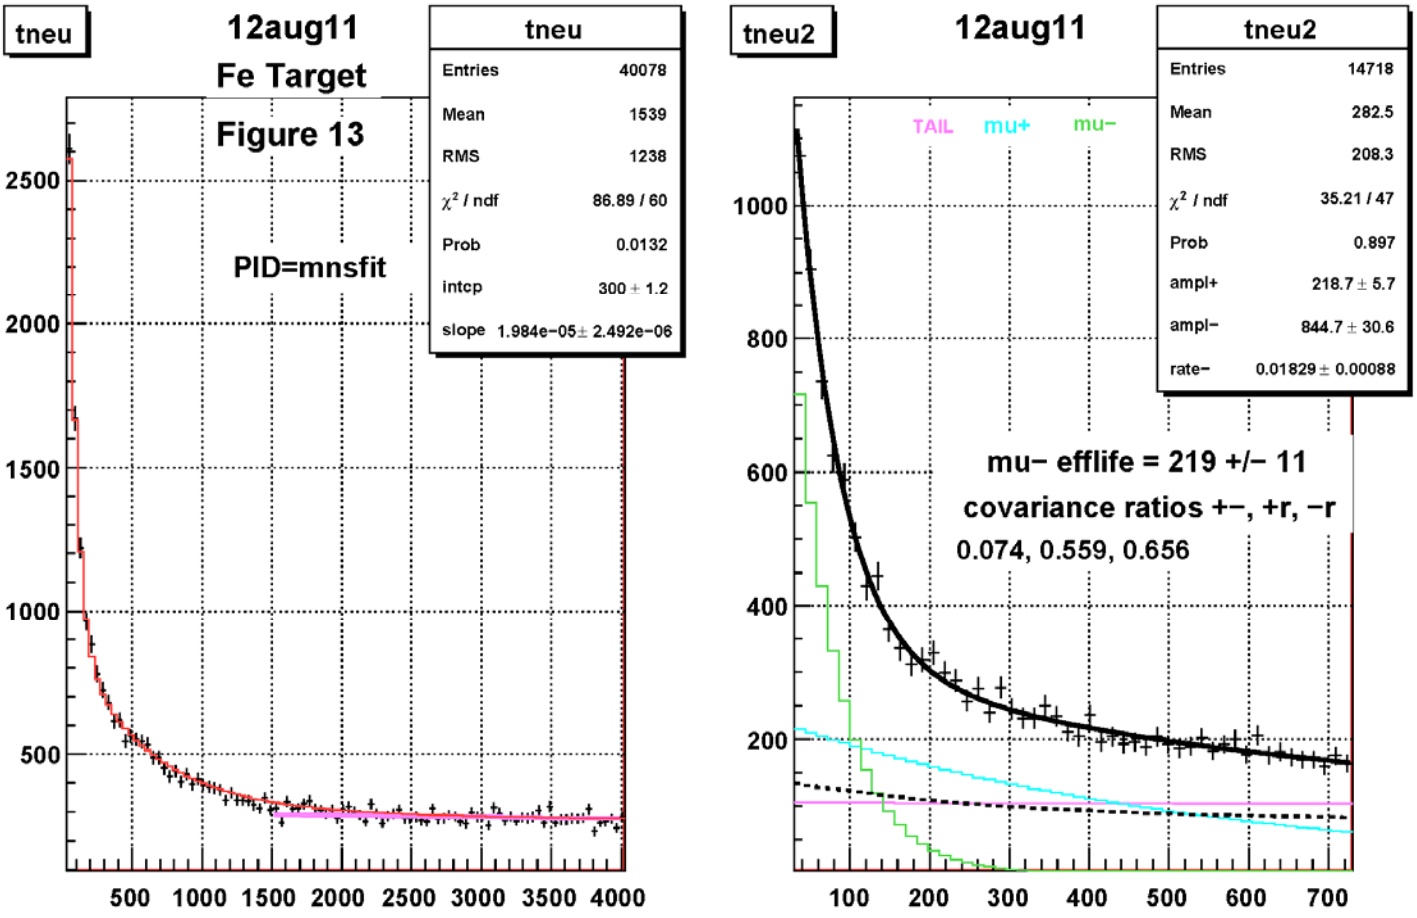

Figure 13: Histograms of fitted data for μ^- experiment with the Fe target. a) shows the TDC distribution on a linear scale from 0 to 16 μs. b) shows a zoom of the same distribution from 0 to 2.8 μs. The various fits are described in the text.

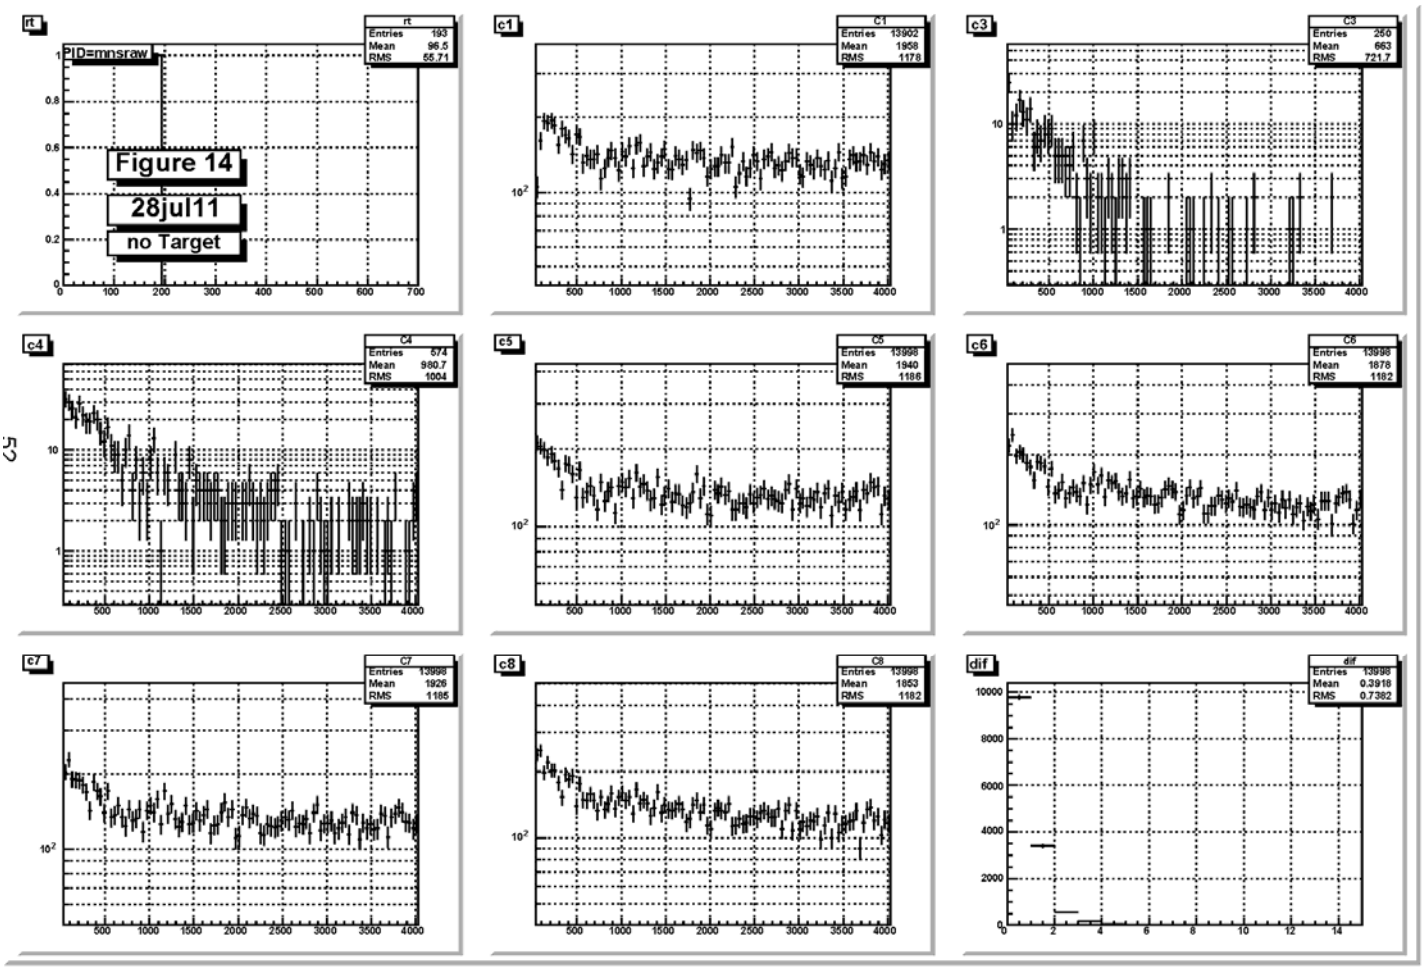

Figure 14: Histograms of raw data for μ^- experiment with the target out. Same distributions as Figure 12, except with no target.

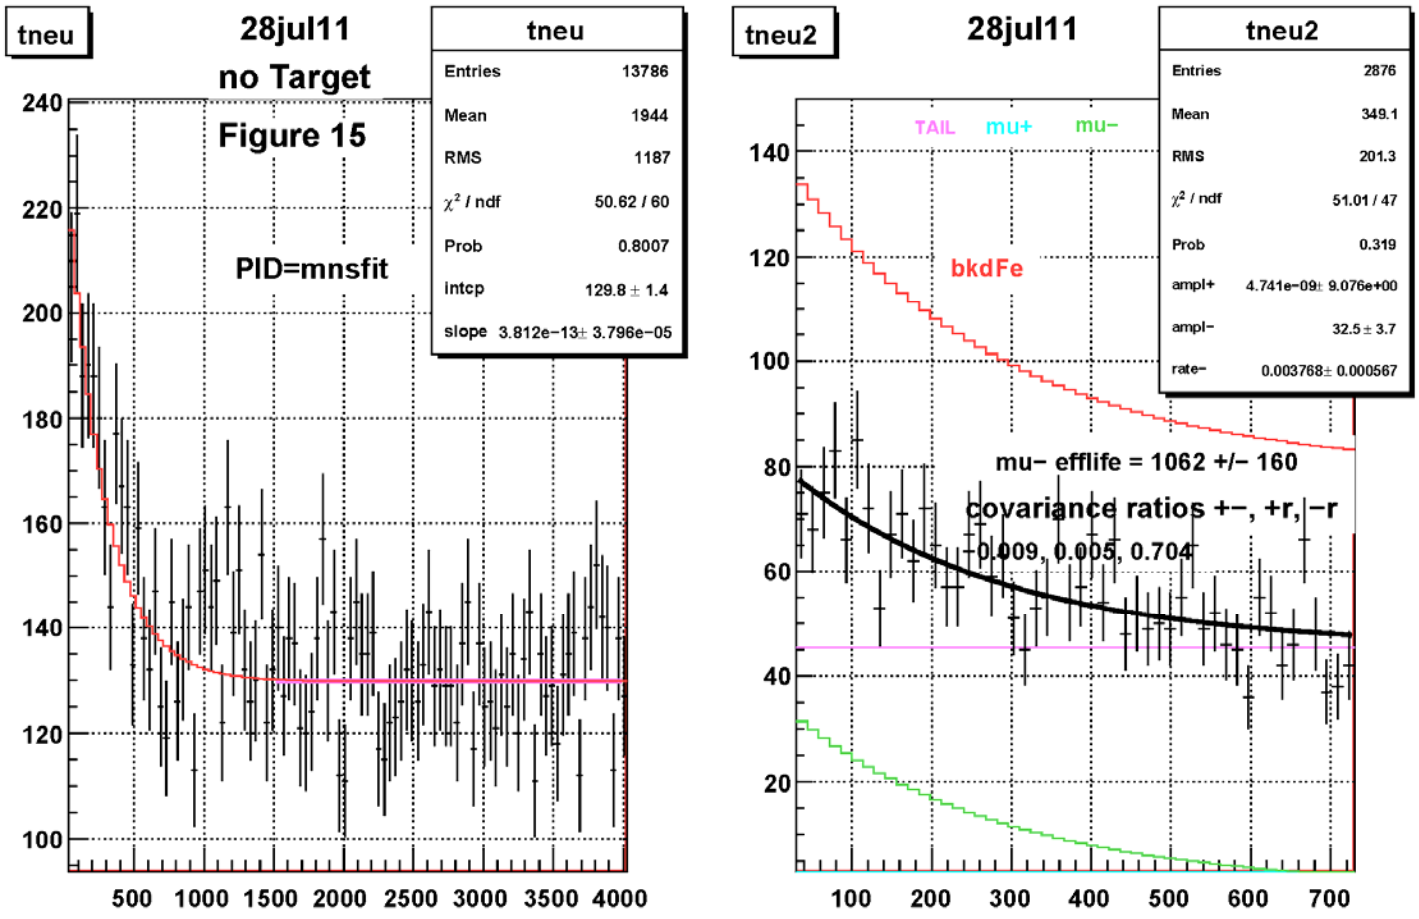

Figure 15: Histograms of fitted data for μ^- experiment for the target-out data. Same distributions as Figure 13, except with no target. The various fits are described in the text.

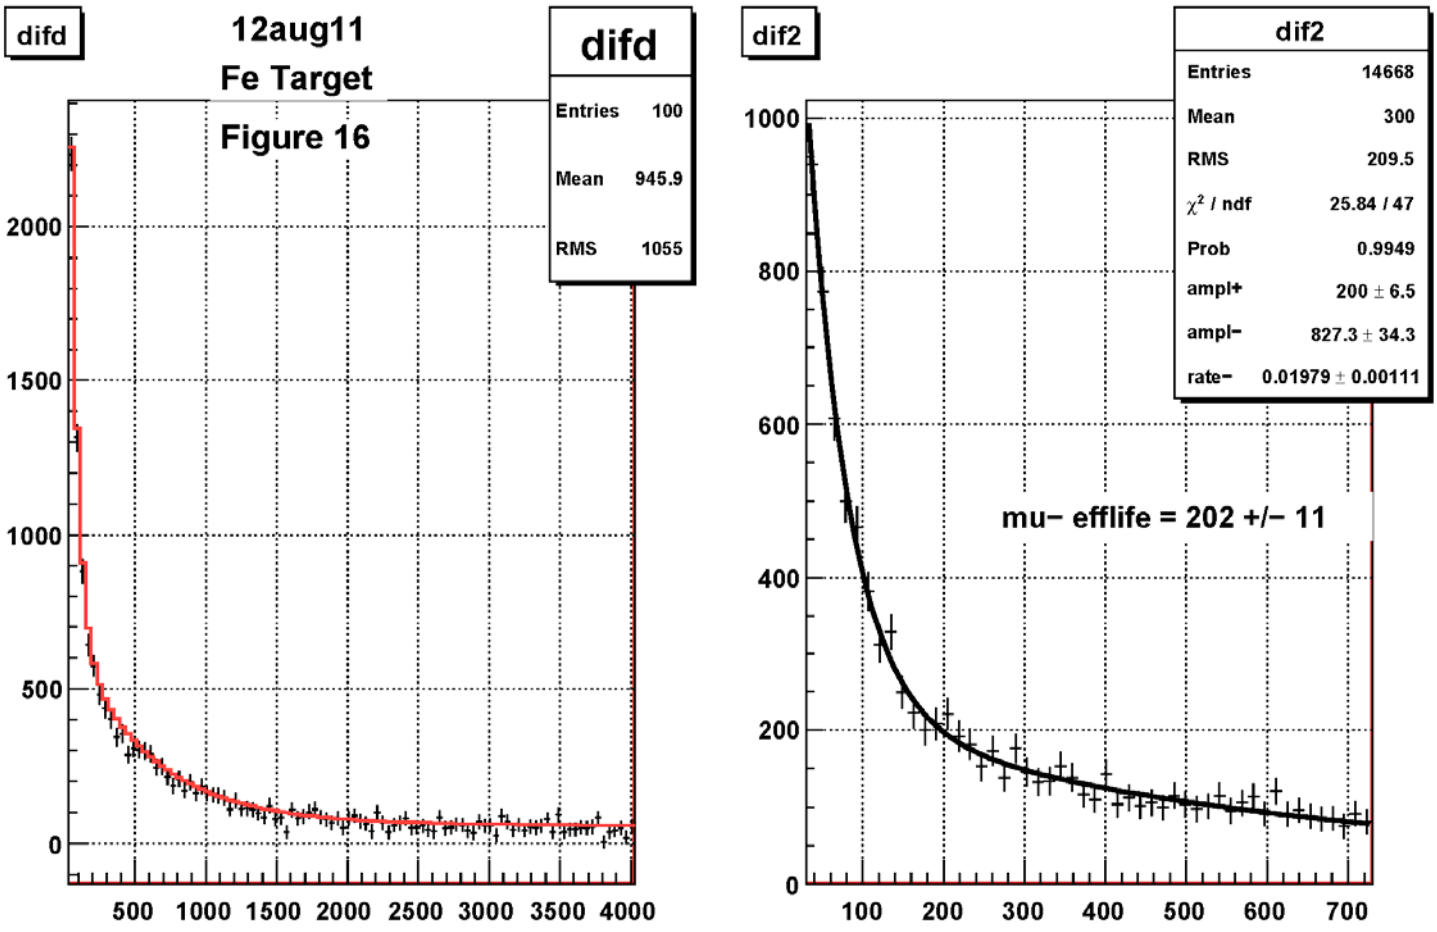

Figure 16: Histograms of fitted data for μ^- experiment for the Fe target, showing the final fits where the target-out data has been statistically subtracted (bin by bin).

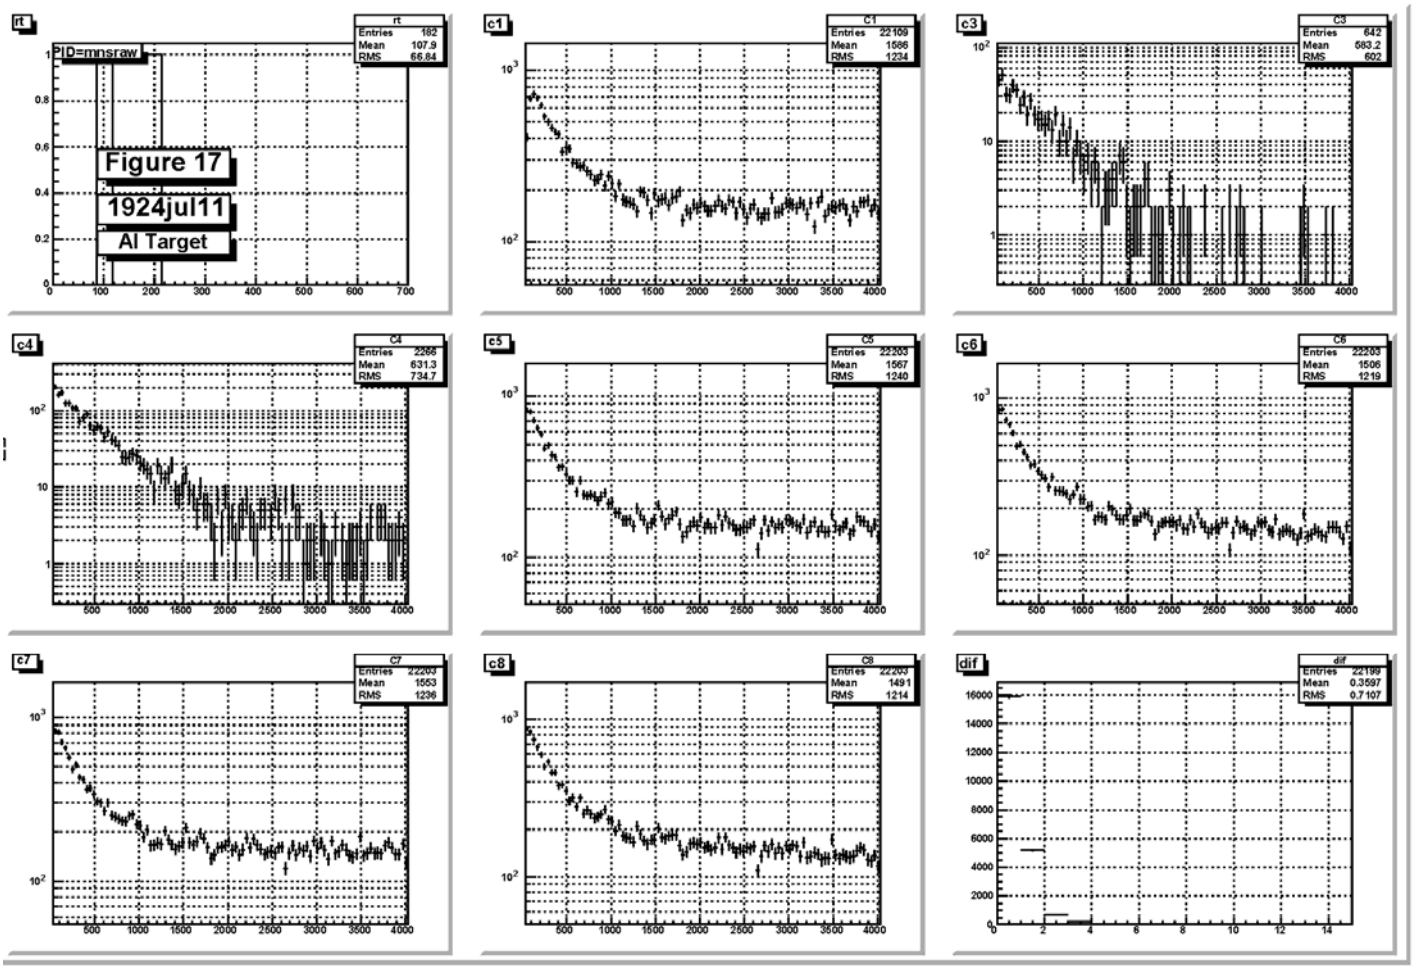

Figure 17: Histograms of raw data for μ^- experiment with the Al target. a) shows the runtime in hours and b) - h) show the TDC distributions for counters C1-C8, respectively.

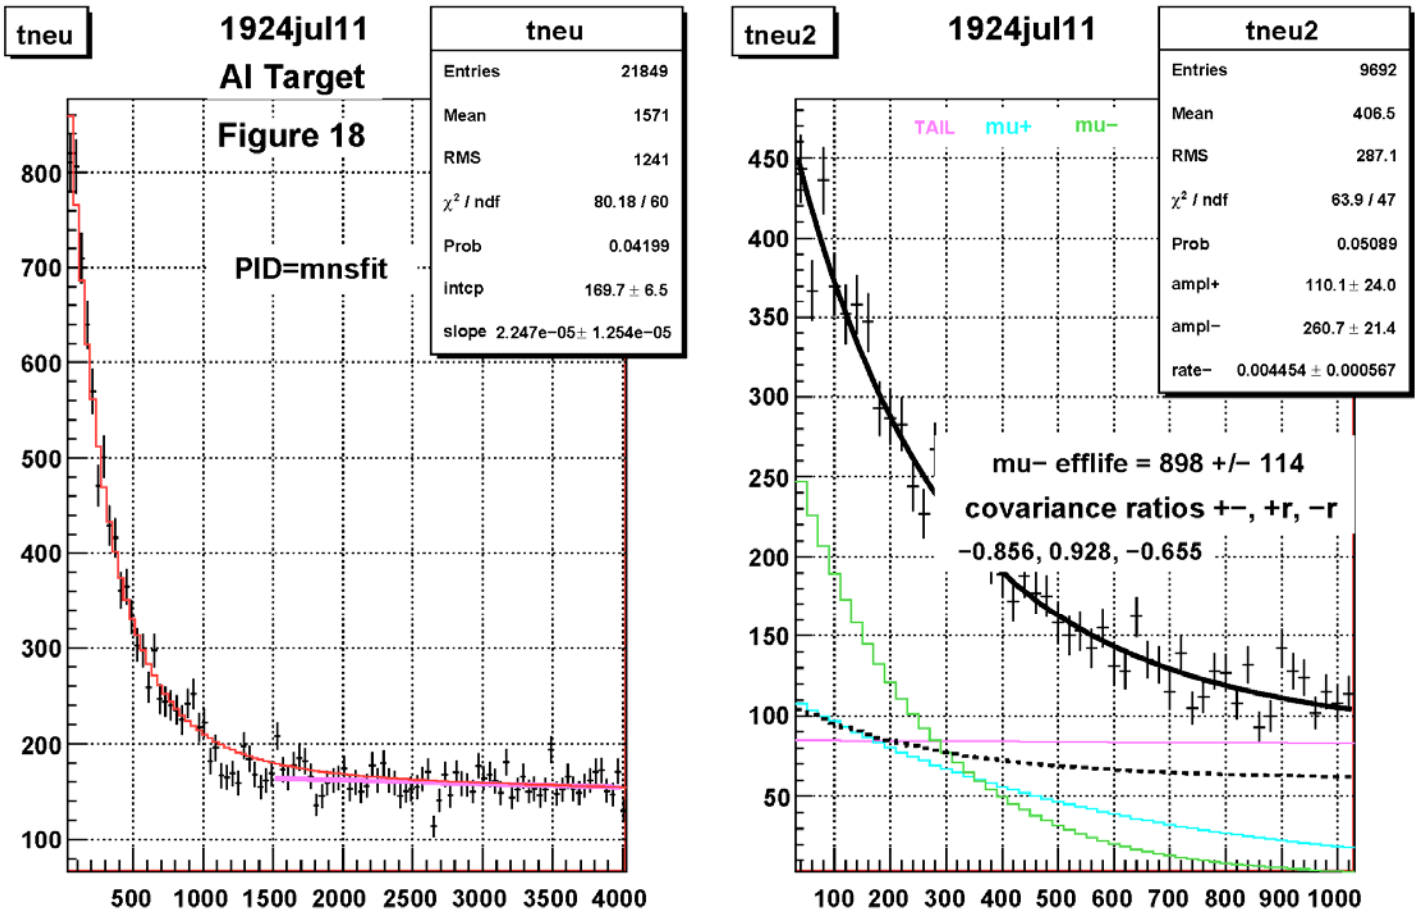

Figure 18: Histograms of fitted data for μ^- experiment with the Al target. a) shows the TDC distribution on a linear scale from 0 to 16 μs. b) shows a zoom of the same distribution from 0 to 4 μs (value of 1000). The various fits are described in the text.

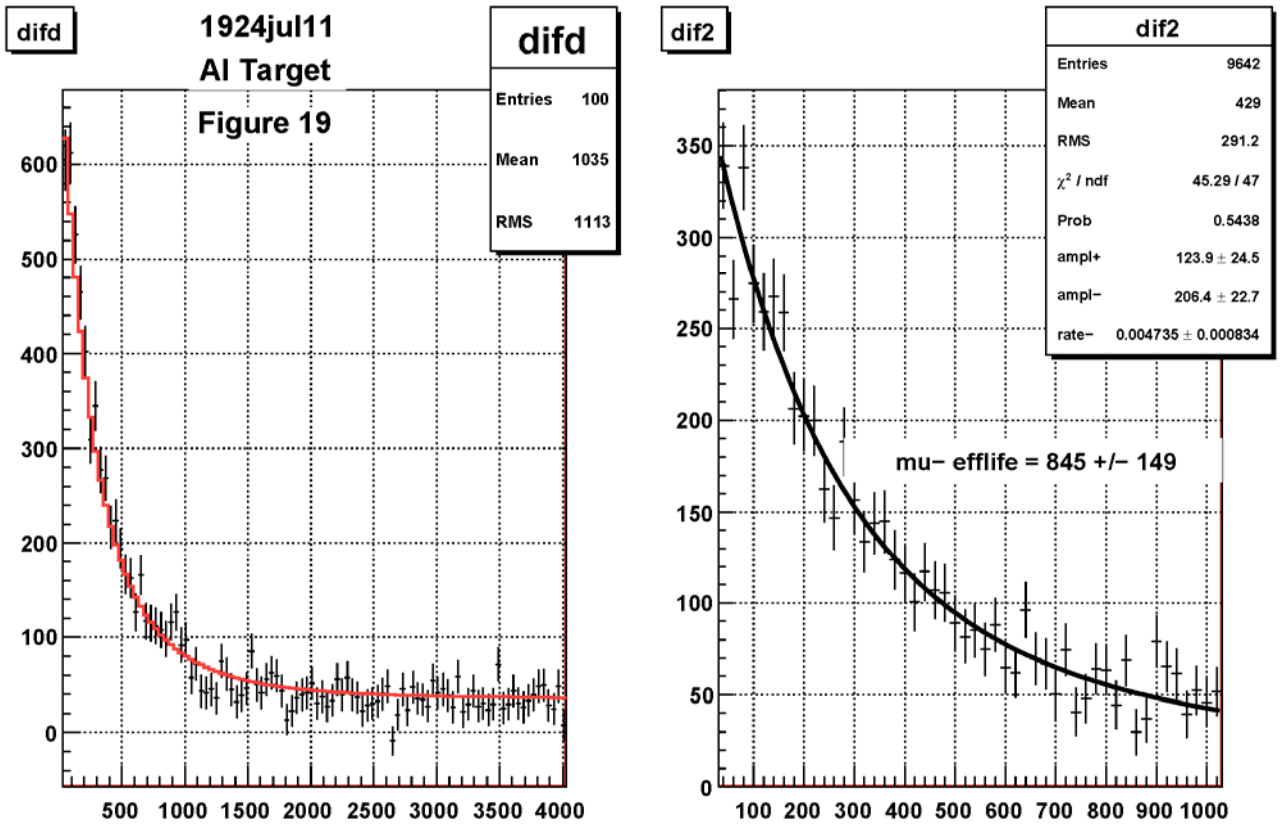

Figure 19: Histograms of fitted data for μ^- experiment for the Al target, showing the final fits where the target-out data has been statistically subtracted (bin by bin).

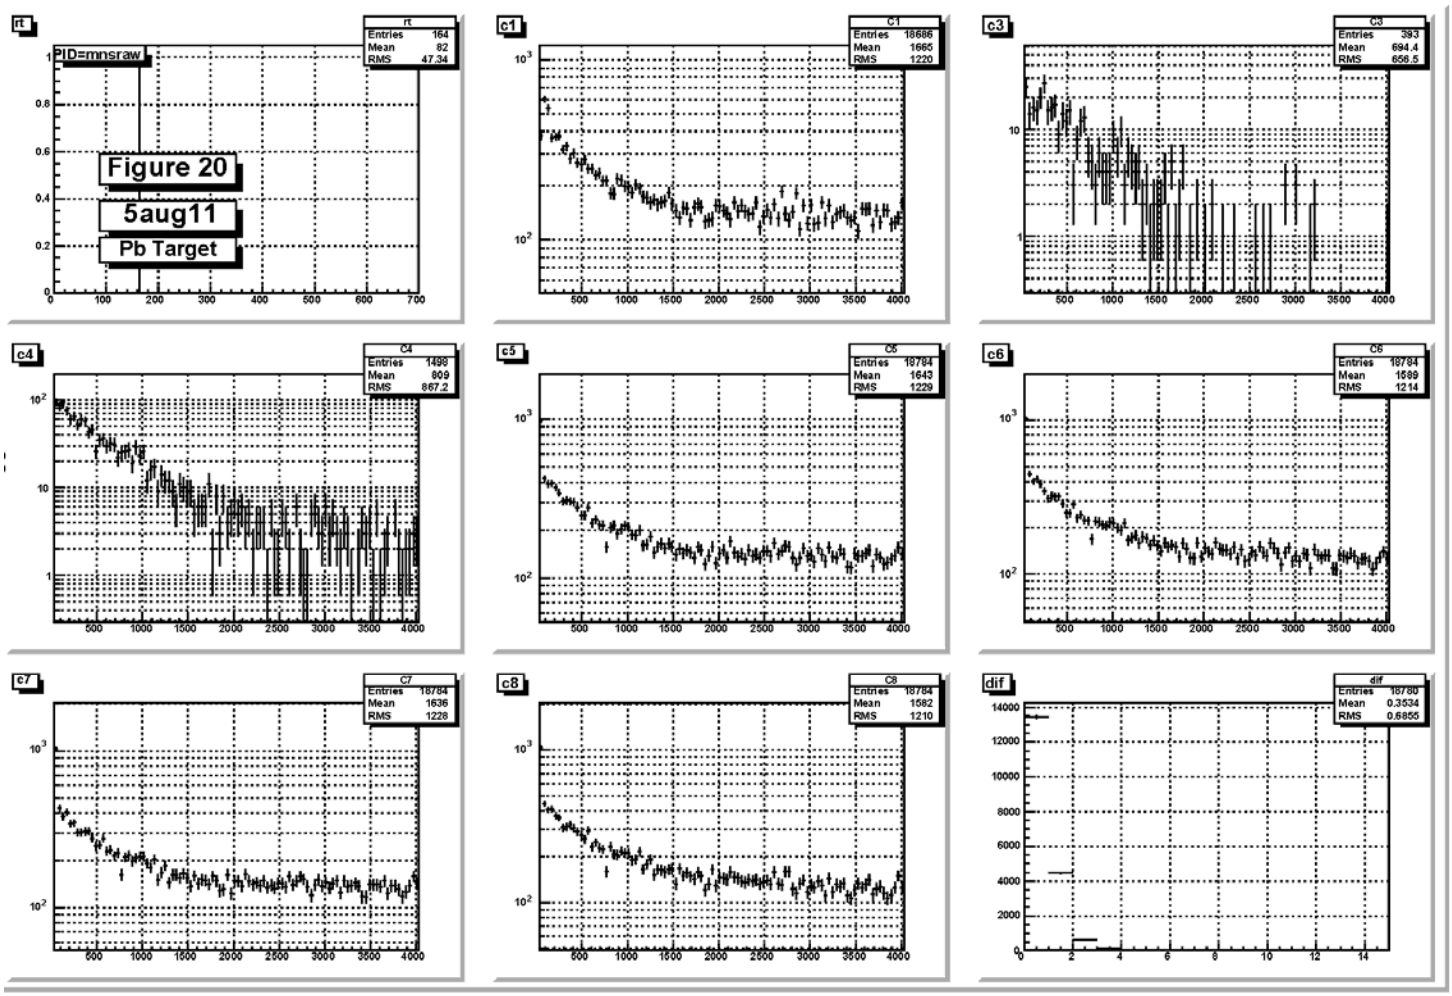

Figure 20: Histograms of raw data for μ^- experiment with the Pb target. a) shows the runtime in hours and b) - h) show the TDC distributions for counters C1-C8, respectively.

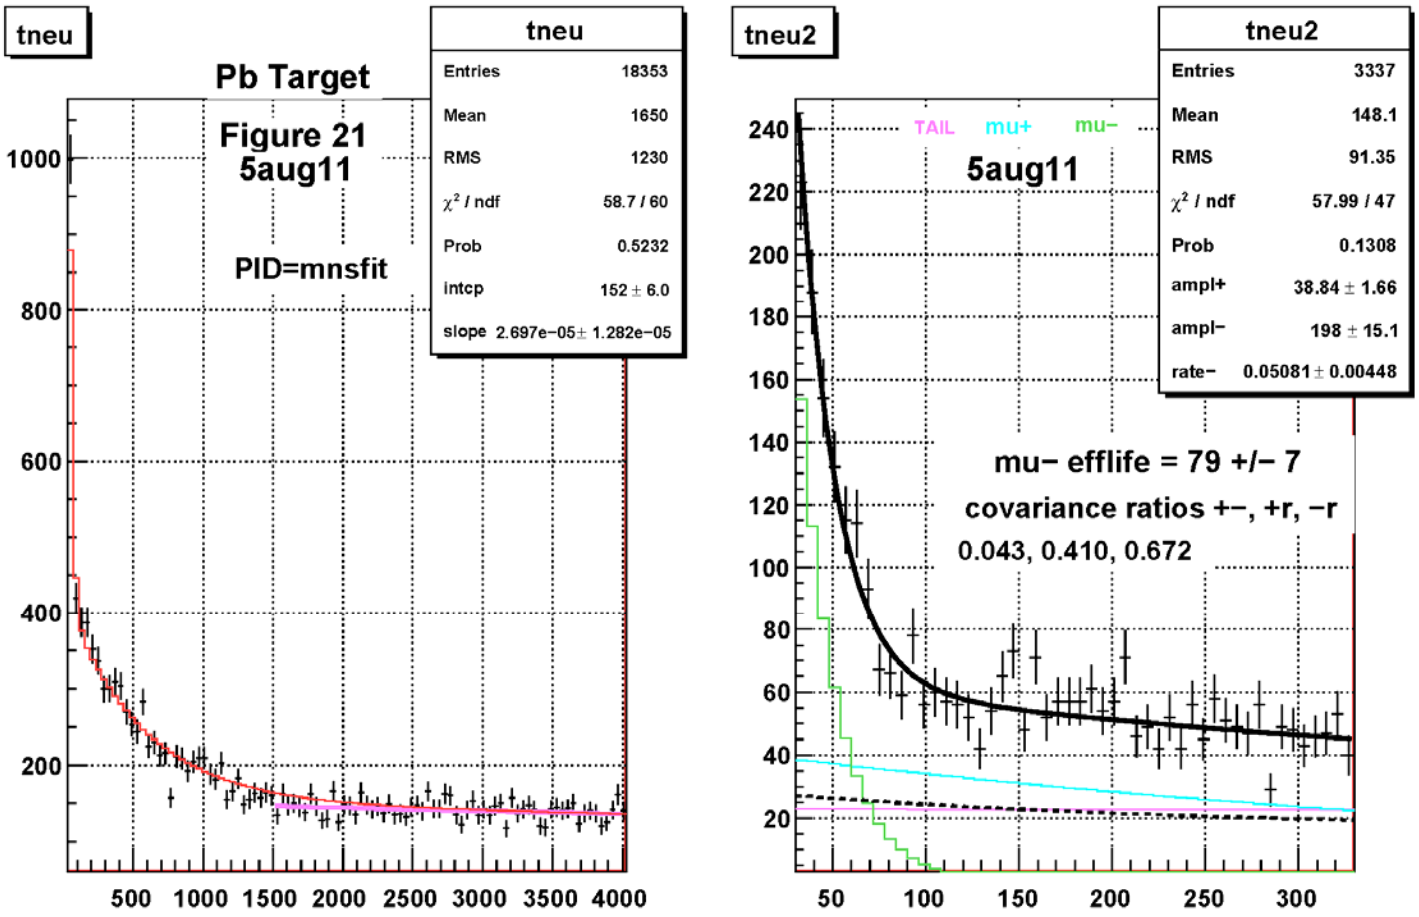

Figure 21: Histograms of fitted data for μ^- experiment with the Pb target. a) shows the TDC distribution on a linear scale from 0 to 16 μs. b) shows a zoom of the same distribution from 0 to 1.2 μs (value of 300). The various fits are described in the text.

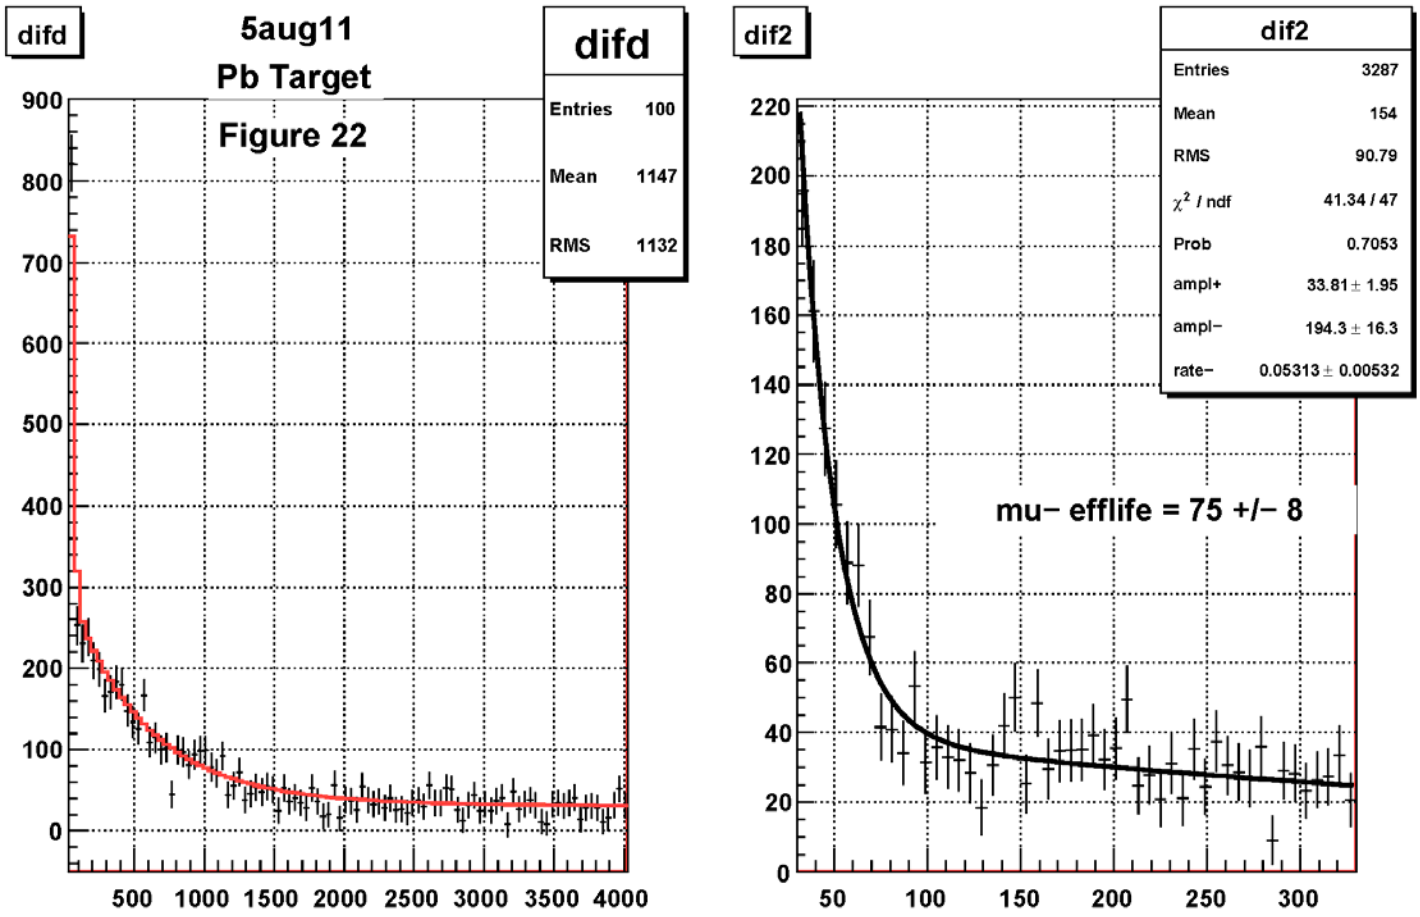

Figure 22: Histograms of fitted data for μ^- experiment for the Pb target, showing the final fits where the target-out data has been statistically subtracted (bin by bin).

C. MCS Experiment

In the MCS experiment, low-momentum muons are selected by a set of scintillation counters (see Figure 14 of the Lab Manual) and the associated coincidence electronics (see Figure 15). When a valid trigger occurs, the spark chamber is pulsed and camera records an image of the resulting sparks. Your data consist of a series of camera images from which you will need to analyze in order to study the multiple Coulomb scattering process. To do the analysis you need to determine the relationship between the coordinates of the camera images (i.e. the pixel values in x and y) and the coordinates of the experiment in real space (i.e. the physical horizontal and vertical coordinates). You can do this by stacking many events on top of one another and the resulting plot will show the positions and dimensions of the spark chambers (or at least those portions of the chambers that are efficiently producing visible sparks). My measuring the positions and dimensions of the chambers with a ruler you can then determine the mapping between image space and physical space. Once you know where the chambers are, you will need to determine the scattering angle on an event-by-event basis. You will need four-chamber events to do this because two points above the scattering plane determine a line, as does two points below the scattering plane. Basically, the scattering angle can be thought of as the deviation of the trajectory of the muon from above the scattering plane to below it. Of course, the angle you will measure is projected into the two-dimensional plane of the camera and so you will need to convert from the angle you measure to the full three-dimensional scattering angle.

There are a variety of ways to go about analyzing the data. A reasonable way to proceed would be to have the following separate programs to: 1) read in the raw image data and store it in ROOT format (ROOT tree), 2) read in a single event and display it using ROOT, 3) stack a large number of events together and plot the aggregate images, and 4) reconstruct the tracks above and below the scattering plane, measure the scattering angle on an event-by-event basis and store that angle in a ROOT histogram for later analysis and plotting. As of 2016, the first three of these programs are provided in basic form in the analysis starter kits uploaded to the class website. Working with the provided programs you should be able to produce your own version of 4) to calculate and histogram the scattering angles.

Note: after you have determined the mapping between the camera coordinates and the physical coordinates, make sure to not bump the camera or adjust its settings.

Sample Histograms

The plots on the following pages show some characteristic distributions produced by an older set of ROOT macros.

Figures 23 and 24 show the results from stacking many images on top of one another. In Figures 23, a two-dimensional scatterplot shows the stacked images as a function of the x and y coordinates of the camera. Figure 24 shows a projection along the x coordinate clearly showing the chamber boundaries in the vertical direction. (Note that different data were used to make Figures 23 and 24).

Figure 25 shows the same data as in Figure 24, except zoomed in and using bins of single pixels; this makes it easier to see the chamber boundaries.

In Figure 26, projections along the y coordinate are shown for the four chambers, each on a separate plot. You can see that sparks cluster at the middle of the chambers; these plots are typical of reasonably good chamber performance.

Figure 27 gives histograms of the actual intensities for each hit pixel in the four chambers; clearly chambers 1 and 4 are providing the brightest sparks.

In Figure 28, a number of individual events are shown, along with reconstructed tracks above and below the scattering plane. Here the dashed line is the (un-scattered) extrapolation of the original track to below the scattering plane.

Figure 29 shows histograms of the measured scattering angle from a sample of 291 good events, on linear (left) and logarithmic (right) scales. The distribution of angles is well-centered on zero and is fit by a sum of two Gaussian distributions, with each Gaussian distribution having a separate normalization, mean and standard deviation.

Figure 30 shows similar histograms as in Figure 29, except now for different amounts of material in the scattering plane.

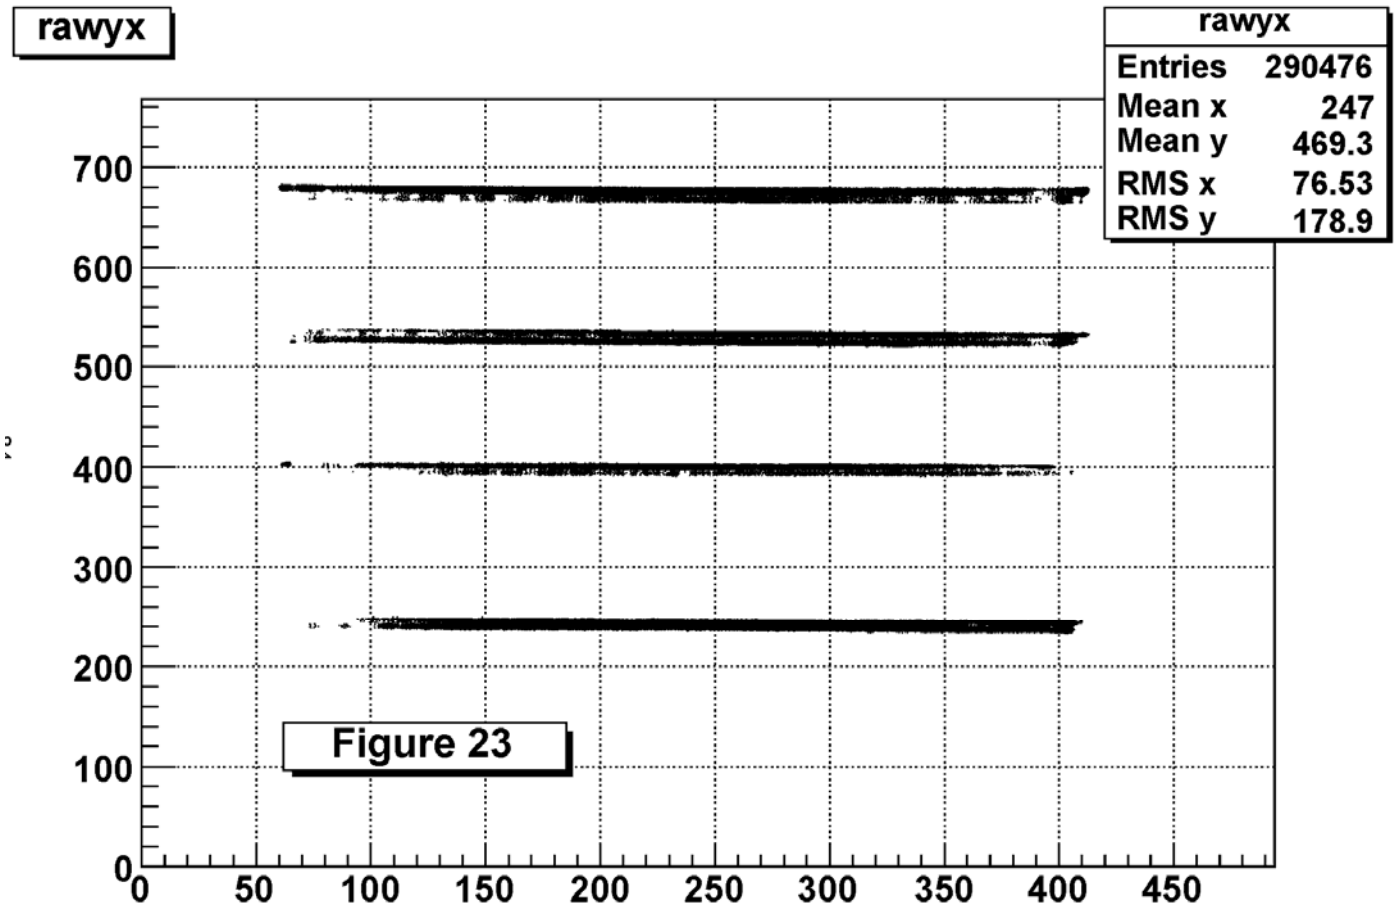

Figure 23: Two-dimensional scatter plot of a set of stacked camera images. The x and y coordinates correspond to the camera dimensions.

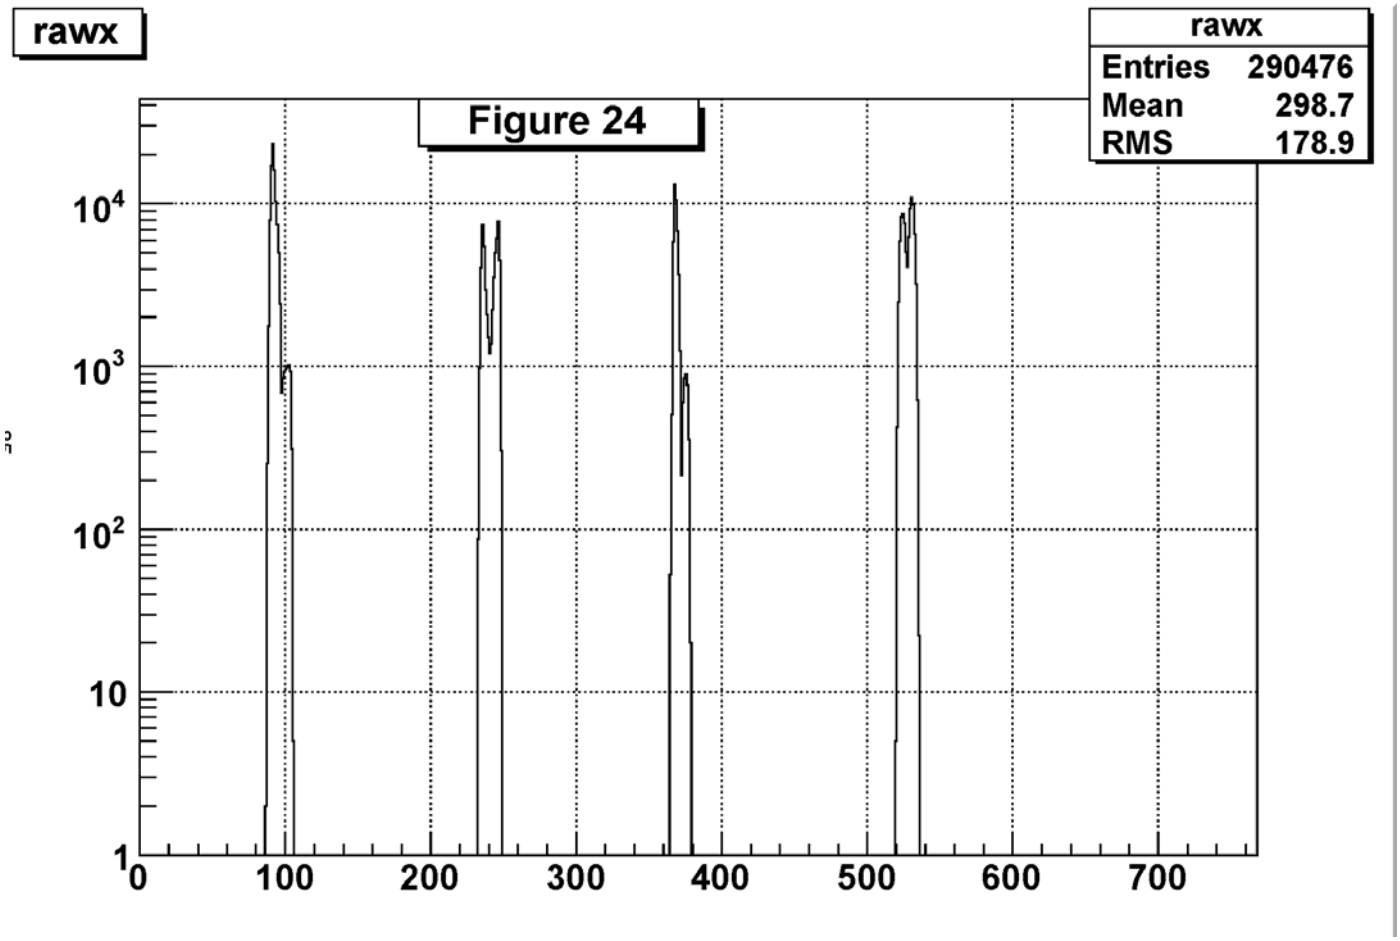

Figure 24: Projection of a set of stacked camera images along the x direction, showing the locations of the spark chamber planes in the vertical (physical) dimension.

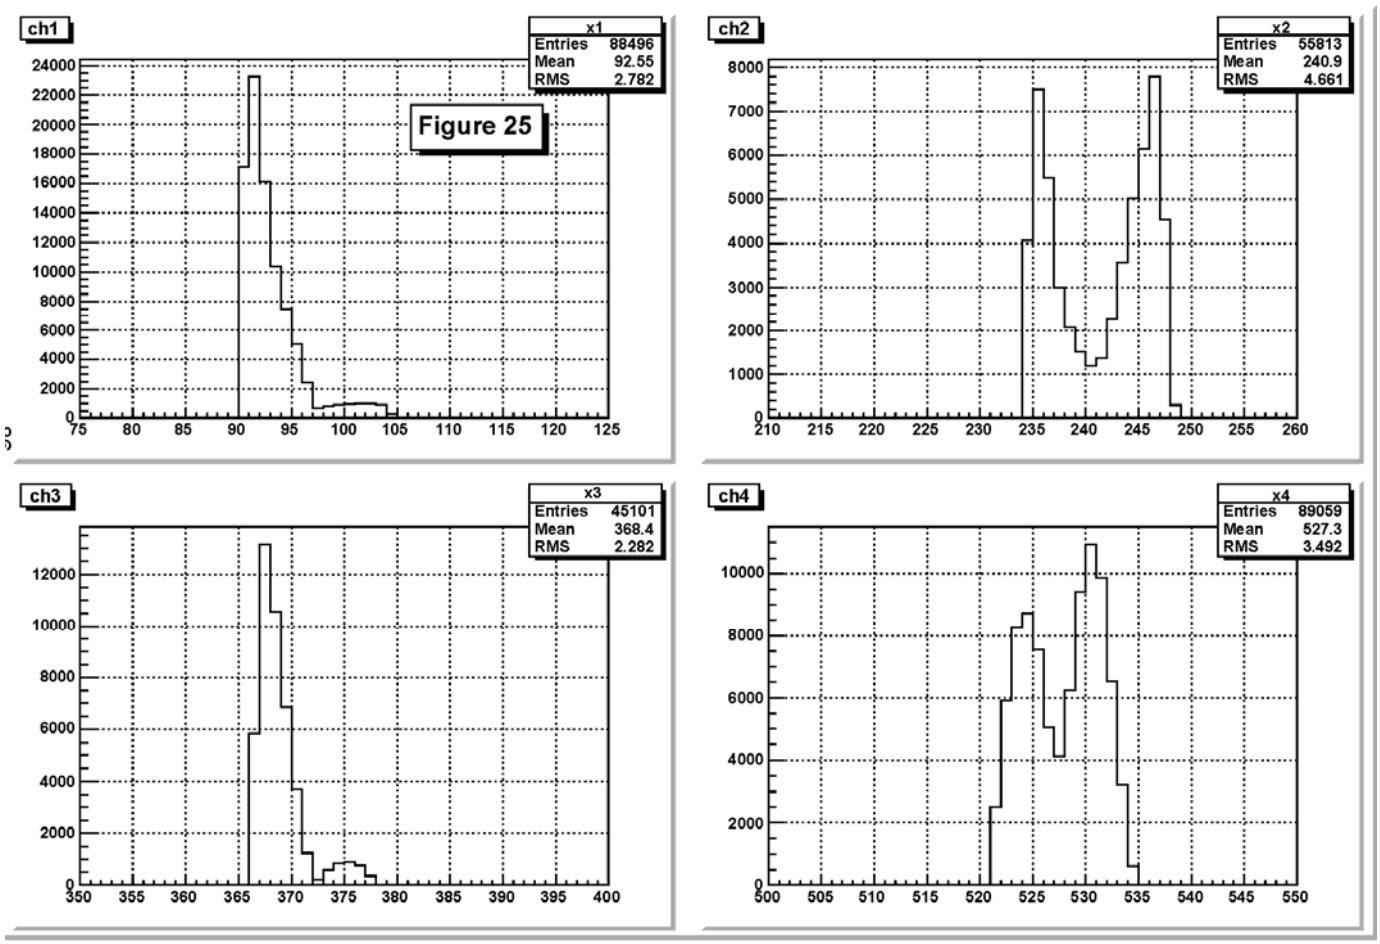

Figure 25: Same plot as Figure 24, except zoomed in and using bins of single pixels; this makes it easier to see the chamber boundaries.

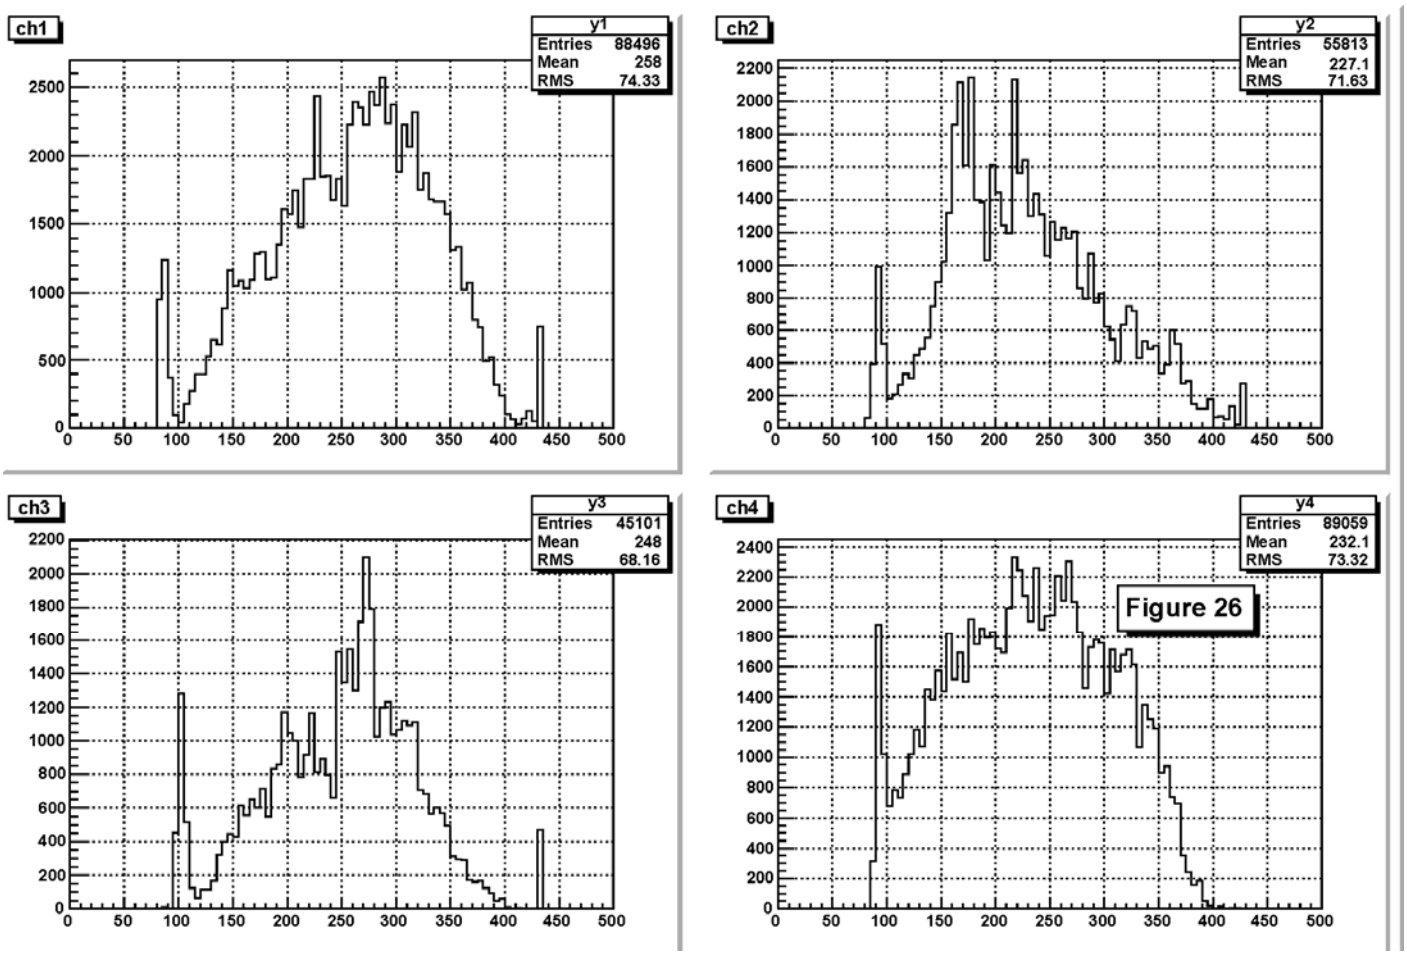

Figure 26: : Projections of a set of stacked camera images along the y direction, showing the locations of the sparks along the horizontal direction in each of the chambers.

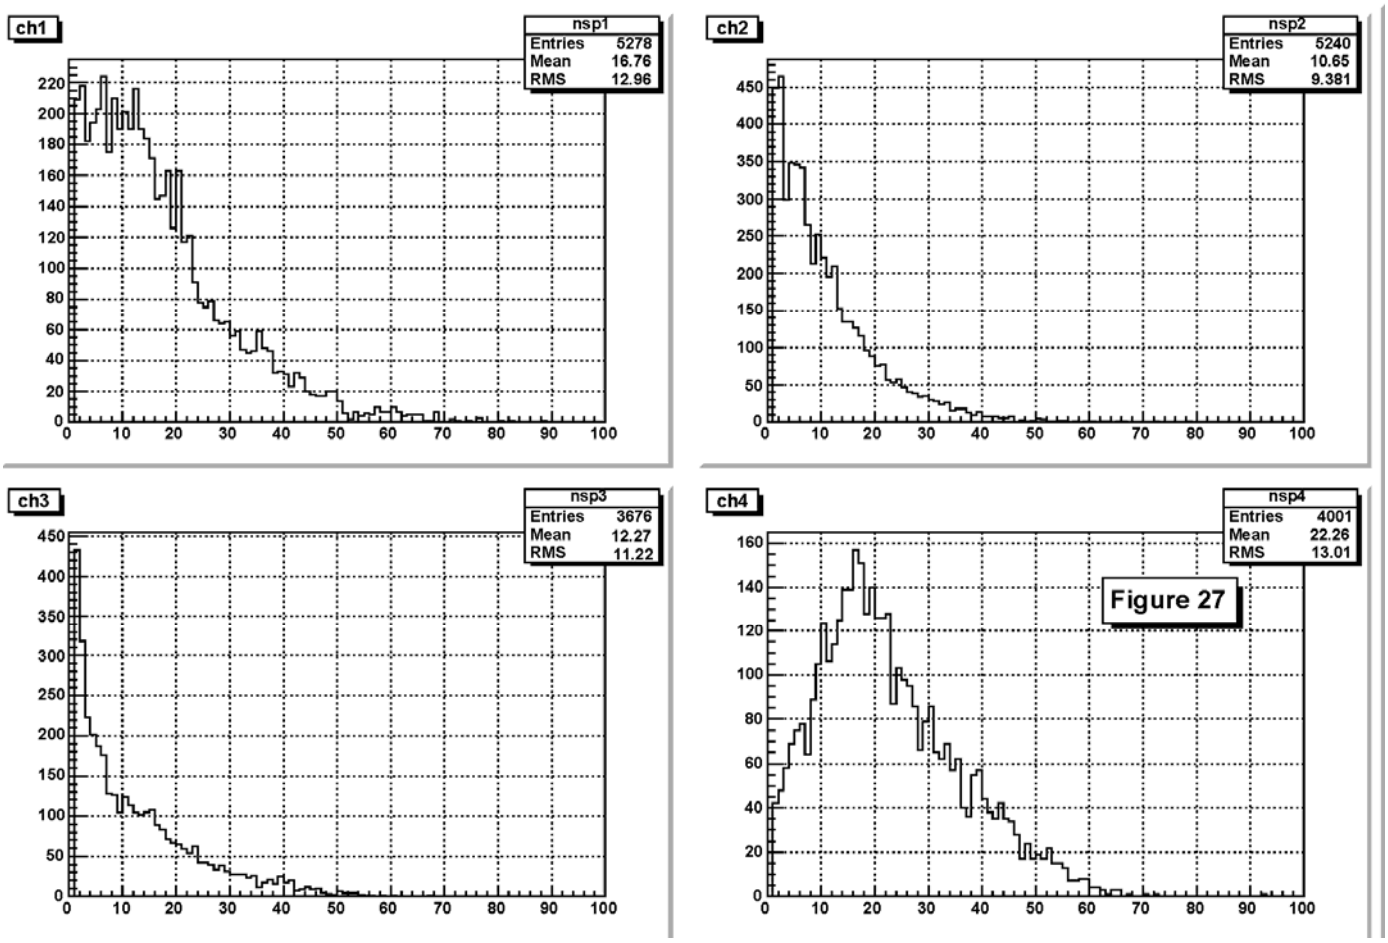

Figure 27: Histograms of the intensities for each hit pixel in the four chambers for many events.

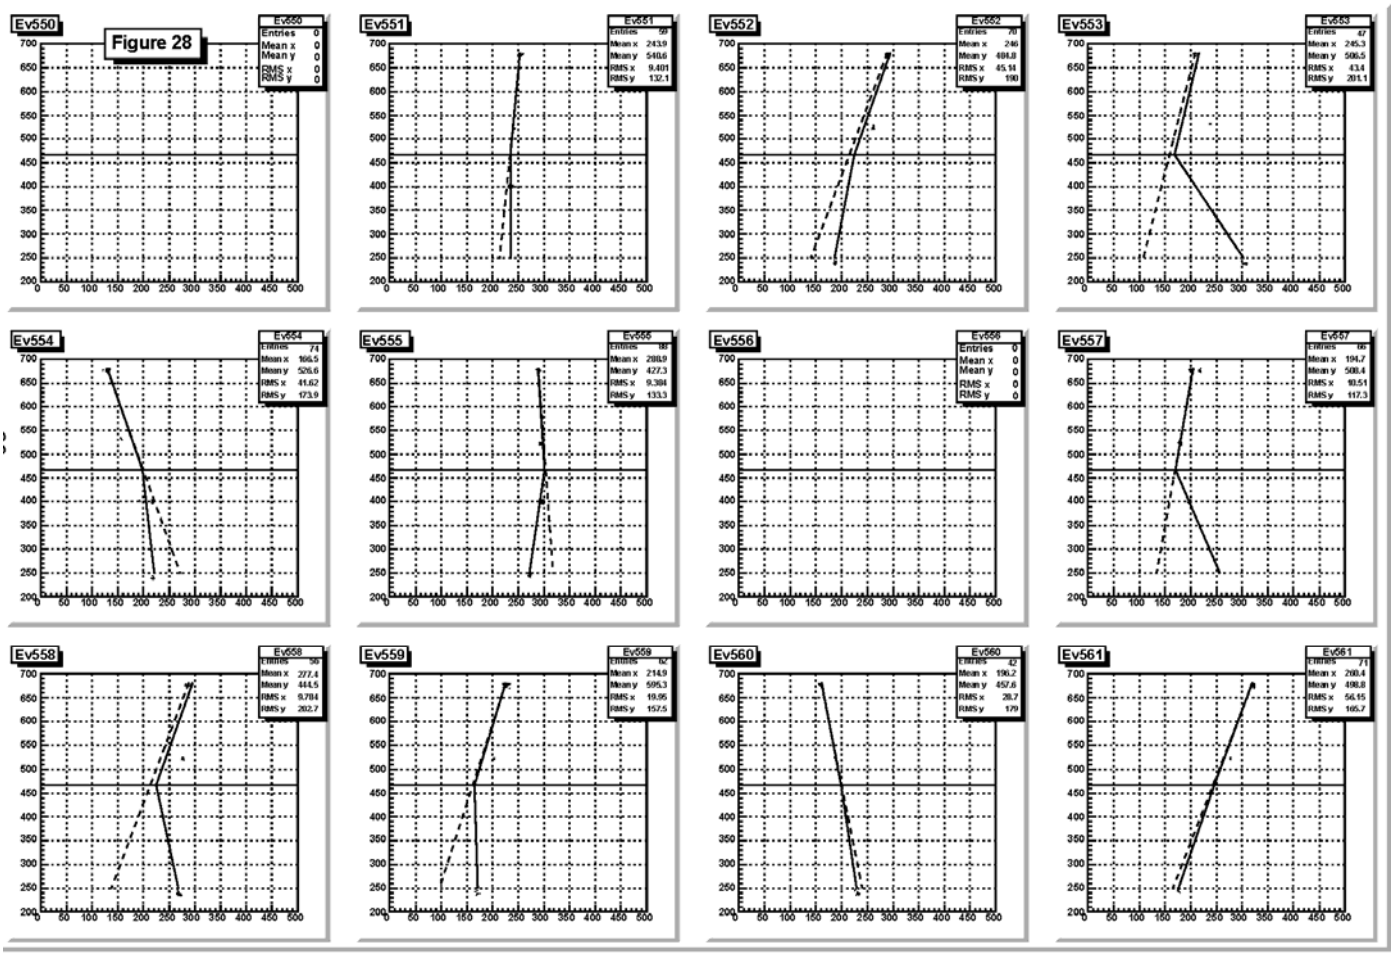

Figure 28: Scatterplots of a set of individual events Shown are the reconstructed tracks above and below the scattering plane. The dashed line is the (un-scattered) extrapolation of the original track to below the scattering plane.

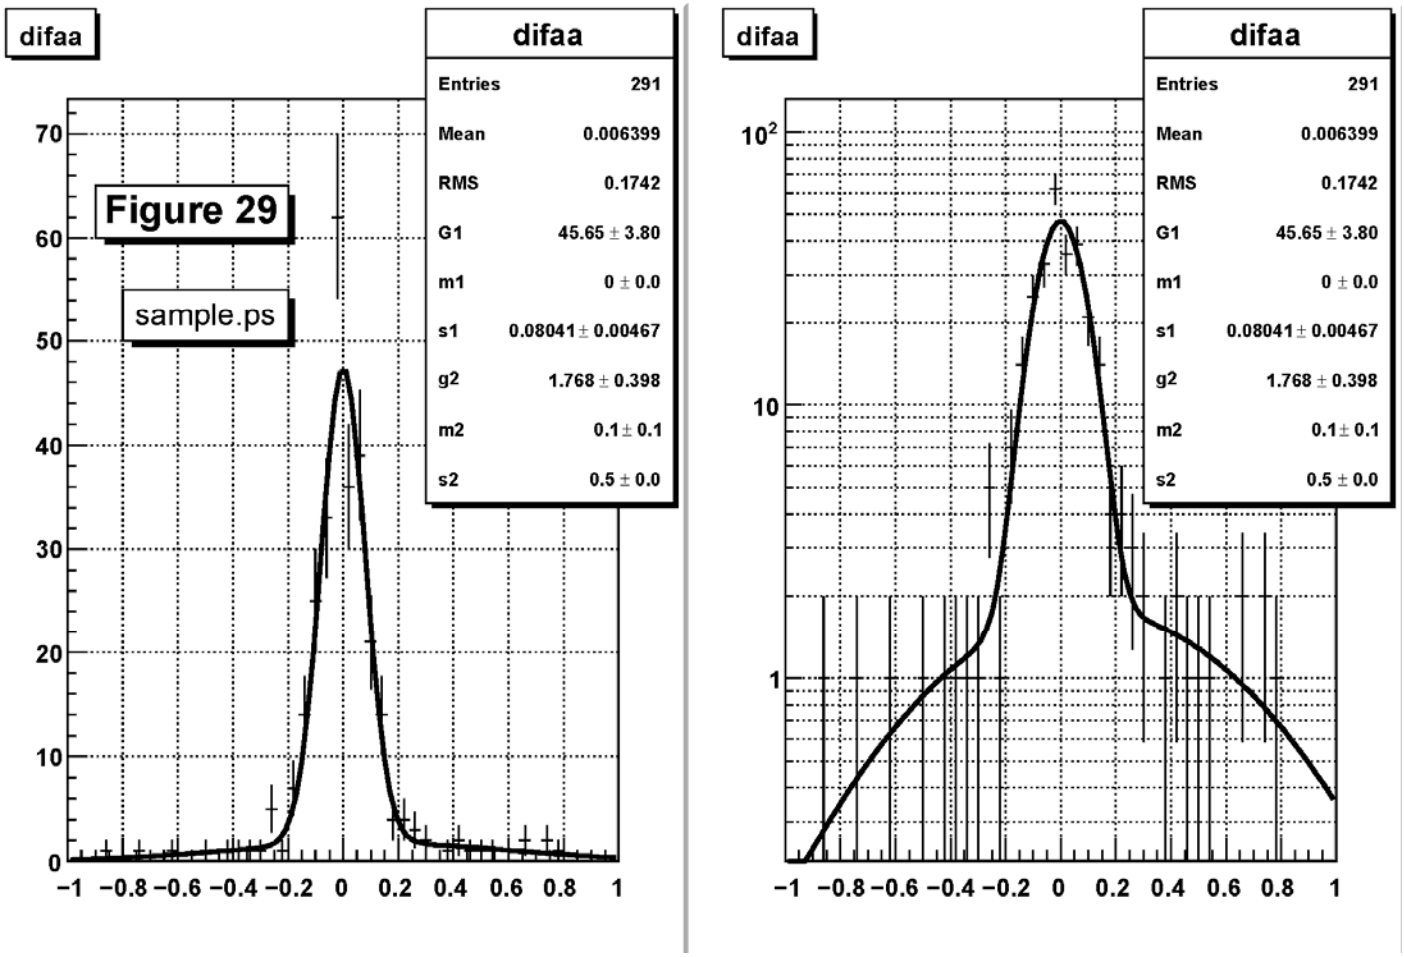

Figure 29: Histograms of the measured scattering angle from a sample of good events, on linear (left) and logarithmic (right) scales. The distribution of angles is fit by a sum of two Gaussian distributions, with each Gaussian distribution having a separate normalization, mean and standard deviation.

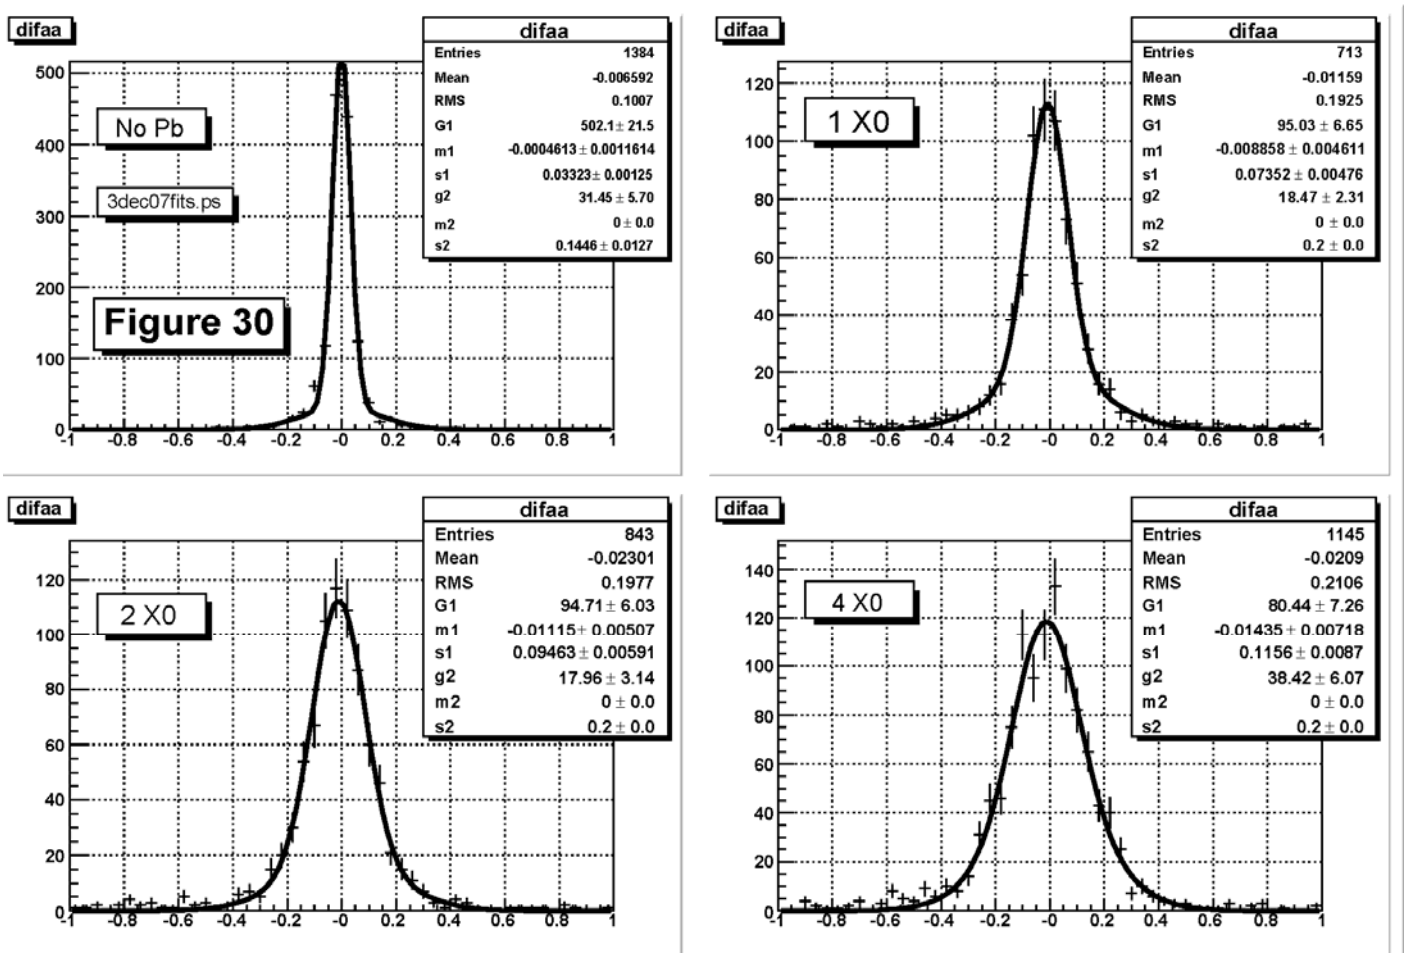

Figure 30: Histograms of the measured scattering angle for four different amounts of lead in the scattering plane, as indicated by the captions. The distribution of angles is fit by a sum of two Gaussian distributions, with each Gaussian distribution having a separate normalization, mean and standard deviation.