lifetime of the muon

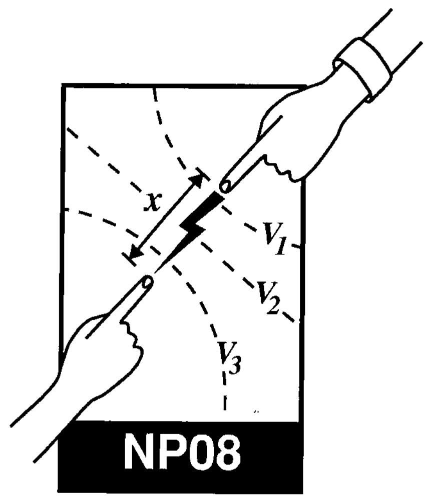

revised MGB October 2006

1 What is a muon?

The muon is a massive copy of the electron – it has negative charge, spin $\frac{1}{2}\hbar$ and is point-like at the level of 10^{-18} m. Its mass is 105.66 MeV, or 206.77 m_e. It interacts with matter through its charge and through the weak interactions which are also responsible for nuclear beta decay. It has been provided with its own species of neutrino and it decays into the much lighter electron via the process

$$\mu^- \rightarrow \nu_\mu \; e^- \; \bar{\nu}_e$$

with a mean lifetime of 2.197×10^{-6} s.

This weak decay process is of just the kind which leads to

$$n \rightarrow p \; e^- \; \bar{\nu}_e$$

with a mean lifetime of $8.87\times10^{+2}$ s.

How can such disparate lifetimes result from the same underlying process? All else being equal, the decay rates of two particles are in the ratio of the number of final states available and for decays into *three* particles the number of states available is (approximately) proportional to the *fifth* power of the maximum (electron) energy which is 0.78 MeV for free neutron decay and 52.5 MeV for muon decay [PROBLEM: Where do these numbers come from?] You are going to measure the lifetime of the muon.

The c and s quarks [see NP 10] are associated with the μ and ν_μ and nature has done it again. The τ is another heavy electron (mass 1777 MeV or 3477 me, mean life 290×10^{-15} s) which has its own neutrino and is associated with the t and b quarks. Why has nature done things in threes? I wish I knew.

2 Where do muons come from?

$\mu^+ \mu^-$ pairs can be produced by high energy gammas, but because the μ is massive the rate is much lower than $e^+ e^-$ pair production. Muons are produced prolifically as a result of nuclear collisions – in the strong interactions quanta of the pion field are shaken off and

$$\pi^- \rightarrow \mu^- \; \bar{\nu}_\mu \qquad \pi^+ \rightarrow \mu^+ \; \nu_\mu \, .$$

(The mean life of the charged pion is 2.6×10^{-8} s.)

The muons in this laboratory are produced as a result of interaction of the cosmic radiation in the atmosphere, at altitudes ~30 km. Even with time dilation the pions generated do not make it to sea level – but the muons do. (Note that at speed c, 30 km is 10^{-4} s – one of the early pieces of evidence for time dilation was the penetration of muons to sea level.) The angular distribution of cosmic ray muons is $\sim\cos^2\theta$ (where θ is the angle from the zenith) and for relatively small zenith angles the flux of muons at sea level is ~1 cm^{-2} min^{-1} $sterad^{-1}$.

This gentle rain of cosmic ray muons penetrates the overburden of NAPL, passes through you and your apparatus and penetrates hundreds of meters into the planet beneath you. However, the muons have an energy spectrum and are losing energy by ionisation [see NP 01] and some come to rest in your apparatus – and in you. They then stick around for microseconds....

3 Principles of the experiment

A charged particle moving through plastic scintillator produces a flash of light. This light falls on the photocathode of a photo- multiplier (PMT) and liberates electrons. These electrons are accelerated in a strong electric field, collide with a dynode and eject more electrons.... $10^6 - 10^7$ electrons arrive at the ANODE of the PMT for each electron liberated at the photocathode. A flash of light in the scintillator is transformed into a voltage pulse from the PMT anode. The modules in the electronics rack turn these electrical signals into intelligible physics. The core of the apparatus is three pieces of plastic scintillator; two slabs and one cylinder

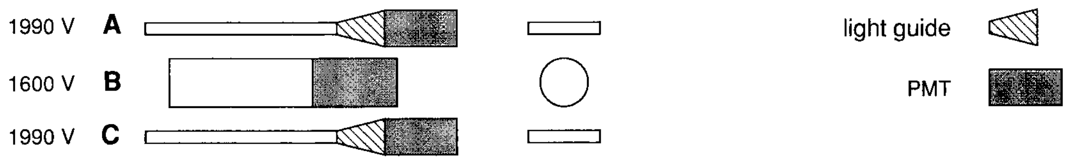

Figure 1

The slabs A and C are thin and do not stop many muons; we rely on the cylinder B for this. If a muon penetrates A and B and C, all three produce electrical signals within an interval of at most a few nanoseconds.

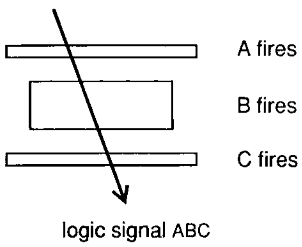

Figure 2a Muon penetrates A and B and C.

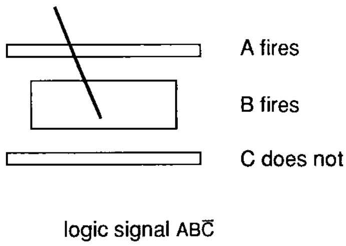

Figure 2b Muon stops in B and then decays.

If a muon stops in B and decays into an electron, the logic signal of such an event is $AB\bar{C}$; B and the two signals are separated by the lifetime of that particular muon.

MEASURE THE TIME SEPARATION BETWEEN $AB\bar{C}$ and B for each $AB\bar{C}$. Histogram the number of events as a function of time. In an ideal situation you would accumulate an exponential distribution, the mean of which would be the mean lifetime of the muon.

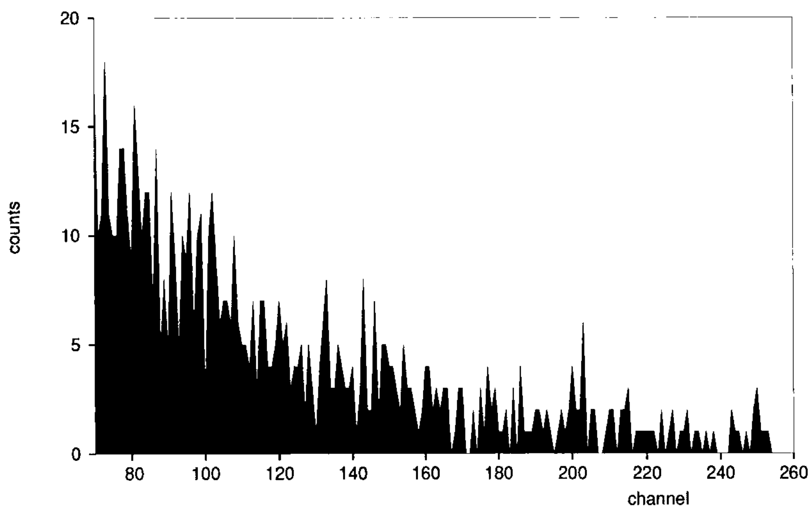

Figure 3

Ideally, this can be done by using $AB\bar{C}$ to START a clock and a subsequent B to stop it. A Time to Analogue Converter (TAC) measures the time interval and generates an output pulse in the range of 0 V (no time difference) to 10 V (maximum time difference for that mode of operation of the TAC – I'd set 10 μs). The output from the TAC is fed to an Analogue to Digital Converter (ADC) followed by a MultiChannel Analyser (MCA); these inhabit a card in the Personal Computer. A time interval of (say) 5 μs is converted into an output pulse of 5 V from the TAC and *one count* is added to (maybe) channel 134. This display of counts against channel number IS the histogram of the number of events as a function of time.

4 Electronic modules

4.1 Logic units

We need a unit which will accept pulses from A and B, generate an output logic pulse if they arrive together in time AND will NOT generate an output if a signal from C arrives at the same time. There are four such units packaged in a single module (755 Quad four-fold logic unit) in the lower shelf of the electronics rack. The face of each unit looks like Figure 4a:

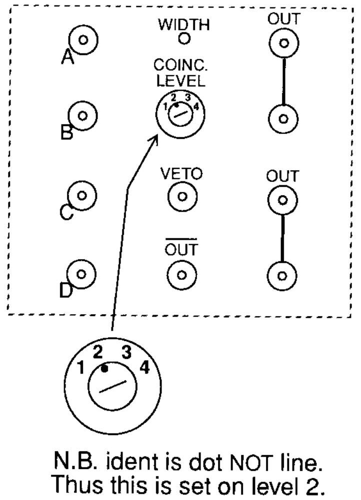

Figure 4a Logic unit panel.

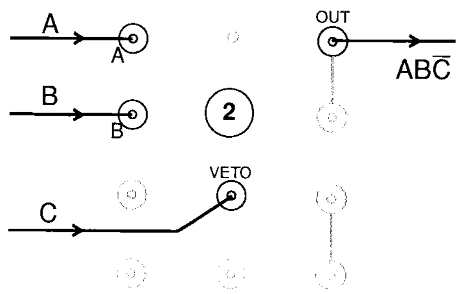

Figure 4b Connections to logic unit.

There are 4 inputs and the outputs produce a negative logic pulse ‾|_|‾. Set on coincidence level 2 and the outputs fire if A and B (or any other input pair) receive in time negative logic pulses. The WIDTH of the output logic pulse can be adjusted with a trim tool. (On coincidence level 1, A OR B will fire the output.) Thus a logic signal generated by A and a logic signal generated by B which are in coincidence will cause an output logic pulse whose value is AB. If a logic pulse from C arrives at the input labelled VETO, in time with A and B, then the output is not generated. (See Figure 4b.)

But wait! The inputs must be negative logic pulses and the signals from the PMT anodes may be negative but are certainly not logic pulses...

4.2 Discriminators

So we feed anode pulses from the PMTs into units which produce a negative logic pulse if the input pulse exceeds a preset threshold.

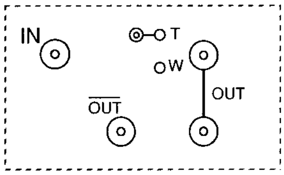

Figure 5

There are 8 of these units in a single module (705 Discriminator), in the lower shelf of the electronics rack. The threshold T above which the input causes an output and the width W of the output pulse can be adjusted with a trim tool. N.B. *Do not* take a trim tool to anything without first consulting a demonstrator.

4.3 Other units

The most important is the TAC, which measures a time interval between START and STOP and delivers an output pulse of magnitude proportional to that time interval. Make sure that you identify correctly the connectors to the TAC START, STOP and OUTPUT. The rack also contains DELAY boxes. These delay signals by 10s of nanoseconds and contain passive cable. A nanosecond is ~30 cm of cable; for a delay of microseconds a DELAY AMPLIFIER (an active unit) is included. The rack contains a pulser for generating test pulses and a timer/scaler for measuring signal rates. This is driven by positive pulses exceeding 0.2V and so the rack also contains an inverter to turn negative logic pulses into positive pulses. Finally, there is a Time Calibrator... see section 7.3

5 Starting the experiment

5.1 Single channels

You are provided with a very fast digital scope which makes everything EASY. The scope has a very high input impedance and so it is best to have a 50 ohm terminator across the input.

You are now in a position to **DO SOMETHING**. Start with the signal cables from detectors A, B, C. Trace them through delay boxes and into the discriminator units (orange lemo cables). The output cables to the $AB\bar{C}$ logic unit are red. If you connect the cable from the anode of A directly to the scope, you will see fast negative pulses looking something like Figure 6

Figure 6 PMT anode signal.

With the scope trigger level turned low, you will see that there are lots of small pulses – mostly noise. Split the anode signal and take one line direct to channel 2 of the scope and the other to the A input in the rack. Take the (red) output from the A discriminator unit to channel 1 on the scope and trigger on channel 1. You are now triggering on the logic pulse generated when an A input signal exceeds the preset discriminator threshold.

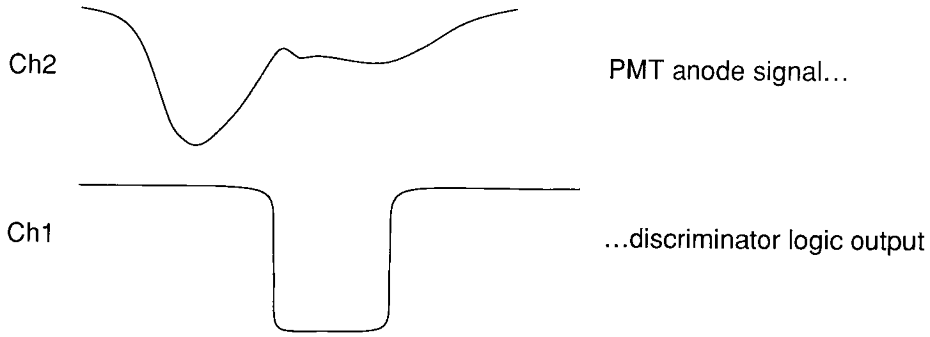

Figure 7

You should see something like Figure 7 and the small anode signals should have disappeared (unless someone has been in with a trim tool and undone all my careful work). Do the same for channels B and C. You will see that the accepted B anode pulses are much bigger than those from A and C. The cylinder of plastic scintillator contains much longer path lengths for a charged particle and so real signals are mostly much bigger.

Having seen and *noted* all this, restore the inputs to their original configuration.

If you take the (red) discriminator outputs one at a time to the inverter and thence to the timer/scaler unit, you can measure the A, B, C rates. Typical values which I found are 1300, 600, 2500 min^{-1}.

5.2 Into the logic unit

The logic unit will only do its stuff if the signals which should arrive in time *do* arrive in time, so you must ensure that this is so. Take the (red) outputs from the A and B discriminator units to the two channels of the scope. You should see a display like Figure 8a with a jitter of a few nanoseconds.

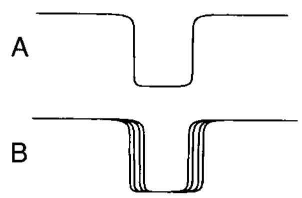

Figure 8a

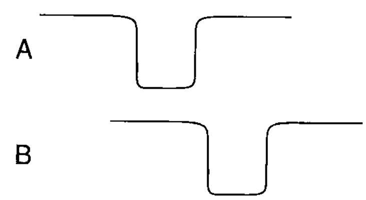

Figure 8b

If the display looks more like Figure 8b then the logic unit will *not* respond (on coincidence level 2) and the delays before input to the discriminators need adjustment. Look at A and C, B and C *before* changing delays, so that you can figure out the best changes to make. The widths of the A and B logic pulses should be ~25 ns; the width of the C veto pulse has been set larger to ensure that ABC is always vetoed, despite the jitter.

Measure the rate of AB by counting pulses from the logic unit with C *removed* from the veto. Then measure the $AB\bar{C}$ rate after putting the C logic pulse back into the logic unit veto. I found typical values ~200, 80 min^{-1}.

6 Measuring the time intervals

You want to START the TAC with a signal $AB\bar{C}$ (which every so often will signal a stopping muon) and STOP the TAC with a subsequent B pulse, generated by the decay electron. So we take the $AB\bar{C}$ logic unit output to the TAC START? We do. And the B discriminator output to the TAC STOP? We do not – this will not work. The reason is that a B pulse was used to form $AB\bar{C}$ and if the other output is taken to STOP then START and STOP arrive almost simultaneously. Even if the TAC works adequately for a negligible time interval, it will have been stopped *before* a decay corresponding to a measurable time interval has occurred. It will ignore a subsequent B pulse and sit there awaiting a fresh START signal from $AB\overline{C}$.

I have thought of two ways round this problem. First, *delay* the $AB\overline{C}$ START pulse by maybe 0.25 μs. The B from $AB\overline{C}$ arrives *before* the START pulse and is ignored. The B pulse from an electron which is produced in a decay occurring MORE than 0.25 μs after the muon stops gives the first STOP which occurs after START and we histogram (t – 0.25 μs), where t is the individual lifetime. The shape of the exponential is not affected. This method works, but I prefer to do it differently.

Take B into a second logic unit set on coincidence level 1 and put $AB\overline{C}$ into the VETO on this logic unit. This unit then only produces an output (B) pulse when B is *not* in coincidence with $AB\overline{C}$.

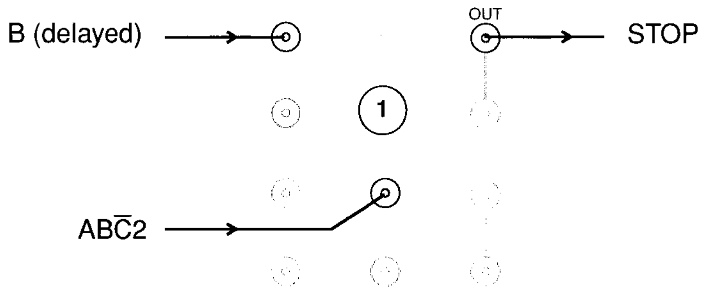

Figure 9

But wait! The $AB\overline{C}$ logic unit takes time to work and therefore $AB\overline{C}$ will arrive after B and will not veto efficiently. We need to delay B by a similar time. Look at the $AB\overline{C}$ logic pulse from the (top) logic unit and the delayed signal from the second B output which is going into the (bottom) logic unit. Adjust this delay until $AB\overline{C}$ embraces this second B signal. Using the fast 2 channel scope it is easy to make sure that the STOP signal is only generated by a B which is *not* part of $AB\overline{C}$. DO IT!

It is useful to see times before time zero and the region around time zero. I therefore took the logic pulse for STOP through the delay amplifier, set to 2 μs. The output from the delay amplifier is not the nice rectangular input pulse (the amplifier cant keep up with the fast switching) so I took the output from the delay amplifier through a discriminator unit (the top one) which generates again a nice fast rectangular pulse suitable to STOP the TAC. The paths of the various signals are laid out below in Figure 10 *phew!*

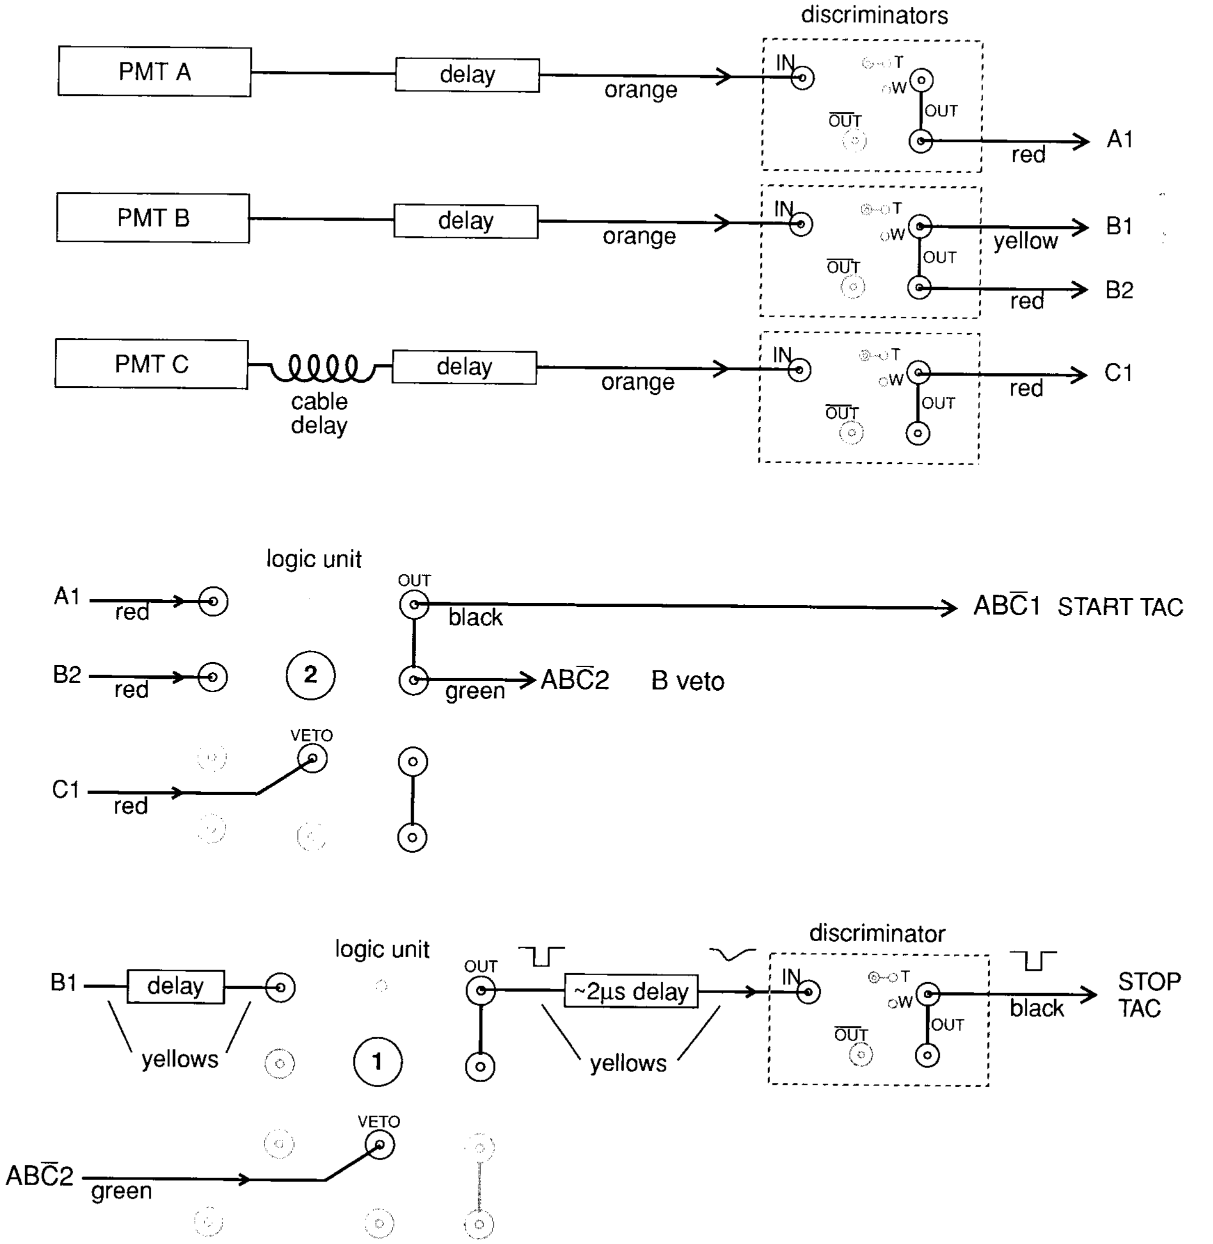

Figure 10

I found it convenient to colour code the lemo cables as shown above. Otherwise it gets too like trying to plait Medusa's hair!

7 Data

7.1 Taking data

With all the connections established correctly and all coincident signals properly timed in, data can be accumulated. All you have to do is hit the big green GO on the computer. I suggest that you first run for a few minutes with ABC̄2 *removed* from the B1 veto. You wont see any decays – the delayed B from the ABC̄ START will STOP the TAC and counts will build in a single MCA channel, at the same rate as ABC̄ counts measured with the timer/scaler. Increasing the delay on the delay amplifier will move this channel and calibrate crudely channels as a function of time. Then replace the ABC̄ veto to B1 and take real data. After ~15 minutes you should see something like the real example shown below.

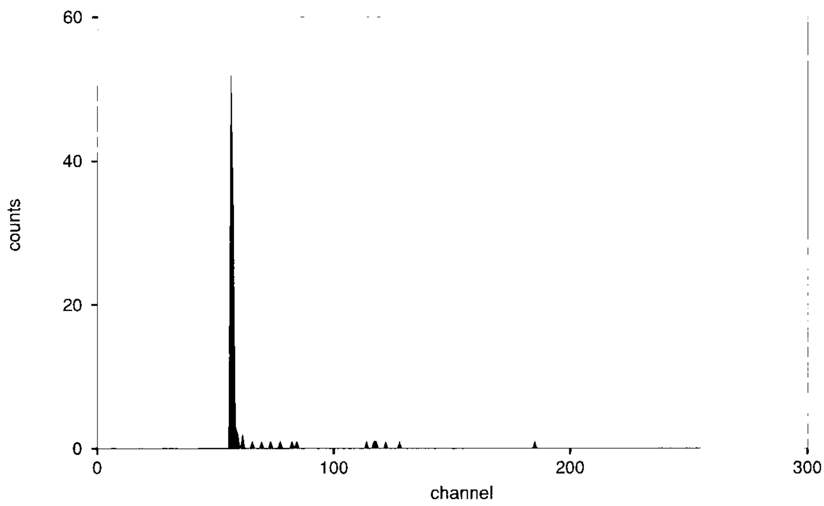

Figure 11

There is a large peak a few channels after the ABC̄ peak which you have already inspected. To the right, the scattering of decay events should correspond to about one detected every minute. It should already be evident that the mean lifetime is ≈2 μs! If the data DO look like this, you can leave the system accumulating for 18 hours or more. If the data DO NOT look like this, something is wrong and you have to start detective work, looking for a faulty cable or shot timing or.....

7.2 Analysing the data

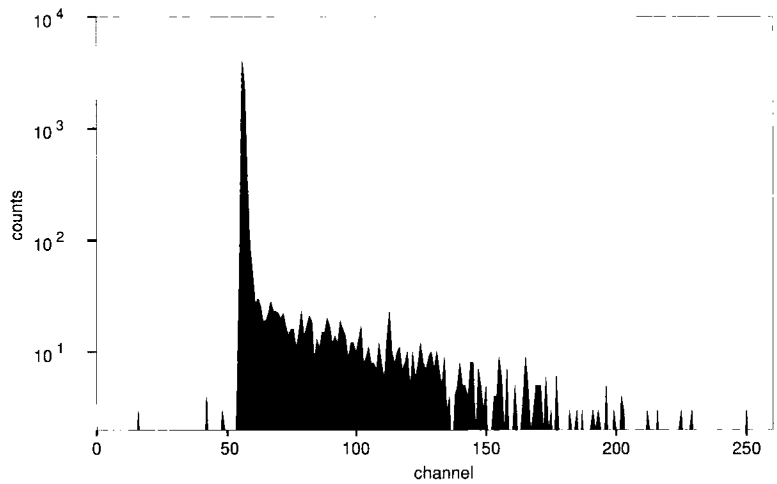

Figure 12

Come the dawn... you should have something like Figure 12 with maybe 1200 real decays. The exponential shape should be clear, to the right of the early peak. SAVE THIS DATA FILE. But where does the horrific early peak come from? If the timing to the B1 logic unit is right, there should be negligible breakthrough from the B in ABC̄ – and the early peak indeed comes a few channels *later* than the ABC̄ peak. Every so often the B channel double pulses, the second following the first by ~120 ns, with some jitter. I haven't been able to get rid of this (and PMTs are known to produce both pre- and after- pulses). See if you can find a way to check that B double pulsing is the reason for the early peak. The existence of the early peak means that you have to start the analysis clear of its influence. With a bit more work we could easily veto out the early peak, but that would not increase the rate of real data and I'd rather know that it is there!

There is a program in the ORIGIN package which will fit an exponential + background to the decay data. Use this – it will yield a lifetime in *channels* and the *error* on this lifetime. It will also draw you a pretty picture of the fitted exponential, superimposed on the data.

However, with only ≈1200 events distributed exponentially over ≈200 channels you will only have ~20 event ch^{-1} early on in the decay exponential and only a few towards the end. I don't trust a commercial package on low statistics and so in Appendix 1 I have suggested two other ways of obtaining the mean lifetime of the muon from the data – you may care to try one or the other or even both.

7.3 Calibration

At this point you have a lifetime and the error on that lifetime, measured in MCA channels on the Computer. You also know that 1 μs corresponds to ~25 channels. You need to check that the TAC – ADC – MCA system is linear over the range you have used *and* make an accurate measurement of the calibration. This is *very easy*, because you have been provided – at vast expense – with an ORTEC 462 Time Calibrator. This produces very fast START and STOP pulses, with precise time intervals which you can dial. Set RANGE on 10.24 μs and PERIOD on 1.28 μs (or 0.64 μs...). The use is obvious. If you take the START output to channel 1 of the fast scope and STOP to channel 2, triggering on channel 1, you can cross check the Time Calibrator against the scope...

8 End

You now have a lifetime for the muon in channels and the conversion factor for channels → microseconds. Your final result for the lifetime and the error on that measurement follows. (You cant do better than ~3% accuracy with ~1000 events.)

There is a subtle correction for a systematic effect which needs to be made. Muons stop in the active target B. NEGATIVE muons drop into the muonic K shell of carbon in the plastic scintillator – which has a radius

$$\frac{1}{207}\frac{1}{Z}a_o$$

where for carbon $Z = 6$ and a_o is the Bohr radius. They can be captured by the process

$$\mu^- + [p] \rightarrow \nu_\mu + n$$

and the rate is $0.4\times10^5\ s^{-1}$ (whereas the free decay rate is $0.45\times10^6\ s^{-1}$). Since the negative muons have two ways to vanish – decay and capture – they disappear a little faster than the positive muons.

The cosmic radiation supplies more μ^+ than μ^- [WHY?] – at these energies $\mu^+:\mu^- \approx 53:47$. Correct your measured lifetime for the effects of nuclear capture in order to obtain the proper mean lifetime of the free muon.

[If you really want to study capture of negative muons, you need to identify μ^+ and μ^- as they come in. Think about how that might be done – then you will understand why we have not attempted it.]

Finally... Capture of negative muons IS a weak process. A negative PION would last about 10^{-22} s after reaching the carbon K shell.

Appendix 1

A lot of fitting programs assume that the number of events N in a bin are normally distributed about a mean $\overline{N}$. In an experiment of this kind we expect the number to follow a Poisson distribution, for which the error is only given by $\sqrt{\overline{N}}$ for $\overline{N}$>>10. I therefore do not trust a fitting program from the supermarket shelf. Here are two relatively simple ways of obtaining the lifetime. Before starting either, calculate the flat background (generated by random events) from the mean number of events/channel to the left of the early peak. Choose a starting channel to the right of the early peak and from each channel to the right subtract the back-ground. What remains should be a pure exponential over a range from perhaps channel 65 to 255. Some bins may now contain negative numbers.

Then $$N(n) = N_0 e^{-n/\overline{n}}$$

where n runs from (say) 0 to 255 - 65 = 190.

A1.1 Method 1

Divide the time range T into two pieces. The first runs from $t = 0$ to T_1; the second from $t = T_1$ to T. Then

$$N_1 = N_o \int_0^{T_1} e^{-t/\tau} dt = N_o \tau \left[1 - e^{-T_1/\tau}\right],$$

and $$N_2 = N_o \int_{T_1}^{T} e^{-t/\tau} dt = N_o \tau \left[e^{-T_1/\tau} - e^{-T/\tau}\right].$$

Thus $$\frac{N_1}{N_2} = \frac{1 - e^{-T_1/\tau}}{e^{-T_1/\tau} - e^{-T/\tau}}. \qquad (1)$$

N_1 and N_2 are sufficiently large that their errors are given quite accurately by and $\sqrt{N_1}$ and $\sqrt{N_2}$.

You can find τ (in channel numbers, $\overline{n}$) by putting sample values into equation (1) and zeroing in on the correct answer. Alternatively, if you *choose* $T_1 = T/2$ then equation (1) becomes

$$\frac{N_1}{N_2} = e^{T/2\tau}; \quad \tau = \frac{T}{2\ln\left[\frac{N_1}{N_2}\right]}. \qquad (2)$$

A1.2 Method 2

The second method is more sophisticated – a poor man's *Maximum Likelihood.* It goes like this....

For $t = 0 \rightarrow \infty$ the probability of a single decay i occurring at time t_i (in interval dt) is

$$P(t_i) = \frac{e^{-t_i/\tau}}{\tau}$$

The likelihood $\mathcal{L}$ is

$$\mathcal{L} = \prod_i \frac{e^{-t_i/\tau}}{\tau}$$

$$\ln\mathcal{L} = \sum_i \left[-\frac{t_i}{\tau}\right] - N\ln\tau. \qquad (3)$$

N_i at t_i are all known; maximise $\ln\mathcal{L}$ with respect to τ and get

$$\tau_{\mathcal{L}} = \frac{\sum t_i}{N} = \bar{t}_i \quad ! \tag{4}$$

You have N (~1200) events between $t = 0$ and $t = T$.

Then
$$P(t_i) = \frac{e^{-t_i/\tau}}{\tau[1 - e^{-T/\tau}]}$$

and with a little manipulation

$$\tau_{\mathcal{L}} = \frac{\sum t_i/N}{1 - \dfrac{(T/\tau_L)\,e^{-T/\tau_L}}{1 - e^{-T/\tau_L}}} \tag{5}$$

(Remember, some t_i bins may contain a negative number of events.)

When the distribution cuts off at time T, the mean underestimates the lifetime. The correction factor in the denominator of equation (5) can be evaluated with an approximate value of τ.

If you want to get really fancy, evaluate $\mathcal{L}$ as a function of τ and plot the values.

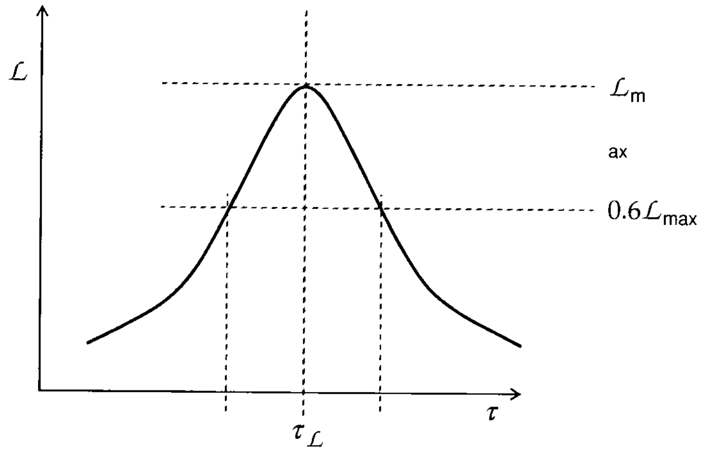

Figure 13

The excursions to $0.6\mathcal{L}_{max}$ give the errors. It is more practical to let $\ln\mathcal{L} \rightarrow \ln\mathcal{L}_{max} - \frac{1}{2}$.