Development of software routines for efficient analysis of high field experiments

In muons, researchers find a very sensitive probe of the atomic-level properties of materials. This has been exploited by condensed matter physicists through the study of magnetism, superconductivity, spin and charge transport, semiconductors and proton conductors, as well by chemists in studying reaction kinetics. The European research community has access to two muon sources: the SμS, located at the Paul Scherrer Institut in Switzerland; and ISIS, operated by the Science and Technology Facilities Council in the UK. Together, they offer researchers access to the full complement of μSR spectroscopic methods.

In order to cater for future research demands, both facilities have been developing instruments optimised for high magnetic field measurements, and the Joint Research Activities in Muons under Framework programmes 6 and 7 has played a pivotal role in these developments. The spectrometers are complementary in the facilities they provide: the ISIS instrument was designed to measure high longitudinal fields, while the instrument at PSI was optimised to measure high transverse fields.

Both instruments are producing datasets that challenge existing muon data analysis packages. For ISIS, there is a demand for an integrated package to simplify processing of Avoided Level Crossing data. At PSI, measurement of very high precession frequencies requires fast timing, and the extended data set require novel computation methods for rapid processing. Development of state-of-the-art data reduction methods is now required to enable a full and efficient analysis of the data from these novel methods. The following routines are discussed in this document:

Contents

Jamie Peck (Jamie.Peck@stfc.ac.uk),
Andreas Suter (andreas.suter@psi.ch),
Stephen Cottrell (Stephen.Cottrell@stfc.ac.uk)

Analysis of Avoided Level Crossing experiments

Background/Context

The Avoided Level Crossing (ALC) μSR technique probes the energy levels of a muoniated radical system, and can be used to elucidate the regiochemistry of muonium addition [1], dynamic processes [2, 3], and reaction kinetics [4], through measurement of the muon and proton hyperfine coupling interactions. During the thermalisation process, a portion of the implanted muons are able to capture electrons to form muonium (μ^+e^-). Muonium adds to centres of unsaturation in a sample (double or triple bonds) to form a muoniated radical species. The spins of the muon, unpaired electron, and protons within the sample interact through the isotropic and anisotropic components of the hyperfine interaction, forming a quantised system, described by a series of discrete energy levels.
In an ALC experiment the magnetic field is incrementally scanned, recording a specified number of positron events at each step. At certain fields, the energy levels in the muon and sample system become nearly degenerate, and are able to interact through the hyperfine coupling interaction. The spins oscillate between the two energy states resulting in a dip in the polarisation, observed as a resonance during the magnetic field scan. The three types of ALC resonance (referred to as Δ_0, Δ_1, and Δ_2 resonances) are characterised by the selection rule ΔM=0, ±1, ±2, where M is the sum of the m_z quantum numbers of the spins of the muon, electron and proton. Isotropic hyperfine coupling interactions manifest as Δ_0 resonances resulting from muon-nuclear spin flip-flop transitions. The Δ_0 resonance field is dependent on the magnitude of both the muon and proton hyperfine interaction (A_μ and A_k, respectively) and can occur in gaseous, liquid, or solid phase samples. The muon spin flip transition that produces the Δ_1 resonance only arises in the presence of anisotropy. Radical systems possessing complete anisotropy produce a single broad resonance and systems with axial or equatorial anisotropy produce an asymmetrical resonance line shape known as a powder pattern. The Δ_2 resonance is also observed in radicals from anisotropic environments. However, these are rarely observed experimentally due to their characteristically weak intensity line shapes. The magnitude of the hyperfine interaction is characteristic of the muon binding site, and can result in a ALC resonance associated with each of the magnetically equivalent nuclei, for each muoniated radical isomer. This can give a complex ALC spectrum in systems with multiple binding sites.

Density Functional Theory (DFT) can be used to calculate hyperfine coupling interactions, and is therefore of significant benefit for the complex interpretation of spectral features. In the relatively simple case of gas phase isobutene, CH_2=C(CH_3)$_2$, muonium adds to the double bond to form the tert-butyl muoniated radical [5, 6]. Two ALC resonances are seen (Figure 1), arising from the two types of protons in the radical – the six equivalent protons in the CH_3 groups and the two equivalent protons in CH_2Mu.

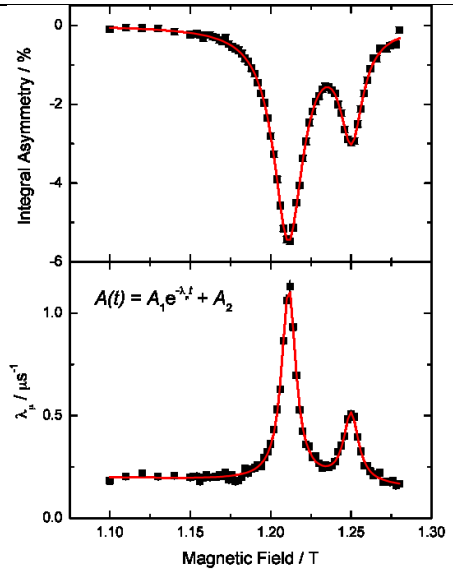

Figure 1: *ALC spectrum of the tert-butyl muoniated radical, formed by implanting muons into gaseous isobutene at 300 K.*

The magnetic field position, the full width at half height (FWHH), and the resonance line shape are the important parameters to be extracted from the ALC spectrum. The field position of a resonance is related to the muon and/or nuclear hyperfine coupling constant. They often show strong temperature dependence and can reveal information regarding the structure of the investigated system. The FWHH of a resonance may indicate any motional dynamics present in the system, and can also be used to determine muonium addition rates. The isotropic environment experienced by the tert-butyl muoniated radical in the example above leads to characteristically sharp, symmetric resonances. The anisotropic environments experienced by radicals in solid samples can produce a variety of 'powder pattern' lineshapes, which are characteristic of the orientation of the effective hyperfine tensors relative to the magnetic field, and can thus indicate any reorientational motion present.

In order to extract these parameters accurately from an ALC spectrum it is first necessary to determine a baseline, perform a baseline subtraction and then fit the peaks. From a user perspective, it is preferable that this can be performed within a single interface that hides the complexity of the underlying program. To help the user explore their data, each step of the analysis should be fully reversible.

References

[1] P.W. Percival et al., Chem. Phys. Lett. 245, (1995), 90.
[2] M. Ricco et al., Phys. Lett. A. 129, (1988), 390.
[3] E. Roduner et al., Chem. Soc. Rev. 22, (1993), 337.
[4] H. Dilger et al., Physica B 374-375, (2006), 317.
[5] D.G. Fleming et al., J. Chem. Phys. 105, (1996), 7517.
[6] P.W. Percival et al., Chem. Phys. 127, (1988), 137-147.

Routine Description

The steps involved in the analysis of a typical ALC spectrum are outline below:

1. ***Reduction of raw data.***
   The evolution of the muon spin is monitored as a function of time by two groups of positron detectors, positioned in the forward (upstream) and backward (downstream) direction relative to the muon momentum. Each detector typically measures decay positrons for 32μs following muon implantation, forming histograms which may be described by:

$$N(t) = N_0 e^{-t/\tau_\mu}[1 + A(t)] + N_{bkg}$$

where N_0 is a normalisation factor related to the total number of positrons counts, τ_μ is the muon lifetime (2.2 μs), and $A(t)$ is the muon asymmetry. If multiple muon decay events occur within a short time period, and are subsequently detected by a single detector, only a single event may be counted. The dead time of a detector is the time after the first event in which a subsequent event will be missed. Dead times can lead to signal distortion, so there is a need to quantify and correct for them using an appropriate model (a suitable algorithm that follows the non-extendable dead time model [1] is already built into Mantid).

A raw data file containing a specified number of decay events is collected for each increment of the field scan. $N_F(t)$ and $N_B(t)$ are histograms representing the time evolution of the muon spin polarisation along the axis defined by the forward (F) and backward (B) detector groupings. A time integral between t_0 and t_1 of the forward and backward counts can be obtained as follows:

$$N_F = \sum_{t=t_0}^{t_1} N_F\,(t_i); \quad N_B = \sum_{t=t_0}^{t_1} N_B\,(t_i)$$

and an integral asymmetry, A, formed over the defined time window:

$$A = \frac{N_F - \alpha N_B}{N_F + \alpha N_B}$$

Where α is a parameter that balances the different sensitivities of the detector groupings. A plot formed using this method is shown in panel 1, Figure 2.

For experiments measuring weak signals baseline shifts can prove problematic. The effects can be reduced by interleaving data acquisition, alternately measuring with the magnetic field at and offset from the set point. In this case, the final dataset consists of a signal acquired into two 'Periods', labelled '1' and '2'. An integral asymmetry can be formed for each period and the signals subtracted to give a differential measurement.

$$A = \frac{N_F^1 - N_B^1}{N_F^1 + N_B^1} - \frac{N_F^2 - N_B^2}{N_F^2 + N_B^2}$$

The routine should provide both options for calculating the integral asymmetry.

2. ***Modelling the baseline***.

The shape of the baseline is a characteristic of the sample size and shape, and may also be affected by systematic errors associated with the spectrometer such as a drift in detector efficiency. The baseline therefore needs to be separately defined for each experiment carried out on the instrument.

Various strategies are possible for defining the form of the baseline, including measurement of a 'blank sample' over a similar field range. In practice, an extended dataset is typically taken on the sample material either side of the ALC resonance peak. Sections of the curve that define the baseline can then be selected and a suitable function used to fit the form of the baseline over the full field range measured. In practice the baseline can generally be modelled by a polynomial function chosen with an appropriate order. The interface should allow users to experiment with different functions and data windows to ensure an optimal fit is achieved. This process is shown in panel 2, Figure 2, where the data recorded in the field range 12500G

to 13250G and 14250G to 15000G is modelling by a 2nd order polynomial.

3. ***Baseline subtraction***.
   Once a function form has been obtained for the baseline it can readily be subtracted to reveal the resonances on a linear background (panel 3, Figure 2). This should be a reversible process, as the user may need to repeat step 2 in case the result is distorted and the baseline needs to be remodelled.

4. ***Peak fitting***.
   A variety of resonance line shapes are possible when measuring an ALC spectrum, and many peaks may be present and are often overlapping one another. Therefore, a choice of fitting functions should be available to optimise their fit. Gaussian and Lorentzian type functions are able to fit the symmetric line shapes of resonances typically found in isotropic liquid systems. However, additional functions are required to model the asymmetric powder patterns that result from anisotropic solid systems; functions could be defined by the user at the time of data analysis or chosen from a fitting library. In the example below (panel 4, Figure 2), the resonances are best described by Lorentzian functions. Key parameters from the fit are the position and FWHH of the peak.

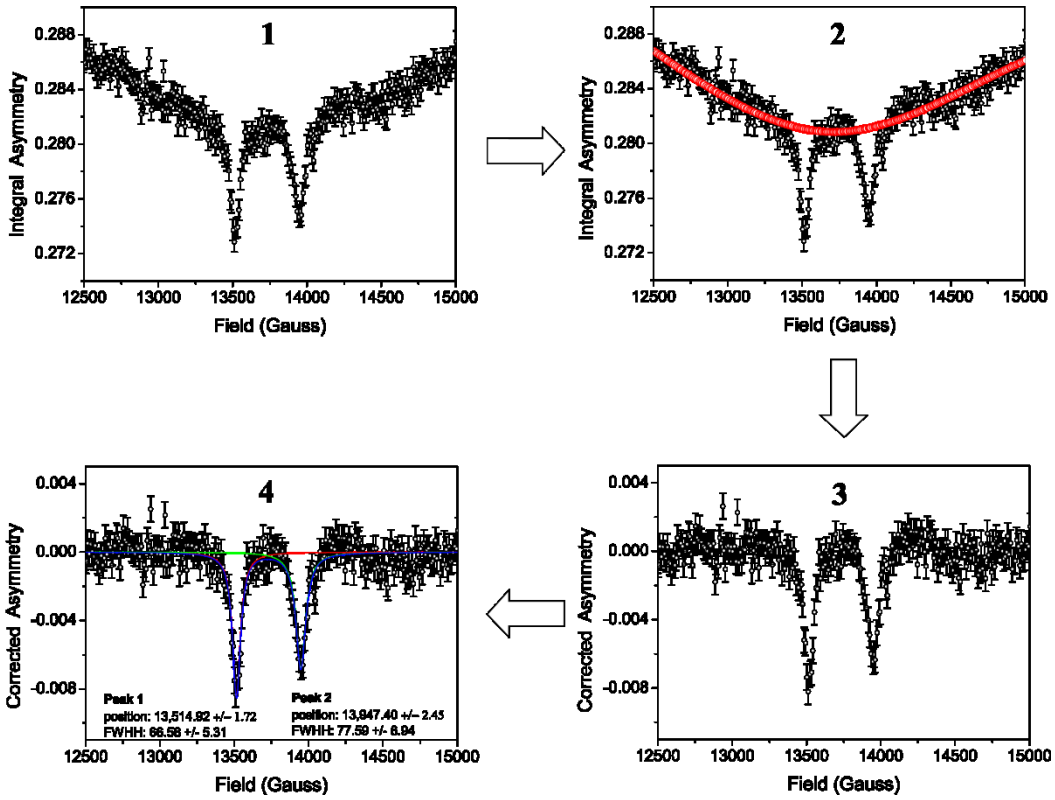

Figure 2: *Steps involved in the analysis of an ALC spectrum: (1) Reduction of the raw data;(2) modelling the baseline; (3) baseline subtraction; (4) peak fitting. Data measured on the HiFi instrument at ISIS.*

References

[1] See W.R. Leo, 'Techniques for Nuclear and Particle Physics Experiments: A How-To Approach', Springer, Pg122.

Implementation Details

At ISIS, the aim would be to incorporate these routines into the supported analysis framework, Mantid

[1], in the form of a single interface. Mantid already has much of the functionality required to follow through the required analysis. For example, the *PlotAysmmetryByLogValue* algorithm reads the raw data and calculates the integral asymmetry as a function of field (step 1), polynomial data fitting is available to model the background (step 2), step 3 can be achieved by manipulating Mantid workspaces and peak fitting (step 4) is available. However, in-depth knowledge of the Mantid package would be required to complete the analysis, and therefore a single user friendly interface, specifically designed for the analysis of ALC data, is preferred. Each analysis step should be reversible, encouraging users to test and refine solutions.

The following should be considered when designing the interface:

- For modelling the baseline shape, limits could be specified interactively on a preview plot (by draggable limit lines) and/or set as numerical values in limit boxes. As an alternative, anchor points might be interactively placed on the ALC curve and then used by the fitter to model the baseline.
- A peak fitting function should provide an interactive method of defining the peak position and peak width, allowing for fine adjustment. These values would be updated into parameter boxes associated with the peak function in real time. This would be performed on each peak. Following completion, it would be advantageous if a plot was generated which included the background subtracted data, the individually fitted peaks, and a cumulatively fitted curve as separate layers.
- Separate workspaces should be generated that include the original raw data, baseline subtracted data, and peak fit curves.
- Results, such as peak position, FWHH and analysis function, should be collected in a tabulated format. Parameters associated with the analysis, such as the integration time window and baseline model, should also be captured.
- The user should be able to reverse each stage of the analysis.

References

[1] For information about Mantid see www.mantidproject.org

Rotating Reference Frame Transform

Background/Context

This routine is required to solve two problems that arise from working with muon spin rotation signals, typically measured when a static magnetic field is applied transverse to the muon polarisation. Firstly, a method will be implemented for combining signals from multiple detector segments to form two output signals in quadrature phase, and secondly, a transformation of these quadrature phase signals to a rotating reference frame will be coded.

A routine to generate a quadrature phase output is particularly important for ISIS, where data from highly segmented detector arrays (containing up to 600 detector elements) needs to be combined to make effective use of the full information available in the dataset. The additional step of making a rotating reference frame transformation provides a useful method of visualising high frequency signals. Since the transformation leaves a comparatively low frequency signal, the data can readily be rebinned to improve counting statistics – this is particularly useful during an experiment to provide an early check on the form of the measured signal. Data analysis (for example least squares fitting) to study the form of the signal relaxation can be completed in either the laboratory or rotating reference frame according to preference.

To date, ad-hoc solutions have been used to by groups requiring these functions. We plan to code a generic and robust solution as a package of one or more subroutines that can be called from the preferred analysis codes.

Routine description

The algorithm for forming the quadrature phase signals and making the rotating reference frame transformation has been adapted for highly segmented detector arrays from the work by Riseman and Brewer [1].

1. ***Forming quadrature phase signals***
   For muon spin rotation measurements the asymmetry signal, $A_i(t)$, measured for the i^{th} detector segment is given by:

$$A_i(t) = \left[\frac{N_i(t) - N_{bkg}}{N_i^0}\right] e^{t/\tau_\mu} - 1$$

   where $N_i(t)$ are the counts recorded as a function of time, N_i^0 a normalisation factor and τ_μ the muon lifetime of ~2.2μs.

   Assuming the detector array is made up of M segments, each with a phase φ_i relative to the incident muon polarisation, the quadrature phase asymmetry in the laboratory reference frame can be formed as follows:

$$A_{lab}(t) = \sum_{i=1}^{M} A_i(t)\cos(\varphi_i) + i\sum_{i=1}^{M} A_i(t)\sin(\varphi_i)$$

   The detector phases, φ_i, may be determined from the spectrometer engineering drawings or, more typically, from calibration measurements. The quadrature phase asymmetry is required as input both for the rotating reference frame transformation (below) and an existing maximum entropy subroutine.

2. The Rotating Reference Frame transformation

Once processed as a complex quadrature phase asymmetry, the signal in the laboratory reference frame, $A_{lab}(t)$, can equivalently be written as:

$$A_{lab}(t) = A_{Re}(t) + iA_{Im}(t)$$

and a transformation of the quadrature phase components to a rotating reference frame (RRF) at frequency ν_0 obtained as follows:

$$A_{\mathrm{Re}}^{\mathrm{RRF}}(\nu_0, t) = +A_{\mathrm{Re}}(t)\cos(2\pi\nu_0 t) + A_{\mathrm{Im}}(t)\sin(2\pi\nu_0 t)$$

$$A_{\mathrm{Im}}^{\mathrm{RRF}}(\nu_0, t) = -A_{\mathrm{Re}}(t)\sin(2\pi\nu_0 t) + A_{\mathrm{Im}}(t)\cos(2\pi\nu_0 t)$$

The choice of frequency is determined both by the apparent frequency, $(\nu - \nu_0)$, required for the signal in the rotating reference frame and the final bin size required.

An example application of the rotating reference frame transformation is shown in Figure 1. Here, a 22.7MHz spin precession signal, recorded during a 90°-180° spin-echo radio frequency pulse sequence (shown top) is presented as a quadrature phase signal in a 20MHz rotating reference frame (bottom). The transformed data has been rebinned to improve counting statistics.

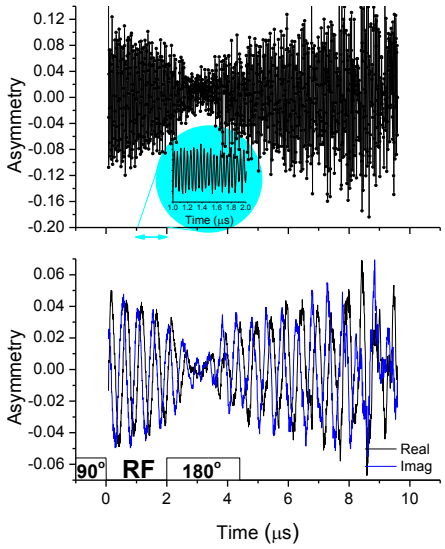

Figure 1. *A 22.7MHz μSR spin precession signal recorded during a 90°-180° pulse sequence (top) and the data presented as a quadrature phase signal in a 20MHz rotating reference frame.*

References

[1] T.M. Riseman and J.H. Brewer, Hyp. Int., 65, (1990), 1107.

Implementation details

This is likely to be implemented as a package of one or more subroutines, coded in an appropriate language, which can be called from the analysis package preferred by the facility.

Fast Data Processing using Multithreading

Background/Context

For high transverse field muon spin rotation experiments a time resolution ≤ 100 ps is required in order to be able to measure all observable muon or muonium precession frequencies present. Typical observation windows are ≥ 10 μsec, and hence the recorded positron histograms per detector have ≈ 10^5 entries or more. For instance, the high field spectrometer at PSI consists of 16 positron detectors, each with 4×10^5 histogram entries. Depending on the underlying physics of the material, data reduction is not possible since it typically leads to a deformation of the line shape which contains valuable physical information. This means that any analysis tool will face the challenge of handling large data sets and being able to fit these data sets within a reasonable time. Fitting these large data sets can easily take one hour when running on a single CPU core. This becomes even more challenging for so called 'global fits' where multiple data sets have to be fitted simultaneously, e.g. to determine a more reliable temperature dependence of fitting parameters within a restricted model.

Modern CPUs typically have a number of cores on which various threads can run in parallel. By using these resources properly the fitting time will scale roughly inversely to the number of threads used in the fitting algorithm.

Routine description

Fitting data typically involves χ^2 or log-max likelihood calculations:

$$\chi^2 = \sum_i \frac{(d_i - t_i)^2}{e_i^2},$$

where d_i is the histogram bin with index i, t_i the theory estimate, and e_i the error estimate. μSR spectra will obey Poisson statistics and hence $e_i^2 = d_i$. Since Poisson statistics apply, the log-max likelihood can be written as

$$\ell = -\sum_i \{d_i \ln t_i - t_i - \ln \Gamma(d_i + 1)\}$$

As can be seen from these formulae, they are forming sums over the data sets and hence are a natural starting point to parallelize the optimization since loops are easily parallizable.

Since the implementations should be open-source, the OpenMP framework [1] is very well suited to attack this problem with minimal overhead. It supports various architectures and the programing languages C, C++, and Fortran. Essentially, compiler directives can be added to code to highlight which parts can be calculated in parallel.

An important point to note: most μSR theory functions (t_i) are light-weight and hence thread-safe from the beginning. However, some specialized μSR functions like the dynamic Kubo-Toyabe LF [2, 3] are currently implemented with a shared memory block in `musrfit` [4]. It will be important to make all these routines thread-safe as well.

References

[1] http://openmp.org
[2] R. S. Hayano *et al.*, Phys. Rev. B **20,** 850 (1979).
[3] P. Dalmas de Réotier and A. Yaouanc, J. Phys.: Condens. Matter **4** 4533 (1992).
[4] A. Suter, and B.M. Wojek, Physics Procedia 30, 69 (2012).

Implementation details

This is likely to be implemented in existing routines, coded in an appropriate language. Specific detailed examples showing how various portion of heavily calculation code can be optimized will help to implement this in the specific packages preferred by the facilities.